

Chapter 3

An example problem: Alarm correlation

In this part we consider a case-study in KDD: the problem of analyzing alarms from a telecommunication network. Telecommunication networks are large systems consisting of such complex components as switches, base stations, radio links, etc. Network elements produce large amounts of alarms about the faults in a network. Fully employing this valuable data is, however, difficult, due to the high volume and the fragmented nature of the information.

This part consists of 4 chapters. In this chapter we introduce the problem and outline one possible framework for knowledge discovery in alarm databases. In Chapter 4 we give a full and domain independent formulation for the framework and present algorithms for the task. In Chapter 5 we present an alternative approach and algorithms. The ideas have been implemented in a system called TASA, for Telecommunication Alarm Sequence Analyzer. We conclude this part in Chapter 6, where we look briefly at the KDD process in this particular case.

3.1 Fault management problems

Telecommunication networks are growing fast, and at the same time the task of identifying and correcting faults is becoming more difficult. The task is critical, since faults that interfere with the services offered by the network are costly for the operator. The quality of service plays also an important role in the growing competition between operators.

Telecommunication networks and alarms A telecommunication network can be viewed as consisting of a number of interconnected components: switches, transmission equipment, etc. Each component can in its turn contain several subcomponents. The number of components depends on the

abstraction level used in viewing the system; a network operated by a local telephone company can be considered to contain 10–1000 components. We use the term *network element* to refer to the components under consideration.

An *alarm* is a message emitted by a network element, typically when a problem is encountered. Telecommunication systems are more or less fault tolerant, so the problem is not necessarily visible to the users. So faults are reported by network elements, and typically before the users are affected, but unfortunately a network element has a very narrow view to the network. A network element can therefore only report the symptoms of the fault from its limited viewpoint. On the other hand, one fault can result in a number of different alarms from several network elements.

Example 3.1 An actual alarm, as shown to the network manager, typically looks something like this:

```
1234 EL1 BTS 940926 082623 A1 Channel missing
```

The first field identifies the type of the alarm and the last field is the name of the alarm type. The second and third fields, **EL1** and **BTS**, identify the network element that emitted the alarm and the type of that network element, respectively. The next two fields are the date and time when the alarm was sent. Finally, the sixth, **A1**, tells the priority level of the alarm on a scale from 1 to 5. □

The information contents of alarms vary a lot. Some alarms concern problems in logical concepts, such as virtual paths, some concern physical devices, such as power supplies. Some alarms report a distinct failure, e.g., that the incoming signal is missing, whereas some only report high error rate without any hint for the cause.

Even in a small telecommunication network there can be hundreds of different types of alarms. The number of alarms produced by a network varies greatly, of course, but typically there can be about 200–10000 alarms a day.

Processing of the alarm flow is a difficult task for the following reasons.

- The size of the networks and the diversity of alarm types mean that there are a lot of different situations that can occur.
- The alarms occur in bursts, and hence there is only little time for network managers to decide what to do with each alarm. However, when lots of alarms occur within a short time the operators should intervene.
- The hardware and software used in telecommunication networks develop fast. As new elements are added to the network or old ones are updated, the characteristics of the alarm sequences are changed. Thus the network managers do not have time to learn what the appropriate response to each situation is.

Alarm correlation A central problem of fault management is to determine which alarms are related, to combine the information in related alarms, and to make inferences about the faults and their locations. It is the task of a *network management center* to correlate the alarms as they are received from the network. *Correlating* alarms means combining the fragmented information they contain and interpreting the flow of alarms as a whole. Alarm correlation systems typically are expert systems performing operations such as (i) *removing* alarms carrying redundant information, (ii) *filtering out* low-priority alarms when higher-priority alarms are present, or (iii) *substituting* a set of alarms by some new information [45]. The goal of alarm correlation is to reduce the amount of information shown to the network managers, improve the usefulness of the information, and ultimately to identify the most probable faults that caused the alarms and to possibly even propose corrective actions. While the prediction of severe faults is a difficult task, the economic benefits that would be obtained from such predictions would be significant.

In addition to the alarms received from a network, several sources of background information are essential in alarm correlation. Knowledge about the topological relationships of network elements is crucial. A lot of other sources are also useful for interpreting alarms. For instance, knowledge about recent problems in the network may help to explain certain alarms.

While the use of alarm correlation systems is quite popular and methods for specifying the correlations are maturing, acquiring all the knowledge necessary for constructing an alarm correlation system for a network and its elements is difficult. The complexity and diversity of network elements and the large variation in the patterns of alarm occurrences pose serious problems for network management experts trying to build a correlation model.

Alarm correlation is typically based on looking at the active alarms within a time window, and interpreting them as a group. We adopt the following formal view to alarm correlation, similar to the one taken, e.g., in [44]. Abstractly, the input to a correlation system is an ordered sequence $(A_1, t_1), (A_2, t_2), \dots$ of alarm occurrences. A *correlation pattern* describes a situation that can be recognized in an alarm sequence within a time window of a given length. Typically, a correlation pattern is an expression on the set of active alarms of, e.g., the last five minutes. If in a given window there is a set of alarms that matches the correlation pattern, then the set is said to be an occurrence of the pattern. Associated with each correlation pattern is a *correlation action*, which is to be executed when there is an occurrence of the corresponding pattern in a window.

Example 3.2 Consider a correlation pattern containing two alarms, “*link alarm* from X with severity 1” (alarm type A) and “*high fault rate* in X ” (alarm type B), where the variable X may be replaced by the same network element in both alarms. Assume that when alarms of types A and B co-occur,

they are known to be followed by fatal problems in the network element X . The correlation action could combine the information in the alarms to a high priority warning message “ X will probably collapse within an hour”, present it to the network manager, and filter out the original alarms. \square

Given correlation patterns and actions, and a sequence s of alarms, an alarm correlation process continuously observes the incoming alarm sequence s , considers the last time window on s , and executes the actions associated with the patterns that occur in the window.

Alarm correlation and filtering systems for telecommunication network management have been presented, e.g., in [29, 30, 44, 45, 47, 71, 9]. Similar approaches have been used successfully also in process control tasks [66].

The problem Building a system for alarm correlation is a difficult task. Networks are large and network elements are complex. The number of correlation patterns can be very large, and acquiring them from technical experts is a tedious task.

Correlations may pass unnoticed by the experts, for different reasons. It can be that an expert knows a correlation but did not come to think about it, or a correlation can be such that the expert did not even know there was a connection between certain alarms. Both networks and network elements evolve quickly over time, so a correlation system is never complete. It also takes time for the experts to learn new correlations and to modify existing ones.

3.2 Sketching a solution

Setting the goal Our high-level goal is to discover useful knowledge from telecommunication alarms, mainly to help in the construction of alarm correlation systems. From a sequence of alarms different types of knowledge can be discovered, for example neural networks, hazard models, or rule-based representations. If the goal would be just to obtain good predictive performance, a neural network could be useful. There is impressive evidence on the wide applicability of neural networks. However, in the current application one important goal is the comprehensibility of the discovered knowledge: the telecommunication operators do not wish to install any “black boxes” into their systems. This rules out the simple-minded use of neural networks.

From the statistical point of view, an alarm sequence can be viewed as a marked point process, and each event (i.e., an alarm) can be considered as a failure of a component. Therefore the hazard rate based methods of analysis of failure-time data for the data can be used. Combined with the Bayesian paradigm, the statistical machinery for these types of models is very powerful. However, these methods require a lot of human effort in the building of the




Figure 3.1: A sequence of alarms.

statistical models, as well as enormous amounts of computational resources. Currently they are not feasible for the analysis of hundreds of different types of alarms and their potential relations.

In this case rather simple rule-based formalisms are useful. Since the correlation systems are often based on looking at a time window of the alarm sequence at a time, it would be useful to use a similar concept in the discovered rules. So we choose as our goal the discovery of rules of the following form: “if certain alarms occur within a time window of a given width, then certain other alarm also occur within the time window”. This type of rules are particularly useful for the following reasons.

- **Comprehensibility:** such rules are easy to understand and verify.
- **“Standards”:** Correlation patterns closely resemble this rule format.
- **Characteristics of the application domain:** such rules can be representations of simple small causal relationships within the domain.

A first solution Let us now outline an exact representation of the problem. The alarms contain a number of fields, and the most important ones are the type of the alarm and the network element that sent the alarm. We first consider only the alarm type. The occurrence time is also essential: in the interest of considering different windows on the alarm sequence, we want to know the times of the alarms. Unfortunately though, the time stamps on the alarms can be inaccurate, due to unproper synchronization of clocks in the network elements. Anyway, we choose to view an occurrence of an alarm abstractly as a pair (A, t) , where A is the *alarm type* and t is the *time* of the alarm. A sequence of alarms can now be represented as a sequence of alarm types, as in Figure 3.1. In the figure, A, B, C, D, E , and F are alarm types, and they have been marked on a time line.

What are the useful rules like more exactly? We want a set of alarms to occur close to each other, i.e., within a time window, as in the correlation patterns. Because of the poor synchronization of clocks, the rules should not be very sensitive to changes in the order of alarms. If we look at a window on the alarm sequence and consider the alarms in it as an (unordered) set, we obtain these goals at least partially. Additionally, the alarm sequence is merged from several sources, and therefore the rules should be insensitive to intervening events. We thus want to search for rules that only test the presence of alarms in a window and that do not require that other alarms are not

present. If we look at an alarm sequence through a sliding window and find out which sets of alarms tend to occur together, we fulfill the requirements.

This task is obviously very similar to the task of finding frequent sets, and we can actually formulate a first solution as a modification of the frequent set discovery task. Consider a window at a given starting position and with a given width as a (multi) set of those alarm types that occur within the window. Denote now the set of all windows of the given width but over all possible starting positions by r , and denote the set of alarm types by R . Now r can be considered a binary database over R , and one can discover the collection $\mathcal{F}(r, \text{min_fr})$ of frequent sets in r with respect to some frequency threshold min_fr . This way one finds all those sets of alarms that occur in a randomly chosen window of the given width with a probability at least min_fr . By discovering association rules one finds the kind of rules we were looking for: $X \Rightarrow Y$ means that “if the set X of alarms occurs within a time window of a given width, then alarms Y also occur in the time window”.

A more general approach Let us have another look at the alarms in the sequence of Figure 3.1 and the kind of patterns that seem to occur frequently. Clearly, an alarm of type E is soon followed by F . One can also make the observation that whenever A and B occur (in either order), C occurs soon. The frequent alarm sets are interesting, but they do not take the order of alarms into consideration at all. What if we know that the order of alarms is important?

Instead of sets of alarms, we can consider partially ordered sets of alarms in general. Let us call such frequently occurring patterns *episodes*. We can draw episodes simply as directed acyclic graphs. Consider, for instance, episodes α , β , and γ in Figure 3.2. Episode α occurs in a sequence only if there are events of types E and F that occur in this order in the sequence. We call such totally ordered episodes *serial*. Episode β is a *parallel* episode: no constraints on the relative order of A and B are given. Episode γ is an example of non-serial and non-parallel episode: it occurs in a sequence if there are occurrences of A and B and these precede an occurrence of C ; no constraints on the relative order of A and B are given.

It would not be very useful to search for all possible episodes with arbitrary partial orders. Even for small episodes there would be a large number of different partial orders, and frequent episodes would most probably contain a lot of redundancy. Instead, we assume that the user restricts the search to a certain class of episodes, in practice to either parallel or serial episodes. Note that even these classes have an infinite number of episodes if episodes are allowed to contain multiple similar events and if the size of episodes is not restricted.

Once the frequent episodes are known, they can be used to obtain rules, e.g., for prediction. For example, if we know that the episode β of Figure 3.2




Figure 3.2: Episodes.

occurs in 4.2 % of the windows and that the superepisode γ occurs in 4.0 % of the windows, we can estimate that after seeing a window with A and B , there is a chance of about 0.95 that C follows in the same window. One can compute such rules and their confidences from the frequencies of episodes, just like confidences are computed for association rules.

Summary of the approach We propose automatic methods for the analysis of alarms, to aid in the knowledge acquisition phase of building an alarm correlation system. Briefly, the scenario for building alarm correlation systems is the following. First, a large database of alarms is analyzed off-line, and episode rules are discovered automatically. Then, for the construction of an alarm correlation system, the network management specialists have a collection of patterns of alarms at hand, and they can build the alarm correlation model from them. In the final step, the correlation rules are applied in real-time fault identification.

The discovery methods, although not directly applicable for on-line analysis themselves, can be used also in network surveillance. The analysis can be rerun or augmented, e.g., every day or every hour. Recent patterns may point to yet unnoticed faults in the network.

To sum up, the knowledge discovery task we consider is the following. Given a class of episodes, an input sequence of events, a window width, and a frequency threshold, find all episodes that are frequent in the windows on the event sequence. In the algorithm for solving this task we can use many ideas from the algorithm for discovering frequent sets. Additionally, episodes can be recognized efficiently in the database pass by “sliding” a window on the input sequence. Windows starting at successive positions have a lot of overlap and are therefore similar to each other. We can take advantage of this similarity: after recognizing episodes in a window, we can make incremental updates in our data structures to determine which episodes occur in the next window.

Chapter 4

Frequent episodes

In Chapter 3 we introduced the problem of analyzing telecommunication alarm sequences, and we briefly outlined a solution based on the idea of discovering frequent alarm sets in alarm sequences. Motivated by Chapter 3, we now give exact and domain-independent formulations for the ideas about finding all frequent episodes.

We define concepts related to event sequences and episodes in Section 4.1. In Section 4.2 we give algorithms for the discovery of all frequent parallel and serial episodes. We touch on the discovery of episodes with other types of partial orders, and on the rule generation process. Illustrative experimental results are presented in Section 4.3. Experiences are discussed in Chapter 6 in the context of the whole KDD process.

4.1 The framework

Let us now work out the approach in more detail. Since the framework is application independent—although motivated by fault management—we talk about events instead of alarms. We start by defining event sequences and windows.

Definition 4.1 Given a set R of *event types*, an *event* is a pair (A, t) , where $A \in R$ is an event type and t is an integer, the (*occurrence*) *time* of the event.

An *event sequence* \mathbf{s} on R is a triple (s, T_s, T_e) , where $T_s < T_e$ are integers, T_s is called the starting time and T_e the ending time, and

$$s = \langle (A_1, t_1), (A_2, t_2), \dots, (A_n, t_n) \rangle$$

is an ordered sequence of events such that $A_i \in R$ and $T_s \leq t_i < T_e$ for all $i = 1, \dots, n$, and $t_i \leq t_{i+1}$ for all $i = 1, \dots, n - 1$. \square

Example 4.2 Figure 4.1 presents graphically the event sequence $\mathbf{s} = (s, 29, 68)$, where

$$s = \langle (E, 31), (D, 32), (F, 33), (A, 35), (B, 37), (C, 38), \dots, (D, 67) \rangle.$$




Figure 4.1: The example event sequence \mathbf{s} and two windows of width 5.

Observations of the event sequence have been made from time 29 to just before time 68. For each event that occurred in the time interval $[29, 68)$, the event type and the time of occurrence have been recorded. \square

Think now of looking at an event sequence through a narrow window, giving a view to the events within a relatively small time period. We define a window as a slice of an event sequence that is seen at any given time. We then discuss the case where one considers an event sequence as a sequence of partially overlapping windows.

Definition 4.3 A *window* on event sequence $\mathbf{s} = (s, T_s, T_e)$ is an event sequence $\mathbf{w} = (w, t_s, t_e)$, where $t_s < T_e, t_e > T_s$, and w consists of those pairs (A, t) from s where $t_s \leq t < t_e$. The time span $t_e - t_s$ is called the *width* of the window \mathbf{w} , and it is denoted $width(\mathbf{w})$. Given an event sequence \mathbf{s} and an integer win , we denote by $\mathcal{W}(\mathbf{s}, win)$ the set of all windows \mathbf{w} on \mathbf{s} such that $width(\mathbf{w}) = win$. \square

There is a small trick in the definition. The first and last windows on a sequence extend outside the sequence, so that the first window only contains the first time point of the sequence, and the last window only contains the last time point. With this definition an event close to either end of a sequence is observed in equally many windows to an event in the middle of the sequence. Given an event sequence $\mathbf{s} = (s, T_s, T_e)$ and a window width win , the number of windows in $\mathcal{W}(\mathbf{s}, win)$ is $T_e - T_s + win - 1$.

Example 4.4 Figure 4.1 shows two windows of width 5 on the sequence \mathbf{s} of the previous example. A window starting at time 35 is shown in solid line, and the immediately following window, starting at time 36, is depicted with a dashed line. The window starting at time 35 is

$$(\langle (A, 35), (B, 37), (C, 38), (E, 39) \rangle, 35, 40).$$

Note that the event $(F, 40)$ that occurred at the ending time is not in the window. The window starting at 36 is similar to this one; the difference is that the first event $(A, 35)$ is missing and there is a new event $(F, 40)$ at the end.

The set of the 43 partially overlapping windows of width 5 constitutes $\mathcal{W}(\mathbf{s}, 5)$; the first window is $(\emptyset, 25, 30)$, and the last is $(\langle (D, 67) \rangle, 67, 72)$.




Figure 4.2: An episode.

Event $(D, 67)$ occurs in 5 windows of width 5, as does, e.g., event $(C, 50)$. If only windows totally within the sequence were considered, event $(D, 67)$ would occur only in window $\langle\langle(A, 65), (D, 67)\rangle\rangle, 63, 68$. \square

We now move on to define episodes formally. We also define when an episode is a subepisode of another; this relation is then used in the candidate generation for frequent episodes.

Definition 4.5 An *episode* α is a triple (V, \leq, g) where V is a set of nodes, \leq is a partial order on V , and $g : V \rightarrow R$ is a mapping associating each node with an event type. The interpretation of an episode is that the events in $g(V)$ have to occur in the order described by \leq . The *size* of α , denoted $|\alpha|$, is $|V|$. Episode α is *parallel* if the partial order \leq is trivial (i.e., $x \not\leq y$ for all $x, y \in V$ such that $x \neq y$). Episode α is *serial* if the relation \leq is a total order (i.e., $x \leq y$ or $y \leq x$ for all $x, y \in V$). Episode α is *injective* if the mapping g is an injection, i.e., no event type occurs twice in the episode. \square

Example 4.6 An episode $\alpha = (V, \leq, g)$ can be drawn as a directed acyclic graph, where $g(x)$ is the label of node $x \in V$, and where edges indicate temporal ordering. Consider episode $\alpha = (V, \leq, g)$ in Figure 4.2. The set V contains two nodes; call them x and y . The mapping g labels these nodes with the event types that are seen in the figure: $g(x) = E$ and $g(y) = F$. An event of type E is supposed to occur before an event of type F , i.e., x precedes y , and we have $x \leq y$. Episode α is injective, since it does not contain duplicate event types; in a window where α occurs there may, however, be multiple events of types E and F . \square

Definition 4.7 An episode $\beta = (V', \leq', g')$ is a *subepisode* of $\alpha = (V, \leq, g)$, denoted $\beta \preceq \alpha$, if there exists an injective mapping $f : V' \rightarrow V$ such that $g'(v) = g(f(v))$ for all $v \in V'$, and for all $v, w \in V'$ with $v \leq' w$ also $f(v) \leq f(w)$. An episode α is a *superepisode* of β if and only if $\beta \preceq \alpha$. We write $\beta \prec \alpha$ if $\beta \preceq \alpha$ and $\alpha \not\preceq \beta$. \square

Example 4.8 Figure 4.3 presents two episodes, β and γ . From the figure we see that we have $\beta \preceq \gamma$ since β is a subgraph of γ . In terms of Definition 4.7,




Figure 4.3: A subepisode and episode.

there is a mapping f that connects the nodes labeled A with each other and the nodes labeled B with each other, i.e., both nodes of β have (disjoint) corresponding nodes in γ . Since the nodes in episode β are not ordered, the corresponding nodes in γ do not need to be ordered, either, but they could be. \square

Consider now what it means that an episode occurs in a sequence. The nodes of the episode need to have corresponding events in the sequence such that the event types are the same and the partial order of the episode is respected. Below we formalize this. We also define the frequency of an episode as the fraction of windows in which the episode occurs.

Definition 4.9 An episode $\alpha = (V, \leq, g)$ occurs in an event sequence

$$\mathbf{s} = (\langle (A_1, t_1), (A_2, t_2), \dots, (A_n, t_n) \rangle, T_s, T_e),$$

if there exists an injective mapping $h : V \rightarrow \{1, \dots, n\}$ from nodes to events, such that $g(x) = A_{h(x)}$ for all $x \in V$, and for all $x, y \in V$ with $x \neq y$ and $x \leq y$ we have $t_{h(x)} < t_{h(y)}$. \square

Example 4.10 The window $(w, 35, 40)$ of Figure 4.1 contains events A , B , C , and E , in that order. Both episodes β and γ of Figure 4.3 occur in the window. \square

Definition 4.11 Given an event sequence \mathbf{s} and a window width win , the frequency of an episode α in \mathbf{s} is

$$fr(\alpha, \mathbf{s}, win) = \frac{|\{\mathbf{w} \in \mathcal{W}(\mathbf{s}, win) \mid \alpha \text{ occurs in } \mathbf{w}\}|}{|\mathcal{W}(\mathbf{s}, win)|},$$

i.e., the fraction of windows on \mathbf{s} in which α occurs.

Given a frequency threshold min_fr , α is frequent if $fr(\alpha, \mathbf{s}, win) \geq min_fr$. The collection of episodes that are frequent in \mathbf{s} with respect to win and min_fr is denoted $\mathcal{F}(\mathbf{s}, win, min_fr)$. The collection of frequent episodes of size l is denoted $\mathcal{F}_l(\mathbf{s}, win, min_fr)$. \square

Algorithm 4.13

Input: A set R of event types, an event sequence \mathbf{s} over R , a set \mathcal{E} of episodes, a window width win , and a frequency threshold min_fr .

Output: The collection $\mathcal{F}(\mathbf{s}, win, min_fr)$ of frequent episodes.

Method:

1. compute $\mathcal{C}_1 := \{\alpha \in \mathcal{E} \mid |\alpha| = 1\}$;
 2. $l := 1$;
 3. **while** $\mathcal{C}_l \neq \emptyset$ **do**
 4. // Database pass (Algorithms 4.19 and 4.21):
 5. compute $\mathcal{F}_l(\mathbf{s}, win, min_fr) := \{\alpha \in \mathcal{C}_l \mid fr(\alpha, \mathbf{s}, win) \geq min_fr\}$;
 6. $l := l + 1$;
 7. // Candidate generation (Algorithm 4.14):
 8. compute $\mathcal{C}_l := \{\alpha \in \mathcal{E} \mid |\alpha| = l, \text{ and } \beta \in \mathcal{F}_{|\beta|}(\mathbf{s}, win, min_fr) \text{ for all } \beta \in \mathcal{E} \text{ such that } \beta \prec \alpha \text{ and } |\beta| < l\}$;
 9. **for all** l **do** output $\mathcal{F}_l(\mathbf{s}, win, min_fr)$;
-

We can now give an exact formulation of the discovery task at hand: given an event sequence \mathbf{s} , a set \mathcal{E} of episodes, a window width win , and a frequency threshold min_fr , find $\mathcal{F}(\mathbf{s}, win, min_fr)$.

4.2 Algorithms

In the algorithms of this section we assume that the search space \mathcal{E} consists of all parallel or serial episodes constructed from the events in a given set R . Our main algorithm makes only the assumption that there are candidates and frequent episodes of each size up to the size of the largest frequent episodes; the more detailed algorithms for candidate generation and database pass, in turn, assume that episodes are either parallel or serial. Towards the end of this section we briefly consider the discovery of episodes with more general partial orders.

4.2.1 Main algorithm

Algorithm 4.13 computes the collection $\mathcal{F}(\mathbf{s}, win, min_fr)$ of frequent episodes. The algorithm has the familiar structure of alternation between candidate generation and database pass phases. The crucial point in the candidate generation is given by the following immediate lemma.

Lemma 4.12 If an episode α is frequent in an event sequence \mathbf{s} , then all subepisodes $\beta \preceq \alpha$ are frequent.

4.2.2 Generation of candidate episodes

We present in detail a candidate generation method which is a generalization of the candidate generation for frequent sets. The method can be adapted to

Algorithm 4.14**Input:** A sorted array \mathcal{F}_l of frequent parallel episodes of size l .**Output:** A sorted array of candidate parallel episodes of size $l + 1$.**Method:**

```

1.   $\mathcal{C}_{l+1} := \emptyset$ ;
2.   $k := 0$ ;
3.  if  $l = 1$  then for  $h := 1$  to  $|\mathcal{F}_l|$  do  $\mathcal{F}_l.block\_start[h] := 1$ ;
4.  for  $i := 1$  to  $|\mathcal{F}_l|$  do
5.       $current\_block\_start := k + 1$ ;
6.      for ( $j := i$ ;  $\mathcal{F}_l.block\_start[j] = \mathcal{F}_l.block\_start[i]$ ;  $j := j + 1$ ) do
7.          //  $\mathcal{F}_l[i]$  and  $\mathcal{F}_l[j]$  have  $l - 1$  first event types in common,
8.          // build a potential candidate  $\alpha$  as their combination:
9.          for  $x := 1$  to  $l$  do  $\alpha[x] := \mathcal{F}_l[i][x]$ ;
10.          $\alpha[l + 1] := \mathcal{F}_l[j][l]$ ;
11.         for  $y := 1$  to  $l - 1$  do
12.             // Build and test subepisodes  $\beta$  that do not contain  $\alpha[y]$ :
13.             for  $x := 1$  to  $y - 1$  do  $\beta[x] := \alpha[x]$ ;
14.             for  $x := y$  to  $l$  do  $\beta[x] := \alpha[x + 1]$ ;
15.             if  $\beta$  is not in  $\mathcal{F}_l$  then continue with the next  $j$  at line 6;
16.         // All subepisodes are in  $\mathcal{F}_l$ , store  $\alpha$  as candidate:
17.          $k := k + 1$ ;
18.          $\mathcal{C}_{l+1}[k] := \alpha$ ;
19.          $\mathcal{C}_{l+1}.block\_start[k] := current\_block\_start$ ;
20. output  $\mathcal{C}_{l+1}$ ;

```

deal with parallel episodes (i.e., multisets of items), serial episodes (ordered multisets), and injective parallel and serial episodes (sets and ordered sets).

Algorithm 4.14 computes the candidates for parallel episodes. In the algorithm, an episode $\alpha = (V, \leq, g)$ is represented as a lexicographically sorted array of event types. The array is denoted by the name of the episode and the items in the array are referred to using square brackets. For example, a parallel episode α with events of types A, C, C , and F is represented as an array α with $\alpha[1] = A, \alpha[2] = C, \alpha[3] = C$, and $\alpha[4] = F$. Collections of episodes are also represented as lexicographically sorted arrays, i.e., the i th episode of a collection \mathcal{F} is denoted by $\mathcal{F}[i]$.

Since the episodes and episode collections are sorted, all episodes that share the same first event types are consecutive in the episode collection. In particular, if episodes $\mathcal{F}_l[i]$ and $\mathcal{F}_l[j]$ of size l share the first $l - 1$ events, then for all k with $i \leq k \leq j$ we have that $\mathcal{F}_l[k]$ shares also the same events. A maximal sequence of consecutive episodes of size l that share the first $l - 1$ events is called a *block*. Potential candidates can be identified by creating all combinations of two episodes in the same block. For the efficient identification of blocks, we store in $\mathcal{F}_l.block_start[j]$ for each episode $\mathcal{F}_l[j]$ the index i such that $\mathcal{F}_l[i]$ is the first episode in the block.

Theorem 4.15 Algorithm 4.14 works correctly.

Proof The crucial claim is that in the algorithm the pairs $\mathcal{F}_l[i]$ and $\mathcal{F}_l[j]$ of

episodes generate all candidates. For a moment assume that for each episode $\mathcal{F}_l[j]$ the value of $\mathcal{F}_l.\text{block_start}[j]$ is the index i such that $\mathcal{F}_l[i]$ is the first episode in the block. We show later that this assumption holds. —In the following we identify an episode with its index in the collection.

In the outer loop (line 4) variable i iterates through all episodes in \mathcal{F}_l , and in the inner loop (line 6) variable j iterates through those episodes in \mathcal{F}_l that are in the same block with i but are not before i . Consider now any block b of episodes in \mathcal{F}_l . Variables i and j obviously iterate through all (unordered) pairs of episodes in block b , including the case where $i = j$.

Since i and j are in the same block, they have the same $l - 1$ first event types. Conceptually we construct a new *potential candidate* α as the union of episodes (multisets) i and j . We build α by taking first the common $l - 1$ events and the l th event from episode i (both done on line 9), and finally the event number $l + 1$ from episode j (line 10). Then the events of a potential candidate are lexicographically sorted. Since the iteration of episodes proceeds in lexicographical order (over the sorted collection \mathcal{F}_l), the collection of candidates is also constructed in lexicographical order.

Next we show that the collection of potential candidates α contains all valid candidates γ of size $l + 1$. All subepisodes of γ are frequent, and in particular those two subepisodes δ_1 and δ_2 of size l that contain all but the last and the second last events of γ , respectively. Since δ_1 and δ_2 are in \mathcal{F}_l and they have $l - 1$ items in common, they are in the same block. At some time in the algorithm we have $\mathcal{F}_l[i] = \delta_1$ and $\mathcal{F}_l[j] = \delta_2$, and γ is considered as a potential candidate in the algorithm.

We need to show that no false candidates are output. An episode of size $l + 1$ has $l + 1$ subepisodes β of size l , and for all of these we make sure that they are in \mathcal{F}_l . We obtain all these subepisodes by leaving out one of the events in α at a time (line 11). Note that the two subepisodes that were used for constructing α do not need to be checked again. Only if all subepisodes of size $l - 1$ are in \mathcal{F}_l , is α correctly output as a candidate.

Finally we show that we have the correct value $\mathcal{F}_l.\text{block_start}[j] = i$ for all j , i.e., the index i such that $\mathcal{F}_l[i]$ is the first episode in the block. For $l = 1$ the structure is built on line 3: all episodes of size 1 have at least 0 common events, so they are all in the same block, and $\mathcal{F}_1.\text{block_start}[h] = 1$ for all h . For $l \geq 1$ and \mathcal{F}_{l+1} , a block b of \mathcal{F}_{l+1} (or \mathcal{C}_{l+1}) has the property that all episodes in the block have been generated from the same episode $\mathcal{F}_l[i]$. This is due to the simple fact that the first l events have been copied directly from $\mathcal{F}_l[i]$ (line 9). We save for each i the index of the first candidate generated from it (line 5), and then use the saved value to set $\mathcal{C}_{l+1}.\text{block_start}[k]$ correctly for all candidates k in the block (line 19). \square

Algorithm 4.14 can be easily modified to generate candidate serial episodes. Now the events in the array representing an episode are in the order imposed by a total order \leq . For instance, a serial episode β with events

of types C, A, F , and C , in that order, is represented as an array β with $\beta[1] = C$, $\beta[2] = A$, $\beta[3] = F$, and $\beta[4] = C$. Collections of episodes are still stored as lexicographically sorted arrays. The only change to the algorithm is to replace line 6.

Theorem 4.16 With the line

6. **for** ($j := \mathcal{F}_l.block_start[i]; \mathcal{F}_l.block_start[j] = \mathcal{F}_l.block_start[i];$
 $j := j + 1$) **do**

Algorithm 4.14 works correctly for serial episodes.

Proof The proof is similar to the proof for Theorem 4.15; now, however, i and j iterate over all *ordered* pairs of episodes in each block. The (potential) candidates are ordered sequences of event types, not sorted arrays as before, but the candidate collection is still constructed in lexicographical order. The same arguments for the correctness of the candidate collection and the structure $\mathcal{F}_l.block_start$ hold. \square

There are further options with the algorithm. If the desired episode class consists of parallel or serial injective episodes, i.e., no episode should contain any event type more than once, simply add one line.

Theorem 4.17 With the line

6b. **if** $j = i$ **then** continue with the next j at line 6;

inserted after line 6, Algorithm 4.14 works correctly for injective parallel episodes (or injective serial episodes with the change of Theorem 4.16).

Proof Clearly, the effect of the inserted line is that some candidates are not generated. Consider now those excluded candidate episodes. First note that only candidates α that contain some event type at least twice are excluded. Either a candidate is excluded explicitly because $i = j$, or it is not generated because some of its subepisodes is not in \mathcal{F}_l . If α is excluded explicitly, then it contains the event type $\alpha[l] = \alpha[l + 1]$ twice. If, on the other hand, some tested subepisode β is not in the collection \mathcal{F}_l , then there must be a subepisode $\gamma \preceq \beta$ that has been excluded explicitly. Then α contains twice the event type $\gamma[|\gamma|]$.

Now note that no episode α with at least two occurrences of an event type is generated. Let A be an event type that occurs at least twice in α . Then for the episode γ of size 2 such that $\gamma[1] = A$ and $\gamma[2] = A$ we have $\gamma \preceq \alpha$, and thus α cannot be a candidate unless γ is frequent. However, γ has been excluded explicitly by the inserted line in an earlier iteration, and thus α is not a candidate. \square

The time complexity of Algorithm 4.14 is polynomial in the size of the collection of frequent episodes and it is independent of the length of the event sequence.

Theorem 4.18 Algorithm 4.14 (with any of the above variations) has time complexity $\mathcal{O}(l^2 |\mathcal{F}_l|^2 \log |\mathcal{F}_l|)$.

Proof The initialization (line 3) takes time $\mathcal{O}(|\mathcal{F}_l|)$. The outer loop (line 4) is iterated $\mathcal{O}(|\mathcal{F}_l|)$ times and the inner loop (line 6) $\mathcal{O}(|\mathcal{F}_l|)$ times. Within the loops, a potential candidate (lines 9 and 10) and $l - 1$ subcandidates (lines 11 to 14) are built in time $\mathcal{O}(l + 1 + (l - 1)l) = \mathcal{O}(l^2)$. More importantly, the $l - 1$ subsets need to be searched for in the collection \mathcal{F}_l (line 15). Since \mathcal{F}_l is sorted, each subcandidate can be located with binary search in time $\mathcal{O}(l \log |\mathcal{F}_l|)$. The total time complexity is thus $\mathcal{O}(|\mathcal{F}_l| + |\mathcal{F}_l| |\mathcal{F}_l| (l^2 + (l - 1) l \log |\mathcal{F}_l|)) = \mathcal{O}(l^2 |\mathcal{F}_l|^2 \log |\mathcal{F}_l|)$. \square

In practical situations the time complexity is likely to be close to $\mathcal{O}(l^2 |\mathcal{F}_l| \log |\mathcal{F}_l|)$, since the blocks are typically small.

4.2.3 Recognizing episodes in sequences

Let us now consider the implementation of the database pass. We give algorithms which recognize episodes in sequences in an incremental fashion. For two windows $\mathbf{w} = (w, t_s, t_s + win)$ and $\mathbf{w}' = (w', t_s + 1, T_s + win + 1)$, the sequences w and w' of events are similar to each other. We take advantage of this similarity: after recognizing episodes in \mathbf{w} , we make incremental updates in our data structures to achieve the shift of the window to obtain \mathbf{w}' .

The algorithms start by considering the empty window just before the input sequence, and they end after considering the empty window just after the sequence. This way the incremental methods need no other special actions at the beginning or end. For the frequency of episodes, only the windows correctly on the input sequence are, of course, considered.

Parallel episodes Algorithm 4.19 recognizes candidate parallel episodes in an event sequence. The main ideas of the algorithm are the following. For each candidate parallel episode α we maintain a counter $\alpha.event_count$ that indicates how many events of α are present in the window. When $\alpha.event_count$ becomes equal to $|\alpha|$, indicating that α is entirely included in the window, we save the starting time of the window in $\alpha.inwindow$. When $\alpha.event_count$ decreases again, indicating that α is no longer entirely in the window, we increase the field $\alpha.freq_count$ by the number of windows where α remained entirely in the window. At the end, $\alpha.freq_count$ contains the total number of windows where α occurs.

To access candidates efficiently, they are indexed by the number of events of each type that they contain: all episodes that contain exactly a events of type A are in the list $contains(A, a)$. When the window is shifted and the contents of the window change, the episodes that are affected are updated. If, for instance, there is one event of type A in the window and a second

Algorithm 4.19

Input: A collection \mathcal{C} of parallel episodes, an event sequence $\mathbf{s} = (s, T_s, T_e)$, a window width win , and a frequency threshold min_fr .

Output: The episodes of \mathcal{C} that are frequent in \mathbf{s} with respect to win and min_fr .

Method:

```

1. // Initialization:
2. for each  $\alpha$  in  $\mathcal{C}$  do
3.   for each  $A$  in  $\alpha$  do
4.      $A.count := 0$ ;
5.     for  $i := 1$  to  $|\alpha|$  do  $contains(A, i) := \emptyset$ ;
6. for each  $\alpha$  in  $\mathcal{C}$  do
7.   for each  $A$  in  $\alpha$  do
8.      $a :=$  number of events of type  $A$  in  $\alpha$ ;
9.      $contains(A, a) := contains(A, a) \cup \{\alpha\}$ ;
10.   $\alpha.event\_count := 0$ ;
11.   $\alpha.freq\_count := 0$ ;
12. // Recognition:
13. for  $start := T_s - win + 1$  to  $T_e$  do
14.   // Bring in new events to the window:
15.   for all events  $(A, t)$  in  $s$  such that  $t = start + win - 1$  do
16.      $A.count := A.count + 1$ ;
17.     for each  $\alpha \in contains(A, A.count)$  do
18.        $\alpha.event\_count := \alpha.event\_count + A.count$ ;
19.       if  $\alpha.event\_count = |\alpha|$  then  $\alpha.inwindow := start$ ;
20.   // Drop out old events from the window:
21.   for all events  $(A, t)$  in  $s$  such that  $t = start - 1$  do
22.     for each  $\alpha \in contains(A, A.count)$  do
23.       if  $\alpha.event\_count = |\alpha|$  then
24.          $\alpha.freq\_count := \alpha.freq\_count - \alpha.inwindow + start$ ;
25.          $\alpha.event\_count := \alpha.event\_count - A.count$ ;
26.          $A.count := A.count - 1$ ;
27. // Output:
28. for all episodes  $\alpha$  in  $\mathcal{C}$  do
29.   if  $\alpha.freq\_count / (T_e - T_s + win - 1) \geq min\_fr$  then output  $\alpha$ ;
```

one comes in, all episodes in the list $contains(A, 2)$ are updated with the information that both events of type A they are expecting are now present.

Theorem 4.20 Algorithm 4.19 works correctly.

Proof We consider the following two invariants. (1) For each event type A that occurs in any episode, the variable $A.count$ correctly contains the number of events of type A in the current window. (2) For each episode α , the counter $\alpha.event_count$ equals $|\alpha|$ exactly when α occurs in the current window.

The first invariant holds trivially for the empty window starting at $T_s - win$, since counters $A.count$ are initialized to zero on line 4. Assume now that the counters are correct for the window starting at $start - 1$, and consider the computation for the window starting at $start$, i.e., one iteration of the loop starting at line 13. On lines 15 – 16, the counters are updated for each new

event in the window; similarly, on lines 21 and 26, the counters are updated for events no longer in the window.

For the second invariant, first note that each set $\text{contains}(A, a)$ consists of all episodes that contain exactly a events of type A : the lists are initialized to empty on line 5, and then filled correctly for each event type in each episode on line 9. Now consider the counter $\alpha.\text{event_count}$ for any episode α . In the beginning, the counter is initialized to zero (line 10). Given an event type A , denote by a the number of events of type A in α . The effect of lines 18 and 25 is that $\alpha.\text{event_count}$ is increased by a exactly for the time when there are at least a events of type A in the window. Thus $\alpha.\text{event_count} = |\alpha|$ exactly when there are enough events of each type of α in the window.

Finally note that at the end $\alpha.\text{freq_count}$ is correct. The counter is initialized to zero (line 11). Given any number of consecutive windows containing α , by the invariant the index of the first window is stored in $\alpha.\text{inwindow}$ on line 19. After the last window of these, i.e., in the first window not containing α , the counter $\alpha.\text{freq_count}$ is increased by the number of the consecutive windows containing α (line 24). Since the last window considered is the empty window immediately after the sequence, occurrences in the last windows on the sequence are correctly computed. On the last lines the frequent episodes are output. \square

Serial episodes Serial candidate episodes are recognized in an event sequence by using state automata that accept the candidate episodes and ignore all other input. The idea is that there is an automaton for each serial episode α , and that there can be several instances of each automaton at the same time, so that the active states reflect the (disjoint) prefixes of α occurring in the window. Algorithm 4.21 implements this idea.

We initialize a new instance of the automaton for a serial episode α every time the first event of α comes into the window; the instance is removed when the same event leaves the window. When an automaton instance for α reaches its accepting state, indicating that α is entirely included in the window, and if there are no other instances for α in the accepting state already, we save the starting time of the window in $\alpha.\text{inwindow}$. When an automaton instance in the accepting state is removed, and if there are no other instances for α in the accepting state, we increase the field $\alpha.\text{freq_count}$ by the number of windows where α remained entirely in the window.

It is useless to have multiple automaton instance in the same state, as they would only make the same transitions and produce the same information. It suffices to maintain the one that reached the common state last since it will be also removed last. There are thus at most $|\alpha|$ automaton instances for an episode α . For each instance we need to know when it should be removed. We can thus represent all the automaton instances for α with one array of size $|\alpha|$: the value of $\alpha.\text{initialized}[i]$ is the latest initialization time of an

Algorithm 4.21

Input: A collection \mathcal{C} of serial episodes, an event sequence $\mathbf{s} = (s, T_s, T_e)$, a window width win , and a frequency threshold min_fr .

Output: The episodes of \mathcal{C} that are frequent in \mathbf{s} with respect to win and min_fr .

Method:

```

1. // Initialization:
2. for each  $\alpha$  in  $\mathcal{C}$  do
3.   for  $i := 1$  to  $|\alpha|$  do
4.      $\alpha.initialized[i] := 0$ ;
5.      $waits(\alpha[i]) := \emptyset$ ;
6. for each  $\alpha \in \mathcal{C}$  do
7.    $waits(\alpha[1]) := waits(\alpha[1]) \cup \{(\alpha, 1)\}$ ;
8.    $\alpha.freq\_count := 0$ ;
9. for  $t := T_s - win$  to  $T_s - 1$  do  $beginsat(t) := \emptyset$ ;
10. // Recognition:
11. for  $start := T_s - win + 1$  to  $T_e$  do
12.   // Bring in new events to the window:
13.    $beginsat(start + win - 1) := \emptyset$ ;
14.    $transitions := \emptyset$ ;
15.   for all events  $(A, t)$  in  $s$  such that  $t = start + win - 1$  do
16.     for all  $(\alpha, j) \in waits(A)$  do
17.       if  $j = |\alpha|$  and  $\alpha.initialized[j] = 0$  then  $\alpha.inwindow := start$ ;
18.       if  $j = 1$  then
19.          $transitions := transitions \cup \{(\alpha, 1, start + win - 1)\}$ ;
20.       else
21.          $transitions := transitions \cup \{(\alpha, j, \alpha.initialized[j - 1])\}$ ;
22.          $beginsat(\alpha.initialized[j - 1]) :=$ 
23.            $beginsat(\alpha.initialized[j - 1]) \setminus \{(\alpha, j - 1)\}$ ;
24.          $\alpha.initialized[j - 1] := 0$ ;
25.          $waits(A) := waits(A) \setminus \{(\alpha, j)\}$ ;
26.   for all  $(\alpha, j, t) \in transitions$  do
27.      $\alpha.initialized[j] := t$ ;
28.      $beginsat(t) := beginsat(t) \cup \{(\alpha, j)\}$ ;
29.     if  $j < |\alpha|$  then  $waits(\alpha[j + 1]) := waits(\alpha[j + 1]) \cup \{(\alpha, j + 1)\}$ ;
30.   // Drop out old events from the window:
31.   for all  $(\alpha, l) \in beginsat(start - 1)$  do
32.     if  $l = |\alpha|$  then  $\alpha.freq\_count := \alpha.freq\_count - \alpha.inwindow + start$ ;
33.     else  $waits(\alpha[l + 1]) := waits(\alpha[l + 1]) \setminus \{(\alpha, l + 1)\}$ ;
34.      $\alpha.initialized[l] := 0$ ;
35. // Output:
36. for all episodes  $\alpha$  in  $\mathcal{C}$  do
    if  $\alpha.freq\_count / (T_e - T_s + win - 1) \geq min\_fr$  then output  $\alpha$ ;

```

automaton instance that has reached its i th state. Recall that α itself is represented by an array containing its events; this array can be used to label the state transitions.

To access and traverse the automata instances efficiently they are organized in the following way. For each event type $A \in R$, the instances that would next accept A are linked together to a list $waits(A)$. The list contains entries of the form (α, x) meaning that episode α is waiting for its x th event. When an event (A, t) enters the window during a shift, the list $waits(A)$ is traversed. If an automaton instance reaches a common state with another

instance, the earlier entry in the array $\alpha.initialized[]$ is simply overwritten.

The transitions made during one shift of the window are stored in a list *transitions*. They are represented in the form (α, x, t) meaning that episode α got its x th event, and the latest initialization time of the prefix of length x is t . Updates regarding the old states of the automaton instances are done immediately, but updates for the new states are carried out only after all transitions have been identified, in order to not overwrite any useful information. For easy removal of automaton instances when they go out of the window, the instances initialized at time t are stored in a list *beginsat*(t). Note that only lists *beginsat*($start - 1$) – *beginsat*($start + win - 1$) are needed at any particular moment, and lists *beginsat*(j) with $j < start - 1$ can be removed.

Theorem 4.22 Algorithm 4.21 works correctly.

Proof Let α be a serial episode in \mathcal{C} , j an integer such that $1 \leq j \leq |\alpha|$, and A an event type, and consider a window on the input sequence. Denote by $mpt(\alpha, j)$ the maximal time t in the window such that the prefix of length j of α occurs within the subsequence starting at time t and ending at where the window ends. Consider the following invariants.

1. We have $\alpha.initialized[j] = 0$, if the prefix does not occur in the window at all, or if $j < |\alpha|$ and $mpt(\alpha, j) = mpt(\alpha, j + 1)$. Otherwise $\alpha.initialized[j] = mpt(\alpha, j)$.
2. For each time t in the window, we have $(\alpha, j) \in beginsat(t)$ if and only if $\alpha.initialized[j] = t$.
3. The list *waits*(A) consists of entries (α, j) such that $\alpha[j] = A$ and either $j = 1$ or $\alpha.initialized[j - 1] \neq 0$.

The first invariant holds trivially for the empty window in the beginning, as the data structures are initialized to zeros on line 4. Assume now that the data structures are correct for the window starting at $start - 1$, and consider the computation for the window starting at $start$. We show by induction that the computations are correct for all j . First, consider the case $j = 1$. When a new event comes to the window, it is always the latest prefix of length $j = 1$ for all episodes that start with the event type. The value of $\alpha.initialized[1]$ is correctly set to $start + win - 1$ for all such episodes α on lines 19 and 26. Assume now that $j > 1$, that $\alpha[j]$ comes into the window, and that $\alpha.initialized[k]$ is correct for all $k < j$. Now $mpt(\alpha, j)$ clearly equals the old value of $mpt(\alpha, j - 1)$; the correct updates are done on lines 21 and 26 for $\alpha.initialized[j]$ and on line 23 for $\alpha.initialized[j - 1]$. Note that the value of $\alpha.initialized[j - 1]$ is set to non-zero later if $mpt(\alpha, j - 1) > mpt(\alpha, j)$. Note also that when a prefix of length l is not in the window anymore, $\alpha.initialized[l]$ is correctly set to zero on line 33.

The second invariant holds also trivially in the beginning (line 9). Assuming that the data structures are correct for the window starting at $start - 1$, the correct additions to $beginsat$ are done on line 27, and correct removals on line 22. (Removing lists $beginsat(t)$ with $t < start$ is not necessary.)

The third invariant holds for $j = 1$ for the whole algorithm: the $waits$ lists are set correctly on line 7, and they are not altered during the algorithm. For larger prefixes correct additions to the $waits$ lists are made on lines 19, 21, and 28, and correct removals are made when $\alpha.initialized[j - 1]$ becomes zero (lines 24 and 32).

Based on these invariants, the index of the window is correctly stored in $\alpha.inwindow$ for the first of consecutive windows containing α (line 17), and $\alpha.freq_count$ is correctly increased after the last of consecutive windows containing α (line 31). Finally, the frequent episodes are correctly output on the last lines of the algorithm. \square

Analysis of time complexity For simplicity, suppose that the class of event types R is fixed, and assume that exactly one event takes place every time unit. Assume candidate episodes are all of size l , and let n be the length of the sequence. We go back to parallel episodes, and start the analysis from Algorithm 4.19.

Theorem 4.23 The time complexity of Algorithm 4.19 is $\mathcal{O}((n + l^2) |\mathcal{C}|)$.

Proof Initialization takes time $\mathcal{O}(|\mathcal{C}| l^2)$. Consider now the number of the operations in the innermost loops, i.e., accesses to $\alpha.event_count$ on lines 18 and 25. In the recognition phase there are $\mathcal{O}(n)$ shifts of the window. In each shift, one new event comes into the window, and one old event leaves the window. Thus, for any episode α , $\alpha.event_count$ is accessed at most twice during one shift. The cost of the recognition phase is thus $\mathcal{O}(n |\mathcal{C}|)$. \square

In practice the size l of episodes is very small with respect to the size n of the sequence, and the time required for the initialization can be safely neglected. For injective episodes we have the following tighter result.

Theorem 4.24 The time complexity of recognizing injective parallel episodes in Algorithm 4.19 (excluding initialization) is $\mathcal{O}(\frac{n}{win} |\mathcal{C}| l + n)$.

Proof Consider win successive shifts of one time unit. During such sequence of shifts, each of the $|\mathcal{C}|$ candidate episodes α can undergo at most $2l$ changes: any event type A of α can have $A.count$ increased to 1 and decreased to 0 at most once. This is due to the fact that after an event of type A has come into the window, we have $A.count \geq 1$ for the next win time units. Reading the input takes time n . \square

Compare this to a trivial non-incremental method where the sequence is pre-processed into windows, and then frequent sets are searched for. The

time requirement for recognizing $|\mathcal{C}|$ candidate sets in n windows, plus the time required to read in n windows of size win , is $\mathcal{O}(n|\mathcal{C}|l + n \cdot win)$, i.e., larger by a factor of win .

Theorem 4.25 The time complexity of Algorithm 4.21 is $\mathcal{O}(n|\mathcal{C}|l)$.

Proof The initialization takes time $\mathcal{O}(|\mathcal{C}|l + win)$. In the recognition phase, again, there are $\mathcal{O}(n)$ shifts, and in each shift one event comes into the window and one event leaves the window. In one shift, the effort per an episode α depends on the number of automaton instances accessed; there are a maximum of l instances for each episode. The worst-case time complexity is thus $\mathcal{O}(|\mathcal{C}|l + win + n|\mathcal{C}|l) = \mathcal{O}(n|\mathcal{C}|l)$ (note that win is $\mathcal{O}(n)$). \square

The input sequence consists in the worst case of events of only one event type, and the candidate serial episodes consist only of events of that particular type. Every shift of the window results now in an update in every automaton instance. This worst-case complexity is close to the complexity of the trivial non-incremental method $\mathcal{O}(n|\mathcal{C}|l + n \cdot win)$. In practical situations, however, the time requirement is considerably smaller, and we approach the savings obtained in the case of injective parallel episodes.

Theorem 4.26 The time complexity of recognizing injective serial episodes in Algorithm 4.21 (excluding initialization) is $\mathcal{O}(n|\mathcal{C}|)$.

Proof Each of the $\mathcal{O}(n)$ shifts can now affect at most two automaton instances for each episode: when an event comes into the window there can be a state transition in at most one instance, and at most one instance can be removed because the initializing event goes out of the window. \square

4.2.4 General partial orders

So far we have only discussed serial and parallel episodes. We now discuss briefly the use of other partial orders in episodes. The recognition of an arbitrary episode can be reduced to the recognition of a hierarchical combination of serial and parallel episodes. For example, episode γ in Figure 4.4 is a serial combination of two episodes: δ' , a parallel episode consisting of A and B , and δ'' , an episode consisting of C alone. The occurrence of an episode in a window can be tested using such hierarchical structure: to see whether episode γ occurs in a window one checks (using a method for serial episodes) whether δ' and δ'' occur in this order; to check the occurrence of δ' one uses a method for parallel episodes to verify whether A and B occur.

There are, however, some complications one has to take into account. First, it is sometimes necessary to duplicate an event node to obtain a decomposition to serial and parallel episodes. Consider, for instance, the episode on the left in Figure 4.5. There is no hierarchical composition consisting only




Figure 4.4: Recursive composition of a complex episode.




Figure 4.5: Recursive composition of a complex episode.

of serial and parallel episodes. In the composite episode on the right, the node labeled B has been duplicated. Such duplication works with injective episodes, but non-injective episodes need more complex methods. Another important aspect is that composite events have a duration, unlike the elementary events in R .

A practical alternative is to handle all episodes basically like parallel episodes, and to check the correct partial ordering only when all events are in the window. Parallel episodes can be located efficiently; after they have been found, checking the correct partial ordering is relatively fast.

Another interesting approach to the recognition of episodes is to use inverse structures. That is, for each frequent episode we store the identifiers of the windows in which the episode occurs. Then, in the recognition phase, for a candidate episode α we can compute the set of windows in which α occurs as the intersection of the sets of windows for two subepisodes of α . This holds for all but serial episodes, for which some additional information is needed.

4.2.5 Rule generation

Once the frequent episodes and their frequencies are known, one can generate episode rules, similar to association rules. An episode rule can state, for instance, that if there are events of types A and B in a window, then an

event of type C is also in the window (parallel episode rule), or that if there are events E and G in the window and in that order, then an event of type F is between them (serial episode rule). Serial episode rules can point forward or backward in time and, as illustrated by the example, the left-hand side can also have places that are filled by corresponding events on the right hand side. Episode rules can be generated from frequent episodes in the same way that association rules are generated from frequent sets. Extending the association rule methods to deal with multisets and ordered sets is fairly straightforward.

4.3 Experiments

We present experimental results obtained with two telecommunication network fault management databases. The first database s_1 is a sequence of 73 679 alarms covering a time period of 7 weeks. The time granularity is one second. There are 287 different types of alarms with very diverse frequencies and distributions. On the average there is an alarm every minute. The alarms tend, however, occur in bursts: in the extreme cases there are over 40 alarms in one second. We present results from experiments with serial episodes and injective parallel episodes, i.e., the opposite extreme cases of the complexity of the recognition phase.

Performance overview Tables 4.1 and 4.2 give an overview of the discovery of frequent episodes. In Table 4.1, serial episodes and injective parallel episodes have been discovered in s_1 with a fixed frequency threshold 0.003 and a varying window width; in Table 4.2, episodes have been discovered with a fixed window width of 60 seconds and a varying frequency threshold. These ranges for the parameter values have been given by experts of the alarm correlation domain.

The experiments show that the approach is efficient. Running times are between 5 seconds and 8 minutes, in which time hundreds of frequent episodes could be found. The methods are robust in the sense that a change in one parameter only adds or removes some frequent episodes, but does not replace any.

Quality of candidate generation Table 4.3 shows the number of candidate and frequent serial episodes per iteration, with frequency threshold 0.003, and averaged over test runs with window widths 10, 20, 40, 60, 80, 100, and 120 seconds.

In the first iteration, for size 1, all 287 event types have to be checked. The larger the episodes become, the more combinatorial information there exists to take advantage of. From size 4 up, over one half of the candidates turned out to be frequent.

Window width (s)	Serial episodes		Injective parallel episodes	
	Count	Time (s)	Count	Time (s)
10	16	31	10	8
20	31	63	17	9
40	57	117	33	14
60	87	186	56	15
80	145	271	95	21
100	245	372	139	21
120	359	478	189	22

Table 4.1: Results of experiments with s_1 using a fixed frequency threshold of 0.003 and a varying window width.

Frequency threshold	Serial episodes		Injective parallel episodes	
	Count	Time (s)	Count	Time (s)
0.1	0	7	0	5
0.05	1	12	1	5
0.008	30	62	19	14
0.004	60	100	40	15
0.002	150	407	93	22
0.001	357	490	185	22

Table 4.2: Results of experiments with s_1 using a fixed window width of 60 s and a varying frequency threshold.

As can be seen from the table, a possible practical improvement is to combine iterations by generating candidate episodes for several iterations at once, and thus avoid reading the input sequence so many times. This pays off in the later iterations, where there are otherwise only few candidates to recognize, and where the match is good.

Scale-up We performed scale-up tests with 1 to 8 fold multiples of the sequence s_1 , i.e., sequences with approximately 74 000 to 590 000 events. The results in Figure 4.6 show that the time requirement is linear with respect to the length of the input sequence, as could be expected from the analysis.

Incremental recognition We also tested the efficiency of the database pass, in particular the effect of the incremental recognition. Figure 4.7 presents the ratio of times needed for trivial vs. incremental recognition of candidate episodes. The time required to generate the windows for the trivial

Episode size l	Number of episodes 287^l	Number of candidate episodes	Number of frequent episodes	Match
1	287	287.0	30.1	11 %
2	82 369	1 078.7	44.6	4 %
3	$2 \cdot 10^7$	192.4	20.0	10 %
4	$7 \cdot 10^9$	17.4	10.1	58 %
5	$2 \cdot 10^{12}$	7.1	5.3	74 %
6	$6 \cdot 10^{14}$	4.7	2.9	61 %
7	$2 \cdot 10^{17}$	2.9	2.1	75 %
8	$5 \cdot 10^{19}$	2.1	1.7	80 %
9	$1 \cdot 10^{22}$	1.7	1.4	83 %
10-		17.4	16.0	92 %

Table 4.3: Number of candidate and frequent serial episodes in s_1 with frequency threshold 0.003 and averaged over window widths 10, 20, 40, 60, 80, 100, and 120 s.




Figure 4.6: Scale-up results for serial episodes (dotted line) and injective parallel episodes (solid line) in s_1 with window width 60 s and frequency threshold 0.01.




Figure 4.7: Ratio of times needed for trivial vs. incremental recognition methods in s_1 for serial episodes (dotted line) and injective parallel episodes (solid line) as functions of window width.

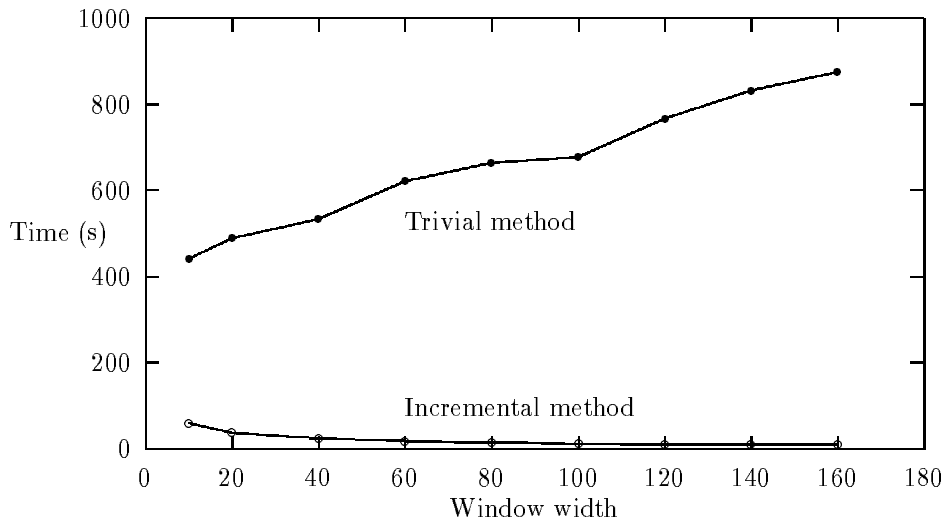
method has been excluded from the results. The figure shows that the incremental methods are faster by a factor of 1–20, roughly linearly with respect to the window width of 10–120 seconds. This is consistent with the analysis of Algorithm 4.19: for injective parallel episodes the worst-case analysis gave a difference of a factor of win . The results indicate that the incremental recognition method is useful in practice also for non-injective serial episodes.

To analyze the effect of the incremental recognition in more detail we conducted the following more controlled tests. We used an alarm database s_2 from a different network; this sequence contains 5000 events covering a time period of 6 days. We ignored the actual times of the events and assumed instead that one alarm had arrived in a time unit. There are 119 event types and the number of their occurrences ranges from 1 to 817. For these tests we considered only injective parallel episodes.

Table 4.4 presents results of test runs with different window widths and frequency thresholds. Results about the efficiency with respect to the number of frequent episodes and candidates are similar to the ones obtained with sequence s_1 . The frequency thresholds have in these experiments been higher since the data is dense: there is an event every time unit.

A central outcome of the analysis of the windowing methods was the effect of window width win on the time complexity. We examined it with tests where all other factors, in particular the candidate collection, were fixed. Our

Window width	Frequency threshold	Candidate episodes	Frequent episodes	Time (s)
10	0.05	444	84	16
20	0.10	463	161	27
40	0.20	632	346	71
60	0.30	767	488	84
80	0.40	841	581	112
100	0.50	755	529	90
120	0.60	578	397	49
160	0.70	633	479	64

Table 4.4: Results of experiments with sequence s_2 .Figure 4.8: Time for the database pass over s_2 as a function of the window width.

collection of candidates consists of 385 episodes of sizes 1 to 11. Figure 4.8 presents total times for recognizing these candidates in the input sequence s_2 . Within window widths of 10 to 160 time units, the total time with the trivial method doubles from 400 to 800 seconds. With the incremental method the time is in turn cut from 60 down to 10 seconds. The running time of the trivial method is approximately $3win + 420$, and for the incremental method $700/win + 5$. These results match the time complexity analysis given earlier. In particular, the time complexity of the trivial method is greater by a factor of the window width win ; the approximating functions give a factor of $0.6win$. The efficiency ratio was in these experiments better than in the experiments described earlier: the ratio ranges from 6 up to 80.

Chapter 5

Minimal occurrences of episodes

In this chapter we describe an alternative approach to the discovery of episodes. Instead of looking at the windows and only considering whether an episode occurs in a window or not, we now look at the exact occurrences of episodes and the relationships between those occurrences. One of the advantages of this approach is that focusing on the occurrences of episodes allows us to more easily find rules with two window widths, one for the left-hand side and one for the whole rule, such as “if A and B occur within 15 seconds, then C follows within 30 seconds”.

We give an introduction to the approach in Section 5.1. Sections 5.2 and 5.3 present the main ideas of the discovery of all frequent episodes and the rule generation, respectively. Finally, experiments are described in Section 5.4. In this chapter our aim is to convey the main ideas of an alternative approach, and we skip detailed proofs and algorithms.

5.1 Outline of the approach

The approach is based on minimal occurrences of episodes. Besides the new rule formulation, the use of minimal occurrences gives rise to the following new method for the recognition of episodes in the input sequence. For each frequent episode we store information about the locations of its minimal occurrences. In the recognition phase we can then compute the locations of minimal occurrences of a candidate episode α as a temporal join of the minimal occurrences of two subepisodes of α . In addition to being simple and efficient, this formulation has the advantage that the confidences and frequencies of rules with a large number of different window widths can be obtained quickly, i.e., there is no need to rerun the analysis if one only wants to modify the window widths. In the case of complicated episodes, the time needed for recognizing the occurrence of an episode can be significant; the use





Figure 5.1: Episodes.

Figure 5.2: The example event sequence \mathbf{s} .

of stored minimal occurrences of episodes eliminates unnecessary repetition of the recognition effort.

We identify minimal occurrences with their time intervals in the following way. Given an episode α and an event sequence \mathbf{s} , we say that the interval $[t_s, t_e)$ is a *minimal occurrence* of α in \mathbf{s} , if (1) α occurs in the window $\mathbf{w} = (w, t_s, t_e)$ on \mathbf{s} , and if (2) α does not occur in any proper subwindow on \mathbf{w} , i.e., not in any window $\mathbf{w}' = (w', t'_s, t'_e)$ on \mathbf{s} such that $t_s \leq t'_s$, $t'_e \leq t_e$, and $\text{width}(\mathbf{w}') < \text{width}(\mathbf{w})$. The set of (intervals of) minimal occurrences of an episode α in a given event sequence is denoted by $mo(\alpha)$: $mo(\alpha) = \{ [t_s, t_e) \mid [t_s, t_e) \text{ is a minimal occurrence of } \alpha \}$.

Example 5.1 Consider the episodes in Figure 5.1 (reproduced from Figure 3.2) and the event sequence \mathbf{s} in Figure 5.2 (reproduced from Figure 4.1). The parallel episode β consisting of event types A and B has four minimal occurrences in \mathbf{s} : $mo(\beta) = \{ [35, 38), [46, 48), [47, 58), [57, 60) \}$. The partially ordered episode γ has the following three minimal occurrences: $[35, 39), [46, 51), [57, 62)$. \square

An *episode rule* is an expression $\beta [win_1] \Rightarrow \alpha [win_2]$, where β and α are episodes such that $\beta \preceq \alpha$, and win_1 and win_2 are integers. The interpretation of the rule is that if episode β has a minimal occurrence at interval $[t_s, t_e)$ with $t_e - t_s \leq win_1$, then episode α occurs at interval $[t_s, t'_e)$ for some t'_e such that $t'_e - t_s \leq win_2$. Formally this can be expressed in the following way. Given win_1 and β , denote $mo_{win_1}(\beta) = \{ [t_s, t_e) \in mo(\beta) \mid t_e - t_s \leq win_1 \}$. Further, given α and an interval $[u_s, u_e)$, define $occ(\alpha, [u_s, u_e)) = \text{true}$ if and only if there exists a minimal occurrence $[u'_s, u'_e) \in mo(\alpha)$ such that $u_s \leq u'_s$

and $u'_e \leq u_e$. The confidence of an episode rule $\beta[win_1] \Rightarrow \alpha[win_2]$ is now

$$\frac{|\{[t_s, t_e] \in mo_{win_1}(\beta) \mid occ(\alpha, [t_s, t_s + win_2])\}|}{|mo_{win_1}(\beta)|}.$$

Example 5.2 Continuing Example 5.1, we have, e.g., the following rules and confidences. For the rule $\beta[3] \Rightarrow \gamma[4]$ we have three minimal occurrences $[35, 38), [46, 48), [57, 60)$ of β of width at most 3 in the denominator. Only one of them, $[35, 38)$, has an occurrence of α within width 4, so the confidence is $1/3$. For the rule $\beta[3] \Rightarrow \gamma[5]$ the confidence is 1. \square

Since in a rule $\beta[win_1] \Rightarrow \alpha[win_2]$ episode β is a subepisode of α , the rule right-hand side α contains information about the relative location of each event in the rule. Thus the “new” events in the rule right-hand side can actually be required to be positioned, e.g., between events in the left-hand side. There is also a number of possible definitions for the temporal relationship between the intervals. For instance, rules that point backwards in time can be defined in a similar way. For brevity, we only consider this one case.

In the previous chapter we defined the frequency of an episode as the fraction of windows that contain the episode. While frequency has a nice interpretation as the probability that a randomly chosen window contains the episode, the concept is not very useful with minimal occurrences: (1) there is no fixed window size, and (2) a window may contain several minimal occurrences of an episode. Instead of frequency we use the concept of *support*, the number of minimal occurrences of an episode: the support of an episode α in a given event sequence \mathbf{s} is $|mo(\alpha)|$. Similarly to a frequency threshold, we now use a threshold for the support: given a support threshold min_sup , an episode α is *frequent* if $|mo(\alpha)| \geq min_sup$.

The current episode rule discovery task can now be stated as follows. Given an event sequence \mathbf{s} , a class \mathcal{E} of episodes, and a set W of time bounds, find all frequent episode rules of the form $\beta[win_1] \Rightarrow \alpha[win_2]$, where $\beta, \alpha \in \mathcal{E}$ and $win_1, win_2 \in W$.

5.2 Finding minimal occurrences of episodes

In this section we describe algorithms that locate the minimal occurrences of frequent serial and parallel episodes. Let us start with some observations about the basic properties of episodes. Lemma 4.12 still holds: the subepisodes of a frequent episode are frequent. Thus we can use the main algorithm (Algorithm 4.13) and the candidate generation (Algorithm 4.14) as they are. We also have results about the minimal occurrences of an episode containing minimal occurrences of its subepisodes.

Lemma 5.3 Assume α is an episode and $\beta \preceq \alpha$ is its subepisode. If $[t_s, t_e) \in mo(\alpha)$, then β occurs in $[t_s, t_e)$ and hence there is an interval $[u_s, u_e) \in mo(\beta)$ such that $t_s \leq u_s \leq u_e \leq t_e$.

Lemma 5.4 Let α be a serial episode of size k , and let $[t_s, t_e) \in mo(\alpha)$. Then there are subepisodes α_1 and α_2 of α of size $k - 1$ such that $[t_s, t_e^1) \in mo(\alpha_1)$ for some $t_e^1 < t_e$ and $[t_s^2, t_e) \in mo(\alpha_2)$ for some $t_s^2 > t_s$.

Lemma 5.5 Let α be a parallel episode of size k , and let $[t_s, t_e) \in mo(\alpha)$. Then there are subepisodes α_1 and α_2 of α of size $k - 1$ such that $[t_s^1, t_e^1) \in mo(\alpha_1)$ and $[t_s^2, t_e^2) \in mo(\alpha_2)$ for some $t_s^1, t_e^1, t_s^2, t_e^2 \in [t_s, t_e]$, and furthermore $t_s = \min\{t_s^1, t_s^2\}$ and $t_e = \max\{t_e^1, t_e^2\}$.

The minimal occurrences of a candidate episode α are located in the following way. In the first iteration of the main algorithm, $mo(\alpha)$ is computed from the input sequence for all episodes α of size 1. In the rest of the iterations, the minimal occurrences of a candidate α are located by first selecting two suitable subepisodes α_1 and α_2 of α , and then computing a temporal join between the minimal occurrences of α_1 and α_2 , in the spirit of Lemmas 5.4 and 5.5.

To be more specific, for serial episodes the two subepisodes are selected so that α_1 contains all events except the last one and α_2 in turn contains all except the first one. The minimal occurrences of α are then found with the following specification:

$$mo(\alpha) = \{ [t_s, u_e) \mid \text{there are } [t_s, t_e) \in mo(\alpha_1) \text{ and} \\ [u_s, u_e) \in mo(\alpha_2) \text{ such that } t_s < u_s, \\ t_e < u_e, \text{ and } [t_s, u_e) \text{ is minimal} \}.$$

For parallel episodes, the subepisodes α_1 and α_2 contain all events except one; the omitted events must be different. See Lemma 5.5 for the idea of how to compute the minimal occurrences of α .

The minimal occurrences of a candidate episode α can be found in a linear pass over the minimal occurrences of the selected subepisodes α_1 and α_2 . The time required for one candidate is thus $\mathcal{O}(|mo(\alpha_1)| + |mo(\alpha_2)| + |mo(\alpha)|)$, which is $\mathcal{O}(n)$, where n is the length of the event sequence. To optimize the running time, α_1 and α_2 can be selected so that $|mo(\alpha_1)| + |mo(\alpha_2)|$ is minimized.

The space requirement of the algorithm is $\sum_i \sum_{\alpha \in \mathcal{F}_i} |mo(\alpha)|$, assuming the minimal occurrences of all frequent episodes are stored, or as $\max_i (\sum_{\alpha \in \mathcal{F}_i \cup \mathcal{F}_{i+1}} |mo(\alpha)|)$, if only the current and next levels of minimal occurrences are stored. The size of $\sum_{\alpha \in \mathcal{F}_1} |mo(\alpha)|$ is bounded by n , the number of events in the input sequence, as each event in the sequence is a minimal occurrence of an episode of size 1. In the second iteration, an event

in the input sequence can start at most $|\mathcal{F}_1|$ minimal occurrences of episodes of size 2. The space complexity of the second iteration is thus $\mathcal{O}(|\mathcal{F}_1|n)$.

While minimal occurrences of episodes can be located quite efficiently from minimal occurrences, the size of the data structures can be even larger than the original database, especially in the first couple of iterations. A practical solution is to use in the beginning other pattern matching methods, e.g., similar to the ones described in Section 4.2, to locate the minimal occurrences.

5.3 Finding confidences of rules

We now show how the information about minimal occurrences of frequent episodes can be used to obtain confidences for various types of episode rules without looking at the data again.

Recall that we defined an episode rule as an expression $\beta[win_1] \Rightarrow \alpha[win_2]$, where β and α are episodes such that $\beta \preceq \alpha$, and win_1 and win_2 are integers. To find such rules, first note that for the rule to be frequent, the episode α has to be frequent. So rules of the above form can be enumerated by looking at all frequent episodes α , and then looking at all subepisodes β of α . The evaluation of the confidence of the rule $\beta[win_1] \Rightarrow \alpha[win_2]$ can be done in one pass through the structures $mo(\beta)$ and $mo(\alpha)$, as follows.

For each $[t_s, t_e) \in mo(\beta)$ with $t_e - t_s \leq win_1$, locate the minimal occurrence $[u_s, u_e)$ of α such that $t_s \leq u_s$ and $[u_s, u_e)$ is the first interval in $mo(\alpha)$ with this property. Then check whether $u_e - t_s \leq win_2$.

The time complexity of the confidence computation for a given episode and given time bounds win_1 and win_2 is $\mathcal{O}(|mo(\beta)| + |mo(\alpha)|)$. The confidences for all win_1, win_2 in the set W of time bounds can be found, using a table of size $|W|^2$, in time $\mathcal{O}(|mo(\beta)| + |mo(\alpha)| + |W|^2)$. For reasons of brevity we omit the details.

<i>min_sup</i>	Number of frequent episodes and informative rules							
	Time bounds W (s)							
	15, 30		30, 60		60, 120		15, 30, 60, 120	
50	1131	617	2278	1982	5899	7659	5899	14205
100	418	217	739	642	1676	2191	1676	3969
250	111	57	160	134	289	375	289	611
500	46	21	59	49	80	87	80	138

Table 5.1: Experimental results: number of episodes and rules

<i>min_sup</i>	Execution times (s)			
	Time bounds <i>W</i> (s)			
	15, 30	30, 60	60, 120	15, 30, 60, 120
50	158	210	274	268
100	80	87	103	104
250	56	56	59	58
500	50	51	51	52

Table 5.2: Experimental results: execution times

5.4 Experimental results

We have experimented with the methods using as test data a part of the WWW server log from the Department of Computer Science at the University of Helsinki. The log contains requests to see WWW pages at the department's server; such requests can be made by WWW browsers at any host in the Internet.

An event in the log can be seen as consisting of the attributes *page*, *host*, and *time*. The number of events in our data set is 116308, and it covers three weeks in February and March, 1996. In total, 7634 different pages are referred to from 11635 hosts. Requests for images have been excluded from consideration. For simplicity, we only considered the *page* and *time* attributes; we used relatively short time bounds to reduce the probability of unrelated requests contributing to the minimal occurrences.

We experimented with support thresholds *min_sup* between 50 and 500, and with time bounds between 15 s and 2 min. In three cases we used two time bounds, and in one case we searched simultaneously for all combinations of four time bounds in *W*. Episode rules discovered with these parameters should reveal the paths through which people navigate when they know where they want to go.

Table 5.1 shows the number of frequent episodes and the number of informative rules with confidence at least 0.2. (A rule $\beta[win_1] \Rightarrow \alpha[win_2]$ is considered informative if its confidence is higher than the confidence of all the rules $\beta[win'_1] \Rightarrow \alpha[win'_2]$ with $win'_1 > win_1$, or $win'_1 = win_1$ and $win'_2 < win_2$.)

The number of frequent episodes is in the range from 40 to 6000, and it seems to grow rather fast when the support threshold becomes lower. Our data is relatively dense, and therefore the effect of the time bounds on the number of frequent episodes is roughly linear. The largest frequent episodes consist of 7 events. Note that the method is robust in the sense that a change in one parameter extends or shrinks the collection of frequent episodes but does not replace any.

Table 5.2 shows the execution times for the experiments on a PC (90 MHz Pentium, 32 MB memory, Linux operating system). The data resided in a 3.0 MB flat text file. The execution times are between 50 s and 5 min. Note, in particular, that searching for episodes with several different time bounds (the right-most columns in the tables) is as fast as searching for episodes with only the largest time bound. Minimal occurrences are thus a very suitable representation for queries with different time bounds.

Following are some examples of the episode rules found (we use the titles of the pages here, or their English translations, which should be self-explanatory). All these rules show users going down in the hierarchy or pages.

- “Department Home Page”, “Spring term 96” [15 s] \Rightarrow “Classes in spring 96” [30 s] (confidence 0.83). In other words, in 83 % of the cases where the departmental home page and the spring term page had been accessed within 15 seconds, the classes page was requested within 30 seconds (that is, within 30 seconds from the request for the departmental home page).
- “Research at the department” \Rightarrow “Staff of the department” [2 min] (confidence 0.29). (There is no time bound for the left-hand side since there is only one event.)
- “Department Home Page”, “Department Home Page in Finnish”, “Classes in spring 96”, “Basic courses” [15 s] \Rightarrow “Introduction to Document Preparation (IDP)”, “IDP Course Description”, “IDP Exercises” [2 min] (confidence 0.42).

Experiments with the alarm data set s_1 of Chapter 4 show that—with comparable parameters—the present method is as fast or faster. The new method has, however, two important advantages: the rule formalism is more useful, and rules with several different time bounds can be found with the same effort.

5.5 Bibliographic notes

Most data mining and machine learning techniques are adapted towards the analysis of unordered collections of data. However, there are important application areas where the data to be analyzed has an inherent sequential structure. Examples of such data are telecommunications network alarms, user interface actions, crimes committed by a person, occurrences of recurrent illnesses, etc. Recently, interest in knowledge discovery from sequences of events has increased: see, e.g., [17, 56, 87, 69, 8]. Episodes were introduced in [64], but the first formulations and results on the discovery of parallel episodes were presented already in [63]. Applications of the methods have been described in [38, 39]. A similar approach to the automatic acquisition of

network management knowledge from the existing data has been presented in [31].

Many ideas, for instance the candidate generation method, stem from the discovery of frequent sets and association rules. Various extensions to association rules apply directly or with minor modifications to episodes, too. For instance, these methods can be extended with an event taxonomy by a direct application of the similar extensions to association rules [37, 40, 81]. See Section 2.7 for extensions and work related to association rules.

Technical problems related to the recognition of episodes have been researched in several fields. Taking advantage of the slowly changing contents of the group of recent events has been studied, e.g., in artificial intelligence, where a similar problem in spirit is the many pattern/many object pattern match problem in production system interpreters [24]. Also, comparable strategies using a sliding window have been used, e.g., to study the locality of reference in virtual memory [16]. Our setting differs from these in that our window is a queue with the special property that we know in advance when an event will leave the window; this knowledge is used in the recognition of serial episodes. In the algorithm utilizing minimal occurrences, we take advantage of the fact that we know where subepisodes of candidates have occurred.

The recent work on sequence data in databases (see [78]) provides interesting openings towards the use of database techniques in the processing of queries on sequences. A problem similar to the computation of frequencies occurs also in the area of active databases. There triggers can be specified as composite events, somewhat similar to episodes. In [28] it is shown how finite automata can be constructed from composite events to recognize when a trigger should be fired. This method is not practical for episodes since the deterministic automata could be very large.

The methods for matching sets of episodes against a sequence have some similarities to the algorithms used in string matching (e.g., [32]). In particular, recognizing serial episodes in a sequence can be seen as locating all occurrences of subsequences, or matches of patterns with variable length don't care symbols, where the length of the occurrences is limited by the window width. Learning from a set of sequences has received considerable interest in the field of bioinformatics, where an interesting problem is the discovery of patterns common to a set of related protein or amino acid sequences. The classes of patterns differ from ours; they can be, e.g., substrings with fixed length don't care symbols [46]. Closer to our patterns are those considered in [87]. The described algorithm finds patterns that are similar to serial episodes; however, the patterns have a given minimum length, and the occurrences can be within a given edit distance. Recent results on the pattern matching aspects of recognizing episodes can be found in [12].

The work most closely related to ours is perhaps [4]. There multiple sequences are searched for patterns that are similar to the serial episodes with

some extra restrictions and an event taxonomy. More recently, the pattern class has been extended with windowing, some extra time constraints, and an event taxonomy [83].—For a survey on patterns in sequential data, see [56].

In stochastics, event sequence data is often called a marked point process [48]. It should be noted that traditional methods for analyzing marked point processes are ill suited for the cases where the number of event types is large.

There are also some interesting similarities between the discovery of frequent episodes and the work done on inductive logic programming, ILP (see, e.g., [70]); a noticeable difference is caused by the sequentiality of the underlying data model, and the emphasis on time-limited occurrences. Some steps towards convergence have been published recently [14, 15, 60]. The problem of looking for one occurrence of an episode can, in turn, be viewed as a constraint satisfaction problem.

The class of patterns discovered can be easily modified in several directions. First, the methods can be used to analyze several sequences. If the sequences are short and windowing is not meaningful, simpler database passes are sufficient. If windowing is used there is actually a variety of choices for the definition of frequency of an episode in a set of sequences. Second, other windowing strategies could be used, e.g., considering only windows starting every win' time units for some win' , or windows starting from every event, or for a serial episode with event types A and B , in this order, only windows starting with an event of type A could be taken into account. The use of minimal occurrences is actually an extension similar to this last idea. Third, other patterns could be searched for, e.g., substrings with fixed length don't care symbols.

Perhaps the most important extensions to the framework are facilities for rule querying and compilation, i.e., methods by which the user could specify the episode class in high-level language and the definition would automatically be compiled into a specialization of the algorithm that would take advantage of the restrictions on the episode class. Other open problems include the combination of episode techniques with intensity models.

Chapter 6

The knowledge discovery process

The process of knowledge discovery aims at the discovery of useful and interesting knowledge. In the alarm analysis, the task of finding frequent episodes in the alarm data is only a part of the KDD process, although a central one. In this chapter we discuss the KDD process in the context of this particular application. We also look at TASA, Telecommunication Alarm Sequence Analyzer, a system for discovering knowledge from telecommunication network alarm databases, and see how it supports the process.

6.1 A KDD framework

In the field of data mining or exploratory data analysis the goal is to discover previously unknown information. That is why it can be hard, or even impossible, to specify beforehand, what is interesting. This is particularly true with telecommunication alarms, because the networks are continuously updated. We concentrate here on *subjective* interestingness, and describe one way of trying to discover the most useful and interesting patterns. (Some simple objective interestingness measures were briefly considered in Section 2.5.)

A KDD process, at a coarse level, consists of the following phases:

1. Understanding the domain.
2. Collecting and cleaning data (selection, transformations etc.).
3. Discovery of patterns.
4. Presentation of the results.
5. Interpretation and utilization of the results.

We follow these general steps, with two distinctive characteristics:

1. In the pattern discovery phase, we aim to find *all* potentially interesting (= frequent) patterns according to rather loose criteria for frequency and confidence.
2. In the presentation phase, the discovered patterns can be explored iteratively and interactively.

Our motivation for discovering a lot of rules at once is that network management experts' requests for different viewpoints to the data can then be responded very quickly: a new pattern discovery phase is not necessary, but simply a new view to the already discovered patterns. By producing all rules at once different views on the data can be created very efficiently in the presentation phase. The idea is that the frequency and confidence thresholds filter effectively non-interesting patterns, and produce only potentially interesting ones. The decision of what is interesting is for the most part left to the domain expert to explore.

We next discuss the data preparation, presentation, and utilization phases of the KDD process.

6.2 Collecting and cleaning the data

The first step in the KDD process is collecting and cleaning the data. In some domains this step can take up to 80 % of the total time needed. In our application the information is already collected to the alarm log, and the data is usually of high quality.

Some problems still remain. One is related to the fact that the time stamps of the alarms are not reliable: there can be differences of up to 3–5 minutes in the synchronization of the clocks. As our goal is to locate regularities that are intimately connected to the temporal aspects of alarms, such errors are problematic. Parallel episodes are reasonably robust with respect to changes in the order of the alarms, but a more general solution would be to use preliminary data analysis for locating which components of the network are likely to have erroneous clocks.

Data transformations take a good deal of time: alarms need to be exported from a fault management database, the appropriate attributes have to be selected, and the data needs to be saved in a form suitable for input for TASA. Selection of relevant attributes is not such a trivial task as it might seem at first. Searching for frequent episodes consisting of alarm types is interesting, but there are other alternatives one should consider. For instance, one can derive a new set of event types by concatenating the information about the alarm type and the network element that sent the alarm. The new events can then look something like *1234_EL1*, meaning that element EL1 sent an alarm of type 1234. Frequent episodes now show connections

between alarms from individual elements, e.g., that an alarm 1234 from EL1 is followed by an alarm 3333 from network element EL43. It is useful to try several such variations of the data set, as they give different views to the data and result in partially different findings.

We use the term *alarm predicate* to refer to the (properties of) alarms. Episodes can be discovered in sequences of alarm predicates, where the predicate considered can express the type of the alarm, or some other (combinations) of properties. Association rules can be searched for in unordered collections of alarms, if alarms are seen as rows and alarm predicates as items. For episode rules, the type of the alarm and the sender of the alarm are the most typical predicates. For association rules we consider also predicates such as the priority of the alarm, the day of the week, whether the alarm occurred during office hours or not, etc.

The task of acquiring the required background knowledge is in our case fairly easy. The background knowledge consists mostly of information about the network topology: how the network elements are connected and which subelements they contain. Additionally, the types of the network elements form an inheritance hierarchy that is useful in classifying and presenting the rules. This information is readily available from the telecommunication operators.

6.3 Presentation of results using templates

The presentation of discovered knowledge is a main part of this methodology. In this phase the interesting patterns should be located in large collections of potentially interesting patterns. But what is interesting? How to define it? Many of the patterns discovered in the alarm data are trivial or uninteresting for the network managers:

- **A rule can correspond to prior knowledge or expectations.** For instance, we might know from the network implementation that if an element sends an alarm *A*, it will also send an explanatory notice *B*.
- **A rule can refer to uninteresting attributes or attribute combinations.** If the user is trouble shooting a particular network element, then rules about unrelated parts of the network are probably not useful.
- **Rules can be redundant.** Rules may contain alarms of different abstraction levels but actually referring to the same fault.

For the most part, what is interesting depends on the case, and is highly based on the user's personal aims and perspective. Knowledge trivial to one expert may not be trivial to another, but with proper tools each expert may filter the rule collection based on his personal background knowledge, and to correspond to his current needs.

We now present methods for exploring large sets of association and episode rules. In TASA, the user can manipulate the set of patterns using selection and ranking operations, as well as more complex operations for including or excluding certain classes of rules. TASA supports the following types of operations:

1. *Focusing*: presentation only of a subset of rules, according to the templates the user specifies (see below).
2. *Sorting or ranking* of rules according to simple objective measures.
3. *Clustering*: grouping of rules into clusters of rules that have similar effects in the analyzed data set.

In TASA, rules can be selected or rejected from the view by templates, simple but powerful pattern expressions.

Definition 6.1 We define templates as regular expressions that describe, in terms of alarm predicates, the form of rules that are to be shown or not shown. More formally, a *template* is an expression $A_1, \dots, A_k \Rightarrow A_{k+1}, \dots, A_l$, where each A_i is either an alarm predicate, the name of an alarm predicate collection, or an expression $C+$ or $C*$, where C is a collection name. Here $C+$ and $C*$ correspond to one or more and zero or more instances of the collection C , respectively. A rule $B_1, \dots, B_m \Rightarrow B_{m+1}, \dots, B_n$ matches a template if the rule can be considered to be an instance of the pattern.

With templates, the user can explicitly specify both what is interesting and what is not. To be interesting, a rule has to match a *selective template*. If a rule, however, matches a *rejective template*, it is considered uninteresting. To be presented to the user, a rule must be considered interesting—i.e. match one of the selective templates—and it must not be uninteresting—i.e. not match with any of the unselective templates. \square

Figure 6.1 shows the user interface to templates as implemented in TASA. Templates are specified using separate fields for the rule left-hand and right-hand sides. In addition to templates, additional bounds can be set, e.g., for frequency, confidence, and number of alarms in one rule. We next describe some illustrative scenarios that utilizing these ideas.

Example 6.2 Focus can be set to, e.g., day-time alarms by selecting only association rules that contain the predicate “office hours = yes”. Or, episode rules containing alarms from separate subnetworks can be obtained by using templates that reject all rules where the senders are in the same subnetwork.

The template concept can be combined with thresholds for rule frequency, confidence, and significance (see Section 2.5). The user may state restrictions such as “rule frequency must be between 5 % and 30 %”, “rule confidence must be at least 80 %”, and “rule significance must be over 0.95”. In this

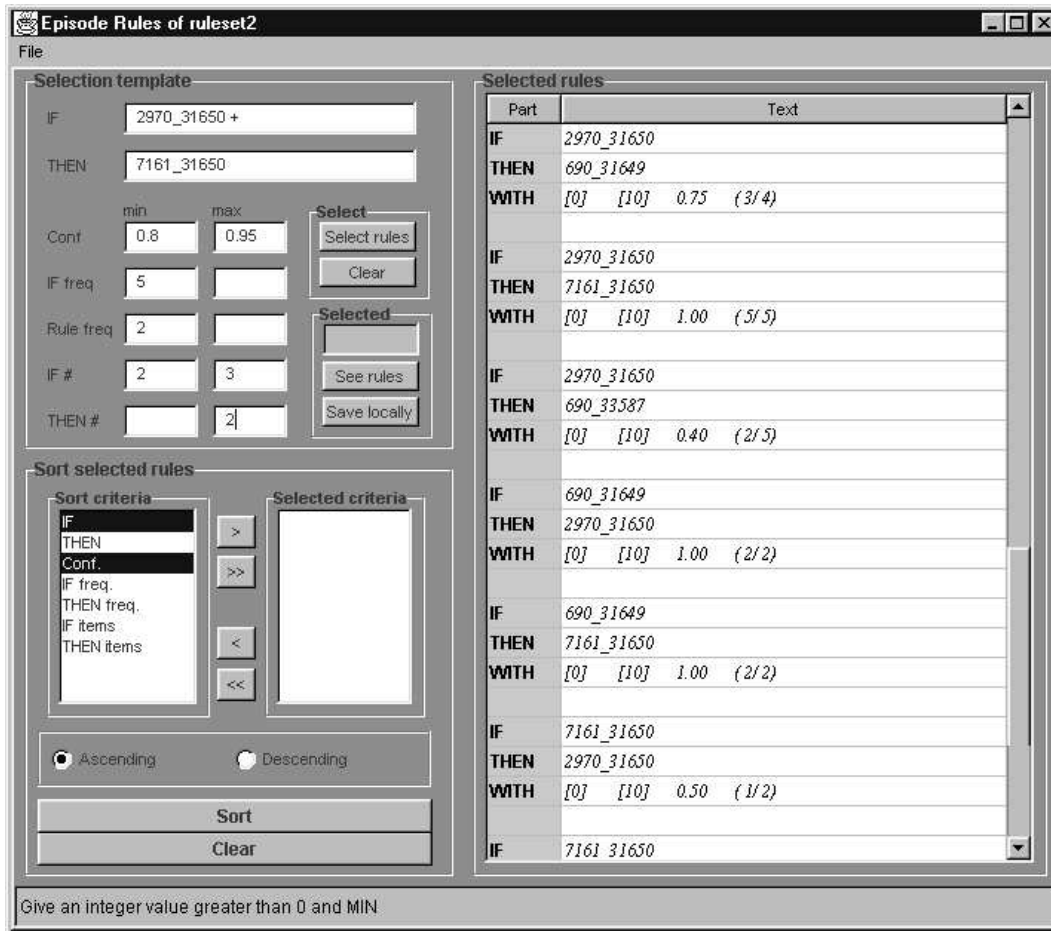


Figure 6.1: Rule Viewing window of TASA: template specification panel on the left, selected rules on the right.

case the user filters out very rare and reasonably frequent rules, and further on selects only those that are both strong and significant. \square

Example 6.3 As an example of how the system can be used for off-line network surveillance, consider the following typical scenario. Assume the network manager has used TASA to discover association rules for the current month. First he might want to see what the alarms have been like during the current week, say week 30, so he uses a template to select rules with the predicate “week = 30” as the left-hand side.

The number of selected rules is still very large. The network manager decides to restrict the rule right-hand side to only contain one predicate, and he also sorts the rules by their confidences.

Looking at the selected rules, he sees the rule “if week = 30 then alarm type = *connection failure*” with confidence 0.12, and he infers that an un-

usually large fraction of alarms during the week has been of type *connection failure*. To see in more detail what the alarms have been like, he refines the template and selects rules with “week = 30 and alarm type = *connection failure*” as the left-hand side.

Looking at the new set of selected rules, the network manager sees that a lot of rules concern the network element *EL1*. That reminds him of maintenance undertaken in the beginning of the week that explains those rules. To remove the rules, he applies a rejective template with the predicate “network element = *EL1*”.

The resulting set of rules shows nothing special, but just to make sure the network manager wants to compare the rules with the corresponding rules from some previous week. He opens a copy of the window, and changes the first template to “week = 29”. If there is anything special or interesting, the viewing criteria can be refined or altered again. □

Clustering of rules aims at giving a larger picture of the behavior of the alarm sequence. In the data there are often various explanations for the occurrence of a particular alarm type, say *path unavailable*. Clustering methods can be used to assign rules to groups so that two rules with the right-hand side *path unavailable* belong to the same cluster if they often explain or predict *path unavailable* in similar situations. This can be useful in pointing out potentially related rules.

In addition to looking and manipulating the discovered knowledge, the users want to be able to use several types of views into the data. They want to see the discovered knowledge, but they also want to be able to see how that knowledge is actually supported by the original data. TASA links rules, alarms, and data together by hyperlinks.

Visualization of information is obviously an important part of KDD applications. For this the TASA system offers some simple facilities. As an example, consider the confidence of a rule. It is only a crude measure of how well the rule manages to predict the occurrence of the right-hand side, and a more complete picture of the interaction between the left and right-hand sides of a rule can be obtained by simply drawing a histogram showing the distance from each occurrence of the left-hand side to the nearest occurrence of the right-hand side. Such histograms are valuable guides for locating possible periodic relationships between the left and right-hand sides, as is demonstrated by Figure 6.2.

6.4 Experiences with TASA, episodes, and associations

Different versions of TASA have been in prototype use in four telecommunication companies since the beginning of 1995. TASA has been found useful

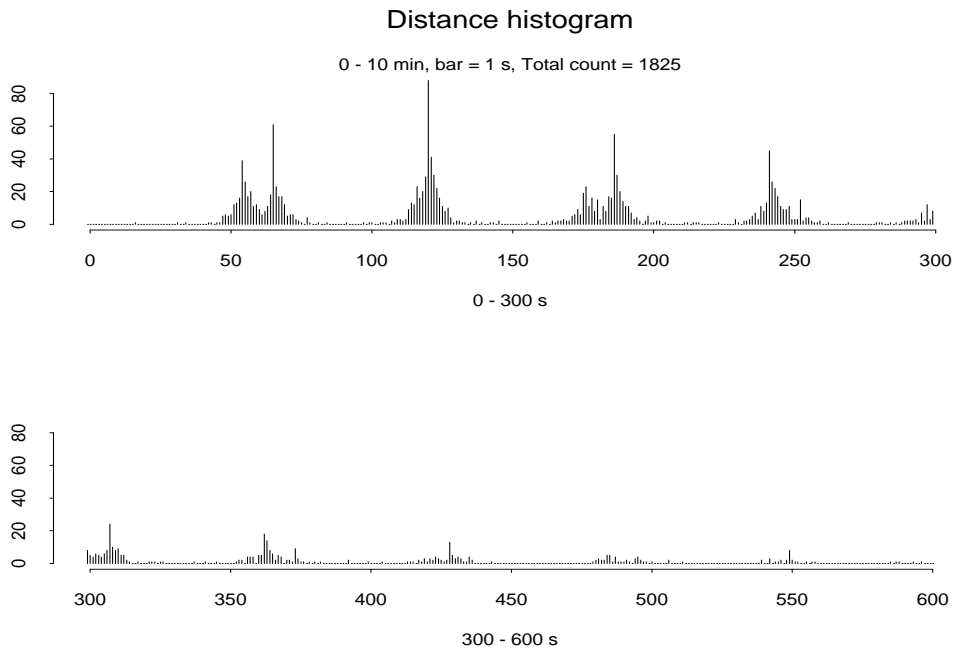


Figure 6.2: Distance histogram

in, e.g., finding long-term, rather frequently occurring dependencies, creating an overview of a short-term alarm sequence, and evaluating the alarm data base consistency and correctness.

Unexpected dependencies have been found, e.g., between network elements which are not closely connected in the network topology. An example of such a dependency is that when a remote device sends alarms, the fault is reflected to another corner of the network through several devices, and not always necessarily via the same routes and devices. So, just analyzing the neighboring devices might not reveal any strong relationships. However, when a larger region is analyzed, such a relationship can be detected. Beginning from the first tests, discovered rules have been integrated into alarm correlation systems.

On the other hand, many of the rules discovered by TASA are deemed trivial by the network managers. Some of the rules correspond to the knowledge that the network managers have about the behavior of the network, and some other rules reflect the assumed functioning of network devices. Luckily, much of the trivial knowledge can be expressed and removed with templates. Templates are also useful since the knowledge trivial to one expert may not be trivial to another, and with templates each expert may filter the rule collection based on his/her personal background knowledge.

The usability of discovery tools has an essential, often perhaps under-

estimated role. The usability of an early version of TASA was tested in the usability laboratory of the Helsinki University of Technology. The tests contained, e.g., user tests taken by four fault management experts from telecommunication companies. In the tests, TASA was generally acknowledged as appealing. On the other hand, first-time users were unfamiliar with many concepts from the knowledge discovery field. Despite these problems with the terminology, the system as a whole got encouraging comments.

Overall, TASA has been considered useful. Episode rules are being used as first drafts of correlation rules, whereas association rules are more typically used for creating short-term overviews in off-line network surveillance. Telecommunication operators are integrating these methods to their alarm analysis and surveillance systems.

6.5 Bibliographic notes

The problem of locating a small set of truly interesting information is a generic problem in data mining (see, e.g., [74]): it is hard to know which aspects of the discovered knowledge really interest the user. Kloesgen [53], for instance, defines a number of criteria for interestingness, following the ones presented in [25]: evidence (statistical significance), non-redundancy, novelty (deviation from prior knowledge), simplicity (syntactical complexity of a finding), and generality (the fraction of the population the finding refers to).

The idea of discovering a large number of potentially interesting patterns can be contrasted with numerous methods, e.g., in machine learning, which are more focused and produce one or at most few patterns that match the given problem specification. These methods usually require that the searched or learned subject is quite carefully described in advance, and they leave any other potentially interesting phenomena hidden. The advantage of these systems is that the patterns they find are more expressive than the relatively simple association and episode rules, and focusing the pattern discovery is thus more important.

TASA uses association and episode rules, but the basic idea—iteration in the pattern presentation phase—can be applied to any formalisms that have some similar properties as association and episode rules:

- There is an algorithm that produces lots of potentially interesting patterns.
- The time requirement for discovering all potentially interesting patterns is not considerably longer than if the discovery was focused to a small subset of the potentially interesting patterns.
- The desired focus is not known definitely in advance.

This approach which allows the user to set the focus, and where the user has a total, explicit control over the resulting rule set, is in many ways similar to the declarative language bias used by some ILP (Inductive Logic Programming) systems, e.g., Claudien [76]. In ILP systems, however, the focus (bias) is usually set before the discovery.

Templates were introduced in [52]. In the Explora system, Hoschka and Kloesgen [41] already have used patterns similar to templates for defining what is interesting, and their ideas have influenced the work on templates. Their approach is based on few fixed statement types and a partial ordering of attributes, whereas templates are closer to regular expressions.

Sometimes considerable amounts of rules remain, even when the user has found the desired focus with the described methods. Automatic pruning, ordering, and structuring methods should at this point be available for invocation by the user, especially for removal of redundancy.

A method for pruning a set of strong association or episode rules by removing redundancy has been presented in [85]. The method is based on computing a *rule cover*. A rule cover is a subset of rules that has predictive power equal to the original set of rules, in the sense that whenever a rule of the original set matches a row, then there is a rule in the cover that matches the row.

An approximate order of interestingness could be assigned to the discovered rules by giving weights to templates: positive weights to selective templates and negative weights to rejective templates. The ranking of rules would then correspond to the sum of weights of matched templates.

Clustering has been proposed as a method for structuring a set of association or episode rules [85]. The idea is to take rules with the same right-hand side, and to group the rules so that rules in the same cluster tend to match the same cases in the database. Clustering of rules has a remote connection to clustering of database rows.

The approaches taken by other KDD systems include, e.g., the following. Explora [41, 53] finds interesting instances of statistical patterns. In Explora, the pattern discovery phase is focused by the user. The system selects and presents the best patterns to the user, and, based on the results, the user can change the focus and repeat the pattern discovery. The patterns discovered by 49er [89] are contingency tables, equations, and logical equivalences. The user can interactively change the focus, e.g., independent and dependent variables, and require for a new pattern discovery. The Key Finding Reporter (Kefir) [65, 75] discovers and explains deviations, and gives recommendations for corrective actions. Applications of Kefir are tailored with a lot of domain knowledge to be aware of the interestingness criteria, corrective actions, etc., of the domain. Given a database from the domain, a Kefir-based application produces a report of the deviations without iteration.

Chapter 7

Discovery of all frequent patterns

In this part we consider a generalization of the problems of discovering all frequent sets or all frequent episodes. The problem we look at is the following: given a set of patterns, a selection criterion, and a database, find those patterns that satisfy the criterion in the database. In this chapter we present a framework for this problem and give a generic levelwise algorithm for solving it, e.g., in the case where the selection criterion is based on the frequency of patterns. We then analyze the algorithm and the problem in Chapter 8.

We start this chapter by describing the discovery task in Section 7.1. We give an algorithm for this task in Section 7.2. In Section 7.3 we give examples of the setting in various knowledge discovery tasks, and we show that, instead of frequency, other criteria can be used for selecting rules. Finally, in Section 7.4, we outline an extension of the setting and the algorithm for discovery in several database states.

7.1 The discovery task

We start by defining the knowledge discovery setting we consider in this chapter. Given a set of patterns, i.e., a class of expressions about databases, and a predicate to evaluate whether a database satisfies a pattern, the task is to determine which patterns are satisfied by a given database.

Definition 7.1 Assume that \mathcal{P} is a set and q is a predicate $q : \mathcal{P} \times \{\mathbf{r} \mid \mathbf{r} \text{ is a database}\} \rightarrow \{\text{true}, \text{false}\}$. Elements of \mathcal{P} are called *patterns* and q is a *selection criterion* over \mathcal{P} . Given a pattern φ in \mathcal{P} and a database \mathbf{r} , we say that φ is *selected* if $q(\varphi, \mathbf{r})$ is true. Since the selection criterion is often based on the frequency of the pattern, we use the term *frequent* as a

synonym for “selected”. Given a database \mathbf{r} , the *theory* $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ of \mathbf{r} with respect to \mathcal{P} and q is $\mathcal{T}(\mathcal{P}, \mathbf{r}, q) = \{\varphi \in \mathcal{P} \mid q(\varphi, \mathbf{r}) \text{ is true}\}$. \square

Example 7.2 The problem of finding all frequent item sets can be described as a task of discovering frequent patterns in a straightforward way. Given a set R , a binary database r over R , and a frequency threshold min_fr , the set \mathcal{P} of patterns consists of all item sets, i.e., $\mathcal{P} = \{X \mid X \subseteq R\}$, and for the selection criterion we have $q(\varphi, r) = \text{true}$ if and only if $fr(\varphi, r) \geq min_fr$. \square

Note that we do not specify any satisfaction relation for the patterns of \mathcal{P} in \mathbf{r} : this task is taken care of by the selection criterion q . For some applications, “ $q(\varphi, \mathbf{r})$ is true” could mean that φ occurs often enough in \mathbf{r} , that φ is true or almost true in \mathbf{r} , or that φ defines, in some way, an interesting property or subgroup of \mathbf{r} . Obviously, the task of determining the theory of \mathbf{r} is not tractable for arbitrary sets \mathcal{P} and predicates q . If, for instance, \mathcal{P} is infinite and $q(\varphi, \mathbf{r})$ is true for infinitely many patterns, an explicit representation of $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ cannot be computed.

In the discovery tasks considered here the aim is to find all patterns that are selected by a relatively simple criterion—such as exceeding a frequency threshold—in order to efficiently identify a space of potentially interesting patterns. Other criteria can then be used for further pruning and processing of the patterns. Consider as an example the discovery of association rules: first frequent sets are discovered, then all rules with sufficient frequency are generated, and a confidence threshold is used to further prune the rules.

The task of discovering frequent sets has two noteworthy properties. First, all frequent sets are needed for the generation of association rules. It is not sufficient to know just the largest frequent sets, although they determine the collection of all frequent sets. The second important property is that the selection criterion, i.e., frequency, is monotone decreasing with respect to expansion of the set. We consider only the situation where the predicate q is monotone with respect to a given partial order on the patterns.

Definition 7.3 Let \mathcal{P} be a set of patterns, q a selection criterion over \mathcal{P} , and \preceq a partial order on the patterns in \mathcal{P} . If for all databases \mathbf{r} and patterns $\varphi, \theta \in \mathcal{P}$ we have that $q(\varphi, \mathbf{r})$ and $\theta \preceq \varphi$ imply $q(\theta, \mathbf{r})$, then \preceq is a *specialization relation* on \mathcal{P} with respect to q . If we have $\theta \preceq \varphi$, then φ is said to be *more special* than θ and θ to be *more general* than φ . If $\theta \preceq \varphi$ and not $\varphi \preceq \theta$ we write $\theta \prec \varphi$. \square

Example 7.4 The set inclusion relation \subseteq is a specialization relation for frequent sets. For instance, if the set $\{A, B, C\}$ is frequent, then its subset $\{A, C\}$ must also be frequent.

In more general, given two item sets $X, Y \subseteq R$, the set X is more general, $X \preceq Y$, if and only if $X \subseteq Y$. That is, $X \preceq Y$ implies that if the more specific set Y is frequent then the more general set X is frequent, too. \square

For practical purposes the specialization relation has to be computable, i.e., given patterns φ and θ in \mathcal{P} , it must be possible to determine whether $\varphi \preceq \theta$. Typically, the specialization relation \preceq is a restriction of the converse of the semantic implication relation: if $\theta \preceq \varphi$, then φ implies θ . If the predicate q is defined in terms of, e.g., statistical significance, then the semantic implication relation is not a specialization relation with respect to q : a pattern can be statistically significant even when a more general pattern is not. Recall that the predicate q is not meant to be the only way of identifying the interesting patterns; a threshold for the statistical significance can be used to further prune patterns found using q .

7.2 The generic levelwise algorithm

In this section we present an algorithm for the task of discovering all frequent patterns in the case where there exists a computable specialization relation between patterns. We use the following notation for the relative speciality of patterns.

Definition 7.5 Given a specialization relation \preceq on patterns in \mathcal{P} , the *level* of a pattern φ in \mathcal{P} , denoted $level(\varphi)$, is 1 if there is no θ in \mathcal{P} for which $\theta \prec \varphi$. Otherwise $level(\varphi)$ is $1 + L$, where L is the maximum level of patterns θ in \mathcal{P} for which $\theta \prec \varphi$. The collection of frequent patterns of level l is denoted by $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q) = \{\varphi \in \mathcal{T}(\mathcal{P}, \mathbf{r}, q) \mid level(\varphi) = l\}$. \square

Algorithm 7.6, analogical to Algorithm 2.14, finds all frequent patterns. It works in a levelwise or breadth-first manner, starting with the set \mathcal{C}_1 of the most general patterns, and then generating and evaluating more and more special candidate patterns. The algorithm prunes those patterns that cannot be frequent given all the frequent patterns obtained in earlier iterations.

The algorithm is generic: details depending on the specific types of patterns and data are left open, and instances of the algorithm must specify these. The levelwise algorithm aims at minimizing the number of evaluations of q on line 5. As with the frequent set discovery algorithm, the computation to determine the candidate collection does not involve the database at all.

Theorem 7.7 Algorithm 7.6 works correctly.

Proof We show by induction on l that $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q)$ is computed correctly for all l . For $l = 1$, the collection \mathcal{C}_1 contains all patterns of level one (line 1), and collection $\mathcal{T}_1(\mathcal{P}, \mathbf{r}, q)$ is then correctly computed (line 5).

For $l > 1$, assume the collections $\mathcal{T}_i(\mathcal{P}, \mathbf{r}, q)$ have been computed correctly for all $i < l$. Note first that $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q) \subseteq \mathcal{C}_l$. Namely, consider any pattern φ in $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q)$: we have $level(\varphi) = l$ and thus for all patterns $\theta \prec \varphi$ we have $level(\theta) < l$. Since $\mathcal{T}_i(\mathcal{P}, \mathbf{r}, q)$ has been computed for each $i < l$, each $\theta \prec \varphi$

Algorithm 7.6

Input: A database schema \mathbf{R} , a database \mathbf{r} over \mathbf{R} , a finite set \mathcal{P} of patterns, a computable selection criterion q over \mathcal{P} , and a computable specialization relation \preceq on \mathcal{P} .

Output: The set $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ of all frequent patterns.

Method:

1. compute $\mathcal{C}_1 := \{\varphi \in \mathcal{P} \mid level(\varphi) = 1\}$;
 2. $l := 1$;
 3. **while** $\mathcal{C}_l \neq \emptyset$ **do**
 4. // Database pass:
 5. compute $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q) := \{\varphi \in \mathcal{C}_l \mid q(\varphi, \mathbf{r})\}$;
 6. $l := l + 1$;
 7. // Candidate generation:
 8. compute $\mathcal{C}_l := \{\varphi \in \mathcal{P} \mid level(\varphi) = l \text{ and } \theta \in \mathcal{T}_{level(\theta)}(\mathcal{P}, \mathbf{r}, q) \text{ for all } \theta \in \mathcal{P} \text{ such that } \theta \prec \varphi\}$;
 9. **for all** l **do** output $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q)$;
-

is correctly in $\mathcal{T}_{level(\theta)}(\mathcal{P}, \mathbf{r}, q)$, and so φ is put into \mathcal{C}_l (line 8). The collection $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q)$ is then computed correctly on line 5.

Finally note that for every $\varphi \in \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ there is an iteration where the variable l has value $level(\varphi)$. By the definition of level, there are more general patterns $\theta \prec \varphi$ on every level less than $level(\varphi)$, and since \preceq is a specialization relation they are all frequent, so the ending condition $\mathcal{C}_l = \emptyset$ is not true with $l \leq level(\varphi)$. \square

The input specification of the algorithm states that the set \mathcal{P} is finite. Actually, it does not always need to be finite: the algorithm works correctly as long as the number of candidate patterns is finite. There are some desirable properties for the specialization relation \preceq . An efficient method for accessing the more specific and more general patterns on neighboring levels is useful, or otherwise finding the collection of valid candidates may be expensive.

7.3 Examples

We now look at the applicability of Algorithm 7.6 for some knowledge discovery tasks. We consider three problems as tasks of discovering patterns that are selected by a given predicate: the discovery of association rules and frequent episodes, and the discovery of exact database rules.

7.3.1 Association rules

Recall that the task of discovering association rules can be split into two: the discovery of frequent sets, and the generation of rules. We summarize below the necessary specifications for discovering frequent sets with Algorithm 7.6.

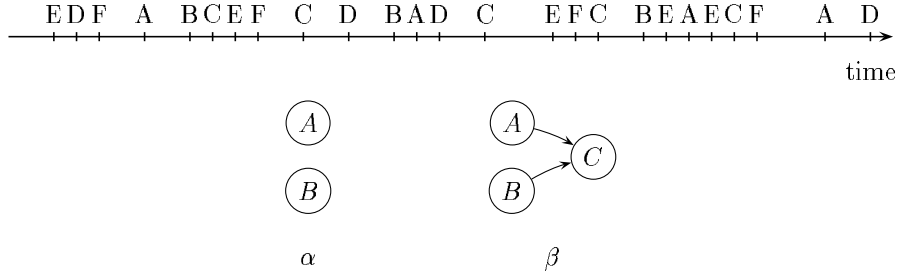


Figure 7.1: An event sequence and two episodes.

Given are a set R , a binary database r over R , and a frequency threshold min_fr . We want to find the frequent sets, so the set \mathcal{P} of patterns consists of all item sets: $\mathcal{P} = \{X \mid X \subseteq R\}$. Our selection criterion q is based on the frequency of an item set, so that q is true for a set $X \in \mathcal{P}$ if and only if $\text{fr}(X, r) \geq \text{min_fr}$. Obviously the set inclusion relation \subseteq is monotone with respect to frequency, so it can be used as a specialization relation \preceq .

With these specifications Algorithm 7.6 proceeds as Algorithm 2.14. Note in particular that candidate generation on line 8 of Algorithm 7.6 gives a result equivalent to the candidate collection of Definition 2.13.

7.3.2 Episodes

Consider the problem of recognizing frequent episodes in sequences of events. Figure 7.1 depicts a sequence of events and two episodes that occur several times in the sequence. Episode α contains two events, A and B , but does not specify any order for them. Episode β contains additionally an event C , and states that C must occur after both A and B .

Recall from Chapter 4 that the task is to discover all episodes whose frequency exceeds a given threshold min_fr . Given a window width win , the frequency of α in a given event sequence \mathbf{s} is defined as the fraction of windows of length win on \mathbf{s} that contain an instance of α .

The discovery of all frequent episodes can be solved with the levelwise algorithm. The subepisode relation \preceq is a specialization relation on episodes. That is, for $\alpha = (V, \leq, g)$ and $\beta = (V', \leq', g')$ we have $\beta \preceq \alpha$, if and only if (1) $V' \subseteq V$ (after possible renaming of nodes), (2) for all $v \in V'$ we have $g'(v) = g(v)$, and (3) for all $v, w \in V'$ with $v \leq' w$ also $v \leq w$. The relation $\alpha \preceq \beta$ holds for the episodes in Figure 7.1. In the case of arbitrary partial orders, Algorithm 7.6 is more conservative than Algorithm 4.13. Algorithm 7.6 considers a totally ordered episode consisting of events A and B as candidate only, if the trivially ordered episode consisting of A and B is frequent. Algorithm 4.13, in turn, was specified to increase the size of episodes by one in each iteration.

7.3.3 Exact database rules

Exact database rules [74] are a rule formalism that is somewhat more general than association rules: numerical and categorical attributes are considered. On the other hand, the confidence of exact rules must be 1; a small variation gives *strong rules* which can have a confidence less than 1. Before introducing exact database rules formally, we define the notion of a taxonomy on an attribute.

Definition 7.8 Given an attribute A , a *taxonomy* on A is a set $T(A)$ such that $Dom(A) \subseteq T(A)$, where $Dom(A)$ is the domain of A , and such that there is a partial order *is-a* on $T(A)$. We assume that *is-a* is reflexive and that there is a special member *any* in the taxonomy such that for all $a \in T(A)$ we have a *is-a* *any*. \square

Example 7.9 Consider an attribute *department* for which the domain $Dom(\textit{department})$ is $\{\textit{dept}_1, \textit{dept}_2, \dots, \textit{dept}_{15}\}$. Now a taxonomy $T(\textit{department})$ could consist of $Dom(\textit{department}) \cup \{\textit{management_dept}, \textit{production_dept}, \textit{sales_dept}, \textit{any}\}$, i.e., of names of departments and their types. The partial order *is-a* could then classify each department to its type by defining that \textit{dept}_1 *is-a* $\textit{management_dept}$, \textit{dept}_2 *is-a* $\textit{management_dept}$, that for $i = 3, \dots, 11$ \textit{dept}_i *is-a* $\textit{production_dept}$, and that for $i = 12, \dots, 15$ \textit{dept}_i *is-a* $\textit{sales_dept}$. Additionally, for every \textit{dept} in $T(\textit{department})$ we have \textit{dept} *is-a* \textit{dept} and \textit{dept} *is-a* *any*. \square

Definition 7.10 Let r be a relation over a relation schema R . Assume taxonomies are given for the non-numerical attributes in R . A *simple condition* on a row t in r is either of the form $a_1 \leq t[A] \leq a_2$, where $A \in R$ is a numerical attribute and $a_1, a_2 \in Dom(A)$, or of the form $t[A]$ *is-a* a , where $A \in R$ is non-numerical and has a taxonomy $T(A)$, and $a \in T(A)$.

An *exact database rule* is now an expression of the form $C_1 \Rightarrow C_2$, where both C_1 and C_2 are simple conditions. The rule $C_1 \Rightarrow C_2$ *holds* in r if C_2 is true on every row t of r that C_1 is true on. \square

Example 7.11 Assume a relation r represents employees. Consider only attributes *department*, as in the previous example, and *age*, and assume that the domain of *age* is $\{18, 19, \dots, 65\}$. The exact database rule

$$t[\textit{department}] \textit{ is-a } \textit{management_dept} \Rightarrow 43 \leq t[\textit{age}] \leq 65$$

tells that employees in the management departments are at least 43 years old. \square

Given a relation r and taxonomies for non-numerical attributes, the collection of exact database rules that hold in r is the theory $\mathcal{T}(\mathcal{P}, r, q)$, where the set \mathcal{P} of patterns consists of all possible exact rules φ , and the predicate

$q(\varphi, r)$ is true if and only if φ holds in r . Next we show what is a specialization relation for exact rules.

Theorem 7.12 The following relation \preceq is a specialization relation with respect to the set \mathcal{P} of all possible exact database rules:

$$(C_1 \Rightarrow C_2) \preceq (C'_1 \Rightarrow C'_2) \text{ if and only if } C'_1 \sqsubseteq C_1 \text{ and } C_2 \sqsubseteq C'_2,$$

where \sqsubseteq is a partial order on simple conditions defined as follows:

$$(a_1 \leq t[A] \leq a_2) \sqsubseteq (b_1 \leq t[B] \leq b_2) \text{ if and only if} \\ A = B \text{ and } [b_1, b_2] \subseteq [a_1, a_2]$$

and

$$(t[A] \text{ is-a } a) \sqsubseteq (t[B] \text{ is-a } b) \text{ if and only if } A = B \text{ and } b \text{ is-a } a.$$

Proof Denote by $\mathcal{M}(C)$ the set of rows on which condition C is true. By the definition, the relation \sqsubseteq on simple conditions has the following property: if C_1 is true on a row t , then every $C_2 \sqsubseteq C_1$ is true on t , i.e., $\mathcal{M}(C_1) \subseteq \mathcal{M}(C_2)$, and \sqsubseteq is actually a specialization relation on simple conditions.

Assume the exact database rule $C'_1 \Rightarrow C'_2$ holds, i.e., $\mathcal{M}(C'_1) \subseteq \mathcal{M}(C'_2)$. Consider now any more general rule $(C_1 \Rightarrow C_2) \preceq (C'_1 \Rightarrow C'_2)$. From the properties of \sqsubseteq it follows that $\mathcal{M}(C_1) \subseteq \mathcal{M}(C'_1)$ and $\mathcal{M}(C'_2) \subseteq \mathcal{M}(C_2)$. Thus $\mathcal{M}(C_1) \subseteq \mathcal{M}(C_2)$, i.e., the rule $(C_1 \Rightarrow C_2)$ holds. \square

The proof shows that the specialization relation \preceq is a restriction of the converse of the semantic implication: for any two patterns φ and θ , if we have $\varphi \preceq \theta$ then θ implies φ . Intuitively, the specialization relation means here that once we have an exact database rule that holds, we know that a modified rule where the left-hand side only matches a subset of rows must hold as well, and that if the right-hand side matches a superset of rows, the modified rule must also hold.

Algorithm 7.6 would start with those rules that are most likely to hold, and then loosen the conditions on the left-hand sides while tightening the conditions on the right-hand sides.

Example 7.13 Assume the relation r represents employees. Consider only attributes *department*, as in the previous example, and *age*, and assume that the domain of *age* is $\{18, 19, \dots, 65\}$.

The most general patterns considered by Algorithm 7.6 are such as

$$39 \leq t[age] \leq 39 \Rightarrow t[department] \text{ is-a any}$$

and

$$t[department] \text{ is-a dept.7} \Rightarrow 18 \leq t[age] \leq 65.$$

These and a number of other obvious rules hold. Later, when more meaningful rules are dealt with, the specialization relation prunes rules from consideration in the following way. If, for instance, the rule

$$t[\textit{department}] \textit{ is-a dept.2} \Rightarrow 18 \leq t[\textit{age}] \leq 40$$

does not hold, then rules such as

$$t[\textit{department}] \textit{ is-a dept.2} \Rightarrow 18 \leq t[\textit{age}] \leq 39$$

and

$$t[\textit{department}] \textit{ is-a management_dept} \Rightarrow 18 \leq t[\textit{age}] \leq 40$$

cannot hold. □

Note that the task of discovering exact database rules cannot be split into two phases like the discovery of association rules, where frequent sets, i.e., the rule components, are discovered first. Namely, in the case of exact rules there is no minimum threshold for the frequency of rules.

When strong, i.e., almost always correct rules are searched for, the “almost always correctness” needs to be carefully defined, or the partial order \preceq given above is not a specialization relation. The following example demonstrates this.

Example 7.14 Consider the discovery of strong rules, and the use of a confidence threshold *min_conf*, defined as with association rules, as a means for determining whether a rule is strong or not. If the rule

$$t[\textit{department}] \textit{ is-a dept.7} \Rightarrow 40 \leq t[\textit{age}] \leq 50$$

has a confidence close to but below the threshold *min_conf*, then the rule

$$t[\textit{department}] \textit{ is-a any} \Rightarrow 40 \leq t[\textit{age}] \leq 50$$

might actually be strong, e.g., if all employees in other departments than *dept.7* are between 40 and 50 years old. □

Algorithm 7.6 considers in each database pass a collection of candidate rules, where all the rules are on the same level. The KID3 algorithm [74] for discovering exact database rules, in turn, considers in one iteration all rules with the same attribute on the left-hand side. KID3 does not directly evaluate q on all those rules; instead, it stores some summary information from which rules that hold can be extracted. Both approaches have their drawbacks. The space requirement of the summaries in KID3 is in the worst case linear in the database size. Algorithm 7.6, in turn, does not take any advantage of the fact that rules close to each other in the specialization relation are similar, and could be evaluated efficiently together.

Almost always lots of redundant exact and strong rules hold. For exact rules, for instance, giving the most specific rules that hold would be sufficient, since the rest of the rules are implied by these. Recall, again, that the purpose is to find the rules that hold, and then use other methods to select the most useful ones—the specificity is certainly one criterion, but not the only one [74].

7.4 Discovery in several database states

Knowledge discovery is sometimes criticized for analyzing just one database state. The critics say that it does not give reliable information: it is often impossible to know if a regularity exists in the analyzed database only by chance, or if it is true in most database states. Next we describe how Algorithm 7.6 can be adopted to discover those patterns that are selected by the given criterion in most of the given database states. We define the global selection criterion of a pattern φ to depend on the number of database states where φ is selected.

Definition 7.15 Given a selection criterion q over a set \mathcal{P} of patterns and a *frequency threshold* min_fr , the *global selection criterion* Q is a predicate

$$Q : \mathcal{P} \times \{\mathbf{r} \mid \mathbf{r} \text{ is a set of databases}\} \rightarrow \{\text{true, false}\},$$

such that for any set $\mathbf{r} = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$ of databases we have $Q(\varphi, \mathbf{r}) = \text{true}$ if and only if $|\{i \mid q(\varphi, \mathbf{r}_i)\}| \geq min_fr \cdot n$. The theory $\mathcal{T}(\mathcal{P}, \mathbf{r}, Q)$ is called the *almost always true theory* of \mathbf{r} with respect to \mathcal{P} , q , and min_fr . \square

Note that any partial order \preceq on \mathcal{P} that is a specialization relation with respect to q is also a specialization relation with respect to Q . We have the following theorem.

Theorem 7.16 Let \mathcal{P} , q , and Q be as in Definition 7.15, and let \preceq be a specialization relation with respect to q . Then \preceq is a specialization relation also with respect to the global selection criterion Q .

Proof By definition the relation \preceq is a specialization relation with respect to Q , if for all sets \mathbf{r} of databases and all $\varphi, \theta \in \mathcal{P}$ we have that $Q(\varphi, \mathbf{r})$ and $\theta \preceq \varphi$ imply $Q(\theta, \mathbf{r})$. To see that this is the case, consider a pattern φ for which $Q(\varphi, \mathbf{r})$ holds. For each \mathbf{r}_i in \mathbf{r} for which $q(\varphi, \mathbf{r}_i)$ holds, $q(\theta, \mathbf{r}_i)$ must hold for all more general patterns $\theta \preceq \varphi$, and thus $Q(\theta, \mathbf{r})$ must hold. \square

Since \preceq is a specialization relation with respect to the global selection criterion Q , Algorithm 7.6 can be applied directly for knowledge discovery from several database states; just use the global selection criterion Q instead of q . The evaluation of Q on \mathbf{r} consists now of evaluating q on the individual

database states $\mathbf{r}_i \in \mathbf{r}$. It turns out that we can use here the specialization relation \preceq both locally and globally.

Consider Algorithm 7.6 running with a set $\mathbf{r} = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$ of database states and the global selection criterion Q as its inputs. Candidate patterns φ in the algorithm will be such that all more general patterns than φ are globally selected. Such patterns are then evaluated in each database state, in order to find out if they are selected in by q in sufficiently many database states. However, it is possible that in some database state $\mathbf{r}_i \in \mathbf{r}$ patterns more general than φ are not selected by q , and correspondingly φ cannot be selected by q in \mathbf{r}_i . A key to a more efficient evaluation of the global selection criterion is thus to generate candidates also locally in each database state, and only to evaluate patterns that are candidates both globally and locally.

To be more specific, at level l the global candidate collection \mathcal{C}_l contains those patterns that are potentially selected by Q in \mathbf{r} . The local candidate collections, denoted by \mathcal{C}_l^i , contain for each database state \mathbf{r}_i those patterns that are potentially selected by q in \mathbf{r}_i . During the evaluation of Q , for each database state \mathbf{r}_i we evaluate the predicate q on the intersection $\mathcal{C}_l \cap \mathcal{C}_l^i$ of global and local candidates.

By using information about the local candidates we can further eliminate evaluations of q . Namely, the global candidate collection \mathcal{C}_l may contain such patterns that are not candidates in sufficiently many collections \mathcal{C}_l^i . This can be the situation for a pattern $\varphi \in \mathcal{C}_l$ when the more general patterns $\theta \prec \varphi$ are selected too often in disjoint database states. Such useless evaluation of candidates can be avoided by a simple check: a pattern $\varphi \in \mathcal{C}_l$ needs not to be evaluated if $|\{i \mid \mathcal{C}_l^i\}| < \text{min_fr} \cdot n$. A similar check can be applied after each failed evaluation of $q(\varphi, \mathbf{r}_i)$, in order to prune a candidate as soon as it turns out that it cannot be globally selected.

In summary, using only information about the global selection criterion of patterns we would at level l investigate the patterns in \mathcal{C}_l against each database state \mathbf{r}_i . Looking at each database state \mathbf{r}_i locally would enable us to investigate the patterns in \mathcal{C}_l^i . Combining local and global information, we see that one has to investigate at most the patterns in $\mathcal{C}_l \cap \mathcal{C}_l^i$.

This method could be used to analyze, e.g., the same database over time, in order to see what regularities hold in most of the database states, or to analyze several similar databases, for instance to find out which association rules hold in most of the stores of a supermarket chain.

Chapter 8

Complexity of finding frequent patterns

We now analyze the complexity of finding all frequent patterns, and we also derive results for the complexity of discovering all frequent sets. In Section 8.1 we introduce the concept of the border between frequent and non-frequent patterns. This notion turns out to be useful in the analysis of the generic algorithm in Section 8.2. Inspired by this, we give in Section 8.3 a guess-and-correct algorithm for the task of finding all frequent patterns. We then analyze the task in Section 8.4. In Section 8.5 we return to the concept of border, and show that it has strong connections to transversals on hypergraphs.

8.1 The border

Consider the theory $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ of some set \mathcal{P} of patterns. The whole theory can be specified by giving only the maximally specific patterns in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$: every pattern more general than any of those is selected by q , and the rest are not. The collection of maximally specific patterns in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ and, correspondingly, the collection of *minimally* specific (i.e., maximally general) patterns *not* in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ are useful in the analysis of the generic algorithm and the problem. For this purpose we introduce the notion of border.

Definition 8.1 Let \mathcal{P} be a set of patterns, \mathcal{S} a subset of \mathcal{P} , and \preceq a partial order on \mathcal{P} . Further, let \mathcal{S} be closed downwards under the relation \preceq , i.e., if $\varphi \in \mathcal{S}$ and $\gamma \preceq \varphi$, then $\gamma \in \mathcal{S}$. The *border* $Bd(\mathcal{S})$ of \mathcal{S} consists of those patterns φ such that all more general patterns than φ are in \mathcal{S} and no pattern more specific than φ is in \mathcal{S} :

$$Bd(\mathcal{S}) = \{ \varphi \in \mathcal{P} \mid \text{for all } \gamma \in \mathcal{P} \text{ such that } \gamma \prec \varphi \text{ we have } \gamma \in \mathcal{S}, \text{ and} \\ \text{for all } \theta \in \mathcal{P} \text{ such that } \varphi \prec \theta \text{ we have } \theta \notin \mathcal{S} \}.$$

Those patterns φ in $\mathcal{Bd}(\mathcal{S})$ that are in \mathcal{S} are called the *positive border* $\mathcal{Bd}^+(\mathcal{S})$,

$$\mathcal{Bd}^+(\mathcal{S}) = \{\varphi \in \mathcal{S} \mid \text{for all } \theta \in \mathcal{P} \text{ such that } \varphi \prec \theta \text{ we have } \theta \notin \mathcal{S}\},$$

and those patterns φ in $\mathcal{Bd}(\mathcal{S})$ that are not in \mathcal{S} are the *negative border* $\mathcal{Bd}^-(\mathcal{S})$,

$$\mathcal{Bd}^-(\mathcal{S}) = \{\varphi \in \mathcal{P} \setminus \mathcal{S} \mid \text{for all } \gamma \in \mathcal{P} \text{ such that } \gamma \prec \varphi \text{ we have } \gamma \in \mathcal{S}\}.$$

□

In other words, the positive border consists of the most specific patterns in \mathcal{S} , the negative border consists of the most general patterns outside \mathcal{S} , and the border is the union of these two sets. Note that a set \mathcal{S} that is closed downwards can be described by giving just the positive or the negative border. Consider, e.g., the negative border. No pattern θ such that $\varphi \preceq \theta$ for some φ in the negative border is in \mathcal{S} , while all other patterns are in \mathcal{S} .

A theory $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ is always closed downwards with respect to a specialization relation, and the concept of border can be applied on the set of all frequent patterns.

Example 8.2 Consider the discovery of frequent sets with items $R = \{A, \dots, F\}$. Assume the collection \mathcal{F} of frequent sets is

$$\{\{A\}, \{B\}, \{C\}, \{F\}, \{A, B\}, \{A, C\}, \{A, F\}, \{C, F\}, \{A, C, F\}\}.$$

The negative border of this collection contains now sets that are not frequent, but whose all subsets are frequent, i.e., minimal non-frequent sets. The negative border is thus

$$\mathcal{Bd}^-(\mathcal{F}) = \{\{D\}, \{E\}, \{B, C\}, \{B, F\}\}.$$

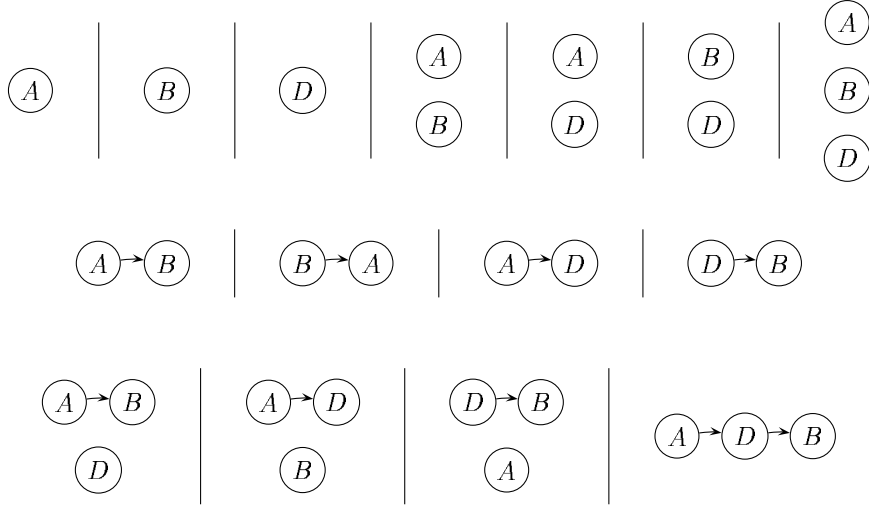

The positive border, in turn, contains the maximal frequent sets, i.e.,

$$\mathcal{Bd}^+(\mathcal{F}) = \{\{A, B\}, \{A, C, F\}\}.$$

□

Example 8.3 Consider the discovery of frequent episodes in a sequence over events A, \dots, D . Assume that arbitrary partial orders are allowed in the episodes, and that the episodes in Figure 8.1 constitute the theory $\mathcal{F}(\mathbf{s}, \text{win}, \text{min_fr})$ of frequent episodes in some given event sequence \mathbf{s} .

There are only two maximally specific episodes, those presented in Figure 8.2. All other episodes in $\mathcal{F}(\mathbf{s}, \text{win}, \text{min_fr})$ are included in either of these two that constitute the positive border $\mathcal{Bd}^+(\mathcal{F}(\mathbf{s}, \text{win}, \text{min_fr}))$. The negative border $\mathcal{Bd}^-(\mathcal{F}(\mathbf{s}, \text{win}, \text{min_fr}))$ is depicted in Figure 8.3. It consists of the maximally general episodes that are not frequent. □

Figure 8.1: A collection $\mathcal{F}(\mathbf{s}, win, min_fr)$ of frequent episodes.Figure 8.2: The positive border $\mathcal{B}d^+(\mathcal{F}(\mathbf{s}, win, min_fr))$.Figure 8.3: The negative border $\mathcal{B}d^-(\mathcal{F}(\mathbf{s}, win, min_fr))$.

8.2 Complexity of the generic algorithm

Consider the complexity of discovering all frequent patterns, in terms of the number of evaluations of the selection criterion q . The trivial method for finding $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ is to test all patterns of \mathcal{P} , and hence use $|\mathcal{P}|$ evaluations of q . Algorithm 7.6 evaluates only patterns in the result, i.e., frequent patterns, and patterns in the negative border of the collection of frequent patterns.

Theorem 8.4 Let \mathcal{P} , \mathbf{r} , and q be as in Algorithm 7.6. Algorithm 7.6 evaluates the predicate q exactly on the patterns in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q) \cup \mathcal{B}d^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$.

Proof First note that no pattern is evaluated more than once: each pattern has a unique level and can therefore be at most in one collection \mathcal{C}_l . Recall now line 8 of Algorithm 7.6, the specification of the candidate collection:

$$8. \quad \text{compute } \mathcal{C}_l := \{\varphi \in \mathcal{P} \mid \text{level}(\varphi) = l \text{ and } \theta \in \mathcal{T}_{\text{level}(\theta)}(\mathcal{P}, \mathbf{r}, q) \text{ for all } \theta \in \mathcal{P} \text{ such that } \theta \prec \varphi\};$$

We show that $\bigcup_l \mathcal{C}_l = \mathcal{T}(\mathcal{P}, \mathbf{r}, q) \cup \mathcal{Bd}^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$. First note that every pattern φ in \mathcal{C}_l is in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q) \cup \mathcal{Bd}^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$. If φ is selected by q , it is in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$. If φ is not selected by q , it is in $\mathcal{Bd}^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$, since all patterns more general than φ are in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$, and φ itself is not.

Now note that every pattern φ in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q) \cup \mathcal{Bd}^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$ is in $\mathcal{C}_{\text{level}(\varphi)}$. When $l = \text{level}(\varphi)$, all more general patterns than φ have been evaluated in earlier iterations since their levels are less than $\text{level}(\varphi)$. All more general patterns $\theta \prec \varphi$ are in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ and thus in $\mathcal{T}_{\text{level}(\theta)}(\mathcal{P}, \mathbf{r}, q)$. So when $l = \text{level}(\varphi)$, the pattern φ will be in $\mathcal{C}_{\text{level}(\varphi)}$.

Finally, it can be shown, as in the proof of Theorem 7.7, that $\mathcal{C}_{\text{level}(\varphi)}$ is constructed for each φ in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q) \cup \mathcal{Bd}^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$. \square

What the candidate generation step of Algorithm 7.6 basically does is to compute the negative border of frequent patterns found so far. Line 8 actually equals the following specification:

$$8. \quad \text{compute } \mathcal{C}_l := \mathcal{Bd}^-(\bigcup_{i < l} \mathcal{T}_i(\mathcal{P}, \mathbf{r}, q)) \setminus \bigcup_{i < l} \mathcal{C}_i;$$

That is, in each iteration the candidate collection generated is exactly the negative border of the patterns selected so far, minus patterns already found not to be selected. We can now use the concept of negative border to restate the complexity of the frequent set discovery algorithm in a compact form.

Corollary 8.5 Given a set R , a binary database r over R , and a frequency threshold min_fr , Algorithm 2.14 evaluates the frequency of sets in $\mathcal{F}(r, \text{min_fr}) \cup \mathcal{Bd}^-(\mathcal{F}(r, \text{min_fr}))$.

Proof The claim follows directly from the fact that Algorithm 2.14 is an instance of Algorithm 7.6 and from Theorem 8.4. \square

Example 8.6 Consider the discovery of frequent sets in a random relation over 20 attributes, where the probability that $t[A]$ has value 1 is p for all rows t and all attributes A . Table 8.1 presents the sizes of the theory and its positive and negative borders in experiments with two random relations, with $p = 0.2$ and $p = 0.5$, and with 1000 rows. In such relations the size of the border seems to be roughly 2 to 4 times the number of frequent sets.

$\mathcal{Bd}(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$ can be small for a large theory $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$. An experiment with frequent sets in a real database gave the following results. We discovered all frequent sets in a course enrollment database where there is a row per student, a column per each course offered, and a row has value 1 in a column

p	min_fr	$ \mathcal{T}(\mathcal{P}, \mathbf{r}, q) $	$ \mathcal{B}d^+(\mathcal{T}(\mathcal{P}, \mathbf{r}, q)) $	$ \mathcal{B}d^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q)) $
0.2	0.01	469	273	938
0.2	0.005	1291	834	3027
0.5	0.1	1335	1125	4627
0.5	0.05	5782	4432	11531

Table 8.1: Experimental results with random data sets.

min_fr	$ \mathcal{T}(\mathcal{P}, \mathbf{r}, q) $	$ \mathcal{B}d^+(\mathcal{T}(\mathcal{P}, \mathbf{r}, q)) $	$ \mathcal{B}d^-(\mathcal{T}(\mathcal{P}, \mathbf{r}, q)) $
0.08	96	35	201
0.06	270	61	271
0.04	1028	154	426
0.02	6875	328	759

Table 8.2: Experimental results with a real data set.

if the corresponding student took the corresponding course. There are 127 courses, and a student has taken 4.3 courses on average; the number of students is 4734. Table 8.2 shows that the size of the border behaves nicely with respect to the size of the theory. \square

8.3 The guess-and-correct algorithm

Algorithm 7.6 starts by evaluating the selection predicate q on the most general sentences, and moves gradually to more specific sentences. As q is assumed to be monotone with respect to the specialization relation, this approach is safe in the sense that no statement satisfying q will be overlooked. However, the approach can be quite slow, if there are frequent patterns that are far from the bottom of the specialization relation, i.e., if there are patterns φ that are selected by q , but which appear in the candidate set \mathcal{C}_i only for a large i . As every iteration of the algorithm requires an investigation of the database, this means that such sentences will be discovered slowly.

An alternative is to start the process of finding $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ from an initial guess $\mathcal{S} \subseteq \mathcal{P}$, and then correcting the guess by looking at the database. The guess can be obtained, e.g., from computing the set $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$, where $\mathbf{s} \subseteq \mathbf{r}$ is a sample of \mathbf{r} . Algorithm 8.7 computes $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ as follows. To begin with, the algorithm is given a guess $\mathcal{S} \subseteq \mathcal{P}$. The algorithm first evaluates the sentences in the positive border $\mathcal{B}d^+(\mathcal{S})$ and removes from \mathcal{S} those that are not interesting. These evaluation and removal steps are repeated until the positive border only contains sentences satisfying q , and

Algorithm 8.7 The *guess-and-correct algorithm* for finding all potentially interesting sentences with an initial guess \mathcal{S} .

Input: A database \mathbf{r} , a language \mathcal{P} with specialization relation \preceq , a selection predicate q , and an initial guess $\mathcal{S} \subseteq \mathcal{P}$ for $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$. We assume \mathcal{S} is closed under generalizations..

Output: The set $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$..

Method:

1. $\mathcal{C}^* := \emptyset$;
 // correct \mathcal{S} downward:
 2. $\mathcal{C} := \mathcal{B}d^+(\mathcal{S})$;
 3. **while** $\mathcal{C} \neq \emptyset$ **do**
 4. $\mathcal{C}^* := \mathcal{C}^* \cup \mathcal{C}$;
 5. $\mathcal{S} := \mathcal{S} \setminus \{\varphi \in \mathcal{C} \mid q(\mathbf{r}, \varphi) \text{ is false}\}$;
 6. $\mathcal{C} := \mathcal{B}d^+(\mathcal{S}) \setminus \mathcal{C}^*$;
 7. **od**;
 - // now $\mathcal{S} \subseteq \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$; expand \mathcal{S} upwards:
 8. $\mathcal{C} := \mathcal{B}d^-(\mathcal{S}) \setminus \mathcal{C}^*$;
 9. **while** $\mathcal{C} \neq \emptyset$ **do**
 10. $\mathcal{C}^* := \mathcal{C}^* \cup \mathcal{C}$;
 11. $\mathcal{S} := \mathcal{S} \cup \{\varphi \in \mathcal{C} \mid q(\mathbf{r}, \varphi) \text{ is true}\}$;
 12. $\mathcal{C} := \mathcal{B}d^-(\mathcal{S}) \setminus \mathcal{C}^*$;
 13. **od**;
 14. output \mathcal{S} ;
-

thus we have $\mathcal{S} \subseteq \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$. The variable \mathcal{C}^* is used to avoid evaluating sentences twice. Then the algorithm expands \mathcal{S} upwards, like Algorithm 7.6: it evaluates such sentences in the negative border $\mathcal{B}d^-(\mathcal{S})$ that have not been evaluated yet, and adds those that satisfy q to \mathcal{S} . Again, the algorithm repeats the evaluation and addition steps until there are no sentences to evaluate. Finally, the output is $\mathcal{S} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$.

The following results are straightforward.

Lemma 8.8 Algorithm 8.7 works correctly.

Theorem 8.9 Algorithm 8.7 uses at most

$$|(\mathcal{S} \Delta \mathcal{T}) \cup \mathcal{B}d(\mathcal{T}) \cup \mathcal{B}d^+(\mathcal{S} \cap \mathcal{T})|$$

evaluations of q , where $\mathcal{T} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$.

How to obtain good original guesses \mathcal{S} ? One fairly widely applicable method is sampling. Take a small sample \mathbf{s} from \mathbf{r} , compute $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ and use it as \mathcal{S} . Applied to association rules this method produces extremely good results. Basically, with a high probability one can discover the association rules holding in a database using only a single pass through the database.

Another method for computing an initial approximation can be derived from the Partition algorithm [77]. The idea is to divide \mathbf{r} into small datasets

\mathbf{r}_i which can be handled in main memory, and to compute $\mathcal{S}_i = \mathcal{T}(\mathcal{P}, \mathbf{r}_i, q)$. In the case of frequent sets, use as the guess \mathcal{S} the union $\bigcup_i \mathcal{S}_i$. The guess is a superset of $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$, and executing the first half of Algorithm 8.7 suffices.

8.4 Problem complexity

Let us now analyze the complexity of the problem of discovering all frequent patterns. Consider first the following verification problem: assume somebody gives a set $\mathcal{S} \subseteq \mathcal{P}$ and claims that $\mathcal{S} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$. How many evaluations of q are necessary for verifying this claim? The following theorem shows that the whole border $\mathcal{Bd}(\mathcal{S})$ must be inspected.

Theorem 8.10 Let \mathcal{P} and $\mathcal{S} \subseteq \mathcal{P}$ be sets of patterns, \mathbf{r} a database, q a selection criterion, and \preceq a specialization relation. If the database \mathbf{r} is accessed only using the predicate q , then determining whether $\mathcal{S} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ (1) requires in the worst case at least $|\mathcal{Bd}(\mathcal{S})|$ evaluations of q , and (2) can be done in exactly $|\mathcal{Bd}(\mathcal{S})|$ evaluations of q .

Proof We show that it is sufficient and in the worst case necessary to evaluate the border $\mathcal{Bd}(\mathcal{S})$. The claims follow then from this.

First assume that patterns in the border are evaluated. If and only if every pattern in $\mathcal{Bd}^+(\mathcal{S})$ and no pattern in $\mathcal{Bd}^-(\mathcal{S})$ is selected by q , then $\mathcal{S} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$, by the definition of the border. If \mathcal{S} and $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ do not agree on the border, then clearly $\mathcal{S} \neq \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$.

Now assume that less than $|\mathcal{Bd}(\mathcal{S})|$ evaluations have been made, all consistent with the claim $\mathcal{S} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$; then there is a pattern φ in the border $\mathcal{Bd}(\mathcal{S})$ for which q has not been evaluated. Now there is no way of knowing whether φ is in $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ or not. The evaluations made for other patterns give no information about φ : since they were consistent with $\mathcal{S} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ and φ is in the border, all more general patterns $\theta \prec \varphi$ are selected by the definition of border, none of the more specific patterns is selected, and the rest are irrelevant with respect to φ . In other words, any set in the negative border can be swapped to the positive border, and vice versa, without changing the truth of q for any other set. \square

Corollary 8.11 Let \mathcal{P} be a set of patterns, \mathbf{r} a database, q a selection criterion, and \preceq a specialization relation. Any algorithm that computes $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ and accesses the data only with the predicate q must evaluate q on the patterns in $\mathcal{Bd}(\mathcal{T}(\mathcal{P}, \mathbf{r}, q))$.

Proof The claim follows directly from the proof of Theorem 8.10. \square

Example 8.12 Continuing Example 8.2 consider the discovery of frequent sets in a relation over attributes $R = \{A, \dots, F\}$. Assume a sample or some

person tells us that the collection of frequent sets is

$$\mathcal{S} = \{\{A\}, \{B\}, \{C\}, \{F\}, \{A, B\}, \{A, C\}, \{A, F\}, \{C, F\}, \{A, C, F\}\}.$$

To verify whether the claim is true we have to evaluate, according to Corollary 8.11, the border $\mathcal{Bd}(\mathcal{S})$. We thus have to evaluate $\mathcal{Bd}^+(\mathcal{S}) = \{\{A, B\}, \{A, C, F\}\}$ and $\mathcal{Bd}^-(\mathcal{S}) = \{\{D\}, \{E\}, \{B, C\}, \{B, F\}\}$ to determine whether \mathcal{S} is the collection of frequent sets. \square

We have further corollaries about the complexity of discovery tasks. Consider the discovery of frequent sets. Recall that Algorithm 2.14 evaluates also the frequency of the candidates that are not frequent, i.e., sets in $\mathcal{Bd}^-(\mathcal{F}(r, \text{min_fr}))$. The following corollary shows that in a limited but reasonable model of computation the evaluation of the non-frequent candidates is inevitable.

Corollary 8.13 Given a set R , a binary database r over R , and a frequency threshold min_fr , finding the collection $\mathcal{F}(r, \text{min_fr})$ using only queries of the form “Is $X \subseteq R$ frequent in r ” requires that sets in the negative border $\mathcal{Bd}^-(\mathcal{F}(r, \text{min_fr}))$ are evaluated.

Proof The claim follows directly from Corollary 8.11. \square

Algorithm 2.14 actually evaluates the whole theory $\mathcal{F}(r, \text{min_fr})$, not only its border. For the discovery of association rules this is in general necessary, since the exact frequencies are needed in the rule generation phase. Algorithm 2.14 is thus optimal under the simplifying restriction that the only way of controlling the algorithm is to use the information whether a set is frequent or not.

The advantage of Corollary 8.11 is that the border $\mathcal{Bd}(\mathcal{S})$ can be small even for large \mathcal{S} . The drawback is that it can be difficult to determine the border. We return to this issue in Section 8.5 where we show a connection between the problem of finding the border and the hypergraph transversal problem.

8.5 Computing the border

We now return to the verification problem: given \mathcal{P} , \mathbf{r} , q , and a set $\mathcal{S} \subseteq \mathcal{P}$, verify whether $\mathcal{S} = \mathcal{T}(\mathcal{P}, \mathbf{r}, q)$. By Corollary 8.11 the border $\mathcal{Bd}(\mathcal{S})$ must be inspected. Given \mathcal{S} , we can compute $\mathcal{Bd}^+(\mathcal{S})$ without looking at the data \mathbf{r} at all: simply find the most specific patterns in \mathcal{S} . The negative border $\mathcal{Bd}^-(\mathcal{S})$ is also defined by \mathcal{S} , and can be determined without looking at the data, but finding the most general patterns in $\mathcal{P} \setminus \mathcal{S}$ can be difficult. We now show how minimal transversals of hypergraphs can be used to determine the negative border.

Definition 8.14 Let R be a set. A collection \mathcal{H} of subsets of R is a *simple hypergraph* on R , if no element of \mathcal{H} is empty and if $X, Y \in \mathcal{H}$ and $X \subseteq Y$ imply $X = Y$. The elements of \mathcal{H} are called the *edges* of the hypergraph, and the elements of R are the *vertices* of the hypergraph. \square

Definition 8.15 Given a simple hypergraph \mathcal{H} on a set R , a *transversal* T of \mathcal{H} is a subset of R intersecting all the edges of \mathcal{H} : T is a transversal if and only if $T \cap X \neq \emptyset$ for all $X \in \mathcal{H}$. A *minimal transversal* of \mathcal{H} is a transversal T such that no $T' \subset T$ is a transversal. We denote the collection of minimal transversals of \mathcal{H} by $Tr(\mathcal{H})$. \square

For our purposes, hypergraphs and transversals apply almost directly on the pattern class of frequent sets. Let the items in R be the vertices and the *complements* of the sets in the positive border be the edges of a simple hypergraph \mathcal{H} . So, for each set X in the positive border we have the set $R \setminus X$ as an edge in \mathcal{H} . Consider now a set $Y \subseteq R$. If there is an edge $R \setminus X$ such that $Y \cap (R \setminus X) = \emptyset$, then $Y \subseteq X$, and Y is frequent. On the other hand, if there is no such edge that the intersection is empty, then Y cannot be frequent. That is, Y is not frequent if and only if Y is a transversal of \mathcal{H} . Minimal transversals are now the minimal non-frequent sets, i.e., the negative border.

In general, for using hypergraphs and transversals to determine the negative border, we need to represent patterns in \mathcal{P} as sets. Frequent sets are such a representation themselves; next we give the requirements for the general case.

Definition 8.16 Let \mathcal{P} be a set of patterns, \preceq a specialization relation, and R a set. A function $f : \mathcal{P} \rightarrow \{X \mid X \subseteq R\}$ is a *set representation* of \mathcal{P} and \preceq , if f is bijective, f and its inverse are computable, and for all $\theta, \varphi \in \mathcal{P}$ we have $\theta \preceq \varphi$ if and only if $f(\theta) \subseteq f(\varphi)$.

For notational convenience, given a collection \mathcal{S} of sets, we write $f(\mathcal{S}) = \{f(X) \mid X \in \mathcal{S}\}$. \square

As was described above for frequent sets, minimal transversals of a suitably constructed hypergraph constitute the negative border.

Definition 8.17 Let \mathcal{P} and $\mathcal{S} \subseteq \mathcal{P}$ be sets of patterns and let f be a set representation of \mathcal{P} . We denote by $\mathcal{H}(\mathcal{S})$ the simple hypergraph on R that contains as edges the complements of sets $f(\varphi)$ for $\varphi \in \mathcal{B}d^+(\mathcal{S})$, i.e., $\mathcal{H}(\mathcal{S}) = \{R \setminus f(\varphi) \mid \varphi \in \mathcal{B}d^+(\mathcal{S})\}$. \square

Now $\mathcal{H}(\mathcal{S})$ is a hypergraph corresponding to the set representation of \mathcal{S} , and $Tr(\mathcal{H}(\mathcal{S}))$ is the set representation of the negative border. The inverse function f^{-1} maps the set representations of the negative border to the patterns in the negative border. That is, the set $f^{-1}(Tr(\mathcal{H}(\mathcal{S})))$ is the negative border. Next we show this formally.

Theorem 8.18 Let \mathcal{P} and $\mathcal{S} \subseteq \mathcal{P}$ be sets of patterns, and let f be a set representation of \mathcal{P} . Then $f^{-1}(\text{Tr}(\mathcal{H}(\mathcal{S}))) = \mathcal{B}d^-(\mathcal{S})$.

Proof We prove the claim in two steps. First we show that X is a transversal of $\mathcal{H}(\mathcal{S})$ if and only if $f^{-1}(X) \notin \mathcal{S}$:

$$\begin{aligned}
& X \text{ is a transversal of } \mathcal{H}(\mathcal{S}) \\
\Leftrightarrow & X \cap Y \neq \emptyset \text{ for all } Y \in \mathcal{H}(\mathcal{S}) \\
\Leftrightarrow & X \cap (R \setminus f(\varphi)) \neq \emptyset \text{ for all } \varphi \in \mathcal{B}d^+(\mathcal{S}) \\
\Leftrightarrow & X \not\subseteq f(\varphi) \text{ for all } \varphi \in \mathcal{B}d^+(\mathcal{S}) \\
\Leftrightarrow & f^{-1}(X) \not\subseteq \varphi \text{ for all } \varphi \in \mathcal{B}d^+(\mathcal{S}) \\
\Leftrightarrow & f^{-1}(X) \notin \mathcal{S}.
\end{aligned}$$

Next we show that $\text{Tr}(\mathcal{H}(\mathcal{S})) = f(\mathcal{B}d^-(\mathcal{S}))$; the theorem then immediately follows.

$$\begin{aligned}
& \text{Tr}(\mathcal{H}(\mathcal{S})) \\
= & \{X \mid X \text{ is a minimal transversal of } \mathcal{H}(\mathcal{S})\} \\
= & \{X \mid X \text{ is a minimal set such that } f^{-1}(X) \notin \mathcal{S}\} \\
= & \{X \mid f^{-1}(X) \notin \mathcal{S} \text{ and } f^{-1}(Y) \in \mathcal{S} \text{ for all } Y \subset X\} \\
= & \{f(\varphi) \mid \varphi \notin \mathcal{S} \text{ and } \theta \in \mathcal{S} \text{ for all } \theta \prec \varphi\} \\
= & f(\mathcal{B}d^-(\mathcal{S})).
\end{aligned}$$

□

Thus for languages representable as sets, the notions of negative border and the minimal transversals coincide.

Example 8.19 Continuing Example 8.12, the discovery of frequent sets in a relation with attributes $R = \{A, \dots, F\}$, we now compute the set $\mathcal{B}d^-(\mathcal{S})$ using the hypergraph formulation. For frequent sets we simply let f be the identity function: $f(X) = X$. We are given the positive border $\mathcal{B}d^+(\mathcal{S}) = \{\{A, B\}, \{A, C, F\}\}$, so we have

$$\mathcal{H}(\mathcal{S}) = \{R \setminus f(X) \mid X \in \mathcal{B}d^+(\mathcal{S})\} = \{\{C, D, E, F\}, \{B, D, E\}\}.$$

The minimal transversals of $\mathcal{H}(\mathcal{S})$ are $\{D\}$, $\{E\}$, $\{B, C\}$, and $\{B, F\}$, thus

$$\mathcal{B}d^-(\mathcal{S}) = f^{-1}(\text{Tr}(\mathcal{H}(\mathcal{S}))) = \text{Tr}(\mathcal{H}(\mathcal{S})) = \{\{D\}, \{E\}, \{B, C\}, \{B, F\}\}.$$

□

We showed in Section 8.4 that under some simplifying restrictions the negative border must be evaluated when discovering all frequent patterns. In this section we showed that for patterns representable as sets the notion of negative border corresponds to the minimal transversals of a suitably defined hypergraph. The advantage of this is that the wealth of material about transversals (see, e.g., [7]) can be used, e.g., in the design of algorithms or complexity analysis for specific knowledge discovery problems.

8.6 Bibliographic notes

Many of the concepts and ideas of this and the previous chapter are known in different contexts. These chapters are, however, among the first attempts to provide a unified viewpoint to several knowledge discovery tasks and a generic algorithm for those tasks. Parts of the material of these chapters have been published in [61].

Algorithm 7.6 is based on the breadth-first search strategy, and it uses a specialization relation for safely pruning branches of the search tree, both well known search methods. The idea of checking *all* more general patterns when generating candidates has, however, been missed, e.g., in the original algorithm for discovering frequent sets [1] and in the inference of inclusion dependencies [49, 58]. In the area of machine learning, the version spaces of Mitchell [68] are the first systematic use of specialization relations and concepts similar to our border. Mitchell's learning task is different from ours, but conceptually Mitchell's set S of the most specific consistent patterns is the same as our positive border. —A generic viewpoint to knowledge discovery algorithms, similar to ours, has been expressed in [79].

There are several knowledge discovery settings that can be seen as the discovery of all frequent patterns. We very briefly contrast the described approach with two important systems, Explora and Claudien.

Explora [53] is a system for discovering patterns describing outstanding subgroups of a given database. Explora employs specialization relations: it organizes patterns into hierarchies and lets the user specify which are specialization relations with respect to the domain and the user's interests. The algorithm repeatedly evaluates patterns along paths between the most general and the most specific patterns until the border is located. The specialization relations are used to prune from evaluation those patterns for which the truth of selection criterion is already known.

Claudien [13] discovers minimally complete first order clausal theories from databases; in our terms a minimally complete theory roughly corresponds to the positive border. Claudien discovers the positive border by finding those patterns that are not frequent, i.e., by a conversed strategy. Due to the large number of patterns, in particular the number of patterns in individual levels, the implementation of Claudien uses depth-first search in order to save space. Therefore Claudien cannot take full advantage of the specialization relation: it may generate and test candidate clauses that cannot hold given all more general patterns.

In general, the depth-first search strategy may be useful if the collections of frequent patterns on each level are very large, or if computing the selection criterion from the database is cheap. Pratt [46], a system for the discovery of patterns in protein sequences, is a good example of such an application. Since the total size of the analyzed sequences is not large, Pratt can store the data in main memory and even use index structures that are significantly

larger than the original data. Evaluating the selection criterion is thus fast, possibly even faster than evaluating whether a potential candidate is a valid candidate or not.

Methods that use sampling in the discovery of frequent sets have been given in [84].

The problem complexity of these settings has not received much attention. Some lower bounds for the problem of finding all frequent sets are given in [2, 62]. The relevance of transversals to computing the theory of a model has been known in the context of finding functional dependencies [59] and several other specific problems [18]. The complexity of computing the minimal transversals of a hypergraph has long been open; see [34, 67] for recent results. It is, however, known that transversals can be computed in time $\mathcal{O}(n^{\mathcal{O}(\log n)})$, where n is the sum of the sizes of the edges of both the hypergraph and its minimal transversals [26, 50]. The connection of the border and transversals have been recently strengthened [33].

Chapter 9

Sampling large databases for frequent sets

The size of the data collection has an essential role in data mining. Large data sets are necessary for reliable results—unfortunately, however, the efficiency of mining algorithms depends significantly on the database. The time complexity of the frequent set discovery algorithm is linear with respect to the number of rows in the database. However, the algorithm requires multiple passes over the database, and subsequently the database size is the most influential factor in the execution time for large databases.

In this chapter we present algorithms that make only one or sometimes two passes over the database. The idea is to pick a random sample from the input database, use it to determine all sets that possibly are frequent in the whole database, and then to verify the results with the rest of the database. These algorithms thus produce a correct set of association rules in one full pass over the database. In those rare cases where our sampling method does not produce all frequent sets, the missing sets can be found in a second pass. The concept of negative border turns out to be useful in this task.

We describe our sampling approach for discovering association rules in Section 9.1. In Section 9.2 we analyze the goodness of the sampling method, e.g., the relation of sample size to the accuracy of results. In Section 9.3 we give variations of algorithms and experimental results. The results show that the methods reduce the disk activity considerably, making the approach attractive especially for large databases. This chapter is based on [84].

9.1 Sampling in the discovery of frequent sets

An obvious way of reducing the database activity of knowledge discovery is to use only a random sample of the database and to find approximate regularities. In other words, one can trade off accuracy against efficiency. This can be useful: samples small enough to be handled totally in main

Algorithm 9.1

Input: A binary database r over a set R , a frequency threshold min_fr , a sample size s_size , and a lowered frequency threshold $low_fr < min_fr$.

Output: The collection $\mathcal{F}(r, min_fr)$ of frequent sets and their frequencies, or its subset and a failure report.

Method:

1. compute $s :=$ a random sample of size s_size from r ;
 2. // Find frequent sets in the sample:
 3. compute $\mathcal{S} := \mathcal{F}(s, low_fr)$ in main memory using Algorithm 2.14;
 4. // Database pass:
 5. compute $\mathcal{R} := \{X \in \mathcal{S} \cup \mathcal{B}d^-(\mathcal{S}) \mid fr(X, r) \geq min_fr\}$ using Algorithm 2.22;
 6. // Output:
 7. for all $X \in \mathcal{R}$ do output X and $fr(X, r)$;
 8. if $\mathcal{R} \cap \mathcal{B}d^-(\mathcal{S}) \neq \emptyset$ then report that there possibly was a failure;
-

memory can give reasonably accurate results. Or, approximate results from a sample can be used to set the focus for a more complete discovery phase.

It is often important to know the frequencies and confidences of association rules exactly. In business applications, for example for large volumes of supermarket sales data, even small differences can be significant. When relying on results from sampling alone, one also takes the risk of losing valid association rules altogether because their frequency in the sample is below the threshold.

Using a random sample to get approximate results is fairly straightforward. Below we give bounds for sample sizes, given the desired accuracy of the results. We show further that exact frequencies can be found efficiently, by analyzing first a random sample and then the whole database as follows. Use a random sample to locate a superset \mathcal{S} of the collection of frequent sets. A superset can be determined efficiently by applying Algorithm 2.14 on the sample in main memory, and by using a lowered frequency threshold. Then use \mathcal{S} as the collection of candidates, and compute the exact frequencies of the sets from the rest of the database. This approach, when successful, requires only one full pass over the database, and two passes in the worst case.

Algorithm 9.1 presents the principle: search for frequent sets in the sample, but use a lower frequency threshold so that it is unlikely that frequent sets are missed.

The concept of negative border is useful here. As was noted in the previous chapter, the border has to be inspected when discovering frequent sets. It is thus not sufficient to locate a superset \mathcal{S} of $\mathcal{F}(r, min_fr)$ using the sample and then to evaluate \mathcal{S} in r . Rather, the collection $\mathcal{F}(r, min_fr) \cup \mathcal{B}d^-(\mathcal{F}(r, min_fr))$ needs to be checked. Obviously $\mathcal{S} \cup \mathcal{B}d^-(\mathcal{S})$ is a superset of this collection if \mathcal{S} is a superset of $\mathcal{F}(r, min_fr)$, so we check the union. Sometimes, however, it happens that we find out that not all

necessary sets have been evaluated.

Definition 9.2 There has been a *failure* in the sampling if all frequent sets are not found in one pass, i.e., if there is a frequent set X in $\mathcal{F}(r, \text{min_fr})$ that is not in $\mathcal{S} \cup \mathcal{B}d^-(\mathcal{S})$. A *miss* is a frequent set Y in $\mathcal{F}(r, \text{min_fr})$ that is in $\mathcal{B}d^-(\mathcal{S})$. \square

If there are no misses, then the sampling has been successful. Misses themselves are not a problem: they are evaluated in the whole database, and thus they are not actually missed by the algorithm. Misses, however, indicate a potential failure. Namely, if there is a miss Y , then some superset of Y might be frequent but not in $\mathcal{S} \cup \mathcal{B}d^-(\mathcal{S})$. A simple way to recognize a potential failure is thus to check if there are any misses.

Theorem 9.3 Algorithm 9.1 works correctly.

Proof Clearly, on lines 5 and 7, a collection of frequent sets is computed and output. We need to show that if no failure report is given, then all frequent sets are found, and that if all frequent sets are not found, then a failure report is, in turn, given.

If there is no failure report, i.e., if \mathcal{R} and $\mathcal{B}d^-(\mathcal{S})$ are disjoint, then $\mathcal{R} \subseteq \mathcal{S}$, and $\mathcal{B}d^-(\mathcal{R}) \subseteq \mathcal{S} \cup \mathcal{B}d^-(\mathcal{S})$. Thus the whole negative border $\mathcal{B}d^-(\mathcal{R})$ has been evaluated, and all frequent sets are found. If all frequent sets are not found, i.e., if there is a frequent set X that is not in $\mathcal{S} \cup \mathcal{B}d^-(\mathcal{S})$, then there exists a set Y in $\mathcal{B}d^-(\mathcal{S})$ such that $Y \subseteq X$ and Y is frequent. This set Y is thus in $\mathcal{R} \cap \mathcal{B}d^-(\mathcal{S})$, and a failure is reported. \square

Example 9.4 Assume that we have a binary database r with 10 million rows over items A, \dots, F , and that we want to find the frequent sets with the threshold 0.02. Algorithm 9.1 randomly picks a small fraction s of r , say 20 000 rows, and keeps this sample s in main memory. The algorithm can now, without any further database activity, discover efficiently the sets that are frequent in the sample.

To make it very probably that the collection of frequent sets in the sample includes all sets that really are frequent in r , the frequency threshold is lowered to, say, 0.015. So Algorithm 9.1 determines the collection $\mathcal{S} = \mathcal{F}(s, 0.015)$ from the sampled 20 000 rows. Let the maximal sets of \mathcal{S} , i.e., the positive border $\mathcal{B}d^+(\mathcal{S})$, be

$$\{A, D\}, \{B, D\}, \{A, B, C\}, \{A, C, F\}.$$

Since the threshold was lowered, \mathcal{S} is likely to be a superset of the collection $\mathcal{F}(r, 0.02)$ of frequent sets. In the pass over the rest of the database r , the frequency of all sets in \mathcal{S} and $\mathcal{B}d^-(\mathcal{S})$ is evaluated. That is, in addition to

Algorithm 9.5

Input: A binary database r over a set R , a frequency threshold min_fr , a subset \mathcal{S} of $\mathcal{F}(r, min_fr)$, and a subset \mathcal{S}^- of $\mathcal{B}d^-(\mathcal{F}(r, min_fr))$.

Output: The collection $\mathcal{F}(r, min_fr)$ of frequent sets and their frequencies.

Method:

1. **repeat** compute $\mathcal{S} := \mathcal{S} \cup (\mathcal{B}d^-(\mathcal{S}) \setminus \mathcal{S}^-)$ **until** \mathcal{S} does not grow;
 2. compute $\mathcal{R} := \{X \in \mathcal{S} \mid fr(X, r) \geq min_fr\}$;
 3. **for** all $X \in \mathcal{R}$ **do** output X and $fr(X, r)$;
-

the sets that are frequent in the sample, we evaluate also those candidates that were not frequent, i.e., the negative border

$$\{E\}, \{B, F\}, \{C, D\}, \{D, F\}, \{A, B, D\}.$$

The goal is to discover the collection $\mathcal{F}(r, 0.02)$. Let the sets

$$\{A, B\}, \{A, C, F\}.$$

and their subsets be the frequent sets. All frequent sets are in \mathcal{S} , so they are evaluated and their exact frequencies are known after the full database pass. We also know that we have found all frequent sets since also sets

$$\{D\}, \{E\}, \{B, C\}, \{B, F\},$$

i.e., sets in the negative border of $\mathcal{F}(r, 0.02)$, were evaluated and found to be non-frequent.

Now assume a slightly different situation, where the set $\{B, F\}$ turns out to be frequent in r , that is, $\{B, F\}$ is a miss. The set $\{A, B, F\}$ could be frequent in r , since all its subsets are. In this case Algorithm 9.1 reports that there possibly is a failure. \square

The problem formulation is now the following: given a database r and a frequency threshold min_fr , use a random sample s to determine a collection \mathcal{S} of sets such that \mathcal{S} contains with a high probability the collection of frequent sets $\mathcal{F}(r, min_fr)$. For efficiency reasons, a secondary goal is that \mathcal{S} does not contain unnecessarily many other sets.

In the fraction of cases where a possible failure is reported, all frequent sets can be found by making a second pass over the database. Algorithm 9.5 can be used to extend Algorithm 9.1 with a second pass in such a case. The algorithm simply computes the collection of all sets that possibly could be frequent. The parameter \mathcal{S} is the collection of frequent sets found by Algorithm 9.1, and \mathcal{S}^- is the collection of non-frequent candidates of Algorithm 9.1. The collection $\mathcal{B}d^-(\mathcal{S})$ can be located in a similar way as candidates are generated.

Theorem 9.6 Algorithm 9.5 works correctly.

Proof All sets computed and output on lines 2 and 3 are clearly frequent. To see that all frequent sets are output, consider any frequent set X and assume the contrary: X is not in \mathcal{S} after line 1. Let $Y \subseteq X$ be the smallest subset of X that is not in \mathcal{S} . Then all subsets of Y are in \mathcal{S} , and Y must be in the negative border $\mathcal{B}d^-(\mathcal{S})$. The only possible reason for Y being excluded from \mathcal{S} is that it is in \mathcal{S}^- . This is, however, a contradiction, since Y must be frequent. Thus all frequent sets are output. \square

The number of candidates in the second pass can, in principle, be too large to fit in the main memory and to be handled in one database pass. This can happen when the sample is very bad and gives inaccurate results.

9.2 Analysis of sampling

Let us now analyze the relation of sample size to the accuracy of results. We first consider how accurate the frequencies computed from a random sample are. As has been noted before, samples of reasonable size provide good approximations for frequencies of sets [2, 62]. Related work on using a sample for approximately verifying the truth of arbitrary sentences of relational tuple calculus is considered in [51].

Definition 9.7 Given an item set $X \subseteq R$ and a random sample s of a binary database over R , the *error* $e(X, s)$ is the difference of the frequencies:

$$e(X, s) = |fr(X, s) - fr(X)|,$$

where $fr(X)$ is the frequency of X in the database from which s was drawn. \square

To analyze the error, we consider sampling with replacement. The reason is that we want to avoid making other assumptions of the database size except that it is large. For sampling with replacement the size of the database has no effect on the analysis, so the results apply, in principle, on infinitely large databases. Note also that for very large databases there is practically no difference between sampling with and without replacement.

In the following we analyze the random variable $|\mathcal{M}(X, s)|$, that is, the number of rows in the sample s that contain X . The random variable has binomial distribution $B(|s|, fr(X))$, i.e., the probability of $|\mathcal{M}(X, s)| = c$, denoted $Pr[|\mathcal{M}(X, s)| = c]$, is

$$\binom{|s|}{c} fr(X)^c (1 - fr(X))^{|s|-c}.$$

First consider the necessary size of a sample, given requirements on the size of the error. The following theorem gives a lower bound for the size of

ε	δ	Sample size
0.01	0.01	27 000
0.01	0.001	38 000
0.01	0.0001	50 000
0.001	0.01	2 700 000
0.001	0.001	3 800 000
0.001	0.0001	5 000 000

Table 9.1: Sufficient sample sizes, given ε and δ .

the sample, given an error bound ε and a maximum probability δ for an error that exceeds the bound.

Theorem 9.8 Given an item set X and a random sample s of size

$$|s| \geq \frac{1}{2\varepsilon^2} \ln \frac{2}{\delta}$$

the probability that $e(X, s) > \varepsilon$ is at most δ .

Proof The Chernoff bounds give the result $\Pr[|x - np| > a] < 2e^{-2a^2/n}$, where x is a random variable with binomial distribution $B(n, p)$ [5]. For the probability at hand we thus have

$$\Pr[e(X, s) > \varepsilon] = \Pr[|fr(X, s) - fr(X)| \cdot |s| > \varepsilon |s|] \leq 2e^{-2(\varepsilon |s|)^2 / |s|} \leq \delta.$$

□

Table 9.1 gives values for the sufficient sample size $|s|$, for $\varepsilon = 0.01, 0.001$ and $\delta = 0.01, 0.001, 0.0001$. With the tolerable error ε around 0.01, samples of a reasonable size suffice. For instance, if a chance of 0.0001 for an error of more than 0.01 is acceptable, then a sample of size 50 000 is sufficient. For many applications these parameter values are perfectly reasonable. In such cases, approximate rules can be produced based on a sample, i.e., in constant time independent of the size of r . With tighter error requirements the sample sizes can be quite large.

The result above is for a given set X . The following corollary gives a result for a more stringent case: given a collection \mathcal{S} of sets, with probability $1 - \Delta$ there is no set in \mathcal{S} with error at least ε .

Corollary 9.9 Given a collection \mathcal{S} of sets and a random sample s of size

$$|s| \geq \frac{1}{2\varepsilon^2} \ln \frac{2|\mathcal{S}|}{\Delta},$$

the probability that there is a set $X \in \mathcal{S}$ such that $e(X, s) > \varepsilon$ is at most Δ .

Proof By Theorem 9.8, the probability that $e(X, s) > \varepsilon$ for a given set X is at most $\frac{\Delta}{|S|}$. Since there are $|S|$ such sets, the probability in question is at most Δ . \square

The Chernoff bound is not always very tight, and in practice the exact probability from the binomial distribution or its normal approximation are more useful.

Consider now the proposed approach to finding all frequent sets exactly. The idea was to locate a superset of the collection of frequent sets by discovering frequent sets in a sample with a lower threshold. Consider first the following simple setting: take a sample as small as possible but such that it is likely to contain all frequent sets at least once. What should the sample size be?

Theorem 9.10 Given a set X with $fr(X) \geq min_fr$ and a random sample s of size

$$|s| \geq \frac{1}{min_fr} \ln \frac{1}{\delta},$$

the probability that X does not occur in s is at most δ .

Proof We apply the following inequality: for every $x > 0$ and every real number b we have $(1 + \frac{b}{x})^x \leq e^b$. The probability that a frequent set X does not occur on a given row is at most $1 - min_fr$. The probability that X does not occur on any row is then at most $(1 - min_fr)^{|s|}$, which is further bounded by the inequality by $e^{-|s| min_fr} \leq \delta$. \square

The sample size given by the theorem is small, but unfortunately the approach is not very useful: a sample will include a lot of garbage, i.e., sets that are not frequent nor in the border. For instance, a single row containing 20 items has over a million subsets, and all of them would then be checked from the whole database. Obviously, to be useful the sample must be larger. It is likely that best results are achieved when the sample is as large as can conveniently be handled in the main memory.

We move on to the following problem. Assume we have a sample s and a collection $\mathcal{S} = \mathcal{F}(s, low_fr)$ of sets. What can we say about the probability of a failure? We use the following simple approximation. Assuming that the sets in $\mathcal{B}d^-(\mathcal{S})$ are independent, an upper bound for the probability of a failure is the probability that at least one set in $\mathcal{B}d^-(\mathcal{S})$ turns out to be frequent in r .

This approximation tends to give too large probabilities. Namely, a set X in $\mathcal{B}d^-(\mathcal{S})$ that is frequent in r , i.e., an X that is a miss, does not necessarily indicate a failure at all. In general there is a failure only if the addition of X to \mathcal{S} would add sets to the negative border $\mathcal{B}d^-(\mathcal{S})$; often several additions to \mathcal{S} are needed before there are such new candidates. Note also that the assumption that sets in $\mathcal{B}d^-(\mathcal{S})$ are independent is unrealistic.

An interesting aspect is that this approximation can be computed on the fly when processing the sample. Thus, if an approximated probability of a failure needs to be set in advance, then the frequency threshold low_fr can be adjusted at run time to fit the desired probability of a miss. A variation of Theorem 9.8 gives the following result on how to set the lowered frequency threshold so that misses are avoided with a high probability.

Theorem 9.11 Given a frequent set X , a random sample s , and a probability parameter δ , the probability that X is a miss is at most δ when

$$low_fr \leq min_fr - \sqrt{\frac{1}{2|s|} \ln \frac{1}{\delta}}.$$

Proof Using the Chernoff bounds again—this time for one-sided error—we have

$$Pr[fr(X, s) < low_fr] \leq e^{-2(\sqrt{\frac{1}{2|s|} \ln \frac{1}{\delta}} |s|)^2 / |s|} = \delta.$$

□

Consider now the number of sets checked in the second pass by Algorithm 9.5, in the case of a potential failure. The collection \mathcal{S} can, in principle, grow a lot. Each independent miss can in the worst case generate as many new candidates as there are frequent sets. Note, however, that if the probability that any given set is a miss is at most δ , then the probability of l independent misses can be at most δ^l .

9.3 Experiments

We now describe the experiments we conducted to assess the practical feasibility of the sampling method for finding frequent sets. In this section we also present variants of the method and give experimental results.

Test organization We used three synthetic data sets from [3] in our tests. These databases model supermarket basket data, and they have been used as benchmarks for several association rule algorithms [2, 3, 40, 73, 77]. The central properties of the data sets are the following. There are $|R| = 1\,000$ items, and the average number T of items per row is 5, 10, or 20. The number $|r|$ of rows is approximately 100 000. The average size I of maximal frequent sets, i.e., sets in the positive border, is 2, 4, or 6. Table 9.2 summarizes the parameters for the data sets; see [3] for more details of the data generation.

We assume that the real data collections from which association rules are discovered can be much larger than the test data sets. To make the experiments fair we use sampling with replacement. This means that the

Data set name	$ R $	T	I	$ r $	Size (MB)
T5.I2.D100K	1 000	5	2	97 233	2.4
T10.I4.D100K	1 000	10	4	98 827	4.4
T20.I6.D100K	1 000	20	6	99 941	8.4

Table 9.2: Synthetic data set characteristics (T = row size on average, I = size of sets in the positive border on average).

Frequency threshold	Sample size			
	20 000	40 000	60 000	80 000
0.0025	0.0013	0.0017	0.0018	0.0019
0.0050	0.0034	0.0038	0.0040	0.0041
0.0075	0.0055	0.0061	0.0063	0.0065
0.0100	0.0077	0.0083	0.0086	0.0088
0.0150	0.0122	0.0130	0.0133	0.0135
0.0200	0.0167	0.0177	0.0181	0.0184

Table 9.3: Lowered frequency thresholds for $\delta = 0.001$.

real data collections could have been arbitrary large data sets such that these data sets represent their distributional properties.

We considered sample sizes from 20 000 to 80 000. Samples of these sizes are large enough to give good approximations and small enough to be handled in main memory. Since our approach is probabilistic, we repeated every experiment 100 times for each parameter combination. Altogether, over 10 000 trials were run.

Number of misses and database activity We experimented with Algorithm 9.1 with the above mentioned sample sizes 20 000 to 80 000. We selected the lowered threshold so that the probability of missing any given frequent set X is less than $\delta = 0.001$, i.e., given any set X with $fr(X) \geq min_fr$, we have $Pr[fr(X, s) < low_fr] \leq 0.001$. The lowered threshold depends on the frequency threshold and the sample size. The lowered threshold values are given in Table 9.3; in the computations of the lowered thresholds we used the exact probabilities obtained from the binomial distribution, not the Chernoff bounds.

Figure 9.1 shows the number of database passes for the three different types of algorithms: Algorithm 2.14, Partition, and the sampling Algorithm 9.1. The Partition algorithm [77] was discussed already shortly in Chapter 2. It is based on the idea of partitioning the database to several parts, each small enough to be handled in main memory. The algorithm




Figure 9.1: The number of database passes made by frequent set algorithms.

almost always makes two passes over the database: in the first pass, it uses a variant of Algorithm 2.14 to find frequent sets in each partition, and in the second pass it checks in the whole database all sets that were frequent in at least one partition. Each of the data points in the results shown for Algorithm 9.1 is the average value over 100 trials.

Explaining the results is easy. Algorithm 2.14 makes $L(+1)$ passes over the database, where L is the size of the largest frequent set. The Partition algorithm makes two passes over the database whenever there are any frequent sets. For Algorithm 9.1, the fraction of trials with misses is expected to be larger than $\delta = 0.001$, depending on how many frequent sets have a frequency relatively close to the threshold and are thus likely misses in a sample. The algorithm has succeeded in finding all frequent sets in one pass in almost all cases. The number of database passes made by Partition algorithm is practically twice that of Algorithm 9.1, and the number of passes of Algorithm 2.14 is up to six times that of Algorithm 9.1.

Frequency threshold	Sample size			
	20 000	40 000	60 000	80 000
0.0025	0	1	0	0
0.0050	0	1	0	1
0.0075	0	0	0	0
0.0100	0	0	0	0
0.0150	0	0	0	0
0.0200	0	0	0	0

Frequency threshold	Sample size			
	20 000	40 000	60 000	80 000
0.0050	0	2	1	1
0.0075	0	1	1	1
0.0100	1	0	1	1
0.0150	0	2	0	0
0.0200	0	0	0	0

Frequency threshold	Sample size			
	20 000	40 000	60 000	80 000
0.0100	0	0	0	0
0.0150	1	1	1	0
0.0200	0	1	0	2

Table 9.4: Number of trials with misses.

Table 9.4 shows the number of trials with misses for each data set, sample size, and frequency threshold. In each set of 100 trials, there have been zero to two trials with misses. The overall fraction of trials with misses was 0.0038. We repeated the experiment with $\delta = 0.01$, i.e., so that the probability of missing a given frequent set is at most 0.01. This experiment gave misses in fraction 0.041 of all the trials. In both cases the fraction of trials with misses was larger than δ , but of the same magnitude.

The actual amount of reduction in the database activity depends on the database storage structures. For instance, if the database has 10 million rows, a disk block contains on average 100 rows, and the sample size is 20 000, then the sampling phase could read up to 20 % of the database. An alternative for randomly drawing each row in separation is, of course, to draw whole blocks of rows to the sample. Depending on how randomly the rows have been assigned to the blocks, this method can give good or bad results. For the design and analysis of sampling methods see, e.g, [72]. The related problem of sampling for query estimation is considered in [35].

Frequency threshold	Sample size				Algorithm 2.14
	20 000	40 000	60 000	80 000	
0.0050	382 282	368 057	359 473	356 527	318 588
0.0075	290 311	259 015	248 594	237 595	188 024
0.0100	181 031	158 189	146 228	139 006	97 613
0.0150	52 369	40 512	36 679	34 200	20 701
0.0200	10 903	7 098	5 904	5 135	3 211

Table 9.5: Number of item sets considered for data set T10.I4.D100K.

The reduction in database activity is achieved at the cost of considering some item sets that Algorithm 2.14 does not generate and check. Table 9.5 shows the average number of sets considered for the data set T10.I4.D100K with different sample sizes and the number of candidate sets of Algorithm 2.14. The largest absolute overhead occurs with low thresholds, where the number of item sets considered has grown from 318 588 by 64 694 in the worst case. This growth is not significant for the total execution time since the item sets are handled entirely in main memory. The relative overhead is larger with higher thresholds, but since the absolute overheads are small the effect is negligible. Table 9.5 indicates that larger samples cause less overhead (with equally good results), but that for sample sizes from 20 000 to 80 000 the difference in the overhead is not significant.

To obtain a better picture of the relation of δ and the experimental number of trials with misses, we conducted the following test. We took 100 samples (for each frequency threshold and sample size) and determined the lowered frequency threshold that would have given misses in one out of the hundred trials. Figure 9.2 presents these results (as points), together with lines showing the lowered thresholds with $\delta = 0.01$ or 0.001 , i.e., the thresholds corresponding to miss probabilities of 0.01 and 0.001 for a given frequent set. The frequency thresholds that give misses in fraction 0.01 of cases approximate surprisingly closely the thresholds for $\delta = 0.01$. Experiments with a larger scale of sample sizes give comparable results. There are two explanations for the similarity of the values. One reason is that there are not necessarily many potential misses, i.e., not many frequent sets with frequency relatively close to the threshold. Another reason that contributes to the similarity is that the sets are not independent.

In the case of a possible failure, Algorithm 9.5 generates iteratively all new candidates and makes another pass over the database. In our experiments the number of frequent sets missed—when any were missed—was one or two for $\delta = 0.001$, and one to 16 for $\delta = 0.01$. The number of candidates checked on the second pass was small compared to the total number of item sets checked.

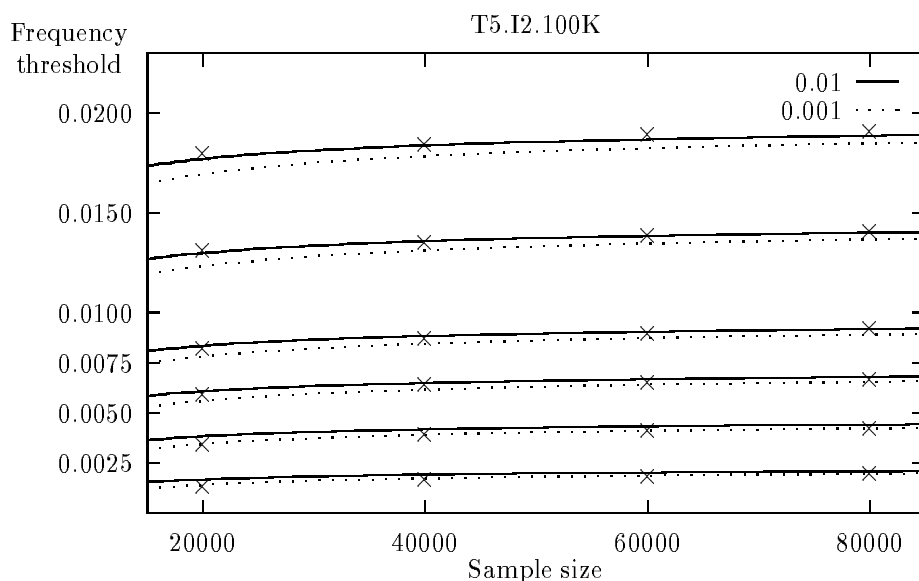


Figure 9.2: Frequency thresholds giving misses in 0.01 cases (points) and lowered thresholds with $\delta = 0.01$ and 0.001 (lines).

Approximate $1 - \Delta$ success probability Setting the lowered threshold for Algorithm 9.1 is not trivial: how to select it so that the probability of a failure is low but there are not unnecessarily many sets to check? An automatic way of setting the parameter would be desirable. Consider, for instance, an interactive mining tool. It would be useful to know in advance how long an operation will approximately take—or, in the case of mining association rules, how many database passes there will be.

We now present two algorithms that find all frequent sets in approximately fraction $1 - \Delta$ of the cases, where Δ is given by the user. Under the assumption that sets in the negative border are independent, the algorithms are actually guaranteed to find the correct sets at least in fraction $1 - \Delta$ of the cases. The first algorithm uses a simple greedy principle to find the optimal lowered threshold under the independence assumption. The other algorithm is not as optimal, but its central phase is almost identical to Algorithm 2.14 and it is therefore easy to incorporate into existing implementations. We present experimental results with this latter algorithm.

The greedy Algorithm 9.12 starts with an empty set \mathcal{S} , and it then decreases the probability of failure by adding the most probable misses to \mathcal{S} until the approximated probability of a potential failure is at most Δ .

Theorem 9.13 Algorithm 9.12 works correctly.

Algorithm 9.12

Input: A binary database r , a frequency threshold min_fr , a sample size s_size , and a miss probability Δ .

Output: The collection $\mathcal{F}(r, min_fr)$ of frequent sets and their frequencies at least in fraction $1 - \Delta$ of the cases (assuming that the frequencies of any two sets X, Y are independent whenever $X \not\subseteq Y$ and $Y \not\subseteq X$), and a subset of $\mathcal{F}(r, min_fr)$ and a failure report in the rest of the cases.

Method:

1. compute $s :=$ random sample of size s_size from r ;
 2. $\mathcal{S} := \emptyset$;
 3. // Find frequent sets in the sample:
 4. **while** the estimated probability of a miss is larger than Δ **do**
 5. select $X \in \mathcal{B}d^-(\mathcal{S})$ with the highest probability of being a miss;
 6. $\mathcal{S} := \mathcal{S} \cup \{X\}$;
 7. // Database pass:
 8. compute $\mathcal{R} := \{X \in \mathcal{S} \cup \mathcal{B}d^-(\mathcal{S}) \mid fr(X, r) \geq min_fr\}$;
 9. **for** all $X \in \mathcal{R}$ **do** output X and $fr(X, r)$;
 10. **if** $\mathcal{R} \cap \mathcal{B}d^-(\mathcal{S}) \neq \emptyset$ **then** report that there possibly was a failure;
-

Proof See the proof of Theorem 9.3 for the correctness of the output and failure reports. From the assumption of independence of sets it follows, in particular, that sets in $\mathcal{B}d^-(\mathcal{S})$ are independent. Then the probabilities can be easily computed on lines 4 and 5, and the algorithm fails in less than fraction Δ of cases. Note also that a single miss does not always indicate a failure, and therefore the probability of a miss is an upper bound for the probability of a failure. \square

The assumption of independence of sets in the algorithm is unrealistic. For this reason Δ in practice only approximates (an upper bound of) the fraction of failures. Algorithm 9.14 is a simple variation of Algorithm 2.14. It utilizes also the failure probability approximation described in Section 9.2: it monitors the approximated probability of a miss and keeps the probability small by lowering the frequency threshold low_fr , when necessary, for the rest of the algorithm. When using Algorithm 2.14 for the discovery of frequent sets in a sample with the dynamic adjustment of the lowered threshold, the only modification concerns the phase where candidates are either added to the collection of frequent sets or thrown away. Every time there is a candidate X that is not frequent in the sample, compute the probability p that X is frequent. If the total probability P of a miss increases too much (see below), then lower the frequency threshold low_fr to the frequency of X in the sample for the rest of the algorithm. Thus X is eventually considered frequent in the sample, and so are all following candidate sets that would increase the overall probability of a miss at least as much as X .

We use the following heuristic to decide whether the possibility of a miss increases too much. Given a parameter γ in $[0, 1]$, the frequency threshold

Algorithm 9.14

Input: A binary database r over a set R , a sample size s_size , a frequency threshold min_fr , a miss probability Δ , and $\gamma \in [0, 1]$.

Output: The collection $\mathcal{F}(r, min_fr)$ of frequent sets and their frequencies at least in fraction $1 - \Delta$ of the cases (assuming that the frequencies of any two sets X, Y are independent whenever $X \not\subseteq Y$ and $Y \not\subseteq X$), and a subset of $\mathcal{F}(r, min_fr)$ and a failure report in the rest of the cases.

Method:

1. compute $s :=$ random sample of size s_size from r ;
 2. $P := 0$;
 3. $low_fr := min_fr$;
 4. // Find frequent sets in the sample:
 5. $\mathcal{C}_1 := \{\{A\} \mid A \in R\}$;
 6. $l := 1$;
 7. **while** $\mathcal{C}_l \neq \emptyset$ **do**
 8. $\mathcal{R}_l := \emptyset$;
 9. **for all** $X \in \mathcal{C}_l$ **do**
 10. **if** $fr(X, s) < low_fr$ **then**
 11. $p := Pr[X \text{ is frequent in } r]$;
 12. **if** $p/(\Delta - P) > \gamma$ **then** $low_fr := fr(X, s)$
 13. **else** $P := 1 - (1 - P)(1 - p)$;
 14. **if** $fr(X, s) \geq low_fr$ **then** $\mathcal{R}_l := \mathcal{R}_l \cup \{X\}$;
 15. $l := l + 1$;
 16. compute $\mathcal{C}_l := \mathcal{C}(\mathcal{R}_{l-1})$;
 17. // Database pass:
 18. compute $\mathcal{R} := \{X \in \bigcup_l \mathcal{C}_l \mid fr(X, r) \geq min_fr\}$;
 19. **for all** $X \in \mathcal{R}$ **do** output X and $fr(X, r)$;
 20. **if** $\mathcal{R} \cap (\bigcup_l (\mathcal{C}_l \setminus \mathcal{R}_l)) \neq \emptyset$ **then** report that there possibly was a failure;
-

is lowered if the probability p is larger than fraction γ of the “remaining error reserve” $\Delta - P$. More complex heuristics for changing the frequency threshold could be developed, e.g., by taking into account the number of candidates on the level and whether the number of frequent sets per level is growing or shrinking. The observations made from Figure 9.2 hint that the lowered threshold can be set in the start-up to roughly correspond to the desired probability of a miss, i.e., for $\Delta = 0.01$ the lowered threshold could be set as for $\delta = 0.01$.

Theorem 9.15 Algorithm 9.14 works correctly.

Proof See the proof of Theorem 9.3 for the correctness of the output and failure reports. Consider the invariant that the variable P is the probability that any of the sets in the negative border is a miss (under the assumption that sets in the negative border are independent), and that $0 \leq P \leq \Delta$.

The invariant holds trivially in the beginning, where P is initialized to zero and no sets have been considered yet. We now show that the invariant continues to hold during the algorithm. On line 13, P is updated for each set X in the negative border in the sample, i.e., for each potential miss, to

T5.I2.D100K

Frequency Threshold	Sample size			
	20 000	40 000	60 000	80 000
0.005	3	3	0	2
0.010	0	0	0	0

T10.I4.D100K

Frequency threshold	Sample size			
	20 000	40 000	60 000	80 000
0.0075	1	4	2	1
0.0150	0	2	4	1

T20.I6.D100K

Frequency threshold	Sample size			
	20 000	40 000	60 000	80 000
0.01	2	1	1	1
0.02	1	3	1	3

Table 9.6: Number of trials with misses with $\Delta = 0.1$.

correctly correspond to the total probability of a miss. Given non-negative real numbers p and $\Delta - P$ such that $p/(\Delta - P) \leq \gamma \leq 1$, we have $p \leq \Delta - P$. Thus the new value of P is $1 - (1 - P)(1 - p) \leq 1 - (1 - P)(1 - (\Delta - P)) = \Delta + (P - \Delta)P$, and this is further bounded by Δ by the invariant itself. Thus the invariant continues to hold, and the probability of a miss when the program exits is $P \leq \Delta$. \square

Remember, again, that a single miss does not necessarily indicate a failure, and thus P is an upper bound for the probability of a failure. Since sets in the negative border are not necessarily independent, the upper bound is only an approximation.

We tested Algorithm 9.14 with maximum miss probability $\Delta = 0.1$ and dynamic adjustment parameter $\gamma = 0.01$ for two frequency thresholds for each data set. The number of trials with misses is shown in Table 9.6. The number successfully remained below $100\Delta = 10$ in each set of experiments. As Table 9.6 shows, the number of cases with misses was actually less than half of 10. The reason is that with a small γ the algorithm tends to be conservative and keeps a lot of space for the probability of a miss in reserve. This is useful when there can be very many candidates. The negligible trade-off is that the algorithm may consider unnecessarily many sets as frequent in the sample.

To summarize briefly, the experiments show that the proposed approach works well in practice: all frequent sets can actually be found in almost one

pass over the database. For the efficiency of mining association rules in large databases the reduction of disk I/O is significant.

Bibliography

- [1] R. Agrawal, T. Imielinski, and A. Swami. Mining association rules between sets of items in large databases. In P. Buneman and S. Jajodia, editors, *Proceedings of ACM SIGMOD Conference on Management of Data (SIGMOD'93)*, pages 207 – 216, Washington, D.C., USA, May 1993. ACM.
- [2] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. I. Verkamo. Fast discovery of association rules. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 307 – 328. AAAI Press, Menlo Park, CA, 1996.
- [3] R. Agrawal and R. Srikant. Fast algorithms for mining association rules in large databases. In *Proceedings of the Twentieth International Conference on Very Large Data Bases (VLDB'94)*, pages 487 – 499, Sept. 1994.
- [4] R. Agrawal and R. Srikant. Mining sequential patterns. In *Proceedings of the Eleventh International Conference on Data Engineering (ICDE'95)*, pages 3 – 14, Taipei, Taiwan, Mar. 1995.
- [5] N. Alon and J. H. Spencer. *The Probabilistic Method*. John Wiley Inc., New York, NY, 1992.
- [6] T. Anand. Opportunity explorer: Navigating large databases using knowledge discovery templates. *Journal of Intelligent Information Systems*, (4):27–37, 1995.
- [7] C. Berge. *Hypergraphs. Combinatorics of Finite Sets*. North-Holland Publishing Company, Amsterdam, 1989.
- [8] C. Bettini, X. S. Wang, and S. Jajodia. Testing complex temporal relationships involving multiple granularities and its application to data mining. In *Proceedings of the Fifteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database and Knowledgebase Systems (PODS'96)*, pages 68 – 78, Montreal, Canada, June 1996.

- [9] A. T. Bouloutas, S. B. Calo, A. Finkel, and I. Katzela. Distributed fault identification in telecommunication networks. *Journal of Network and Systems Management*, 3(3):295 – 312, 1995.
- [10] R. J. Brachman, T. Khabaza, W. Kloesgen, G. Piatetsky-Shapiro, and E. Simoudis. Mining business databases. *Communications of the ACM*, pages 42–48, Nov. 1996.
- [11] M. C. Burl, U. M. Fayyad, P. Perona, P. Smyth, and M. P. Burl. Automating the hunt for volcanoes on venus. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 302–309, Los Alamitos, CA, USA, June 1994. IEEE Computer Society Press.
- [12] G. Das, R. Fleischer, L. Gasieniec, D. Gunopulos, and J. Kärkkäinen. Episode matching. In *Proceedings of the 8th Annual Symposium on Combinatorial Pattern Matching (CPM '97)*, pages 12 – 27, Aarhus, Denmark, June 1997.
- [13] L. De Raedt and M. Bruynooghe. A theory of clausal discovery. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, pages 1058 – 1053, Chambéry, France, 1993. Morgan Kaufmann.
- [14] L. Dehaspe and L. De Raedt. Mining association rules in multiple relations. In N. Lavrač and S. Džeroski, editors, *Proceedings of the 7th International Workshop on Inductive Logic Programming*, volume 1297 of *Lecture Notes in Artificial Intelligence*, pages 125–132. Springer-Verlag, 1997.
- [15] L. Dehaspe, H. Toivonen, and R. D. King. Finding frequent substructures in chemical compounds. In *Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining (KDD'98)*, pages 30 – 36, New York, NY, Aug. 1998. AAAI Press.
- [16] P. J. Denning. The working set model of program behavior. *Communications of the ACM*, 11(5):323 – 333, 1968.
- [17] C. Dousson, P. Gaborit, and M. Ghallab. Situation recognition: Representation and algorithms. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, pages 166 – 172, Chambéry, France, Aug. 1993.
- [18] T. Eiter and G. Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing*, 24(6):1278 – 1304, Dec. 1995.

- [19] J. Elder IV and D. Pregibon. A statistical perspective on knowledge discovery in databases. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 83 – 113. AAAI Press, Menlo Park, CA, 1996.
- [20] U. Fayyad, D. Haussler, and P. Stolorz. Mining scientific data. *Communications of the ACM*, pages 51–57, Nov. 1996.
- [21] U. M. Fayyad, S. G. Djorgovski, and N. Weir. Automating the analysis and cataloging of sky surveys. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 471 – 494. AAAI Press, Menlo Park, CA, 1996.
- [22] U. M. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. From data mining to knowledge discovery: An overview. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 1 – 34. AAAI Press, Menlo Park, CA, 1996.
- [23] U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors. *Advances in Knowledge Discovery and Data Mining*. AAAI Press, Menlo Park, CA, 1996.
- [24] C. L. Forgy. Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence*, 19:17 – 37, 1982.
- [25] W. J. Frawley, G. Piatetsky-Shapiro, and C. J. Matheus. Knowledge discovery in databases: An overview. In G. Piatetsky-Shapiro and W. J. Frawley, editors, *Knowledge Discovery in Databases*, pages 1 – 27. AAAI Press, Menlo Park, CA, 1991.
- [26] M. Fredman and L. Khachiyan. On the complexity of dualization of monotone disjunctive normal forms. Technical Report LCSR-TR-225, Rutgers University, 1994.
- [27] T. Fukuda et al. Mining optimized association rules for numeric attributes. In *Proceedings of the Fifteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database and Knowledgebase Systems (PODS'96)*, 1996.
- [28] N. Gehani, H. Jagadish, and O. Shmueli. Composite event specification in active databases. In *Proceedings of the Eighteenth International Conference on Very Large Data Bases (VLDB'92)*, pages 327 – 338, Vancouver, Canada, Aug. 1992.

- [29] J. Goldman, P. Hong, C. Jeromnimon, G. Louit, J. Min, and P. Sen. Integrated fault management in interconnected networks. In B. Meandzija and J. Westcott, editors, *Integrated Network Management, I*, pages 333 – 344. Elsevier, 1989.
- [30] R. M. Goodman, B. E. Ambrose, H. W. Latin, and C. T. Ulmer. NOAA – An expert system managing the telephone network. In *Integrated Network Management IV*, pages 316 – 327. Chapman & Hall, London, 1995.
- [31] R. M. Goodman and H. Latin. Automated knowledge acquisition from network management databases. In I. Krishnan and W. Zimmer, editors, *Integrated Network Management, II*, pages 541 – 549. Elsevier Science Publishers B.V (North-Holland), Amsterdam, The Netherlands, 1991.
- [32] R. Grossi and F. Luccio. Simple and efficient string matching with k mismatches. *Information Processing Letters*, 33:113 – 120, 1989.
- [33] D. Gunopulos, R. Khardon, H. Mannila, and H. Toivonen. Data mining, hypergraph transversals, and machine learning. In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database and Knowledgebase Systems (PODS'97)*, pages 209 – 216, Tucson, Arizona, May 1997. ACM.
- [34] V. Gurvich and L. Khachiyan. On generating the irredundant conjunctive and disjunctive normal forms of monotone boolean functions. Technical Report LCSR-TR-251, Rutgers University, 1995.
- [35] P. J. Haas and A. N. Swami. Sequential sampling procedures for query size estimation. In *Proceedings of ACM SIGMOD Conference on Management of Data (SIGMOD'92)*, pages 341 – 350, San Diego, CA, June 1992.
- [36] T. Hagerup and C. Rüb. A guided tour of Chernoff bounds. *Information Processing Letters*, 33:305 – 308, 1989/90.
- [37] J. Han and Y. Fu. Discovery of multiple-level association rules from large databases. In U. Dayal, P. M. D. Gray, and S. Nishio, editors, *Proceedings of the 21st International Conference on Very Large Data Bases (VLDB'95)*, pages 420 – 431, Zürich, Switzerland, Sept. 1995. Morgan Kaufmann.
- [38] K. Hättönen, M. Klemettinen, H. Mannila, P. Ronkainen, and H. Toivonen. Knowledge discovery from telecommunication network alarm databases. In S. Y. Su, editor, *Proceedings of the 12th International Conference on Data Engineering (ICDE'96)*, pages 115 – 122, New Orleans, Louisiana, USA, Feb. 1996. IEEE Computer Society Press.

- [39] K. Hätönen, M. Klemettinen, H. Mannila, P. Ronkainen, and H. Toivonen. TASA: Telecommunication alarm sequence analyzer, or "How to enjoy faults in your network". In *Proceedings of the 1996 IEEE Network Operations and Management Symposium (NOMS'96)*, pages 520 – 529, Kyoto, Japan, Apr. 1996. IEEE.
- [40] M. Holsheimer, M. Kersten, H. Mannila, and H. Toivonen. A perspective on databases and data mining. In U. M. Fayyad and R. Uthurusamy, editors, *Proceedings of the First International Conference on Knowledge Discovery and Data Mining (KDD'95)*, pages 150 – 155, Montreal, Canada, Aug. 1995. AAAI Press.
- [41] P. Hoschka and W. Kloesgen. A support system for interpreting statistical data. In G. Piatetsky-Shapiro and W. J. Frawley, editors, *Knowledge Discovery in Databases*, pages 325 – 345. AAAI Press, Menlo Park, CA, 1991.
- [42] M. Houtsma and A. Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, October 1993.
- [43] T. Imielinski and H. Mannila. A database perspective on knowledge discovery. *Communications of the ACM*, 39(11):58 – 64, Nov. 1996.
- [44] G. Jakobson and M. Weissman. Real-time telecommunication network management: Extending event correlation with temporal constraints. In *Integrated Network Management IV*, pages 290 – 301. Chapman & Hall, London, 1995.
- [45] G. Jakobson and M. D. Weissman. Alarm correlation. *IEEE Network*, 7(6):52 – 59, Nov. 1993.
- [46] I. Jonassen, J. F. Collins, and D. G. Higgins. Finding flexible patterns in unaligned protein sequences. *Protein Science*, 4(8):1587 – 1595, 1995.
- [47] J. F. Jordaan and M. E. Paterok. Event correlation in heterogenous networks using OSI management framework. In H.-G. Hegering and Y. Yemini, editors, *Integrated Network Management, III*, pages 683 – 695. Elsevier Science Publishers B.V (North-Holland), Amsterdam, The Netherlands, 1993.
- [48] J. D. Kalbfleisch and R. L. Prentice. *The Statistical Analysis of Failure Time Data*. John Wiley Inc., New York, NY, 1980.
- [49] M. Kantola, H. Mannila, K.-J. Räihä, and H. Siirtola. Discovering functional and inclusion dependencies in relational databases. *International Journal of Intelligent Systems*, 7(7):591 – 607, Sept. 1992.

- [50] R. Khardon. Translating between Horn representations and their characteristic models. *Journal of Artificial Intelligence Research*, 3:349–372, 1995.
- [51] J. Kivinen and H. Mannila. The power of sampling in knowledge discovery. In *Proceedings of the Thirteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'94)*, pages 77 – 85, Minneapolis, MN, USA, May 1994. ACM.
- [52] M. Klemettinen, H. Mannila, P. Ronkainen, H. Toivonen, and A. I. Verkamo. Finding interesting rules from large sets of discovered association rules. In *Proceedings of the Third International Conference on Information and Knowledge Management (CIKM'94)*, pages 401 – 407, Gaithersburg, MD, USA, Nov. 1994. ACM.
- [53] W. Kloesgen. Efficient discovery of interesting statements in databases. *Journal of Intelligent Information Systems*, 4(1):53 – 69, 1995.
- [54] A. Krogh, M. Brown, I. S. Mian, K. Sjölander, and D. Haussler. Hidden Markov models in computational biology: Applications to protein modeling. *Journal of Molecular Biology*, 235:1501–1531, Feb. 1994.
- [55] A. Krogh, I. S. Mian, and D. Haussler. A hidden Markov model that finds genes in *e. coli* DNA. *Nucleic Acids Research*, 22:4768–4778, 1994.
- [56] P. Laird. Identifying and using patterns in sequential data. In K. Jantke, S. Kobayashi, E. Tomita, and T. Yokomori, editors, *Algorithmic Learning Theory, 4th International Workshop*, pages 1 – 18, Berlin, 1993. Springer-Verlag.
- [57] D. W. Loveland. Finding critical sets. *Journal of Algorithms*, 8:362 – 371, 1987.
- [58] H. Mannila and K.-J. Räihä. *Design of Relational Databases*. Addison-Wesley Publishing Company, Wokingham, UK, 1992.
- [59] H. Mannila and K.-J. Räihä. Algorithms for inferring functional dependencies. *Data & Knowledge Engineering*, 12(1):83 – 99, Feb. 1994.
- [60] H. Mannila and H. Toivonen. Discovering generalized episodes using minimal occurrences. In E. Simoudis, J. Han, and U. Fayyad, editors, *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD'96)*, pages 146 – 151, Portland, Oregon, USA, Aug. 1996. AAAI Press.
- [61] H. Mannila and H. Toivonen. On an algorithm for finding all interesting sentences. In *Cybernetics and Systems, Volume II, The Thirteenth European Meeting on Cybernetics and Systems Research*, pages

- 973 – 978, Vienna, Austria, Apr. 1996. Austrian Society for Cybernetic Studies.
- [62] H. Mannila, H. Toivonen, and A. I. Verkamo. Efficient algorithms for discovering association rules. In U. M. Fayyad and R. Uthurusamy, editors, *Knowledge Discovery in Databases, Papers from the 1994 AAAI Workshop (KDD'94)*, pages 181 – 192, Seattle, Washington, USA, July 1994. AAAI Press.
- [63] H. Mannila, H. Toivonen, and A. I. Verkamo. Finding association rules efficiently in sequential data. Technical Report C-1994-40, University of Helsinki, Department of Computer Science, P.O. Box 26, FIN-00014 University of Helsinki, Finland, July 1994.
- [64] H. Mannila, H. Toivonen, and A. I. Verkamo. Discovering frequent episodes in sequences. In U. M. Fayyad and R. Uthurusamy, editors, *Proceedings of the First International Conference on Knowledge Discovery and Data Mining (KDD'95)*, pages 210 – 215, Montreal, Canada, Aug. 1995. AAAI Press.
- [65] C. J. Matheus, G. Piatetsky-Shapiro, and D. McNeill. Selecting and reporting what is interesting. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 495 – 515. AAAI Press, Menlo Park, CA, 1996.
- [66] R. Milne, C. Nicol, M. Ghallab, L. Trave-Massuyes, K. Bousson, C. Dousson, J. Quevedo, J. Aguilar, and A. Guasch. TIGER: Real-time situation assessment of dynamic systems. *Intelligent Systems Engineering*, 3(3):103 – 124, 1994.
- [67] N. Mishra and L. Pitt. On bounded-degree hypergraph transversal. Manuscript, 1995.
- [68] T. M. Mitchell. Generalization as search. *Artificial Intelligence*, 18:203 – 226, 1982.
- [69] R. A. Morris, L. Khatib, and G. Ligozat. Generating scenarios from specifications of repeating events. In *Second International Workshop on Temporal Representation and Reasoning (TIME-95)*, Melbourne Beach, Florida, Apr. 1995.
- [70] S. Muggleton. *Inductive Logic Programming*. Academic Press, London, 1992.
- [71] Y. A. Nygate. Event correlation using rule and object based techniques. In *Integrated Network Management IV*, pages 278 – 289. Chapman & Hall, London, 1995.

- [72] F. Olken and D. Rotem. Random sampling from B^+ trees. In *Proceedings of the Fifteenth International Conference on Very Large Data Bases (VLDB'89)*, pages 269 – 277, Amsterdam, Aug. 1989.
- [73] J. S. Park, M.-S. Chen, and P. S. Yu. An effective hash-based algorithm for mining association rules. In *Proceedings of ACM SIGMOD Conference on Management of Data (SIGMOD'95)*, pages 175 – 186, San Jose, California, May 1995.
- [74] G. Piatetsky-Shapiro. Discovery, analysis, and presentation of strong rules. In G. Piatetsky-Shapiro and W. J. Frawley, editors, *Knowledge Discovery in Databases*, pages 229 – 248. AAAI Press, Menlo Park, CA, 1991.
- [75] G. Piatetsky-Shapiro and C. J. Matheus. The interestingness of deviations. In U. M. Fayyad and R. Uthurusamy, editors, *Knowledge Discovery in Databases, Papers from the 1994 AAAI Workshop (KDD'94)*, pages 25 – 36, Seattle, Washington, USA, July 1994. AAAI Press.
- [76] L. D. Raedt and L. Dehaspe. Clausal discovery. *Machine Learning*, 26(2):99 – 146, 1997.
- [77] A. Savasere, E. Omiecinski, and S. Navathe. An efficient algorithm for mining association rules in large databases. In *Proceedings of the 21st International Conference on Very Large Data Bases (VLDB'95)*, pages 432 – 444, Zürich, Switzerland, Sept. 1995.
- [78] P. Seshadri, M. Livny, and R. Ramakrishnan. SEQ: Design & implementation of a sequence database system. In *Proceedings of the 22nd International Conference on Very Large Data Bases (VLDB'96)*, pages 99 – 110, Mumbai, India, Sept. 1996.
- [79] A. Siebes. Data surveying, foundations of an inductive query language. In *Proceedings of the First International Conference on Knowledge Discovery and Data Mining (KDD'95)*, pages 269 – 274, Aug. 1995.
- [80] P. Smyth and R. M. Goodman. An information theoretic approach to rule induction from databases. *IEEE Transactions on Knowledge and Data Engineering*, 4(4):301 – 316, Aug. 1992.
- [81] R. Srikant and R. Agrawal. Mining generalized association rules. In U. Dayal, P. M. D. Gray, and S. Nishio, editors, *Proceedings of the 21st International Conference on Very Large Data Bases (VLDB'95)*, pages 407 – 419, Zürich, Switzerland, Sept. 1995. Morgan Kaufmann.
- [82] R. Srikant and R. Agrawal. Mining quantitative association rules in large relational tables. In *Proceedings of ACM SIGMOD Conference on*

- Management of Data (SIGMOD'96)*, pages 1 – 12, Montreal, Canada, June 1996.
- [83] R. Srikant and R. Agrawal. Mining sequential patterns: Generalizations and performance improvements. In *Advances in Database Technology—5th International Conference on Extending Database Technology (EDBT'96)*, pages 3 – 17, Avignon, France, Mar. 1996.
- [84] H. Toivonen. Sampling large databases for association rules. In T. M. Vijayaraman, A. P. Buchmann, C. Mohan, and N. L. Sarda, editors, *Proceedings of the 22nd International Conference on Very Large Data Bases (VLDB'96)*, pages 134 – 145, Mumbai, India, Sept. 1996. Morgan Kaufmann.
- [85] H. Toivonen, M. Klemettinen, P. Ronkainen, K. Hätönen, and H. Mannila. Pruning and grouping of discovered association rules. In G. N. Yves Kodratoff and C. Taylor, editors, *Workshop Notes of the ECML-95 Workshop on Statistics, Machine Learning, and Knowledge Discovery in Databases*, pages 47 – 52, Heraklion, Crete, Greece, Apr. 1995. MLnet.
- [86] J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley Publishing Company, Reading, MA, 1977.
- [87] J. T.-L. Wang, G.-W. Chirn, T. Marr, B. Shapiro, D. Shasha, and K. Zhang. Combinatorial pattern discovery for scientific data: Some preliminary results. In *Proceedings of ACM SIGMOD*, pages 115 – 125, Minneapolis, MN, May 1994.
- [88] K. Yoda, T. Fukuda, and Y. Morimoto. Computing optimized rectilinear regions for association rules. In D. Heckerman, H. Mannila, D. Pregibon, and R. Uthurusamy, editors, *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining (KDD'97)*, pages 96 – 103, Newport Beach, California, USA, Aug. 1997. AAAI Press.
- [89] R. Zembowicz and J. M. Zytkow. From contingency tables to various forms of knowledge in databases. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 329 – 349. AAAI Press, Menlo Park, CA, 1996.

