

Data Mining in the Classroom: Discovering Groups' Strategies at a Multi-tabletop Environment

Roberto Martinez-Maldonado, Kalina Yacef, Judy Kay
School of Information Technologies
The University of Sydney,
NSW 2006, Australia
{roberto, judy, kalina}@it.usyd.edu.au

ABSTRACT

Large amounts of data are generated while students interact with computer based learning systems. These data can be analysed through data mining techniques to find patterns or train models that can help tutoring systems or teachers to provide better support. Yet, how can we exploit students' data when they perform small-group *face-to-face* activities in the classroom? We propose a novel approach that aims to address this by discovering the strategies followed by students working in small-groups at a multi-tabletop classroom. We apply two data mining techniques, sequence and process mining, to analyse the actions that distinguish groups that needed more coaching from the ones that worked more effectively. To validate our approach we analysed data that was automatically collected from a series of *authentic* university tutorial classes. The contributions of this paper are: i) an approach to mine face-to-face collaboration data unobtrusively captured at a classroom with the use of multi-touch tabletops, and ii) the implementation of sequence mining and process modelling techniques to analyse the strategies followed by groups of students. The results of this research can be used to provide real-time or after-class indicators to students; or to help teachers effectively support group learning in the classroom.

Keywords

Collaborative Learning, Sequence Mining, Process Mining, Interactive Tabletop, Classroom

1. INTRODUCTION

Collaborative face-to-face activities can offer particular advantages compared to computer-mediated group work [17]. These include a natural channel for both verbal and non-verbal communication, improved perception of quality of group discussions, and an increased productivity in completing tasks [17, 18]. The classroom is a common environment in which the teacher can foster face-to-face collaboration skills acquisition by making use of small-group activities [8]. However, even in small-group activities, it is challenging for teachers to provide students the attention that they may require and be aware of the process followed by each group or their individual contributions [21]. Commonly, teachers try to identify the groups that work effectively to leave them work more independently and be able to devote time to groups needing their attention.

Multi-user shared devices, such as interactive tabletops, provide an enriched space where students can communicate face-to-face with each other and, at the same time, interact with a large work area that has access to digital content and allows the creation of persistent artefacts [14]. Interactive tabletops may afford new possibilities to support learning but they also introduce

additional challenges for a new space of interaction. In order for these tabletops to be integrated into the classroom, as with any emerging technology, they should provide additional support to teachers compared with what they can currently do without such technology [4]. Currently, these devices are making their way into the classroom in the form of multiple interactive tabletops that have the potential of providing teachers with new ways to control groups [1, 11]; plan and enact authentic collaborative activities [10]; and monitor students' progress [5, 11].

At the same time, the increasing usage of technology for learning and instruction has made it possible to collect students' traces of activity resulting in large amounts of data gathered while they interact with computer based learning systems. These data can be analysed through data mining techniques to find patterns or train models that can help tutoring systems or teachers to provide enhanced support [3]. Although there is substantial research work on mining students' data obtained from individual or online learning systems, there is still little research on automatically exploiting the data generated when learners perform small-group *face-to-face activities in the classroom*.

A slightly hidden potential of interactive tabletops is that they can open new opportunities for capturing learners' *digital traces of activity*, offering teachers and researchers the possibility to inspect the process followed by students and recognise patterns of group behaviour [12]. This paper presents a novel approach that focuses on analysing face-to-face collaboration data to discover the strategies that distinguish groups that need more coaching from the ones that work effectively.

To validate our approach we analysed data that was automatically and unobtrusively collected from authentic tutorials that covered part of the regular curricula of a university subject in the area of Management. The teacher designed a small-group collaborative activity, based on the concept mapping learning technique, using our multi-tabletop classroom environment called MTClassroom [11]. This allows multiple small-groups of students to work around a number of interactive tabletops, perform a series of tasks, discuss a topic and provide a solution to a case proposed by the teacher. The system automatically logs identified students' actions on the shared device and all the steps that groups performed to build a collaborative artefact. We describe the application of two data mining techniques. First, we used a sequential pattern mining technique to look for patterns that can help find differences between groups according to the teacher assessment. Then, we used the Fuzzy Miner tool [6] to discover the processes most often followed by both high and low achieving groups. The main contributions of this paper are: i) an approach to mine face-to-face collaboration data unobtrusively captured at a classroom

with the use of multi-touch tabletops, and ii) the implementation of sequence mining and fuzzy modelling techniques to analyse and discover strategies followed by groups of students.

The paper is structured as follows. The next section describes the state of research on the areas of interactive tabletops in the classroom and data mining for collaborative learning. Then, we present details of the multi-tabletop tutorials and our technical infrastructure. Section 4 presents the motivation and design of study. Section 5 describes the data pre-processing and the methods. Section 6 presents a discussion of the results. Section 7 states the conclusions and the avenues for future research.

2. RELATED WORK

There is a steady growth of the usage of tabletops in education. More specifically, there are a number of research projects that have used *multiple tabletops or shared devices in the classroom*. One of these is Synergynet [1], a multi-tabletop setting that has served to study the ways school kids collaborate and interact to achieve group goals. This project also included the design of tools for the teacher to *control* the classroom activities. Another approach was proposed by Do Lenh [5], who developed a setting for training on logistics, that consisted of four tangible horizontal devices that could be *orchestrated* by the teacher using paper-based commands or through a remote computer. This project also offered minimalist indicators of progress of each small group presented at a wall display. Even though these two previous projects included real students and teachers, they were mostly designed and deployed as experimental scenarios. A different approach was followed by Martinez-Maldonado et al. [10], who presented a multi-tabletop system that permitted teachers to *assess the design and enactment* of their planned classroom activities through the use of analytics tools. This is the only previous work that has focused on exploiting the collected data from a multi-shared device environment to describe the activities that occur in an authentic classroom.

In the case of *data mining applied to collaborative settings*, the closest study to ours was presented by Martinez-Maldonado et al. [12]. It consisted in extracting and clustering frequent sequential patterns to then link them with high level group actions at a pen-based tabletop learning application called Mysteries. One important study, even though not related to tabletops, was performed by Perera et al. [20] who explored the usage of sequence mining alphabets and clustering to find trends of interaction associated with effective group-work behaviours in the context of a software development tool. Moreover, Anaya et al. [2] analysed a computer-mediated learning tool to classify and cluster learners according to their level of collaboration.

The work reported in this paper is the first effort we are aware of that proposes an integrated solution, inspired by authentic needs of the teacher in the classroom, to exploit the students' data that can be captured by multiple tabletops though the application of a data mining technique and a process modelling tool.

3. MULTI-TABLETOP TUTORIALS

This section describes our technical infrastructure that consists of: the multi-tabletop classroom, a teacher's dashboard, the system for capturing identified learners' actions and a learning tool for building concept maps. We also describe the teacher's design of the tutorials.

3.1 Technical Infrastructure

Our multi-tabletop classroom is called MTClassroom [11]. This has a number of interconnected multi-touch interactive tabletops (four in this study). Figure 1 shows an instance of MTClassroom for a demo tutorial. Each tabletop consists of a 26 inch PQlabs overlay placed over a high-definition display that is enriched with Collaid [9]. Collaid is a system that provides an ordinary interactive tabletop the capability of automatically and unobtrusively identifying which person is touching where, based on an over-head depth sensor (www.xbox.com/kinect). Using this system, each tabletop can identify actions performed by each student according to their seating position.

The logging system of each tabletop records the activity logs to a central synchronised repository that can be accessed in real time by other services. One of these is a teacher's dashboard called MTDashboard [11]. This dashboard provides functions for the teacher to orchestrate the tabletops (e.g., blocking the touch input of all tables or moving the class to the next phase) and to see key live-indicators of work progress of each small-group. Figure 2 shows the teacher holding the dashboard, displayed on a tablet device, while she provides feedback to a group. The classroom activity consisted in elaborating collaborative concept maps about a case proposed by the teacher. *Concept mapping* is a technique that promotes learning by allowing students to visually represent their understanding in the form of *concepts* associated by *linking words* that creates statements [16]. We used a minimalist version of a tabletop concept mapping application called Cmate [9]. Cmate provides students with a list of concepts and linking words suggested by the teacher, and also allows them to type their own words, in order to build a concept map that represents their solutions. Prior to the tutorials, the teacher creates a *master concept map* with the crucial concepts and links that learners are expected to include in their maps.

3.2 Tutorials Design

Eight tutorial sessions were organised in the School of Business of the University of Sydney during week 6 of semester 2, 2012 for the course: *Management and organisational ethics*. The teacher designed a case resolution activity to cover the topic of the curricula corresponding to that week. A total of 140 students attended these tutorials (from 15 to 20 students per session) that were organised in groups of 4, 5 or 6 students.

The teacher designed the tutorial script as follows: 1) *Introduction* (10 minutes): the teacher forms groups, explains to students how to use the concept mapping application and



Figure 1. MTClassroom: a multi-tabletop classroom with capabilities for capturing differentiated students' activity.

introduces the first activity. 2) *Activity 1* (10 min.): using the MTDashboard, the teacher cleans up the four tabletops for all groups to start at the same time. Students are instructed to create a concept map that represents how the main actors of the case are associated. 3) *Reflection 1* (5 min.): the teacher blocks the tabletops, leads a short class discussion about partial solutions and introduces Activity 2. 4) *Activity 2* (15 min.): this is for the teacher “*the most important activity of the tutorial from the learning perspective*”. The teacher unblocks the tabletops, and students discuss and focus on representing a final solution to the case in their concept map. 5) *Class sharing and reflection* (10 min.): the teacher asks each group to share their solution with the class. After each group has explained their map, the teacher summarises the outcomes of the tutorial, finishes the session and assesses each group in private. The class time was fixed to 50 minutes. Details of these tutorials can be found in [11].

4. STUDY DESIGN AND DATASET DESCRIPTION

The teacher in the classroom can face a number of challenges related with control, awareness and resources management [22] which depend on a number of factors that may fall out of the scope of what tabletop systems can capture. The tabletop systems are not totally aware of the classroom situation, for example, if a group of students is talking, if they work on-task or if someone needs to leave the class. The teacher can have a better idea of the productivity of students’ discussions within each group, however, one of the main conclusions after finishing the tutorials was that for the teacher it is not easy to know aspects of the final artefacts that students built or their individual contributions [11].

In a post-tutorial interview the teacher expressed her view as follows: “*I don’t want to see a lot of information in the dashboard, this can be distracting. But more information can be provided after the tutorials for assessment, like who did what, when, and the quality of the work*”. These are indeed the aspects of group work that tabletops are aware of in detail. Our system can capture: 1) differentiated students’ action on the tabletop; 2) the sequential actions performed to build the group artefact.

Inspired by the above teacher needs, but framed on what tabletops can actually capture in an authentic classroom, we propose an approach to distinguish strategies followed by groups that either needed more coaching or worked effectively. We analyse three sources of contextual information i) identified individual actions on the tabletop that can occur in parallel, in turn, or on other students’ objects, ii) the quality of students’ actions according to the teachers’ artefact, and iii) the impact of students’ actions on the group artefact. In this paper we focus on the students’ actions performed in Activity 1. This is important because a certain degree of success in Activity 1 is required for Activity 2. This also allows the approach to be applicable in real-time, to provide feedback to teachers before the tutorial is over, so they can target their support during Activity 2.

Table 1. Possible actions on the concept mapping tabletop system

| High impact actions (content and structure) | Low impact actions (layout) | No impact actions |
|--|--------------------------------|---------------------------|
| Add a concept/link | Move a concept/link | Open or close menus |
| Delete a concept/link | Merge two links | Move/scroll menu-concepts |
| Edit a concept/link | | |



Figure 2. A teacher attending a group while holding the MTDashboard

The teacher assessed groups at the end of each tutorial, using one of three possible values: low, medium or high achievement. The teacher specified that the assessment criteria mostly considered the *quality of each group solution* presented at the end of the tutorial and the *quality of their discussions* during the tutorial. We considered the activity data of all the 32 groups divided in two sets: 20 groups that were high achieving and 12 groups that were medium or low achieving.

The *initial raw data* of each group consists of a long sequence of actions in which each element is defined as: {*Resource, ActionType, Author, Owner, Time, Relevance*}, where *Resource* can be: *Conc* (concept), *Link* (proposition) or *Menu*. *ActionType* can be: *Add* (create a concept or link), *Del* (delete), *Mov* (move) links, *Chg* (edit), *Scroll*, *Open* or *Close* (a menu). *Author* is the learner who performed the action, *Owner* is the learner who created an object or owns a menu, *Time* is the timestamp when the action occurred and *Relevance* indicates if the concept or link belongs to the crucial elements of the teacher’s map. Table 1 lists all the possible actions in the dataset grouped by their impact on the group concept map. Some examples of actions are: {*ConceptA, Add, 3, 3, 17:30:02, Crucial*}, when a learner adds a crucial concept to the map; {*LinkY, Move, 2, 6, 17:30:04, Irr*}, when s/he moves a link created by another learner; and {*MenuConcepts, Open, 2, 2, 17:30:07,-*} when s/he opens the list of suggested concepts. The original sequence obtained for each group contained from 74 to 377 *physical actions*.

We address four research questions regarding the strategies and characteristics that can differentiate groups according to their extent of achievement. The formulation of these is based on the triangulation of the nature of the available data (differentiated students’ actions and their impact on their artefact), the teacher’s needs (awareness on students’ participation and quality of their work), and open issues in the study of multi-tabletop classrooms [10]. Our research questions are the following. 1) *Can we distinguish groups by inspecting patterns of parallelism and turn-taking?* As the teacher is interested in the participation of all students in the construction of the group solution [10], we analyse whether it is possible to find differences among groups where students worked at the same time (in parallel or taking turns) or not. 2) *Can we distinguish groups by inspecting students’ interactions on others’ objects?* Other studies inspired this question; these have suggested that interacting with what others’ have done may trigger further discussion that is beneficial for tabletop collaboration [11, 13]. 3) *Can we distinguish groups*

by inspecting students' map quality? This and the next question are directly inspired by teachers' needs, as noted above, and the data captured by our system about the groups' artefacts and the process followed to build them. 4) *Can we distinguish groups by inspecting the process followed by students' actions and their impact on the group artefact?*

5. METHOD

Sequential mining and process mining are techniques that have been used to identify patterns in educational datasets by considering the order of students' actions [7, 12, 19]. We used a sequential pattern mining technique called *differential sequence mining* [7] to distinguish strategies followed by groups that were either high or low achievers and address each of our first three research questions. For these, we analysed two of the sources of contextual information listed in the previous section: i) identified actions on the tabletop and ii) the quality of students' artefact. In order to address the fourth question, and analyse the strategies that distinguish groups according to iii) the impact of students' actions on the group map, we used the *Fuzzy Miner* tool [6]. Next subsections present the motivation for using these tools, the data pre-processing and the implementation of each technique.

5.1 Sequence mining

One of the data mining techniques that has been successfully applied to identify patterns that differentiate high from low achieving students is differential sequence mining (DSM) [7]. In general, a sequential pattern is a consecutive or non-consecutive ordered sub-set of a sequence of *events* that is considered frequent when it meets a minimum *support* threshold. In educational contexts, the *events* commonly correspond to individual or grouped students' actions logged by the learning system. The DSM algorithm extracts frequent consecutive ordered sequences of actions from 2 datasets and performs an analysis of significance to obtain the patterns that differentiate them. The actions can also contain contextual information as defined by an *alphabet*. Alphabets can be used to *encode* each action to a set of concatenated keywords. In our study, each action was encoded to the format *{Resource-ActionType-Context}*. We implemented a DSM solution to investigate the differential patterns in terms of degree of parallelism, actions of students on others' objects and relevance of the links and concepts students use according to the teacher's map. Table 2 presents the keywords of each of our three alphabets. The encoded actions encoded using any alphabet should contain at least one keyword for the *Resource* column and one for the *ActionType* column. We add one keyword of the corresponding contextual information (three rightmost columns in Table 2) according to the Resource type. Alphabet 1 aims to model the differentiated individual actions performed on the tabletop that occur in parallel (with other students' actions, keyword:

Parallel), in turns (when the previous action was performed by a different student, keyword: *Other*), or as a series of actions by the same student (*Same*). Alphabet 2 models the actions that students perform on their own objects (*Own*) or on other students' objects (*NoOwn*). Finally, Alphabet 3 indicates whether the concept or link involved in the action belongs to the crucial objects defined by the teacher (*Cruc* or *NoCruc*).

In a previous study, we found that it is very important to consider the periods of significant inactivity registered by the tabletop [11]. During these periods of inactivity students can be having productive discussions, off-task talking or not working collaboratively at all. In our study, even when we do not perform speech detection, it is important to at least consider the occurrence of inactivity. To define a period of inactivity, we explored the time gap between each action performed on the tabletop. We found that time gaps between actions below one standard deviation from the mean ($< \mu + 1\sigma$) account for the 92% of the set. ($\mu = 4.30$ seconds, $\sigma = 8.62$, $\mu + 1\sigma = 13$ seconds). This means that a period above 13 seconds without logged actions can be considered as a block of inactivity. We defined these blocks as *short* when the gap was between 13 ($\mu + 1\sigma$) and 22 ($\mu + 2\sigma$) seconds, and *long*, for gaps longer than 22 seconds ($\mu + 2\sigma$). We detected from 6 to 19 periods of inactivity in each group.

The output of the DSM algorithm, using the three alphabets, consists of three sets of frequent sequential patterns that differentiate high from low achieving groups according to the teacher's assessment. In this study, we set a minimum support of 0.5 to consider a pattern as frequent and a maximum error of one to allow matching sequences with up to 1 different action, similarly to previous work on educational data exploration [7].

5.2 Process mining

The sequence mining approach presented above can extract patterns of activity that distinguishes groups; however, it does not give insights of the higher level view of the *processes* followed. The Fuzzy miner [6] is a process discovery tool that can generate a meaningful abstraction of a general process, from multiple instances by distinguishing the activities that are important. It is especially suitable to mine *unstructured* processes, like the concept mapping construction in this study. The input of this algorithm is a series of consecutive actions, or group of actions. The result is a directed graph in which each node represents an action, or group of actions, and the edges represent the transitions between these. The nodes and edges that appear in the graph should meet a *conformance threshold* based on the instances that were used to build the model.

The objective of this second analysis is to discover the meaning of the higher level steps that high and low achieving groups performed to build the concept map and the impact of such actions. For this, we performed the following data preparation before using the Fuzzy miner tool.

Table 2. Keywords included in the alphabets for the sequential pattern mining.

| Resource | Action type | | Alphabet 1 <i>Parallelism –turn taking</i> | Alphabet 2 <i>Actions on others' objects</i> | Alphabet 3 <i>Master map distance</i> |
|-----------------------------------|------------------------|-------------------------|---|---|--|
| <i>Concept (Conc)-C</i> | <i>Add -C,L</i> | <i>Delete (Del)-C,L</i> | <i>Parallel</i> | <i>Own</i> | <i>Cruc (C,L)</i> |
| <i>Link -L</i> | <i>Edit (Chg) -C,L</i> | <i>Merge (Move)-L</i> | <i>Other</i> | <i>NoOwn</i> | <i>NoCruc (C,L)</i> |
| <i>Menu -M</i> | <i>Move -C,L,M</i> | <i>Open -M</i> | <i>Same</i> | | |
| | | <i>Close -M</i> | | | |
| <i>Inactivity block (Inact)-B</i> | <i>Short(Shrt) -B</i> | <i>Long -B</i> | | | |

1) *Data grouping.* We grouped the actions into periods of activity in order to generalise similar actions according to their impact on the concept map. First, we explored the number of actions contained in each period of activity between periods of inactivity. Figure 3 illustrates the frequencies of the number of actions within blocks of activity in the dataset ($\mu=12.85$ actions, $\sigma=17.68$). The distribution shows a high frequency of periods with a small number of continuous actions, and a long tail of longer sets of actions. In fact, the 71% of the periods of continuous activity were below the mean size (13 actions) and the 87% of them were below one standard deviation from the mean (30 actions). We considered the mean (13 actions) as a practical threshold for the

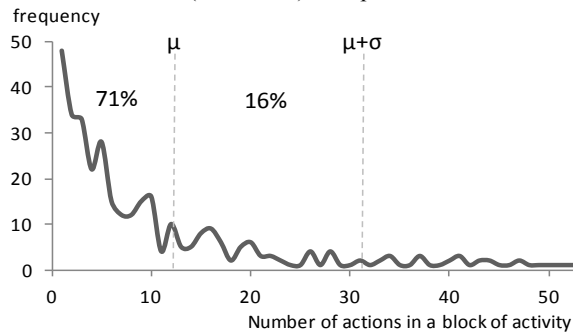


Figure 3. Distribution of the length of the sets of activity in terms of number of actions.

maximum size of a block of activity.

2) *Actions categorisation.* Based on the definition and previous research on concept mapping [15, 16], we categorise students' actions according to their impact on the group map. Actions that make a change in the structure or content of the concept map are categorised as *High-impact actions*. These include actions that modify the quantity or content of concepts and links (Table 1). The second category is *Low-impact actions*, which includes actions that modify the layout of the map, which is important for the activity, but not crucial. These actions include moving concepts and links, or merging links. Finally, actions performed on the menus of the application belong to *No-impact actions*.

3) *Blocks categorisation.* Each block was categorised according to the actions that occurred within that period following the next rules: *HighOnly* for blocks that contained only high-impact actions and some no-impact actions; *HighLow*, if the block contained at least one high-impact and one low-impact actions; *LowOnly*, for blocks that contained only low-impact actions and no-impact actions; and *NoImpact* if the block contained just no-impact actions. Periods of inactivity were categorised as either *InactShort* or *InactLong*, as explained earlier.

4) *Addition of contextual information.* According to our research aim, we highlighted the importance of distinguish the learners who work on their own or on other students' objects. For this, we added the information about who touched which object with the keywords *NoOwn* if most of the actions were performed on others' objects and *Own* if the actions were performed on the same learners' objects.

After performing the data preparation we divided the dataset into two sets, one for high and one for low achieving groups, as we did for the sequential mining. We generated two corresponding fuzzy models using the plugin implemented in the ProM framework (www.processmining.org). Then, we performed two

model analyses: analysis of the number of active learners, and a validation of the models to discriminate groups.

Analysis of number of active learners. We explore whether there is a difference in the number of learners that were actively involved in each of the significant activities that appear in each fuzzy model (the nodes of the model). For the latter, the explored values corresponded to blocks of activity in which only one learner (*1u*), two (*2u*), or more than 2 learners (*+u*) were involved in the actions within a block of activity. This takes into account that all groups had from 4 to 6 group members. No correlation was found between the group size and the level of achievement of each group ($r=0.2$).

Validation of the models. We performed a cross validation of the two models to evaluate if they can be used to effectively differentiate high from low achieving groups. To do this, we calculate, for each group process, the level of conformance of both fuzzy models and validate that the model that fit the most corresponds to the level of achievement of the group.

6. RESULTS AND DISCUSSION

6.1 Sequence mining results

After applying the DSM algorithm on the encoded datasets according to our three alphabets, we selected the patterns whose *instance* support (number of times the pattern is repeated within a group log) differed between the high and low achieving groups ($p \leq 0.10$) and that were composed of at least 2 actions. Table 3 presents the top-4 most frequent sequences for each of the three alphabets explored in this part of the study.

Alphabet 1: focused on parallelism and turn-taking. We obtained a total of 23 *differential patterns* for groups that were either high or low achieving after analysing the first encoded dataset. The top sequences in Table 3 indicate the presence of actions in parallel for *move* events (sequence A) and actions that contain the keyword *Other*, when adding and moving elements of the concept map (sequences B, C and D). These provide evidence that in high achieving groups more than 1 student quite often interacted with the tabletop at the same time. In fact, the keywords *Parallel* and *Other* appeared in 13% and 66% in the frequent patterns of high groups, while in the low achieving groups there were no patterns with the keyword *Parallel* and the keyword *Other* only appeared in the 30% of them.

Alphabet 2: focused on actions on others students' objects. In this case, we obtained a total of 29 *differential patterns*. Table 3 shows that in high achieving groups, students tended to interact with objects created by other students, such as moving and adding links using others' concepts, either followed or preceded by periods of inactivity (keywords *NoOwn* and *Inact* in sequences I, J, and L). The keyword *NoOwn* appeared two times more often in the frequent sequences of the high groups than in the achieving groups (in 42% and 22% of the sequences respectively). The presence of actions on students' own objects (*Own*) was similar in all groups.

Alphabet 3: focused on Master map distance. We obtained 28 *differential patterns* by analysing the encoded dataset. This includes contextual information of the concepts and links that belong to the crucial elements defined by the teacher. The patterns in Table 3 show that in high achieving groups, students

tended to work with more crucial elements than low achieving groups. However, an analysis of all patterns found showed that there was not a large difference in actions performed on crucial elements (keyword *Cruc* was present in 87% and 84% of the patterns of high and low achieving groups respectively). However, the key difference was that high achieving groups interacted with less non-crucial concepts and links (keyword *NoCruc* was in 19% and 73% of the patterns of high and low achieving groups respectively).

The sequences of events extracted using this technique, provides some insights about the strategies followed by groups. Low achieving groups tend to have long periods of inactivity on the tabletop before or after creating links or performing a chain of actions that affect the layout of their concept map (e.g. action *Inact-Long* in patterns G, H, N, O and X). High achieving groups also had periods of inactivity, but these were shorter. Long periods of inactivity appeared two times more in the low achieving groups, followed or preceded by other actions (*Inact-Long* appeared in 48% and 22% of the sequences of high and low achieving groups respectively). There was no difference in the appearance of short periods of inactivity.

These findings suggest that, to discover the strategies followed by groups, this approach offers a limited view of the meaning of the actions. The frequent sequences that were found can be used to build a model or benchmark to ‘detect’ if students’ actions are similar to either high or low achieving groups. However, the patterns themselves do not provide information about the process that groups followed during the activity that would be easily associated with groups’ behaviours.

6.2 Process mining results

Figure 4 shows the resulting fuzzy models after applying the second approach to mine the process of both, high and low achieving groups where the conformance with their corresponding datasets was above 80%. Nodes of the graph represent categories of action blocks of activity and the edges the transitions between these. Each node contains: the name of the block category, the conformance of the block with the dataset, and the rates of active students that were involved in the activities (*1u*, *2u* and *+u*). Nodes with conformance rates below

to 0.1 were not considered in the models to include the majority of the block categories but disregarding the actions that rarely appeared in the data and that would make the graph unnecessarily complex. The numbers next to the edge lines are indicators of conformance of the transitions with the datasets.

By visually comparing both graphs we can highlight that they share the same core blocks of activity. These include: the blocks *Inact-Short* and *Inact-Long* (marked with an orange small square in the top left of the node). We confirmed the results obtained with the sequence mining, where low achieving groups showed more long periods of inactivity compared with high groups (conformance of 0.68 and 0.98 respectively). Both models also have in common the categories *HighLow-NoOwner* and *HighLow-Owner* (blue markers) that represent activity that combined high and low impact actions on the group map (conformance of 1 and around 0.4 respectively). The last similarity, in terms of nodes, corresponds to blocks of low impact actions where students interacted with other students’ objects (*LowOnly-NoOwner*, red markers).

The nodes marked with a yellow star correspond to activity blocks that appear in one model but not in the other. High achieving groups, contrary to the expected, presented more blocks of actions with no impact on the concept map (*NoImpact-Owner/NoOwner*). However, both nodes had the least conformance with the model (0.11 and 0.2 respectively). In contrast, low achieving groups presented blocks of activity with only high impact actions (*HighOnly-Owner/NoOwner*). The conformance of these blocks was not low (conformance of 0.37 and 0.74 respectively).

However, the main difference between the models is in the structure of the transitions. For the model of high achieving groups, there is only one transition between *different* blocks of activity. This was, in addition, not very frequent (0.08 conformance, between *NoImpact-Owner* and *HighLow-Owner*). By contrast, the model of low achieving groups contains 5 transitions between activity nodes with a conformance of up to 0.17 (between *HighLow-Owner* and *LowOnly-NoOwner*). Additionally, we did not find any observable difference in the actions performed on other students objects (*NoOwner*) and students’ own objects (*Owner*).

Table 3. Top-4 most frequent sequences after applying differential sequence mining on each encoded dataset.

| Alphabet 1 | High achieving groups | Low achieving groups |
|--|-----------------------|---|
| A- {Menu-Mov-Same}>{Menu-Mov-Same}>{Menu-Mov-Parallel} | | E- {Link-Add-Same}>{Link-Rem-Same}>{Con-Mov-Same} |
| B- {Con-Mov-Other}>{Link-Add-Same}>{Con-Mov-Same}>{Link-Add-Same} | | F- {Link-Rem-Same}>{Con-Mov-Same}>{Link-Add-Same} |
| C- {Inact-Shrt}>{Con-Mov-Other}>{Link-Add-Same} | | G- {Link-Add-Same}>{Link-Chg-Same}>{Inact-Long} |
| D- {Con-Mov-Other}>{Link-Add-Same}>{Con-Mov-Same} | | H- {Inact-Long}>{Inact-Shrt}>{Con-Mov-Same} |
| Alphabet 2 | High achieving groups | Low achieving groups |
| I- {Con-Mov-NoOwn}>{Con-Mov-NoOwn}>{Link-Add-Own}>{Inact-Shrt} | | M- {Inact-Shrt}>{Con-Mov- NoOwn }>{Link-Add-Own}>{Link-Chg-Own} |
| J- {Inact-Shrt}>{Con-Mov- NoOwn }>{Con-Mov- NoOwn }>{Link-Add-Own} | | N- {Link-Add-Own}>{Link-Chg-Own}>{Inact-Long} |
| K- {Link-Mov- NoOwn }>{Link-Mov- NoOwn }>{Con-Mov- NoOwn } | | O- {Link-Chg-Own}>{Inact-Long} |
| L- {Inact-Shrt}>{Con-Mov- NoOwn }>{Con-Mov- NoOwn } | | P- {Inact-Long}>{Inact-Shrt}>{Con-Mov- NoOwn } |
| Alphabet 3 | High achieving groups | Low achieving groups |
| Q- {Con-Mov-Cruc}>{Link-Add-Cruc}>{Con-Mov-Cruc}>{Link-Add-Cruc} | | U- {Link-Rem-NoCruc}>{Con-Mov-Cruc}>{Link-Add-Cruc}>{Link-Chg-NoCruc} |
| R- {Inact-Shrt}>{Con-Mov-Cruc}>{Con-Mov-Cruc}>{Link-Add-Cruc} | | V- {Link-Chg-NoCruc}>{Link-Chg-NoCruc}>{Inact-Shrt} |
| S- {Link-Add-Cruc}>{Link-Mov-Cruc}>{Con-Mov-Cruc} | | W- {Inact-Shrt}>{Link-Add-Cruc}>{Link-Chg-NoCruc} |
| T- {Link-Chg-Irr}>{Con-Mov-Cruc}>{Link-Add-Cruc} | | X- {Con-Mov-Cruc}>{Link-Add-Cruc}>{Link-Chg-NoCruc}>{Inact-Long} |

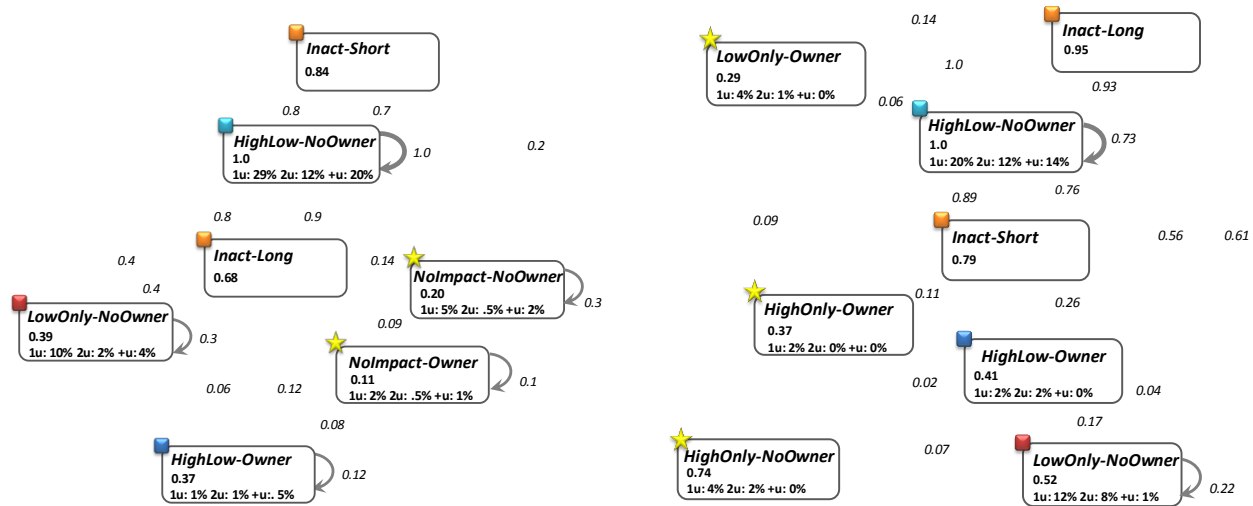


Figure 4. Fuzzy model generated from groups' activity. Left: Fuzzy model of high achieving groups (Conformance: 86% , Cutoff: 0.1). Right: Fuzzy model of low achieving groups (Conformance: 81% , Cutoff: 0.1).

Next, we present the analysis of the number of students involved in the activities and the validation to determine if the observable differences can distinguish high from low achieving groups.

Active learners. Table 4 shows the results of the cumulated distribution of the number of learners involved in the periods of activity for both high and low achieving groups (partial rates displayed in the third line of text inside each node of Figure 4). Both high and low achieving groups presented more than the half of the blocks of activity performed by a single student (54/55%). The main difference found was that high achieving groups presented blocks of activity in which more than two learners were involved in comparison with low achieving groups (+u, 27% and 19% respectively). In low achieving groups most of the blocks of activity were performed by either one or two learners.

Table 4. Distribution of the number of active learners in blocks of activity

| Achievement | One learner (1u) | Two learners (2u) | More learners (+u) |
|-------------|------------------|-------------------|--------------------|
| High | 55% | 18% | 27% |
| Low | 54% | 27% | 19% |

Validation. In order to validate that the two models generated by the fuzzy miner are different and can be used to distinguish the process followed by either high or low achieving groups, we estimated how accurately each model will conform to each group's activity. We performed a cross-validation to compare the level of fit of both models to the data blocks of each group by measuring whether the conformance of the model that corresponded to the level of achievement of the group was higher. Table 5 shows the confusion matrix which layouts the results of this analysis. This indicates that the fuzzy model for low achievement could distinguish the 100% of the low achieving cases, however, three high achieving groups presented a superior conformance to this model. The conformance of the model of high achievement was higher for the high achieving

Table 5. Validation of the fuzzy models

| | | Predicted class | |
|--------------|------|-----------------|-----|
| | | High | Low |
| Actual class | High | 17 | 3 |
| | Low | 0 | 12 |

groups in 17 of the 20 cases. The difference between the levels of fit of each model was statistical significant for high achieving groups (paired $t(23) = 2.46, p = 0.0219$) and very close to statistical significance ($p < .05$) for the model of low achievement (paired $t(7) = 2.16, p = 0.061$).

7. CONCLUSIONS AND FUTURE WORK

This paper described the technological infrastructure and the data mining and process mining techniques used to analyse the strategies that distinguish high from low achieving groups in the classroom. We presented a novel approach to mine traces of collaboration of students working face-to-face on an activity linked with the regular curricula and supported by a number of teacher-orchestrated interactive tablespots. Our goal was to exploit students' data that was unobtrusively captured in an authentic classroom in contrast to a controlled experimental setting. This can make our approach immediately applicable in a real classroom context equipped with the technology required.

Sequential frequent mining was applied to find patterns of activity that differentiate groups. Results revealed interesting patterns that indicated students in high achieving groups worked more often in parallel, interacted with other students' objects and mostly focused on the crucial elements of the problem to solve. The fuzzy miner tool was used to model the process that groups followed by grouping and categorising students' actions. This modelling proved effective in helping distinguish part of the process followed by groups. High achieving groups tended to build their concept map interweaving periods of focused activity with periods of tabletop inactivity. Low achieving groups, by contrast, presented more transitions between different categories of blocks of activity including periods with only actions that caused high impact on the map. We also found that important strategies can be mined from early data. Our analysis was only performed on the data captured from the first activity of the classroom sessions. This gives time for the results of the analysis to be used by facilitators or group members in the classroom.

The knowledge generated by the sequence patterns and the fuzzy models can be used in several valuable ways. Firstly, derived groups' indicators can be displayed in a processed form on the teachers' dashboard to help them adapt in real-time the support to groups that might need closer attention. Secondly, the findings can be used to generate indicators of group learning to be shown

to the teacher for after-class reflection or re-design of the activity or to reflect on students' performance or assessment. Thirdly, this information can be the basis to build student models that can be shown to learners to encourage reflection and self-assessment.

We acknowledge some current limitations of our approach. The first is that the technology to capture students' actions is not yet developed to automatically record verbal interactions in the classroom, which is crucial in collaborative work. However, our approach proved that even modest interaction data can provide insights about their strategies. Regarding the configuration of the data mining method, especially for the Fuzzy process mining, changing some thresholds can produce different results. For example, the size of blocks of activity was set to the mean number of actions between two periods of inactivity (13 actions). We explored the generation of fuzzy models using two more heuristics for the maximum block size: $\mu/2$ and $\mu+\sigma$. We obtained conformance rates as low as 60% for the block size heuristic of $\mu/2$, and very similar fuzzy models and conformance rates for the heuristic $\mu+\sigma$ compared to the one we used in the study. Even when these rates are lower than the ones we obtained using the μ heuristic, a deeper analysis of the configuration of the approach is part of the work in progress.

Our current work includes the exploration of ways to present the results of our approach to the teacher, in real time and for after class analysis. We also aim to connect the students' data that can be captured when they work at the tabletop with other activities that they perform, for example, through online learning systems.

8. REFERENCES

- [1] AlAgha, I., Hatch, A., Ma, L. and Burd, L. Towards a teacher-centric approach for multi-touch surfaces in classrooms. *International Conference on Interactive Tabletops and Surfaces 2010 (ITS 2010)* ACM (2010), 187-196.
- [2] Anaya, A. and Boticario, J. Application of machine learning techniques to analyse student interactions and improve the collaboration process. *Expert Systems with Applications*, 38, 2 (2011), 1171-1181.
- [3] Baker, R. S. and Yacef, K. The State of Educational Data Mining in 2009: A Review and Future Visions. *Journal of Educational Data Mining (JEDM)*, 1, 1 (2009), 3-17.
- [4] Cuban, L., Kirkpatrick, H. and Peck, C. High access and low use of technologies in high school classrooms: Explaining an apparent paradox. *American Educational Research Journal*, 38, 4 (2001), 813-834.
- [5] Do-Lenh, S. *Supporting Reflection and Classroom Orchestration with Tangible Tabletops*. PhD dissertation École Polytechnique Fédérale de Lausanne, 2012.
- [6] Günther, C. and Aalst, W. P. Fuzzy Mining – Adaptive Process Simplification Based on Multi-perspective Metrics. *Business Process Management* Springer Berlin Heidelberg (2007), 328-343.
- [7] Kinnebrew, J. S. and Biswas, G. Identifying Learning Behaviors by Contextualizing Differential Sequence Mining with Action Features and Performance Evolution. *Educational Data Mining* (2012), 57-64.
- [8] Leonard, S., Mary Elizabeth, S. and Samuel, S. D. Effects of Small-Group Learning on Undergraduates in Science, Mathematics, Engineering, and Technology. *Review of educational research*, 69, 1 (1997), 21-51.
- [9] Martinez-Maldonado, R., Collins, A., Kay, J. and Yacef, K. Who did what? who said that? Collaid: an environment for capturing traces of collaborative learning at the tabletop. *International Conference on Interactive Tabletops and Surfaces (ITS2011)* ACM (2011), 172-181.
- [10] Martinez-Maldonado, R., Dimitriadis, Y., Kay, J., Yacef, K. and Edbauer, M.-T. Orchestrating a multi-tabletop classroom: from activity design to enactment and reflection. *International Conference on Interactive Tabletops and Surfaces (ITS 2012)* ACM (2012), 119-128.
- [11] Martinez-Maldonado, R., Dimitriadis, Y., Kay, J., Yacef, K. and Edbauer, M.-T. MTClassroom and MTDashboard: supporting analysis of teacher attention in an orchestrated multi-tabletop classroom. *Submitted to International Conference on Computer Supported Collaborative Learning (CSCL2013)* (2013).
- [12] Martinez-Maldonado, R., Yacef, K., Kay, J., Kharrufa, A. and Al-Qaraghuli, A. Analysing frequent sequential patterns of collaborative learning activity around an interactive tabletop. *International Conference on Educational Data Mining (EDM 2011)* (2011), 111-120.
- [13] Martinez-Maldonado, R., Yacef, K., Kay, J. and Schwendimann, B. An interactive teacher's dashboard for monitoring multiple groups in a multi-tabletop learning environment. *International Conference on Intelligent Tutoring Systems (ITS 2012)* (2012), 482-492.
- [14] Müller-Tomfelde, C. and Fjeld, M. Tabletops: Interactive Horizontal Displays for Ubiquitous Computing. *Computer*, february (2012), 78-81.
- [15] Nousiainen, M. and Koponen, I. T. Concept maps representing knowledge of physics: Connecting structure and content in the context of electricity and magnetism. *Nordic Studies in Science Education*, 6(2010), 155-172.
- [16] Novak, J. Concept mapping to facilitate teaching and learning. *Prospects*, 25, 1 (1995), 79-86.
- [17] Ocker, R. and Yaverbaum, G. Asynchronous Computer-mediated Communication versus Face-to-face Collaboration: Results on Student Learning, Quality and Satisfaction. *Group Decision and Negotiation*, 8, 5 (1999), 427-440.
- [18] Olson, J. S., Teasley, S., Covi, L. and Olson, G. The (currently) unique advantages of collocated work. *Distributed work*(2002), 113-135.
- [19] Pechenizkiy, M., Trcka, N., Vasilyeva, E., Aalst, W. v. d. and Bra, P. D. Process Mining Online Assessment Data. *2nd International Conference On Educational Data Mining* (2009), 279-288
- [20] Perera, D., Kay, J., Koprinska, I., Yacef, K. and Zaiane, O. Clustering and Sequential Pattern Mining of Online Collaborative Learning Data. *IEEE Transactions on Knowledge and Data Engineering* 21, 6 (2009), 759-772.
- [21] Race, P. *A briefing on self, peer & group assessment*. Learning and Teaching Support Network, York 2001.
- [22] Zhang, D., Zhao, J. L., Zhou, L. and Jay F. Nunamaker, J. Can e-learning replace classroom learning? *Communications ACM*, 47, 5 (2004), 75-79.