

©2005 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Emergent Intertransaction Association Rules for Abnormality Detection in Intelligent Environments

Sebastian Lühr, Svetha Venkatesh, Geoff West

Department of Computing, Curtin University of Technology
Kent Street, Bentley, Western Australia, {luhrs, svetha, geoff}@cs.curtin.edu.au

Abstract

This paper is concerned with identifying anomalous behaviour of people in smart environments. We propose the use of emergent transaction mining and the use of the extended frequent pattern tree as a basis. Our experiments on two data sets demonstrate that emergent intertransaction associations are able to detect abnormality present in real world data and that both short and long term behavioural changes can be discovered. The use of intertransaction associations is shown to be advantageous in the detection of temporal association anomalies otherwise not readily detectable by traditional "market basket" intratransaction mining.

1. INTRODUCTION

Population ageing is a modern phenomenon that sees a growing proportion of the global population becoming aged 60 years and older [1]. An expected consequence of this trend is an increase in the number of people choosing to age at home in order to maintain an independent lifestyle for as long as possible [2].

To help accommodate this growing number, we seek to create an intelligent environment that is capable of providing both cognitive support to its occupant and assistance in the event of an accident or a significant change in a person's normal behaviour. This work, the first of a two stage process, aims to identify the existence of anomalous behaviour that may require further investigation in the absence of a representative abnormal data set. We propose the use of emergent intertransaction association rules as a novel approach to this problem.

Intertransaction association rule mining [3] is an extension of the *intratransaction association rule mining* problem [4] to include the discovery of relationships, or itemsets, spanning transactions in one or more arbitrary dimensions. In our application, we wish to find the temporal associations among state-change sensor events recorded in the homes of two volunteer subjects [5], [6]. The emergent intertransaction associations are defined as those rules that display strong growth in a database of new observations given some historical data set. Discovery of such rules indicates the presence of new or unexpected increases in the frequency of behaviours and signals the need for further investigation or intervention by an occupant's supervisor or guardian.

2. RELATED WORK

Our work is similar to Jumping Emerging Patterns (JEPs) [7] and Emerging Patterns (EPs) [8] which seek to find those itemsets that display significant growth between sets of data. Emerging patterns can be used in the detection of trends and have application in classification problems due to their ability to find the discriminative features between classes of data. The difference between the EP mining problem and our work is that the former is concerned only with finding those itemsets that exhibit strong growth whereas we are interested in retaining the associative relationship among the items in our rule base.

A good example of the application of data mining to the detection of abnormality is given in [9]. Here association rules and frequent episodes [10] are used to build a set of audit data and guide the design of a classifier for the detection of intrusive and malicious behaviour in a network environment. In our work, we chose to employ intertransaction associations rather than frequent episodes in order to retain some of the temporal ordering of associated events.

Another approach is to identify unusual patterns in data using domain specific measures of unexpectedness based on a set of beliefs. Here, unexpectedness is both a measure of how the patterns in some new data affect the degree of belief the system has in a set of soft beliefs and the discovery of patterns that violate user defined immutable beliefs [11]. The work of [12], [13], for example, looks at the discovery of unusual patterns where unexpectedness is defined as a logical contradiction in beliefs. Our work, in contrast, looks for previously unseen patterns and significant growth in the frequency of known patterns. We do not consider the affect that new data has on existing patterns nor do we impose assumptions on what constitutes expected behaviour.

3. EMERGENT INTERTRANSACTION ASSOCIATION RULES

Multiple concurrent activities and unexpected interruptions in activity make the recognition of human behaviour as precise sequences of events difficult given the low level sensor data available to us. In our system *intratransaction* associations function as a model of the events that are expected to occur in close temporal proximity while the *intertransaction* associations capture the larger context in which the intratransaction events occur. The temporal component of the rules also provide us with a predictive element that can be used to forecast

$$growth(r|DB_H, DB_N) = \begin{cases} 0, & \text{if } sup(r|DB_H) = 0 \text{ and } sup(r|DB_N) = 0 \\ \infty, & \text{if } sup(r|DB_H) = 0 \text{ and } sup(r|DB_N) \neq 0 \\ \frac{sup(r|DB_N)}{sup(r|DB_H)}, & \text{otherwise} \end{cases} \quad (1)$$

the events we expect to see in a future interval given events already observed.

A. Intertransaction Association Rules

Intertransaction association rules are implications of the form $X \Rightarrow Y$ with the properties [14]

$$X \subseteq E, Y \subseteq E \quad (2)$$

$$\exists a_i^0 \in X \quad (3)$$

$$\exists a_i^d \in Y, d > 0 \quad (4)$$

$$X \cap Y = \emptyset \quad (5)$$

where $E = \{a_1^0, a_2^0, \dots, a_i^d, \dots, a_M^w\}$ is the set of all possible extended items formed by assigning an offset d to items that occur relative to some transaction T_i in a sliding intertransaction window of size w as it is passed over a database DB while forming the extended transaction set. The database $DB = \langle T_1 T_2 \dots T_N \rangle$ is said to consist of transactions T_i ($1 \leq i \leq N$) such that $T_i(x) \in I \forall$ items x in T_i where I is the set of all items $I = \{a_1, a_2, \dots, a_i, \dots, a_M\}$ found in DB .

The support measure for an association rule r is calculated as $sup(r) = \frac{|T_{xy}|}{N}$ where $|T_{xy}|$ is the number of extended transactions containing all items in $X \cup Y$ and N is the number of transactions in DB . The confidence of a rule is calculated as $conf(r) = \frac{|T_{xy}|}{|T_x|}$ where $|T_x|$ is the number of extended transactions containing all items in X . We refer to the intratransaction items $\{a_1^0, a_2^0, \dots, a_i^0, \dots, a_M^0\}$ as intratransaction items and the extended transaction items $\{a_1^1, a_2^1, \dots, a_i^d, \dots, a_M^w\}$ as intertransaction items.

As an example, the rule $A_0, B_0, C_1, D_3 \Rightarrow E_3$ implies that if we encounter A and B in the current transaction interval, C in the next, and D three transactions from now, then E will also occur three transactions from now. If A, B, C, D and E map to the event descriptors “kitchen sink cold water on”, “kitchen sink cold water off”, “dishwasher open”, “dishwasher closed” and “dishwasher on” respectively then one possible interpretation of the rule may be that it is normal to open the dishwasher, close it again and then turn the machine on shortly after having used the cold water faucet in the kitchen sink, presumably to rinse the dishes.

B. Discovery of the Emergent Intertransaction Associations

The discovery of emergent intertransaction associations seeks to find those rules that display significant growth in a database of new observations DB_N over a historical data set DB_H . We say that an association rule r is emergent when its growth, the ratio of its support in DB_N to its support in DB_H , is greater than or equal to some threshold δ and a minimum support threshold α on DB_N has been met. We have adopted Equation 1 [8] as the growth function for this work.

Association rule mining on a database populated with many frequent items can become computationally expensive, especially at the lower support thresholds where exponential growth in the number of rules present may result in a computational explosion. In our application, it is reasonable to expect that our historical data set will be larger and contain a much greater number of associations than will be present in the new data set. It is impractical, therefore, to find the complete set of association rules for both DB_H and DB_N and then compare these sets to find the emergent association rules when the rules in DB_N are expected to be a subset of those in DB_H .

We instead employ the Extended FP-Tree (EFP-Tree) [15], an extension of the Frequent Pattern Tree (FP-Tree) [16], [17], as an intermediate representation of both DB_H and DB_N . The EFP-Tree is a tree structure of descending frequency ordered inraitem nodes that contain zero or one interitem FP-Tree subtrees where the frequency ordering of the subtrees is conditioned on the inraitem parent node. Items are placed into the EFP-Tree such that the frequent items of an arbitrary extended transaction can be restored by traversing the tree.

Intertransaction association rules are retrieved from the EFP-Tree structure using the pattern growth property, a divide and conquer approach, to recursively find the complete set of intertransaction associations without the need for candidate generation.

Item constraints [18] are applied to the mining of DB_H so that only a desired subset of rules are extracted. This provides us with the ability to “query” DB_N on the historical frequency of only those associations found to be frequent in DB_N .

Simple examples for the historical database and the new observations database are given in Table 1 and Table 3 respectively. Table 2 shows the extended itemsets that are retrieved from the historical database using a sliding intertransaction window of size $w = 2$. The extended items retrieved from the new observations database are shown in Table 4. Mining the extended transactions in Table 4 with a minimum support of $\alpha = 0.5$ reveals the associations $D_0 \Rightarrow E_0$, $A_0 \Rightarrow D_1$, $A_0 \Rightarrow E_1$ and $A_0 D_1 \Rightarrow E_1$. We summarise the support of these associations in the new and historical databases with their corresponding growth values in Table 5.

C. Minimal Emergent Rules

Only the set of minimal emergent rules were sought in this work. This was to prevent the discovery process from returning an overwhelming number of rules, the majority of which would be unlikely to offer any valuable information not already present in the minimal set. Finding the set of minimal rules was achieved by ordering the mined associations by ascending rule length and discarding those rules known to contain emergent association subsets using previously detected

TABLE 1: HISTORICAL DATABASE

Trans. ID	Time	Items
100	1	A B E
200	2	E
300	3	A D
400	4	B C D

TABLE 2: EXTENDED TRANSACTIONS RETRIEVED FROM THE HISTORICAL DATABASE IN TABLE 1

Time	Extended transaction items
1	A ₀ B ₀ E ₀ E ₁ A ₂ D ₂
2	E ₀ A ₁ D ₁ B ₂ C ₂ D ₂
3	A ₀ D ₀ B ₁ C ₁ D ₁
4	B ₀ C ₀ D ₀

emergent rules. An emergent association q is said to be non-minimal when $\exists \{r, t\}$ such that $r(t) \subset q(t=0)$ where $r(t) = \{r_1^{d_1+t}, r_2^{d_2+t}, \dots, r_i^{d_i+t}, \dots, r_Z^{d_Z+t}\}$ is a known emergent rule whose intertransaction offsets $\{d_1, d_2, \dots, d_i, \dots, d_Z\}$ have been incremented t intervals. For example, if r is an emergent rule $C_0, D_2 \Rightarrow E_2$ and q is the rule $A_0, B_0, C_1, D_3 \Rightarrow E_3$ then q is also known to be emergent because $r(1) \subset q$ holds true.

Although we applied filtering as a post process, the EFP-Tree offers an opportunity to move the filtering process into the tree mining algorithm so as to guide the mining of both DB_H and DB_N .

Only those associations whose frequency were queried in the historical data set are included in the reported rule counts in Section 4.

Only the associations $D_0 \Rightarrow E_0$, $A_0 \Rightarrow D_1$ and $A_0 \Rightarrow E_1$ from the discovered rules in Table 5 are minimal emergent. $A_0 D_1 \Rightarrow E_1$ is considered to be non-minimal as it is a superset of both $A_0 \Rightarrow D_1$ and $A_0 \Rightarrow E_1$.

4. EXPERIMENTATION AND RESULTS

We have applied our method of emergent intertransaction rule discovery on the real world data described in [5], [6]. This data set comprises state-change sensor event logs recorded over a period of sixteen days in the homes of two volunteer subjects, a working professional and an eighty year old retiree. A total of 77 and 84 sensors were installed in the apartments of the first and second subject respectively. Both participants were living independently at the time the data was gathered.

The event logs were discretised into transactions of five minute intervals and the unique sensor IDs were removed to reduce the information known about each event to include only the room in which the sensor was installed, the appliance or item of furniture it was attached to, and its state. For example, as we are only interested in a subject's general use of a chest of drawers and not the precise drawer that is opened, several sensors attached to a chest of draws are reduced to the events `Bedroom/Drawer true` and `Bedroom/Drawer false`.

The sixteen days worth of event logs from each subject were divided into two weeks, each week containing eight days of events. Each week was in turn used as the historical database

TABLE 3: NEW OBSERVATIONS DATABASE

Trans. ID	Time	Items
100	1	A B
200	2	D E
300	3	A
400	4	C D E

TABLE 4: EXTENDED TRANSACTIONS RETRIEVED FROM THE NEW OBSERVATIONS DATABASE IN TABLE 3

Time	Extended transaction items
1	A ₀ B ₀ D ₁ E ₁ A ₂
2	D ₀ E ₀ A ₁ C ₂ D ₂ E ₂
3	A ₀ C ₁ D ₁ E ₂
4	C ₀ D ₀ E ₀

DB to find the emergent rules present in the other. Discretisation of the event logs from the first subject transformed the log data into two databases of 287 and 375 transactions from the first and second week respectively. The event logs from the second subject were discretised into databases of 440 transactions from the first week of data and 319 transactions from the second.

A. Observations on the Working Professional Data Set

Intertransaction association rules from the first subject were mined using a raw minimum support threshold of $\alpha = 8$ for both weeks of data. This support level was chosen to balance the quality and the quantity of the discovered associations, providing us with an ample number of rules for analysis while remaining resistant to noise.

Association rules discovered in DB_N were classified emergent if a minimum growth of $\delta = 5$ was measured.

1) *Week One:* Mining the first week of data lead to the discovery of 101 emergent intertransaction associations from 1,808 investigated rules. Examining the context of each rule revealed the existence of four distinct groups of associations.

The first group was found to relate to a flurry of kitchen activity on the third day of the event logs. Here the subject appears to be triggering the sensors on a variety of kitchen appliances and furniture over an unusually long period, behaviour that results in significant growth in kitchen related associations. Manual examination of the emergent rules and the log data did not, however, reveal any discernible behavioural trends aside from the unusual length of the activity and the frequency of the events therein.

The second set of rules highlight significant growth in kitchen door associated events. Inspection of the event logs reveals repeated opening and closing of the door on day four of the data. Although it is possible that these readings were caused by the subject frequently entering and leaving the kitchen, we believe that the length of the activity and the regularity with which the readings occur suggests that a momentary glitch with the sensor has been detected.

Analysis of the emergent rules from the third group shows that it is unusual to observe the medicine cabinet and the foyer door being opened or closed in adjacent transactions. This discovery is attributed to insufficient historical data;

TABLE 5: SUMMARY OF THE INTERTRANSACTION RULES DISCOVERED FROM THE EXTENDED ITEMS IN TABLE 4

Intertransaction Rule r	New Support $sup(r DB_N)$	Historical Support $sup(r DB_H)$	Growth $growth(r DB_H, DB_N)$	Minimal Emergent
$D_0 \Rightarrow E_0$	0.5	0	∞	✓
$A_0 \Rightarrow D_1$	0.5	0.25	2	✓
$A_0 \Rightarrow E_1$	0.5	0.25	2	✓
$A_0D_1 \Rightarrow E_1$	0.5	0.25	2	

discretisation of the event logs resulted in the foyer door and the medicine cabinet events being predominately grouped into adjacent intervals in the first week and into single transactions in the second. We see from the event logs that it is quite normal in both weeks of data for the subject to walk through the foyer door shortly before opening the medicine cabinet.

The last set of emergent rules features the subject both flushing the toilet and using the bathroom medicine cabinet in a single interval. These two events occur normally one or two intervals apart in the historical data set but not in the same transaction. This behaviour does not appear to be due to any unusual circumstances and again suggests that a larger historical data set is required.

2) *Week Two*: Only 46 of the 2,437 intertransaction associations tested from the second week were identified as emergent. Inspection of these rules again revealed four groups of contextually related associations.

The first group is a set of two single events revealing that the bathroom door is being opened and shut significantly more often in the second week than in the first. This innocuous new behaviour is evident over several days of data and is confirmed through manual inspection of the event logs.

Twenty four rules associating use of the microwave with use of the fridge, the freezer and the kitchen cabinet within a single transaction period make up the second group of associations. The historical data shows that while we can expect some of these kitchen based events to occur within a single interval, the use of the microwave in addition to two or more of these events all within a single transaction is unusual. Examination of the event logs showed that meal preparations appear to be a more involved process in the second week than in the first.

The third set of emergent rules suggest that it is abnormal for the hot and cold water faucets in the bathroom to be turned both on and off in adjacent intervals. It appears that in the first week the hot and cold water is used in shorter sessions and that these events normally occur within a single transaction. The event logs suggest that the subject does appear to be spending more time in the bathroom and is opening the water faucets more frequently than in the first week. It remains conceivable, however, that this phenomenon is caused by a sensor aberration and not a change in behaviour.

Although the rules in the final set of emergent associations highlight no apparent abnormal behaviour, they do document significant growth in the number of times the bedroom jewellery box is opened and the bedroom light switch is activated in contiguous transactions.

B. Observations on the Retiree Data Set

The minimum support threshold for mining the first week of data was reduced to $\alpha = 6$ in order to provide us with a sufficient number of rules for analysis not otherwise available at higher support levels. The support level for the second week was again set to $\alpha = 8$ and a minimum growth measurement of $\delta = 5$ was required for both weeks.

1) *Week One*: Of the 1,061 intertransaction associations tested, 72 were classified emergent. The emergent rules were again grouped into sets of related associations.

The first of these groups reveal that the subject is watching more television in the first week than in the second. The log files confirm that interaction with the TV is common in the first week yet it is rarely seen in the second.

A second group of associations highlight an anomaly centered around the repeated opening and closing of the kitchen door one and two intervals after it had previously been opened or closed. Delving into the event logs reveals a considerable increase in the number of times the door is opened in this week relative to the historical data set. It appears that the subject is carrying out activities that require frequent access to the kitchen over several days of data and that this is not due to a sensor malfunction as was previously discovered in Section 4-A.1.

The next group of emergent rules highlight an unexpected behaviour prevalent on the first and sixth day. We see from the log data that the bathroom door is frequently opened and closed over a fifteen minute interval on these days yet the historical data suggests that the bathroom door is normally left open or shut for periods of time outside the span of the intertransaction sliding window.

A bevy of emergent rules covering significant growth in the relationships among a variety of objects in the kitchen make up the final set. Emergent rules highlight growth between use of the microwave and the garbage disposal unit, the repeated opening and closing of the kitchen door and repeated access to the fridge and microwave. We found it difficult to locate any anomalous behaviour in the log files as all these events also appear in close proximity in the second week. This suggests that we do not have sufficient historical data to account for all the combinations of intervals and events found here.

2) *Week Two*: The presence of four distinct groups of associations were once again identified amongst the 23 emergent associations discovered from the 758 that were tested.

Two single events from which we learn that the subject accesses their study drawer more frequently in the second week than in the first makes up the first group of emergent associations. Inspection of the logs confirms this find, the study

drawer being opened and closed in seven transactions in week two but only twice in week one.

The rules in the second group imply abnormal behaviour in the form of the shower faucet being repeatedly turned on and off taking place on the first day of this week. Although it is possible that these rules were found due to some new behaviour, the frequency and regularity of the readings over a period of fifteen minutes suggests the possibility that they were the result of an error with the sensors.

The next unusual behaviour is a repeated toggling of the light switch in the butler's pantry on day seven. The emergent rules show that the light switch in the butler's pantry is being turned on for several minutes and then turned off again in a cycle that is repeated several times every few minutes. The expected behaviour associated with the butler's pantry is to see the light switch activated for several minutes up to an hour at a time with longer periods of time passing before the light is turned on again. A single rule in this group highlights frequent use of the microwave twenty five to thirty minutes after the cabinet inside the butler's pantry has been accessed. It is difficult to verify whether these patterns are caused by a new behaviour or not without querying the subject.

The final emergent rules present five new associations related to the use of the kitchen cabinet. The first three of these rules show a new association between the cabinet and the microwave, the latter of which is frequently being opened and closed twenty five to thirty five minutes after the cabinet has been accessed. The last two rules show the closing of the cabinet door, the use of the cold water faucet and the opening and closing of the fridge all within a single transaction is an unexpected but frequent occurrence in the second week. None of these rules suggest a new behaviour that we need to be concerned about.

5. CONCLUSION

This work has introduced a method for the mining of emergent intertransaction association rules and demonstrated their application through the discovery of behavioural changes and sensor aberrations present in the state-change sensor event logs from the homes of two volunteer subjects.

Experimental results show that emergent associations are able to be applied to the discovery of short term abnormalities and to detect changes that occur in a subject's behavioural patterns over a span of several days. Both intratransaction and intertransaction anomalies were detected, attesting to the benefit of incorporating temporal associative relationships into the mining process.

ACKNOWLEDGMENTS

The authors wish to thank Emmanuel Munguia Tapia, Stephen Intille and Kent Larson for the real world data used in this work.

REFERENCES

- [1] United Nations, *World population ageing: 1950-2050*, United Nations Publications, New York, 2002.

- [2] E. D. Mynatt, I. Essa, W. Rogers, "Increasing the opportunities for aging in place", In: *Proc. 2000 Conf. Universal Usability*, Arlington, Virginia, USA, 2000, pp. 65-71.
- [3] H. Lu, L. Feng, J. Han, "Beyond intra-transaction association analysis: mining multi-dimensional inter-transaction association rules", In: *ACM Trans. Information Systems*, Vol. 18, No. 4, 2000, pp. 423-454.
- [4] R. Agrawal, T. Imielinski, A. Swami, "Mining association rules between sets of items in large databases", In: *Proc. ACM-SIGMOD 1993 Int'l Conf. Management of Data*, Washington, District of Columbia, USA, May, 1993, pp. 207-216.
- [5] E. M. Tapia, *Activity Recognition in the Home Setting Using Simple and Ubiquitous Sensors*, Masters thesis, Massachusetts Institute of Technology, 2003.
- [6] E. M. Tapia, S. S. Intille, K. Larson, "Activity recognition in the home using simple and ubiquitous sensors", In: *Int'l Conf. Pervasive Computing*, Vol. 3001 of *Lecture Notes in Computer Science*, 2004, pp. 158-175.
- [7] G. Dong, J. Li, X. Zhang, "Discovering jumping emerging patterns and experiments on real datasets", In: *Int'l Database Conf. Heterogeneous and Internet Databases*, Hong Kong, China, 1999, pp. 155-168.
- [8] G. Dong, J. Li, "Efficient mining of emerging patterns: discovering trends and differences", In: *Int'l Conf. Knowledge Discovery and Data Mining*, San Diego, California, USA, 1999, pp. 43-52.
- [9] W. Lee, S. J. Stolfo, K. W. Mok, "Adaptive intrusion detection: a data mining approach", In: *Artificial Intelligence Review*, Vol. 14, No. 6, 2000, pp. 533-567.
- [10] H. Mannila, H. Toivonen, A. I. Verkamo, "Discovering frequent episodes in sequences", In: *Proc. Int'l Conf. Knowledge Discovery and Data Mining*, Montréal, Canada, 1995, pp. 210-215.
- [11] A. Silberschatz, A. Tuzhilin, "What makes patterns interesting in knowledge discovery systems", In: *IEEE Trans. Knowledge and Data Engineering*, Vol. 8, No. 6, 1996, pp. 970-974.
- [12] B. Padmanabhan, A. Tuzhilin, "A belief-driven method for discovering unexpected patterns", In: *Int'l Conf. Knowledge Discovery and Data Mining*, New York City, New York, USA, August, 1998, pp. 94-100.
- [13] B. Padmanabhan, A. Tuzhilin, "Unexpectedness as a measure of interestingness in knowledge discovery", In: *Decision Support Systems*, Vol. 27, No. 3, 1999, pp. 303-318.
- [14] A. K. H. Tung, H. Lu, J. Han, L. Feng, "Efficient mining of inter-transaction association rules", In: *IEEE Trans. Knowledge and Data Engineering*, Vol. 15, No. 1, 2003, pp. 43-56.
- [15] S. Lühr, G. West, S. Venkatesh, "An extended frequent pattern tree for intertransaction association rule mining", Technical Report 2005/1, Department of Computing, Curtin University of Technology, 2005.
- [16] J. Han, J. Pei, Y. Yin, "Mining frequent patterns without candidate generation", In: *ACM SIGMOD Int'l Conf. Management of Data*, Dallas, Texas, USA, May, 2000, pp. 1-12.
- [17] J. Han, J. Pei, Y. Yin, R. Mao, "Mining frequent patterns without candidate generation: a frequent-pattern tree approach", In: *Data Mining and Knowledge Discovery*, Vol. 8, No. 1, 2004, pp. 53-87.
- [18] R. Srikant, Q. Vu, R. Agrawal, "Mining association rules with item constraints", In: *Int'l Conf. Knowledge Discovery and Data Mining*, Newport Beach, California, USA, August, 1997, pp. 67-73.