

Recognizing episodes in sequences

- first problem: given a sequence and an episode, find out whether the episode occurs in the sequence
- finding the number of windows containing an occurrence of the episode can be reduced to this
- successive windows have a lot in common
- how to use this?
- an incremental algorithm

Parallel episodes

- for each candidate α maintain a counter $\alpha.event_count$: 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
 - save the starting time of the window in $\alpha.inwindow$
- when $\alpha.event_count$ decreases again, increase the field $\alpha.freq_count$ by the number of windows where α remained entirely in the window

Algorithm

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

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

Method:

1. // Initialization:
2. for each α in \mathcal{C} do
3. for each A in α do
4. $A.count := 0$;
5. for $i := 1$ to $|\alpha|$ do $contains(A, i) := \emptyset$;
6. for each α in \mathcal{C} do
7. for each A in α do
8. $a :=$ number of events of type A in α ;
9. $contains(A, a) := contains(A, a) \cup \{\alpha\}$;
10. $\alpha.event_count := 0$;
11. $\alpha.freq_count := 0$;

Algorithm Method:

```
1. // Recognition:
2. for  $start := T_s - win + 1$  to  $T_e$  do
3.   // Bring in new events to the window:
4.   for all events  $(A, t)$  in  $s$  such that  $t = start + win - 1$  do
5.      $A.count := A.count + 1$ ;
6.     for each  $\alpha \in contains(A, A.count)$  do
7.        $\alpha.event\_count := \alpha.event\_count + A.count$ ;
8.       if  $\alpha.event\_count = |\alpha|$  then  $\alpha.inwindow := start$ ;
9.   // Drop out old events from the window:
10.  for all events  $(A, t)$  in  $s$  such that  $t = start - 1$  do
11.    for each  $\alpha \in contains(A, A.count)$  do
12.      if  $\alpha.event\_count = |\alpha|$  then
13.         $\alpha.freq\_count := \alpha.freq\_count - \alpha.inwindow + start$ ;
14.         $\alpha.event\_count := \alpha.event\_count - A.count$ ;
15.       $A.count := A.count - 1$ ;
16. // Output:
17. for all episodes  $\alpha$  in  $\mathcal{C}$  do
18.   if  $\alpha.freq\_count / (T_e - T_s + win - 1) \geq min\_fr$  then output  $\alpha$ ;
```

Theorem 1 *Algorithm 102 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. \square*

Complexity

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.

Theorem 2 *The time complexity of Algorithm 102 is $\mathcal{O}((n + l^2) |\mathcal{C}|)$.*

Proof *Initialization takes time $\mathcal{O}(|\mathcal{C}| l^2)$.*

How many accesses to α .event_count on lines 7 and 14.

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 , α .event_count is accessed at most twice during one shift.

The cost of the recognition phase is thus $\mathcal{O}(n |\mathcal{C}|)$.

□

Serial episodes

- use state automata that accept the candidate episodes
- example: episode A B A B

General episodes

different alternatives

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

Episode size	Number of episodes	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

Experiences in alarm correlation

Useful in

- finding long-term, rather frequently occurring dependencies,
- creating an overview of a short-term alarm sequence, and
- evaluating the consistency and correctness of alarm databases
- discovered rules have been applied in alarm correlation
- lots of rules are trivial

Chapter 5: Minimal occurrences of episodes

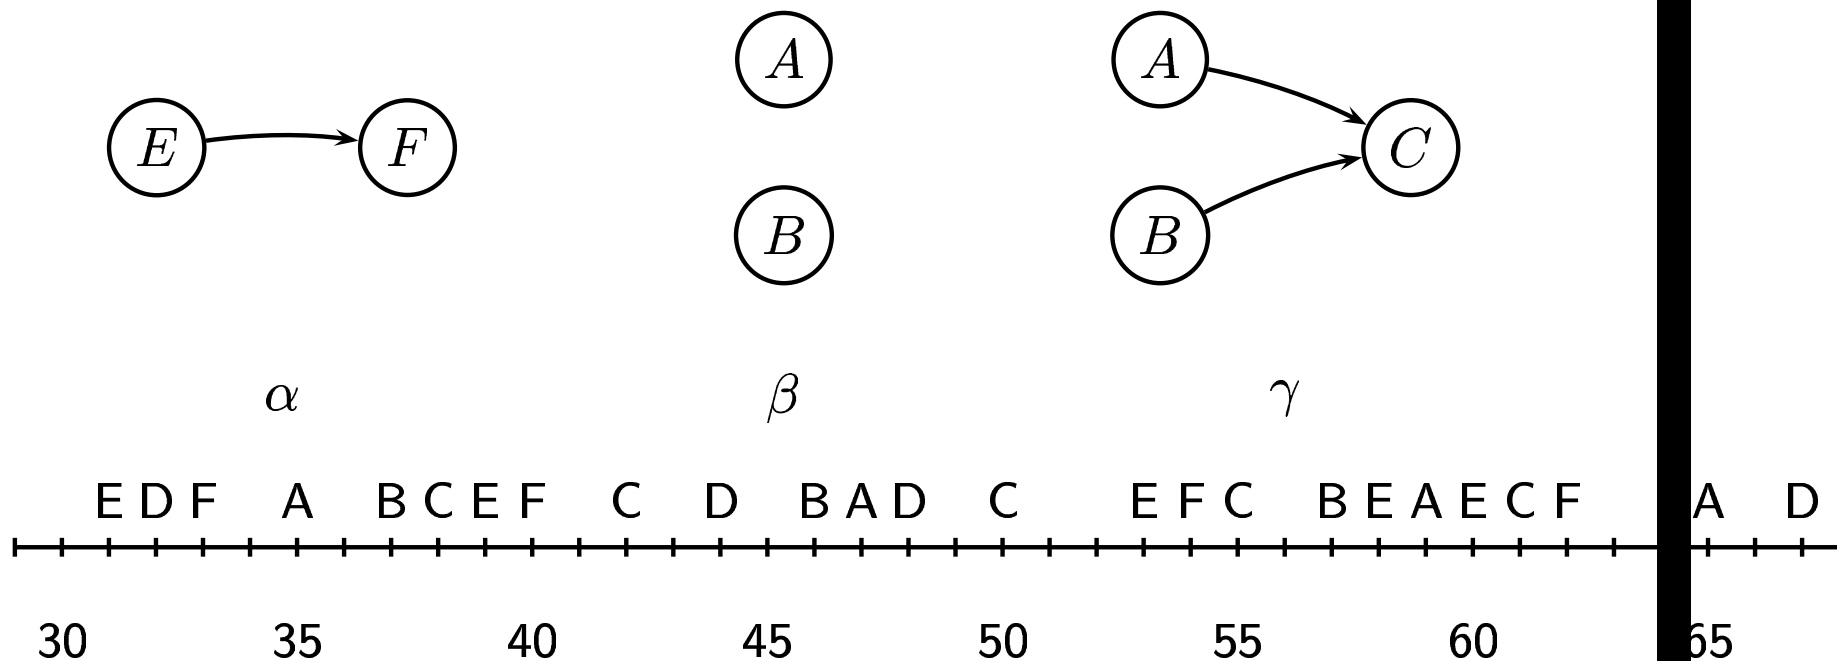
5. Minimal occurrences of episodes

- an alternative approach to discovery of episodes
- no windows
- for each potentially interesting episode, find out the exact occurrences of the episode
- advantages: easy to modify time limits, several time limits for one rule (“if A and B occur within 15 seconds, then C follows within 30 seconds”)
- disadvantages: uses lots of space

Definitions

- an episode α and an event sequence s
- interval $[t_s, t_e)$ is a *minimal occurrence* of α in s , if
 - α occurs in the window $w = (w, t_s, t_e)$ on s
 - α does not occur in any proper subwindow on w
- *set of (intervals of) minimal occurrences* of an episode α :
 $mo(\alpha) = \{ [t_s, t_e) \mid [t_s, t_e) \text{ is a minimal occurrence of } \alpha \}$.

Example



Figures 5.1: Episodes and 5.2: The example event sequence s

$$mo(\beta) = \{[35, 38), [46, 48), [47, 58), [57, 60)\}.$$

$$mo(\gamma) = \{[35, 39), [46, 51), [57, 62)\}.$$

Episodes rules, new version

- *episode rule*: $\beta [win_1] \Rightarrow \alpha [win_2]$,
- β and α are episodes such that $\beta \preceq \alpha$
- win_1 and win_2 are integers
- 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 $t'_e - t_s \leq win_2$
- (old version: $\beta [w] \Rightarrow \alpha [w]$, in windows containing β)

- formally: $mo_{win_1}(\beta) = \{[t_s, t_e) \in mo(\beta) \mid t_e - t_s \leq win_1\}$
- 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, cont.

- $\beta [3] \Rightarrow \gamma [4]$
- three minimal occurrences $[35, 38)$, $[46, 48)$, $[57, 60)$ of β of width at most 3 in the denominator
- Only $[35, 38)$, has an occurrence of α within width 4, so the confidence is $1/3$.
- rule $\beta [3] \Rightarrow \gamma [5]$ the confidence is 1.

Rule forms

- temporal relationships can be complex

Frequency and support

- previously: frequency = fraction of windows containing the episode
- no fixed window size
- several minimal occurrences within a window
- support of an episode: the number of minimal occurrences of an episode, $|mo(\alpha)|$

Rule discovery task

- an event sequence s
- a frequency threshold min_fr
- a class \mathcal{E} of episodes
- a set W of time bounds
- find all frequent episode rules of the form $\beta [win_1] \Rightarrow \alpha [win_2]$
- $\beta, \alpha \in \mathcal{E}$ and $win_1, win_2 \in W$.

Chapter 6: Episode discovery process

6. Episode discovery process

- The knowledge discovery process
- KDD process of analyzing alarm sequences
- Discovery and post-processing of large pattern collections
- TASA, Telecommunication Alarm Sequence Analyzer

The knowledge discovery process

Goal: discovery of useful and interesting knowledge

1. Understanding the domain
2. Collecting and cleaning data
3. Discovery of patterns
4. Presentation and analysis of results
5. Making conclusions and utilizing results

Pattern discovery is only a part of the KDD process (but the central one)

The knowledge discovery process

Questions implied by the KDD process model:

- How to know what could be interesting?
- How to ensure that correct and reliable discoveries can be made?
- How to discover potentially interesting patterns?
- How to make the results understandable for the user?
- How to use the results?

Episode discovery process for alarm sequences

Collecting and cleaning the data

- Can take a lot of time
- Collection of alarms rather easy
- Data cleaning? Inaccuracy of clocks
- Missing data?
- What are the event types?
 - Alarm type? Network element? A combination of the two?
- How to deal with background knowledge: network topology, object hierarchies for network elements
- “Alarm predicates”: properties of alarms

Discovery of patterns

Strategy:

1. Find *all* potentially interesting patterns
⇒ lots of rules
 2. Allow users to explore the patterns iteratively and interactively
-
1. All potentially interesting patterns
 - Episodes: combination of alarms
 - Association rules: what are alarms like
 - Frequency and confidence thresholds
 - Background knowledge coded into alarm predicates in various alternative ways
 - Network topology used to constrain patterns


Presentation and analysis of results

There can be lots of rules

- only a small part is really interesting
- – subjective
 - hard to define in advance
 - can depend on the case
- also expected regularities (or their absence) can be of interest

⇒ iteration is necessary

⇒ support for personal views is needed



Pruning and ordering:

- alarm predicates on the left or right side
- confidence, frequency, statistical significance

Structuring:

- clusters, hierarchies, etc.

TASA: A KDD tool for alarm analysis

Home Page
Produced on Wed Jun 14 19:26:21 1995

Starting time: 00:49:44-05.09.94
Ending time: 04:54:12-19.10.94
Number of alarms: 26796
Duration: 44d-4h:28s
Frequency (avg): 0.007022/s

Attribute Information | Alarm Information | Attribute Associations
Unordered Episodes | Ordered Episodes | Help Information

Unordered Episodes
Produced on Wed Jun 14 19:26:21 1995

Data is demo.seq
Start: 00:49:44-05.09.94 End: 04:54:12-19.10.94 Alarms: 26796

Template: Select Rules [] Usage Help []

Antecedent predicates: []
Consequent predicates: []

Confidence thresholds: min [] max []
Frequency thresholds: min [] max []
Significance thresholds: min [] max []

No Ordering [] Options: Ante Conse Conf. Freq. Sign.

Apply selections: [Apply] Clear selections: [Reset]

Alarms
Produced on Wed Jun 14 19:26:21 1995

Data is demo.seq
Start: 00:49:44-05.09.94 End: 04:54:12-19.10.94 Alarms: 26796

Alarm ID	Count	Frequency	Significance	Order
2064_30895	5	0.019%	2.6*10^-05	1
2064_30896	1	0.0037%	2.6*10^-07	0

- pages are created automatically from analysis results

TASA: Giving an overview of data



Alarm 1234_5678 Description
 Produced on Wed Jun 14 19:26:21 1995

Data is demo.seq
 Start: 00:49:44-05.09.94 End: 04:54:12-19.10.94 Alarms: 26796

Alarm text(s):
 FAILURE_IN_CHANNEL_ACTIVATION_OR_RESTITUTION
Count: 138
Percentage: 0.52%
Frequency: 0.00071
First time: 19.09.94-08:46:09
Last time: 18.10.94-00:10:59
Active time: 28d-15h:24m:50s
Burst: 1

[Alarm Arrival Times](#) [Distances Between Alarm Occurrences](#)

Distances between occurrences of alarm 1234_5678



statistical information, histograms

TASA: Rule presentation

The screenshot displays the TASA web interface in a Netscape browser window. The main page is titled "Unordered Episodes" and shows the following details:

- Produced on:** Wed Jun 14 19:26:21 1995
- Data is:** demo.seq
- Start:** 00:49:44-05.09.94 **End:** 04:54:12-19.10.94 **Alarms:** 26796

The interface includes several interactive elements:

- Template:** A dropdown menu for "Select Rules" and a "Usage Help" button.
- Antecedent predicates:** A text input field containing "1234".
- Consequent predicates:** A text input field containing "44*|11".
- Confidence thresholds (0.1-10%):** Min: 0.80, Max: [input field]
- Frequency thresholds (0.1-10%):** Min: 0.00003, Max: [input field]
- Significance thresholds (0.1-10%):** Min: 0.98, Max: [input field]
- Ordering:** A dropdown menu set to "Descending".
- Options:** Checkboxes for Ante, Conse, Conf, Freq, and Sign.
- Buttons:** "Apply selections: Apply" and "Clear selections: Reset".

The right-hand pane displays a table of discovered rules:

Antecedent	Consequent	Conf	Frequency	Sign
1234_44545	1234_66656	0.59	0.0000157	1.00
1234_44545	1234_11095	0.58	0.0000152	0.99
1234_44545	6789_44545	1.00	0.0000262	0.88
1234_44545	6789_66656	0.59	0.0000157	0.97
1234_44545	6789_11095	0.58	0.0000152	0.99
1234_44545	3245_44545	0.72	0.0000189	0.79
1234_44545	3245_66656	0.35	0.0000094	-
1234_44545	3245_11095	0.34	0.0000089	-
1234_31608	6789_31608	0.77	0.0000367	0.99
1234_31608	3245_31608	0.69	0.0000330	0.88
1234_66656	1234_44545	0.59	0.0000157	0.99
1234_66656	1234_11095	0.58	0.0000152	0.99
1234_66656	6789_44545	0.59	0.0000157	0.89
1234_66656	6789_66656	1.00	0.0000262	0.99
1234_66656	6789_11095	0.58	0.0000152	0.97
1234_66656	3245_44545	0.35	0.0000094	-
1234_66656	3245_66656	0.72	0.0000189	1.00
1234_66656	3245_11095	0.34	0.0000089	-
1234_11095	1234_44545	0.96	0.0000152	0.99
1234_11095	1234_66656	0.96	0.0000152	0.89
1234_11095	6789_44545	0.96	0.0000152	1.00
1234_11095	6789_66656	0.96	0.0000152	1.00

episode and association rules, views, histograms

TASA: Views with templates

Template:

Select Rules

Usage
Help

Antecedent predicates:

Consequent predicates:

Confidence thresholds:

(0.1=10%)

min

max

Frequency thresholds:

(0.1=10%)

min

max

Significance thresholds:

(0.1=10%)

min

max

- select/prune rules by their contents
⇒ iteration!
- criteria: left-hand/right-hand side of the rule, thresholds

Chapter 7: Generalized framework

7. Generalized framework

- given a set of patterns, a selection criterion, and a database
- find those patterns that satisfy the criterion in the database
- what has to be required from the patterns
- a general levelwise algorithm
- analysis in Chapter 8

Relational databases

- a *relation schema* R is a set $\{A_1, \dots, A_m\}$ of *attributes*.
- each attribute A_i has a *domain* $Dom(A_i)$
- a *row* over a R is a sequence $\langle a_1, \dots, a_m \rangle$ such that $a_i \in Dom(A_i)$ for all $i = 1, \dots, m$
- the i th value of t is denoted by $t[A_i]$
- a *relation* over R is a set of rows over R
- a *relational database* is a set of relations over a set of relation schema (the *database schema*)

Discovery task

- \mathcal{P} is a set of *patterns*
- q is a *selection criterion*, i.e., a predicate
 $q : \mathcal{P} \times \{\mathbf{r} \mid \mathbf{r} \text{ is a database}\} \rightarrow \{\text{true, false}\}$.
- φ is *selected* if $q(\varphi, \mathbf{r})$ is true
- *frequent* as a synonym for “selected”.
- give a database \mathbf{r} , the *theory* $\mathcal{T}(\mathcal{P}, \mathbf{r}, q)$ of \mathbf{r} with respect to \mathcal{P} and q is $\mathcal{T}(\mathcal{P}, \mathbf{r}, q) = \{\varphi \in \mathcal{P} \mid q(\varphi, \mathbf{r}) \text{ is true}\}$.

Example

finding all frequent item sets

- a set R a binary database r over R , a frequency threshold min_fr
- $\mathcal{P} = \{X \mid X \subseteq R\}$,
- $q(\varphi, r) = \text{true}$ if and only if $fr(\varphi, r) \geq min_fr$

Selection predicate

- no semantics given for the patterns
- selection criterion takes care of that
- “ $q(\varphi, \mathbf{r})$ is true” can mean different things:
- φ occurs often enough in \mathbf{r}
- φ is true or almost true in \mathbf{r}
- φ defines, in some way, an interesting property or subgroup of \mathbf{r}
- determining the theory of \mathbf{r} is not tractable for arbitrary sets \mathcal{P} and predicates q

Methodological point

- 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
- e.g., association rules or episode rules

Specialization relation

- \mathcal{P} be a set of patterns, q a selection criterion over \mathcal{P}
- \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
- $\theta \preceq \varphi$, then φ is said to be *more special* than θ and θ to be *more general* than φ
- $\theta \prec \varphi$: $\theta \preceq \varphi$ and not $\varphi \preceq \theta$
- the set inclusion relation \subseteq is a specialization relation for frequent sets

Generic levelwise algorithm

- 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 $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q) = \{\varphi \in \mathcal{T}(\mathcal{P}, \mathbf{r}, q) \mid level(\varphi) = l\}$.

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 $\mathcal{T}_l(\mathcal{P}, \mathbf{r}, q)$;

Theorem 7.7 Algorithm 7.6 works correctly.

