Analysing frequent sequential patterns of collaborative learning activity around an interactive tabletop

R. MARTINEZ, K. YACEF, J. KAY,
School of Information Technologies, University of Sydney, Australia
A. AL-QARAGHULI,
School of Computing Science, Newcastle University, UK and
Faculty of Information Science and Technology, UKM, Malaysia
AND
A. KHARRUFA
School of Computing Science, Newcastle University, UK

Electronic traces of activity have the potential to be an invaluable source to understand the strategies followed by groups of learners working collaboratively around a tabletop. However, in tabletop and other co-located learning settings, high amounts of unconstrained actions can be performed by different students simultaneously. This paper introduces a data mining approach that exploits the log traces of a problem-solving tabletop application to extract patterns of activity in order to shed light on the strategies followed by groups of learners. The objective of the data mining task is to discover which frequent sequences of actions differentiate high achieving from low achieving groups. An important challenge is to interpret the raw log traces, taking the user identification into account, and pre-process this data to make it suitable for mining and discovering meaningful patterns of interaction. We explore two methods for mining sequential patterns. We compare these two methods by evaluating the information that they each discover about the strategies followed by the high and low achieving groups. Our key contributions include the design of an approach to find frequent sequential patterns from multiuser co-located settings, the evaluation of the two methods, and the analysis of the results obtained from the sequential pattern mining.

Keywords and Phrases: Collaborative Learning, Sequence Mining, Hierarchical Clustering, Interactive Tabletops

---

1. INTRODUCTION

Recently, the need to explore, share and manipulate tangible data, in situ, has brought forth the development of new user interfaces offering large display areas and multiple input capabilities. These groupware interfaces are becoming available for educational purposes in the form of whiteboards, multi-display settings and horizontal tabletops. Interactive tabletops offer the potential for new ways to support collaborative learning activities by enabling face to face interactions between students and, at the same time, providing a great opportunity to investigate groups' learning processes by capturing their physical actions. This paper reports our work in the context of Digital Mysteries [Kharrufa et al. 2010], a tabletop collaborative learning tool for the development of students' problem-solving skills. When using this tool, students have to examine the information they are provided with and formulate an answer to a posed question (the mystery). The students' cognitive processes become evident through their physical

Authors' addresses: R. Martinez, J. Kay, K. Yacef, School of Information Technologies, The University of Sydney, Australia. E-mail: {roberto,judy,kalina}@it.usyd.edu.au; A. Al-Qaraghuli, Faculty of Information Science and Technology, UKM, Malaysia. E-mail: aalqaraghuli@gmail.com. A. Kharrufa, School of Computing Science, Newcastle University, UK. E-mail: ahmed@diwan.com.

.

manipulation of the information on the tabletop to solve the mystery and thus observable for researchers [Leat and Nichols 2000]. However, when a class of typical size (20 to 30 students) is divided into several small groups working in parallel, it is very difficult for facilitators to keep track of the learning processes followed by all the groups and they usually end up just looking at the final results. This is a problem as it means that the higher level strategies followed by groups are lost. The work described in this paper addresses this problem by mining and analysing frequent sequences of activity and highlighting key differences between high and low achieving groups.

The use of Data Mining techniques in collaborative learning environments has proven successful in getting insights on the interactions within groups that lead to high-quality results in terms of collaboration [Anaya and Boticario 2011; D'Mello et. al. 2011], conflict resolution [Prata et al. 2009], teamwork [Perera et. al. 2009)] and *correctness* of the task [Talavera and Gaudioso 2004]. However, most of these efforts have focused on studying collaboration supported by online learning systems (e.g. chat, forums, wikis, networked ITS's) rather than tackling the context of supporting small groups collaborating around shared devices, for which there is much less research [Jeong and Hmelo-Silver 2010]. In this paper we focus on the latter. We report our work on the analysis of groups' interactions with the resources at the tabletop and the exploration of two different approaches to consider the raw physical touch actions. We detail these on a technical level and then discuss the patterns resulting from each of them.

This paper is organised as follows. Next section describes other studies that have applied machine learning techniques to analyse groups' interactions. Section 3 introduces the tabletop system and dataset. Section 4 explains the data mining methods. We conclude with reflections and future work in sections 5 and 6.

2. RELATED WORK

A number of research projects have studied the collaborative learning processes applying artificial intelligence techniques; however, they have focused mostly on assisting groups in online learning activities. Talavera and Gaudioso [2004] applied clustering in e-learning data to build student profiles based on the interactions with the user interface performed by the students. Anaya and Boticario [2011] acutely described a method to classify learners according to their level of collaboration using clustering and decision trees. Prata et.al. [2009] presented an automated detector of the nature of the utterances written at a math online system in terms of collaboration focusing on the identification of conflict between peers.

Additionally, several researchers have specifically addressed the analysis of collaboration using sequential pattern extraction. Perera et. al. [2009] modelled key aspects of teamwork on groups working with an online project management system. They clustered groups and learners according to quantitative indicators of activity and also proposed the use of alphabets to represent sequential patterns of interactions that can distinguish strong from weak groups. Other techniques have also been used to mine sequential patterns from collaborative data including Hidden Markov Models [Soller and Lesgold 2007], Social Network Analysis [Casillas and Daradoumis 2009] and Process Mining [Reimann et al. 2009].

In terms of co-located collaboration, Martinez et. al. [2011a] proposed a method to discern the extent of collaboration in groups of learners solving an optimisation problem in a multi-display face-to-face setting. The authors also applied a set of techniques to derive a user model of collaboration from a co-located multi-display setting. This also proved give information about the extent of communication and collaboration of students at the tabletop [Martinez et. al. 2011b]. The work reported in this paper is the first effort we are aware of that has made use of data mining techniques to analyse and discover patterns of interaction from data generated by a multi-user tabletop educational application.

3. THE TABLETOP TASK: DIGITAL MYSTERIES

Digital Mysteries is a collaborative learning tool for the development and assessment of students' higher level thinking skills [Kharrufa et al. 2010]. The task provided to the students is to solve a mystery with an open question in any subject such as mathematics, history, or physics. Students are given the question and a number of data slips which may hold direct clues for solving the mystery, background information, or even red-herrings. They are asked to analyse these to formulate their answer to the question. Among the main design concepts behind the original paper-based mysteries tool [Leat and Nichols 2000] is that the students' cognitive processes become evident through their physical manipulation of these data slips to solve the mystery.

Digital Mysteries divides the task of solving a mystery into *three stages* and provides a set of externalisation tools at each of these. i) For the first "*information gathering*" stage, users are provided with 20-26 data slips. Initially, these slips are displayed in a minimised pictorial form to save space at the tabletop. Consequently, users have to expand them to read the contained clues (see Figure 1, right). ii) For the second "*grouping*" stage, students are provided with a tool for creating "named" groups of slips and they are asked to categorise the slips into meaningful groups. Students usually create groups in support of or against a particular claim, or groups containing information related to a particular person, topic, or event. Students move to the next stage after putting all the slips into a minimum of four named groups. iii) For the third and last "*sequencing and webbing*" stage, students are asked to use a sticky tape tool to build a branched structure that reflects cause-and-effect relations and time sequences embodying the students' answer to the question. After completing this stage, students are asked to write down their answer.

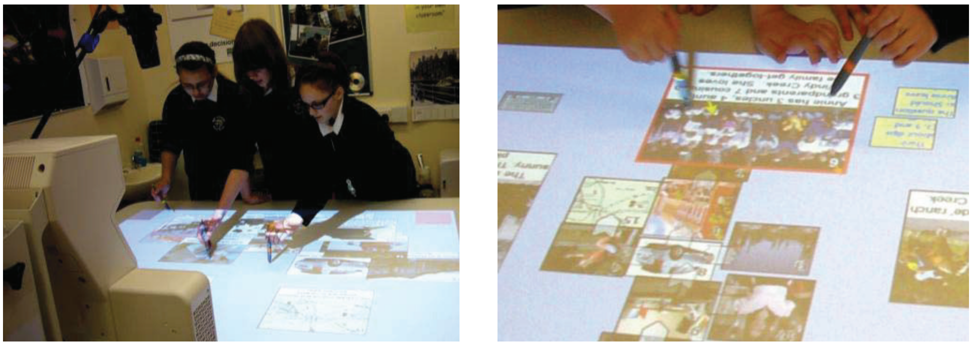

Fig. 1. Left: Three children solving a Digital Mystery. Right: Participants reading a clue

Digital Mysteries was implemented using a prototype of the multi-pen horizontal Promethean Activboard[1]. Using a pen-based tabletop makes it possible to identify the author of each action. In this way, Digital Mysteries captures a rich set of interaction data throughout the mystery solution process that includes user identification or authorship as we will refer to in the rest of this paper.

---

[1] Promethean Interactive Whiteboards: http://www.prometheanworld.com/

Participants and data collection. Every action on the tabletop was logged and all sessions were video recorded. The study involved 18 participants, forming 6 groups of 3 participants each (see Figure 1, left). Some of the groups solved more than one mystery, generating a total of 12 logged sessions. Participants were elementary school students aged between 11 and 14 years. Each group was asked to find the answer to a mystery. They had to read and understand the clues, cluster them into meaningful groups, discuss which clues were related with each other and formalise a response to the mystery. Triads performed between 970 and 2017 actions per session, for a total of 17130 logged actions.

Data exploration. The raw data was coded as a series of Events, where Event= {Time, Author, Action, Object}. The possible actions that can be performed on the data slips are: moving (M), enlarging to maximum size (E), resizing to medium size (N), shrinking (S), Rotating (M), making unions with other data slips (U), add data slip to a group (G) and remove a data slip from a group (R). Out of the 12 sessions, 5 were coded as low achieving groups of students, 5 as high achieving groups and 2 as average groups. The level of achievement was coded considering: the quality of the discussions, the degree of logic thinking and the soundness of the justification for the solution of the mystery. A full report of this analysis can be found in [Kharrufa 2010]. We focus from now on the 10 groups that clearly showed evidence of low orhigh achievement.

4. MINING AND CLUSTERING SEQUENTIAL PATTERNS

From a Data Mining perspective, the dataset collected from our co-located setting poses challenges to general data mining techniques. A first challenge is that there is a diversity of spontaneous actions that can be performed when using a tabletop as opposed to online systems, such as wikis or forums, in which learners have more time to think their actions. As a result, our data might contain more non-relevant human-computer interaction events. The second challenge is the especial importance of the authorship of the low level events performed on Digital Mysteries. To address these issues we have set out to attend two research questions: i) what are the key insights that can be gained from raw and compact logged actions? (e.g. consider N similar actions as a group of actions rather than N individual actions), and ii) what information can be obtained by including authorship information in the post-processing stage of data mining?

The data mining task we set out to solve is to discover sequences of interactions between group members and the data slips at the tabletop that were more frequent in high-achieving groups than in low-achieving ones, and vice-versa. Two important attributes of our data are the sequential order and, as mentioned above, the authorship. One technique that provides insights on the timing of the events is sequential pattern mining. A sequential pattern is a consecutive or non-consecutive ordered sub-set of a sequence of events [Jiang and Hamilton 2003]. However, as noted by Perera et.al. [2009], a frequent pattern of two actions X-Y might not be meaningful if many other events or large gaps of inactivity occur between such actions. We focused on the *consecutive* ordered sub-set of events that can potentially form a pattern. We will refer to these as *frequent sub-sequence sequential patterns*. Our algorithm seeks consecutive and also repeated patterns within the dataset of sequences. A generic flow diagram of our system is shown in Figure 2 (left).

Raw dataset. Our original *raw data* consists of the events performed at the tabletop, along with the authorship information of each of these events. We present a sample excerpt from a group session log in Figure 2 (right). In Digital Mysteries each resource (data slip) provided to solve the mystery is present at the tabletop from the beginning to the end of the session. We took advantage of this to explore how learners interact with the resources at the tabletop. We first broke down each group's long and unique sequence of events into sub-sequences of actions per data-slip. Then, to preserve meaningfulness in the patterns, we broke down these data slips' sub-sequences when a gap of inactivity longer to 120 seconds was detected.

We describe the above with a short scenario: the group decide to read a data-slip D and performs actions to enlarge it (move and enlarge actions), they read the data slip, close it and re-arrange it (more moves and shrink actions); if after this sequence there is a "group action" for the same data slip, but 5 minutes later, we can assume that the "group action" is not directly related with the previous actions. We chose a gap of 120 seconds as a maximum threshold beyond which the set of actions are considered as unrelated. This time frame was chosen based on the observations made on the videos of the sessions and the log files. In summary, the raw dataset we started with as input of step 1 is a dataset of 1618 sequences generated by breaking down the actions of each session in this order: by stage, resource (*data slip*) and long inactivity gap. The length of each sequence obtained was between 4 and 40 elements. In this dataset of sequences, each sequence is related with the session, stage and resource it comes from. Each element within each sequence contains information on timing, authorship and action type.

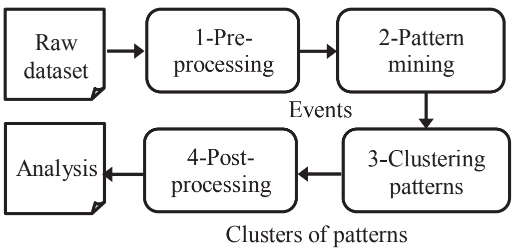

Time	Author	Action	Object
00:10	Mario	M	DataSlip01
00:12	Mario	E	DataSlip01
00:46	Mario	S	DataSlip01
00:47	Alice	M	DataSlip02
00:50	Alice	M	DataSlip01
00:51	Bob	E	DataSlip02
00:53	Mario	M	DataSlip02

Fig. 2. Left: Steps of our data mining approach. Right: Excerpt from the application logs of activity.

Step 1. We explored two pre-processing approaches: the first method consists in going straight into the sequential mining (hence a void step 1). The second method consists in compact similar contiguous actions before applying the sequence mining. Both methods are described in detail in the next section. The output of the first step for both cases is a pre-processed *dataset of sequences*.

Step 2. The *sequence mining* step is generic for both approaches. As mentioned before, our aim is to look for frequent ordered patterns within the action sequences. With the purpose of exploiting not just the frequency but also the redundancy of the patterns we are searching for, we chose an algorithm to extract the frequent sub-sequences from sequences using n-grams [Masataki and Sgisaka 1996]. An n-gram is a subsequence of n items from a given sequence. We set the minimum support threshold to consider a pattern as frequent if this was present in at least one quarter of the total number of data slips. We also set the maximum error in 1 to allow the matching of patterns with sub-sequences if there was an edit distance of 0 (perfect match) or 1 (one different action in the sub-sequence) between them. The output of this step is a list of frequent sequential patterns that meet the minimum given support.

Step 3. The purpose of step 3 is to cluster the patterns found in step 2. Indeed, without further treatment, patterns obtained from step 2 offer limited information to differentiate groups of learners. There can also be many similar patterns. As a result, it is tedious to analyse each pattern distribution across the groups. The patterns were clustered based on their *edit distance*. The edit distance between two patterns was defined as the minimum number of changes needed to convert one pattern of actions into the other, with the allowed operations: insertion, deletion, or substitution of a single action. We used a hierarchical agglomerative clustering technique [Witten and Frank 1999] whose input is a matrix that contains all the edit distances between each pair of patterns. We chose this technique as it has proven successful in mining human-computer interaction data [Fern et al. 2010]. The end result can be visually represented by a dendrogram, showing different levels in which patterns are clustered. These visual representations served to supervise the cluster formation and decide which level of clustering was considered as acceptable.

Step 4. Post-processing and analysis. In the post-processing stage we included the authorship information, by considering the number of students who were involved with the patterns. We also examined the benefits of each method employed at step 1, i.e. the use of raw versus compacted data.

We now describe in detail the specifications of each approach and the results of the data mining outcomes in collaborative learning terms.

4.1. Method 1: Authorship in the post processing

The first method consists in exploring the information that can be obtained by mining the Human-Computer Interaction logs of physical actions without reducing the events.

Pre-processing and sequence mining for method 1. The input data for the sequence mining consisted of a list of sequential raw sequences of events (e.g.{M-E-M-M-S-M-N-G-S-M-R} where M=move, E=enlarge to maximum size, N=resize to normal size, G=add to group, S= shrink and R=remove from group). The output was a list of frequent patterns. Only sequences of at least 4 actions were considered. The final result included 259 frequent patterns found of length varying from 4 to 10 actions.

Post-processing and clustering for method 1. Based on direct observations made on the video recorded sessions and the sequential patterns found, we obtained that many patterns had a similar meaning, although the order or quantity of actions they contained were somewhat different. For example, the sequential patterns S1={M-E-M-M-S} and S2={M-M-M-E-S-M-M} (where M=move, E=enlarge, S=shrink) are both related with the same strategy: read a data-slip, close it and re-arrange it immediately afterwards (presumably to keep the interface organised and tidy). These observations led us to use clustering to group similar patterns. In this part of the process, the input for the hierarchical clustering algorithm was a similarity matrix of 259 x 259 that contained the edit distances of all pairs of sequences. The algorithm produced a dendogram of 4 hierarchies as output. The clusters obtained were supervised to inspect the extent in which the groups were similar. After analysing the dendogram, the second highest level was selected to form eight meaningful clusters. This is the only part of the approach in which the results were manually supervised.

Results of method 1. We examined the results of the clustering by looking at the trends observed between patterns and groups of learners that presented a prominent level of achievement. We found that some sessions (high or low achievers) showed behaviour associated with certain clustered patters. Therefore, we used unpaired student tests *($p \le 0.05$)* to statistically analyse whether there were significant differences between such sessions. Table I summarises the clusters found using this approach and the results of such analysis.

The first two clusters are related with the strategies that learners followed to gather information from the data slips. Cluster 1 contained sequences related with the strategy of reading the slips by enlarging the object and then, after a reasonable time, closing them to keep the interface tidy. Some of these groups positioned the slips in a certain region of the table to indicate they had already read them. On the other hand, Cluster 2 contained sequences of actions in which groups maximised the data slips without closing them. The observations on the videos indicated that some of the groups which followed this behaviour skipped the reading of some slips. We found that high achievers favoured the strategy of *reading, minimising and arranging immediately* (*cluster 1 mean = 124.75, cluster 2 mean = 61.25*). On the contrary, low achievers used both strategies for the information gathering, performing more actions contained in Cluster 2 in which they did not close the slips immediately after reading (*cluster 1 mean = 104.40, cluster 2 mean = 114.80).* This simple change in the strategy for collecting information suggests that reading without re-arranging increases clutter, making the task more difficult to be controlled by the group. Indeed, cluster 3, which contains patterns related with making space actions (moving and shrinking), showed a strong link with low achieving groups (t=2.47, *p= 0.039*). As a result, low achievers spent much more time than the high achievers arranging the elements at the table.

Clusters 6, 7 and 8 contain "union" actions in which learners established links between the data slips they considered to be tightly related. Cluster 6 includes sensible amount of union actions (at most two unions) performed along with arrangement actions. Cluster 7 presented a moderate amount of union actions and cluster 8 presented patterns with an enormous amount of union actions. Low achieving groups favoured clusters 7 and 8 (t=2.97, p=*0.018* and *t=3.98, p=0.0041* respectively). Based on this trend, low achievers created too many unions related to a specific data slip in short periods of time. On the contrary, high achieving groups favoured patterns with modest quantity of unions (t=2.81, *p=0.023*). Clusters 4 and 5 included patterns related with ungroup and group actions. In this case we obtained some differences among sessions. Low achievers made more "corrections" on categorising data slips than high achievers.

Table I. Results for clusters of patterns found by mining the raw events.

Cluster	Example sequence	Favoured Groups	Participants
1- Read and arrange	{M-M-E-M-S-M}	Slightly more in high achievers	Both groups 1-2 authors
2- Read slip	{M-E-M-M}	Slightly more in low achievers	Both groups 1-2 authors
3- Arrangement	{M-M-S-M-M}	Substantially more in low achievers	Low achievers 2-3 authors
4- Ungroup	{M-R-M-G}	Slightly more in low achievers	Both groups 1-2 authors
5- Group	{M-N-M-G-M-S}	Both groups	Both groups 1-2 authors
6- Few unions	{M-M-U-M-M}	Substantially more in high achievers	Low achievers 2-3 authors
7- Moderate unions	{M-U-M-U-M-U}	Substantially more in low achievers	Low achievers 2-3 authors
8- Too many unions	{U-U-U-M-U-U-U}	Substantially more in low achievers	Low achievers 2-3 authors

In regards to authorship, we analysed the way in which participants collectively interacted with the resources in terms of number of authors involved with the data slip in each pattern. For clusters 3, 6, 7 and 8 we obtained a strong statistical difference in the number of participants working together with the same data slip. Low achieving groups presented more sequences in which the *three* authors performed actions sequentially compared with high achieving groups ($p<0.05$ in all cases). For the rest of the clusters there were no significant differences in the number of authors involved with the patterns. For the clusters related with the strategies for gathering information (clusters 1 and 2) and grouping data-slips (clusters 4 and 5) the sequences were performed mostly by one author, and in some cases, by two authors in both high and low achieving groups (see Table I, column Participants).

What we learnt from these findings is that having many hands on the same object at the same time does not imply improved work. In fact, the sequences in which the low achievers have the three participants involved are mostly focused on non-cognitively demanding tasks, such as arranging the elements on the tabletop (cluster 3). In the case of the "union actions" clusters (7 and 8) even when the activity is a cognitively demanding task, we learnt from the analysis described above and from observing the videos that lower achieving groups created a larger number of unions on particular slips that were not necessarily meaningful. Grounding on these results and the video analysis we obtained that the high level groups worked more collaboratively and participants were keener to interact on one data slip at a time, even if they worked in parallel with different objects.

We also explored the possible significant differences between the patterns and the stages in which they appear. As expected, clusters related to gathering information (clusters 1 and 2) are mainly related with stage 1, cluster 3 (re-arrangement) with all the stages, Clusters 4 and 5 with stage 2 (grouping and ungrouping actions) and the clusters related with union actions are evidently related with the third stage (sequencing and webbing). Thus, no further special consideration was put on the staging information.

4.2. Method 2: Authorship in the post processing and generalisation in the pre-processing

The second approach consists of generalising (compacting). Then, we looked at the similarities of this method outcomes with method 1 results.

Pre-processing and sequence mining for method 2. The dataset of sequences was compressed. The aim of the compression was to see how much information will be lost or gained if we generalised the user interface actions that can be attributed to user slips. A simple alphabet was applied which follows a single rule: the sequential actions of the same type (such as the action M in {M-M-M-E-M} or S in {S-U-U-U}) were compacted adding the quantifier for regular expressions + ({M+-E-M} and {S-U+}). The minimum length of the patterns was set to 3 actions, or 2 actions if at least one of the actions contained the quantifier +. In this case the minimum support was also set to one quarter of the data slips. The final result included 261 frequent patterns found of size between 3 and 5 actions each.

Post-processing for method 2. The 261 patterns were clustered following the same process used in method 1. We obtained a dendogram with 5 levels. The first issue found in this method was that patterns were more difficult to cluster accurately as they contained less contextual information (fewer items). The solution was to choose a lower clustering level and manually merge the smaller clusters which contained similar actions. Seven meaningful clusters resulted from the clustering supervision and also 2 extra very small clusters that could not be considered into any other cluster.

Results of method 2. Even though some details in the sequences were lost, we found similar observable tendencies in the presence of patterns of high and low achieving sessions (see table II). This approach provided a deeper difference between the ways the higher and lower achieving groups gather information to solve the problem. For example we found a stronger difference in the strategy of reading without minimising the data slips performed mostly by the low achieving groups (Cluster 2, $t = 2.69$, $p=0.0272$). We also found a significant difference with respect to the strategy "read – close and arrange data slips" favouring the high achieving groups (Cluster 1, $t=3.05$, $p=0.0158$). Results also confirmed that low achieving groups performed a huge amount of unions between data slips in short periods of time (Cluster 7, $t = 3.05$, $p=0.0158$).

For the authoring information, the results from method 1 were also confirmed. The cluster that contains sequences with high amounts of union actions performed by 2 and 3 users at the same time were present mostly in the low achieving groups ($t=2.714$, $p=0.0265$). Some information is lost though; there were no significant differences between groups in any other aspect. In general, this approach confirmed the insights obtained applying method 1 but the quality of the results decreased in some cases.

Table II. Results for clusters of patterns found mining compacted events.

Cluster	Example sequence	Favoured Groups	Participants
1- Read and arrange	{M+-N-S-M+}	Substantially more in high achievers	Both groups 1-2 authors
2- Read slip	{M-E-M+}	Substantially more in low achievers	Both groups 1-2 authors
3- Arrangement	{M+-S-M+}	Slightly more in low achievers	Both groups 1-2 authors
4- Ungroup	{M-R-M+}	Both groups	Both groups 1-2 authors
5- Group	{M+-G-S}	Both groups	Low achievers 2-3 authors
6- Few unions	{M+-U-M-U}	Slightly more in high achievers	Low achievers 2-3 authors
7- Many unions	{M-U+-M-U+-M}	Substantially more in low achievers	Low achievers 2-3 authors

5. DISCUSSION

The design of our approach was motivated by the goal of exploiting the large amounts of data generated from learners' interactions with the interactive tabletop. Our approach shows promise to follow up research on supporting collaborative learning through the use of tabletops and machine learning techniques. Our data mining approach consisted of mining both the raw human computer interactions and the compact logged actions, clustering similar frequent patterns based on edit distance and analyse the proportion of these clusters among group sessions. Both methods we explored produced similar results therefore the compacting method provides very interesting insights even when some details are lost. However, this loss of information impacted negatively on the clustering step, thus this method is unsuitable for being used for automatic support.

Method 1 also requires some human supervision to code the level of achievement of groups. Further research needs to be done on the ways to automatically extract indicators of collaboration. In regards to the educational value of the results, the video analyses confirmed the presence of serial patterns of interaction in the trials. Group members of high achieving groups tried to interact and externalise their thinking. They tended to read all the slips to get clues about the mystery and parallel interactions were clearly observed along with engagement in conversations. The results of our approach do not tell the whole story but are good indicators of desired and undesired patterns of behaviour related with strategies that are followed by groups.

The goal of this line of research is to offer adapted support to groups in the form of direct feedback to students or to their facilitator. The insights obtained in the work reported in this paper are the first steps towards such adapted support that machine learning techniques can offer to the use of tabletop devices. We addressed two questions posed at the beginning of this paper, regarding (i) the key insights that can be acquired from mining raw or compact logged actions and (ii) the information offered by the authorship element of the data logs.

We observed that the results obtained with both methods reflected similar patterns of behaviour, such as the strategies followed by the groups to gather information, arrange resources and the creation of links between data slips. Some elements of the interactions came up by compacting the redundant actions in method 2 (gathering information strategies), but in other elements some information was lost. The most important issue with the compacting method is that more empirical interpretation was needed after the clustering step whilst method 1 offered better clusters. In general, the raw HCI actions offer an adequate degree of detail to obtain meaningful results when studying the interactions *by resource*. In regards to the question related with authorship, results indicated that low achieving groups tended to work sequentially with the same objects. We confirmed from the video analysis that high achieving groups tended to discuss their thoughts and work in parallel with different objects.

6. CONCLUSION AND FUTURE DIRECTIONS

This paper presented an outline of distinctive techniques to extract elements of collaborative interaction at the tabletop. These techniques reveal the importance of the design of specific data mining methods for exploiting traces of collaboration from co-located situations. Our work grounds upon educational data mining research on online collaborative learning and we have proposed a methodology that can be used as a starting point to guide future research on the identification of patterns from educational tabletop settings. An important goal of our work is to mirror useful information about groups to help facilitators and the students themselves to reflect on and improve their learning activity. There are still a number of open questions that we want to address. The next step in this line of research will be the exploration of other ways to analyse sequential actions considering parallel work, looking at the high level problem-solving processes, designing alphabets to include authorship in earlier stages of the data mining.

REFERENCES

ANAYA, A.R. AND BOTICARIO, J.G. 2011. Application of machine learning techniques to analyse student interactions and improve the collaboration process. *Expert Systems with Applications 38*, 1171-1181.

CASILLAS, L. AND DARADOUMIS, T. 2009. Knowledge extraction and representation of collaborative activity through ontology based and Social Network Analysis technologies. *BI and Data Mining 4*, 141-158.

D'MELLO, S., OLNEY, A. AND PERSON, N. 2011. Mining Collaborative Patterns in Tutorial Dialogues. *JEDM - Journal of Educational Data Mining 2*, 1-37.

FERN, X., KOMIREDDY, C., GRIGOREANU, V. AND BURNETT, M. 2010. Mining problem-solving strategies from HCI data. *ACM Trans. Comput.-Hum. Interact. 17*, 1-22.

JEONG, H. AND HMELO-SILVER, C.E. 2010. An Overview of CSCL Methodologies. In *Proceedings of the 9th International Conference of the Learning Sciences*, Chicago, USA2010, 920-921.

JIANG, L. AND HAMILTON, H. 2003. Methods for Mining Frequent Sequential Patterns. In *Advances in Artificial Intelligence*, Y. XIANG AND B. CHAIB-DRAA Eds. Springer Berlin / Heidelberg, 992-992.

KHARRUFA, A.S. 2010. Digital tabletops and collaborative learning. PhD thesis. School of Computing Science, Newcastle University, Newcastle Upon Tyne.

KHARRUFA, A., LEAT, D. AND OLIVIER, P. 2010. Digital mysteries: designing for learning at the tabletop. In *Interactive Tabletops and Surfaces* ACM, 197-206.

LEAT, D. AND NICHOLS, A. 2000. Brains on the Table: Diagnostic and formative assessment through observation. *Assessment in Education: Principles, Policy & Practice 7*, 103 - 121.

MARTINEZ, R., WALLACE, J., KAY, J. AND YACEF, K. 2011a. Modelling and identifying collaborative situations in a collocated multi-display groupware setting. In *Proceedings of the International Conference on Artificial Intelligence in Education* 2011.

MARTINEZ, R., KAY, J., WALLACE, J. AND YACEF, K. 2011b. Modelling symmetry of activity as an indicator of collocated group collaboration. In *Proceedings of the International Conference on User Modeling, Adaptation and Personalization*.

MASATAKI, H. AND SGISAKA, Y. 1996. Variable-order N-gram generation by word-class splitting and consecutive word grouping. In *Proceedings of the Conference on Acoustics, Speech, and Signal Processing, 1996*1996 IEEE Computer Society, 1256487, 188-191.

PERERA, D., KAY, J., KOPRINSKA, I., YACEF, K. AND ZAIANE, O. 2009. Clustering and Sequential Pattern Mining of Online Collaborative Learning Data. *IEEE Tran. on Knowledge and Data Engineering 21*, 759-772.

PRATA, D., BAKER, R., COSTA, E., ROSÉ, C., CUI, Y. AND DE CARVALHO, A. 2009. Detecting and Understanding the Impact of Cognitive and Interpersonal Conflict in Computer Supported Collaborative Learning Environments. In *2nd International Conference On Educational Data Mining*, 131-140.

REIMANN, P., FREREJEAN, J. AND THOMPSON, K. 2009. Using process mining to identify models of group decision making in chat data. In *Proceedings of the international conference on Computer Supported Collaborative Learning* 2009 International Society of the Learning Sciences, 98-107.

SOLLER, A. AND LESGOLD, A. 2007. Modeling the process of collaborative learning. *The Role of Technology in CSCL*, 63-86.

TALAVERA, L. AND GAUDIOSO, E. 2004. Mining Student Data to Characterize Similar Behavior Groups in Unstructured Collaboration Spaces. In *Proceedings of the European Conf. Artificial Intelligence* 2004.

WITTEN, I.H. AND FRANK, E. 1999. *Data mining: Practical machine learning tools and techniques with Java implementations*. Morgan Kaufmann, San Francisco.