Towards Adaptive Learning Support on the Basis of Behavioural Patterns in Learning Activity Sequences

Mirjam Köck
FIM, Johannes Kepler University
Linz, Austria
koeck@fim.uni-linz.ac.at

Alexandros Paramythis
FIM, Johannes Kepler University
Linz, Austria
alpar@fim.uni-linz.ac.at

***Abstract*—Monitoring and interpreting sequential user activities contributes to enhanced, more fine-grained user models in e-learning systems. We present in this paper different behavioural patterns from the domain of problem-solving that can be determined by targeted, ultimately automated clustering. For the identification of these patterns, we apply a new approach – based on the modeling of activity sequences – to real-world learning activity sequence data, monitored via an Intelligent Tutoring System. This paper describes the identified behavioural patterns, explains the process used for their detection, and compares the patterns to related ones in earlier literature. It further discusses implications of the patterns themselves, and of the employed approach, on adaptively supporting individual and group-based collaborative learning.**

***Keywords*-data mining; clustering; problem-solving styles; learning activities; adaptivity;**

I. Introduction

During the past decades, the field of adaptive systems has been steadily gaining attention within the research community. In parallel, the number of potential application domains has been increasing rapidly. One such application domain is the field of e-learning where adaptivity can be used not only for the personalization of learning content and navigation through learning material, but also to actively support the learning process itself for individual users and groups. In the later perspective the learning process is treated as consisting of interconnected learning activities (see, e.g. [1]), rather than the passive or active consumption of knowledge.

A related trend in adaptive systems research has been the increasing utilization of data mining and machine learning techniques for the derivation of user-centered models (see, e.g. [2], [3], [4] or [1]). These techniques can contribute to the enhancement of adaptivity by allowing for a more fine-grained analysis of user activity data, which can subsequently lead to more fine-grained information in the respective user models. This, in turn, can lead to better tailored, personalized support throughout the learning process.

In this context, this paper addresses a specific area of learning, namely problem-solving, and discusses specific behavioral patterns, their significance and their relation to ones reported in recent literature (e.g., [5]), and types of adaptive interventions that can be potentially employed in response to the detected patterns. The approach applied here for the detection of patterns is based on targeted and ultimately automated clustering of users' problem-solving sequences represented by Discrete Markov Models (DMMs). Detection has addressed: (a) predefined concrete problem-solving styles, (b) new styles along known, predefined learning dimensions, and (c) new learning dimensions and related concrete styles. The paper discusses how the findings presented here are, although gained through analysis of individual users' behaviour, also applicable to the enhancement of collaborative learning processes.

The rest of this paper is structured as follows: section II briefly describes the approach used for different levels of pattern detection. Section III explains different types of patterns discovered during the pattern detection process and compares them to related ones. Section IV explains how the detected patterns can be used in adaptive systems, and section V discusses implications for adaptive collaboration support. Finally, section VI summarizes the paper and provides an overview of related work.

II. Multi-Level Clustering

This section provides a brief overview of the activity data used in the presented work, the preprocessing and analysis processes, and the clustering mechanism used for pattern detection. A detailed description of the applied approach can be found in [6].

A. Data Preprocessing and Analysis

The data used in this study consists of problem-solving activity sequences, carried out in (and monitored by) the Intelligent Tutoring System (ITS) Andes [7] and published via the PSLC DataShop [8]. Andes provides students with tasks to be solved, and stores, for every student, the respective activity sequence for a task. This results in a sequence pool, each related to a student/task combination. Our experiments operated with data from a Physics course of the US Naval Academy (2007 - 2009). Our system preprocesses base data and converts the sequences to DMMs that are later used for analysis and clustering. The conversion process starts with an empty DMM, containing default transition- and prior probabilities. Next, the activity sequences are sent through the model, which is updated with every incoming activity. The states in the DMMs correspond to actions a student can perform within the problem-solving process (see Figure 1). The goal is to infer from the sequences different strategies students employ to deal with the tasks. Thus, aggregating the sequences to statistical summaries would bear the risk of losing information about the process that could be indicative of specific types of behavior.

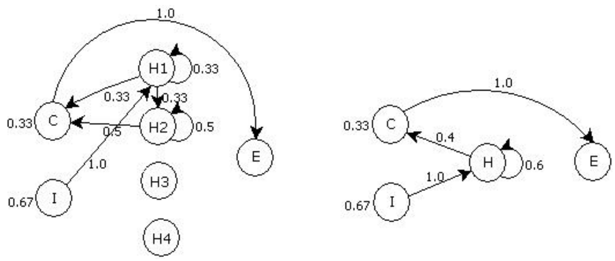

Figure 1. This figure shows two DMMs used for the clustering approach. Notation: C stands for the submission of a correct answer, I for an incorrect answer, $H*$ for a hint request, and E is an artificial end-state indicating that a student finished a sequence. The numbers next to nodes denote their prior probability, the numbers next to transitions denote the transition probability (both showing exemplary values here). As there are different types of hints, we introduced (and used in our experiments) a model distinguishing between different H states (left), and one aggregating them (right).

B. The Clustering Process

Clustering was performed on data sets containing serialized versions of the models created during preprocessing, and additional basic statistical metrics. The goal has been the detection of problem-solving styles, either based on pre-existing descriptions, or with the aim to discover new ones. We introduced different clustering levels that correspond to the different kinds of patterns described in the next section: *Level I (pattern-driven)* uses well-established, predefined problem-solving styles and aims at discovering them in students' behavior. *Level II (dimension-driven)* operates at the level of dimensions rather than concrete styles, i.e., the system tries to identify given dimensions and discover concrete styles along these dimensions. *Level III (open discovery)* does not receive definitions of styles or dimensions as input but aims at the automatic discovery of both.

An inherent (and of fundamental importance) part of the approach described here is the employment of a set of optimization metrics that are applied to determine when an "optimum" cluster setting has been reached (thus, further clustering is not necessary). The "optimum" cluster is the one with the best results according to specific metrics. These metrics can vary and are usually adapted for every different scenario and clustering purpose. Specifically, we evaluate as metrics the distribution of students and problems to clusters with the help of entropy measures, and determine the variance of attribute values in the clusters, as well as the clusters' potential to serve for the prediction of success within a problem-solving sequence. These indices can then be weighted to emphasize specific factors and are, in combination, capable of determining the quality of a setting in the specific domain of problem-solving. A detailed description of the approach used can be found in [6].

III. Pattern Detection at Different Levels

This section describes different patterns that were identified by the different levels of clustering.

A. Level I (Pattern-Driven)

To demonstrate pattern detection at this level, the well-known problem-solving style *Trial and Error* [9] (also referred to as *Trial and Success*) was chosen, describing behaviour that is based on chance at the beginning, and on learning by making mistakes later. In this context, the available attributes (i.e., the features represented by the DMMs and the additional statistical information) were evaluated with regards to their potential to contribute to the identification of *Trial and Error* behaviour. Specifically, we expected a person with this behaviour to have high prior probabilities for incorrect answer submissions and low-to-medium prior probabilities for correct answers. Note, that the probability of guessing a correct answer is lower than the probability of getting it wrong by chance; usually there is only one correct answer as opposed to several incorrect possibilities. Additionally, such a person was expected to have a low hint request rate, a relatively high rate of incorrect submissions, and only a weak tendency to request hints after an incorrect submission. The attributes best corresponding to these characteristics were selected, stored in a new data set, and clustering was performed on this set. The results of a cluster configuration with 8 clusters showed two clusters providing a clear identification of the *Trial and Error* problem-solving style. The procedure as described above can be applied for every other predefined problem-solving style.

B. Level II (Dimension-Driven)

At this level, we are concerned not with the "recognition" of expected behaviour when it occurs, but rather with establishing whether it is possible to identify distinct behavioural patterns in relation to specific semantic dimensions of the activities being analyzed. This translates into performing clustering along known learning dimensions, in order to identify concrete problem-solving styles the learners may exhibit.We have chosen *Help-Seeking* behaviour [10], [11] as a well-known learning dimension and identified the behaviour elements that we expected to be defining here. These included the rate at which learners request help, the probabilities for hint requests following an incorrect answer submission, the prior probability for hint requests, and the probabilities for hint requests that occur in sequences.

Again, clustering showed clear variations of the examined behaviour, indicating that the attributes selected formed a coherent whole, capable of exposing the elements of variability in the learners' behavior. Next, different concrete styles within this learning dimension were analyzed, which successfully led to the detection of the help-related problem-solving styles described in Table I. The four *Help-Seeking* styles identified here can be explained as follows: A problem solver of type H_1 shows *Trial and Error* behaviour and tends to request hints in sequences, whereas type H_2 makes sure not to submit wrong answers but requests a lot of help, even before having tried. This might lead to the assumption that this problem solver uses the help functionality instead of sufficient preparation. Type H_3 does not request help right at the beginning and does not request help too often; when help is requested though, this is done in sequences. This may be indicative, for instance, of a learner that is interested in really understanding a problem before continuing. Type H_4 is very similar to H_2, and in settings with a lower number of clusters these styles might have been combined.

We can compare the results at this level to the help-seeking model discussed in [5]. The authors introduce a taxonomy of "help-seeking bugs" in students' behaviour and list the following categories: *Help Abuse, Help Avoidance, Try-Step Abuse and Miscellaneous Bugs*. The type *Help Abuse* comprises behaviour like clicking through hints or asking for hints even if it would not be necessary because the student would be skilled enough to solve the task without help. The H_2 and H_4 types identified by our system partly correspond to this *Help Abuse* type in that a H_2/H_4 problem solver may also show such behaviour instead of spending more time on understanding the content before. Our model can, however, additionally identify if the student uses help before or after trying to submit an answer first. This undesired kind of behaviour can also be compared to *gaming the system* as explained in [12]. *Try-Step Abuse* can be compared to the *Trial and Error* behaviour as shown by type H_1 who also tends to solve a problem too early even if not sufficiently skilled yet. Parallels can also be drawn between the H_3 type and the *Help Avoidance* style described in [5] concerning the general tendency to keep the amount of requested help low. Yet, the *Help Avoidance* type mainly considers trying unfamiliar steps without help and could thus also be described as a subcategory of the *Try-Step Abuse* type. In our case, H_3 and H_1 are clearly distinct as the behaviour of a problem solver of type H_3 can also be attributed to the desire to avoid the submission of incorrect answers. Thus, level II pattern detection has not only confirmed the taxonomy of "help-seeking bugs" described in [5] but also added some distinct aspects to it.

C. Level III (Open Discovery)

Pattern detection at this level goes one step further than its predecessor and is intended to perform open-ended analysis with the goal of identifying, firstly, potential new dimensions of learning behaviour, and, secondly, concrete patterns within each dimension. Here, this process has again been targeted towards the identification of concrete types of problem-solving behaviour. This level is controlled by the system (excluding the assessment and interpretation of results that need to be performed by a human operator) whose task it is to automatically select feature combinations with high discriminatory capacity, create new data sets containing attributes of one feature combination each, perform clustering on each of the new data sets, and analyze the resulting clusters for significant trends in order to autonomously detect problem-solving styles.

Table I
THIS TABLE SHOWS FOUR PROBLEM-SOLVING STYLES IN THE *Help-Seeking* DIMENSION DISCOVERED BY THE CLUSTERING PROCESS (8 CLUSTERS). THE REMAINING CLUSTERS NOT SHOWN HERE CONTAIN NON-*Help-Seeking* BEHAVIOUR. THE SYNTAX IS TO BE READ AS FOLLOWS: THE PERCENTAGE RESULTS HAVE BEEN ABSTRACTED TO THE FIVE CATEGORIES *very low, low, medium, high, very high*, REPRESENTED BY THE MORE EASY TO READ IDENTIFIERS $--$, $-$, o, $+$, $++$.

Style	**Size**	**P.P. I.A.**	**P.P. H.R.**	**P.H. I.A.**	**S.H.**	**P.I.A.**	**P.H.R.**
H_1	- -	o	o	o	+	o	o
H_2	+	- -	+	-	++	- -	+
H_3	o	-	-	-	+	-	o
H_4	o	- -	+	-	++	- -	+

Size	Number of Instances in the Cluster
P.P.I.A.	Prior Probability for Incorrect Answers
P.P.H.R.	Prior Probability for Hint Requests
P.H.I.A.	Probability for Hint Requests Following Incorrect Answers
S.H.	Subsequent Hints
P.I.A.	Percentage of Incorrect Answers
P.H.R.	Percentage of Hint Requests

The process of selecting subsets of the initial feature set, clustering on them and determining the quality of the set according to, e.g., its discriminatory capacity, is a relatively simple one, but cannot be done by humans, because of the high number of calculations involved. The problem of finding combinations is of exponential complexity, hence we do not compute all possible combinations but limit the number of features in the resulting combination sets.

Using the computed feature combinations, the process continues by creating a "copy" of the original data set containing only the selected features and values. This results in a high amount of data, all depicting different aspects of the same activities. After the subsequent clustering process, the results are compared in terms of their average cluster quality for a specific feature set. The algorithm used here for measuring cluster quality is based on Linear Discriminant Analysis (LDA) as described in [13], maximizing the distance between cluster centroids and minimizing the average distance between the elements within the clusters.

The top ranked feature sets, based on the results of cluster quality evaluation, become the system's recommendation as potentially meaningful dimensions. These recommended feature sets are then analyzed by a human investigator who makes a decision about what sets to pass back to Level II clustering in order to detect concrete problem-solving types. The investigator's decisions are based on semantic reflections a machine would hardly be able to provide.

In the case study discussed here, the top ranked results suggest variations of a *Help-Seeking* dimension similar to the one we manually defined for Level II, based on descriptions in the related literature. Table II shows experimental clustering results using the top ranked feature set of each group (a group contains all feature sets with the same number of features in them). The results listed there are an example of what a human observer would see when applying level II clustering on the dimensions suggested by level III. Of course, a human observer would be provided not only one top ranked feature set but several. The results can be analyzed as follows (treating each feature set as a potential dimension of problem solving behavior, and identifying different types of behavior for each such dimension).

$Rank1, n = 1$**:** This dimension, defined by one single feature, models the users' tendency to request help in sequences. The clusters show a clear distinction between different types of behaviour (e.g., cluster 2 vs. cluster 4). The concrete types along this dimension are $T_{1.1}$ showing a strong tendency to request help in sequences (clusters 1, 3, 4), $T_{1.2}$, not requesting help in sequences (cluster 2), and $T_{1.3}$, occasionally requesting help in sequences (cluster 0).

$Rank1, n = 2$**:** This dimension is defined by two features, modeling the tendency to request help in sequences and to end a problem-solving sequence with a hint request (i.e., in most cases, without having submitted a final solution). Three clusters (1, 2, 4) show similar results and can be summarized as type $T_{2.1}$, tending to request help in sequences and not to conclude a problem with a hint request. The second type identified here, $T_{2.2}$, is described in clusters 0 and 3, where users request help in sequences occasionally and also occasionally end a problem-solving sequence with a hint.

$Rank1, n = 3$**:** This dimension is defined by three features, adding the percentage of help requests to the attributes already explained. Here, we could identify significant types as follows. $T_{3.1}$ learners do not request help in sequences, do not end problem-solving sequences with help requests, and in general request only little help (cluster 2). $T_{3.2}$ learners tend to request help in sequences but do not end problems with help requests and in general request a lot of help (clusters 1 and 3). In cluster 4 we can find the behaviour of $T_{3.3}$, not requesting much help and when so, not in sequences, but showing a strong tendency to end problem-solving sequences with hints. In cluster 0 we can see that this dimension is more expressive than the previous ones. Here, the percentage of requested help steps is low enough to round down to 0.00. Thus, the values indicating tendencies of requesting help in sequences and of ending a sequence with a hint are not relevant. We conclude that the types discovered before are useful, but only in combination with a basic statistical indicator on the general use of help. We define type $T_{3.4}$ behaviour as tending to not use help at all.

$Rank1, n = 4$**:** This dimension adds to the feature pool the probability of submitting a wrong answer directly after a hint request. Cluster 0 behaves as before. Cluster 3 identifies a type of behaviour $T_{4.1}$ that has not been detected by the previous dimensions; learners of this type stop their problem-solving sequence with a hint request in 100% of the cases while not showing a generally very low help request rate. This kind of behaviour is rare and here only affects 1% of the problem-solving sequences. In the clusters 1, 2, and 4 we identify the following types: $T_{4.2}$ shows a very high help request rate, a strong tendency to request help in sequences, and a very low rate of incorrect submissions or quits after a hint. $T_{4.3}$ differs from $T_{4.2}$ only in a slightly lower rate of help requests and a medium rate of quits after a hint, and $T_{4.4}$ is defined by a medium rate of help requests, a medium rate of incorrect attempts or further help requests after a help request, and a low rate of quits after a help request.

$Rank1, n = 5$**:** This dimension comprises, further to the general help rate and the tendency to request help in sequences, the prior probability for help, i.e. when users request help as a first step, and the rate of requested hints directly following a correct attempt. Clusters 1 and 2 define $T_{5.1}$ by a high help rate, a high prior probability for the use of help, and the tendency to request help in sequences. $T_{5.2}$ is derived from cluster 3, defined by a high help sequence rate, a medium overall help request rate and a relatively low prior probability for help. $T_{5.3}$ is similar to $T_{5.2}$ in the prior probability for help requests but differs in other aspects and includes a lower rate of help sequences and, in general, a lower help rate.

$Rank1, n = 6$**:** This dimension adds to the features in $Rank_G = 1, n = 5$ the probability of a wrong answer after a hint request. Clusters 0 and 2 identify type $T_{6.1}$ and show a relatively high prior probability for help requests, a high general help rate and a tendency to help request sequences. Clusters 3 and 4 both show a low help rate, a low prior probability for help requests and a very low probability of help requests after a correct submission. $T_{6.2}$ is characterized by a high help sequence rate, whereas $T_{6.3}$ shows a very low help sequence rate but a relatively high percentage of incorrect attempts after a hint.

$Rank1, n = 7$**:** This dimension adds to the features in $Rank1, n = 6$ the probability of a help request being the last activity in a sequence. Cluster 3 shows a new type of learners ($T_{7.1}$) exhibiting a strong tendency to close a sequence with help, which in most cases is indicative of "giving up" before the problem was solved.

From the above, we can conclude that dimensions with very few features can be indicative of problem-solving types but results are prone to being distorted. Yet, a very high number of features may not allow for the identification of the most significant types but rather suggest a range of "subtypes" many of which could be combined. In order not to fall prey to either of these threats, we suggest a medium number of features for the purpose of dimension detection that lies between a fourth and a third of the overall count.

The results presented here show that the *Help-Seeking* dimension dominates. Yet, if we consider the 5 top-ranked results of every group, we can already discover a different dimension comprising: *the probability for a correct answer as the first action*; *the probability for a hint request as the first action*; *the probability for a hint request directly following a correct answer*; *the probability for an incorrect answer directly following a hint request*; *the probability for subsequent hints directly following a hint request*; and, *the percentage of hint requests*. Clustering along this dimension, we get, among others, student types characterized by: (a) very high probability for correct answers at first attempt, extremely low help request rate, low help sequence rate, and (b) a medium rate of correct answers at first attempt, a low rate of initial help requests, and a relatively low overall help rate. Type (b) is very similar to the *Trial and Error* type in section III-A where we did not operate at the level of dimensions yet but considered a predefined concrete style.

Table II
THIS TABLE SHOWS EXPERIMENTAL CLUSTERING RESULTS (5 CLUSTERS) BASED ON THE TOP RANKED FEATURE SETS (TRFS). THE VALUES FOR THE CLUSTERS DENOTE THE MEAN FOR THE RESPECTIVE FEATURE IN THIS CLUSTER.

TRFS	***Features***	$C0$	$C1$	$C2$	$C3$	$C4$
$n = 1$	*P* for sequential hints	0.25	0.69	0.03	0.52	0.77
$n = 2$	*P* for sequential hints	0.25	0.69	0.75	0.15	0.51
	P for quitting after a hint	0.25	0.06	0.02	0.27	0.01
$n = 3$	*P* for sequential hints	0.25	0.74	0.01	0.60	0.05
	P for quitting after a hint	0.25	0.03	0.00	0.06	0.88
	% of hint requests	0.00	0.67	0.10	0.34	0.28
$n = 4$	*P* for an incorrect answer following a hint	0.25	0.03	0.08	0.00	0.23
	P for sequential hints	0.25	0.74	0.69	0.00	0.34
	P for quitting after a hint	0.25	0.04	0.26	1.00	0.08
	% of hint requests	0.00	0.71	0.48	0.24	0.23
$n = 5$	*P* for a hint request being the first action	0.00	0.76	0.45	0.11	0.09
	P for a hint request following a correct answer	0.00	0.01	0.04	0.01	0.02
	P for sequential hints	0.24	0.72	0.70	0.66	0.18
	P for a hint request being the first action	0.30	0.12	0.14	0.09	0.18
	% of hint requests	0.00	0.72	0.59	0.34	0.09
$n = 6$	*P* for a hint request being the first action	0.75	0.00	0.45	0.11	0.07
	P for a hint request following a correct answer	0.00	0.00	0.05	0.01	0.01
	P for an incorrect answer following a hint	0.04	0.25	0.07	0.10	0.59
	P for sequential hints	0.73	0.25	0.67	0.52	0.01
	P for a hint request being the first action	0.71	0.03	0.42	0.10	0.09
	% of hint requests	0.72	0.00	0.58	0.15	0.07
$n = 7$	*P* for a hint request being the first action	0.00	0.43	0.74	0.31	0.09
	P for a hint request following a correct answer	0.00	0.04	0.04	0.05	0.01
	P for an incorrect answer following a hint	0.25	0.07	0.04	0.01	0.18
	P for sequential hints	0.25	0.71	0.73	0.13	0.48
	P for quitting after a hint	0.25	0.04	0.04	0.83	0.03
	P for a hint request being the first action	0.03	0.41	0.71	0.32	0.08
	% of hint requests	0.00	0.57	0.72	0.28	0.26

IV. POTENTIAL ADAPTIVE BEHAVIOUR

This section outlines possible ways in which the styles and dimensions discussed in section III can be used within adaptive e-learning settings. In general, system interventions can be grouped into means of individual user support (see, e.g., [14]) and means of collaboration support (see, e.g., [15] or [16]). In what follows we describe potential system interventions for individual user support, and their applicability for the different problem solving types discussed before.

A. Hint Tailoring

The system may react to student behaviour by limiting the available help. Example ways in which this can be done include: (a) *reducing or increasing the number of hints*, or (b) *reducing or increasing the granularity of information within the hints*. The first approach would be applicable, for instance, for the types H_2, H_4 (these seem to show a natural aversion for submitting incorrect answers but use a disproportionate amount of help) and the types in the open dimensions $T_{2.*}$, $T_{4.*}$, $T_{6.*}$ and $T_{7.*}$ (e.g., students who tend to quit after a hint request). The second approach might also be applicable for the types H_2 and H_4 and the types in the open dimensions $T_{2.*}$, $T_{4.*}$, and $T_{7.*}$.

B. Hint Withholding

The system may prevent users from accessing hints at a specific time by actively withholding them. This can be necessary when students use hints before having tried to understand the learning content, or where a clearly disproportionate amount of help is used. This kind of intervention is applicable for the types in the open dimension $T_{5.*}$ (students not being well prepared, or trying to over-exploit the help functionality to reduce own efforts) if the respective users show a tendency to request help as a first activity.

C. Proactive Hint Delivery

The system may also try to encourage the use of help by actively offering it, i.e., even without users having requested it. This kind of intervention can be useful if students show a low tendency to use help, or in cases where a user would be able to solve a problem with a hint but does not request one. Proactive hint delivery is applicable for students showing the *Trial and Error* style (here the goal is to prevent students from making uneducated guesses), for the type H_3

in the *Help-Seeking* dimension (these students show a high inhibition threshold regarding help request), and the types in the open dimension $T_{3.*}$ (students that show a tendency not to use help at all).

V. Implications for Collaboration

In addition to individual user support, adaptive systems can actively support collaboration. Adaptive collaboration support can be split into two phases: *adaptive support for collaboration establishment* and *adaptive support during the collaboration process* [17]. Our approach is applicable in both cases; yet, the nature of the information analyzed here is better suited for collaboration establishment, usually based on learners' learning characteristics [17], [18].

The actual way of adaptively supporting collaboration establishment is strongly dependent on the respective learning scenario and underlying teaching concepts and learning theories. Adaptive collaboration establishment support includes encouraging students to cooperate with others, or recommendations of tools to use for collaboration, or partners to collaborate with [19]. Group synthesis recommendations are based on specific rules that may consider users' backgrounds, interaction behavior, etc. In general, it may be desirable for the system to group students that could potentially benefit from cooperation, considering criteria like complementarity or competitiveness [20].

In [20] and [21], the authors analyze the student learning style based on the Felder and Silverman model [22], [23] which categorizes learning styles along five dimensions (*active/reflective, sensing/intuitive, visual/verbal, sequential/global, inductive/deductive*), and conclude that (a) learning styles affect the performance of students when working together, (b) for the dimensions *active/reflective* and *sensing/intuitive*, the mixed pairs tend to work better, (c) heterogeneous groups in general get better results, and (d) students themselves tend to group randomly without respect to their learning styles. Their findings show that it is a worthy goal to use learners' models as a basis for group synthesis recommendations, and that learning styles are a potentially relevant criterion to base grouping algorithms on. We expect effects to be even more pronounced when individuals' problem solving styles are taken into consideration when deciding on the synthesis of groups that are to engage in problem-based collaborative learning.

The second phase of adaptive collaboration support, *adaptive support during the collaboration process* [17] requires, in addition to the analysis of individual users' activities, the same kind of analysis of activities in group settings. The base data, monitored via an arbitrary collaborative environment, may include activities within tools for multi-user communication and cooperation like, for instance, chat, forum, audio- and videoconferencing, wiki, shared resources, etc. In such settings, a statistical analysis can provide a basic model of a user's behaviour, including, for instance, this user's level of activity in the group, tendencies to correct other users' contributions, or to initiate new ones. Of more relevance to the approach discussed in this paper is the analysis of group activity sequences. The related literature proposes computational models for analyzing such sequences to determine metrics such as the *centrality* of group members, or the *cohesion* of a group (see, e.g., [24]). Nevertheless, more "dialogical" forms of analysis have typically been constrained to group discussions, and require annotation of activities by experts, or content analysis of exchanged messages (again, see [24] for an example).

The application of our approach for the analysis of activity sequences in a group is expected to have impact on two fronts: the detection of behavioural patterns of individual learners, with regards to their contact within group learning settings; and, the detection of patterns that emerge in the behaviour of the group as a whole (and may require intervention). The first type of information can potentially enhance the collaboration process, although based on the analysis of individual users' activities, because individual user characteristics become part of the group model and influence the collaboration behaviour. The second type of information can not only be fed back into the user models and become basis for further adaptive collaboration support, but potentially form the basis for novel kinds of adaptation for collaboration. Each of these cases would require a different graph-based representation of activity sequences (replacing the problem solving-oriented DMMs shown here). Furthermore, it may be necessary to revisit the metrics used to drive the clustering process, so as to better capture the semantics of group-based behavior (for instance, defining measurable indices that characterize a group's "success", in analogy to a learner's solving a given problem). A successful application of our approach in this context would arguably be of high value, as recent work has demonstrated that when adverse group behaviour can be detected, there is ample potential in adaptively supporting online communities [25].

VI. Summary and Related Work

In this paper, we discussed how different kinds of patterns in learners' problem-solving behaviour can be automatically discovered. The process is based on a clustering approach applied at three different levels, each tailored to different pattern detection purposes: recognizing predefined concrete problem-solving styles; detecting new styles along known learning dimensions; and, automatically identifying both new learning dimensions and related concrete styles. We described the patterns we detected during our experiments with real-world problem-solving data. Further, we explored possibilities of how the newly gained information in the user model can be used within adaptive systems, aiming at enhancing the (individual and group) learning process, and adaptivity in e-learning systems in general. The interventions proposed here are not regarded as the best possible approaches in the respective cases (something that would depend on the didactic approach employed), but rather demonstrate how the adaptation cycle can be "completed".

Recent work on the discovery of, and reaction to, different learning and problem-solving styles has put emphasis on implications that specific types of behaviour have on the learning process and on concepts of how intelligent systems can tailor their interaction with the user to the respective types. In [11], for instance, the authors analyze help-seeking behaviour in general and describe how help can be designed in intelligent environments in order to support its users individually. The authors state, based on a series of studies, that very often help is used inefficiently. Moreover, they report that generally, students tend to prefer the hints closest to the full answer, which often prohibits use of the best suited type of help. In some cases, the provision of on-demand help might even interfere with productive learning processes. Thus, the provision of help must be very carefully tailored to users' help-seeking behaviour. Another interesting finding is that students without high prior knowledge have better learning gains when they seek help more often. Additionally, students who show a balanced use of different types of help learn more than those who have a tendency to focus on a specific type. The authors also list learner-related factors that have to be considered when deciding on how to provide help: prior knowledge, self-regulation, age and gender, goal-orientation, and epistemological beliefs.

In [5], the authors focus on guidance adapted to students' meta-cognitive abilities and analyze whether and how their approach can help students become better help-seekers and subsequently better learners. The authors evaluate help-seeking behaviour and identify positive and negative ("help-seeking bugs", see section III-B) factors. A pilot study with a small number of students that included a Help Tutor led to the conclusion that such a tutor is generally accepted by the learners, but did not provide findings regarding its influence on students' help-seeking and learning performance. This is discussed in, e.g., [26], [27], or [28]. [26] includes an evaluation of the aforementioned Help Tutor, which showed that the tutor was generally able to improve help-seeking behaviour but was less successful in improving all help-seeking related actions, and in improving learning (maybe a result of suboptimal timing according to the authors).

In [27], the role of feedback in preparation for future learning is discussed, and different types of feedback are compared. Additionally, the authors evaluate guided meta-cognitive feedback and conclude that although directed feedback may allow the student to quickly achieve immediate goals set by the learning environment, guided metacognitive feedback performs better in preparing the student for learning. In [28], the authors discuss the effect of hints and model answers when students experience difficulties in solving applied problems. In this study, the systematic use of hints led to a significant improvement of students' problem-solving skills. Although the concrete results may relate to the field of application (the study was specifically tailored to an application field in secondary physics education) or the target group, they indicate that hints can be a highly relevant factor in the process of problem-solving.

A different perspective on adaptive guidance and help provision is a problem known as "gaming the system" behaviour (see [12] and [29]). "Gaming the system" aims at deceiving the system, regarding, for instance, learning speed, knowledge level, or contents looked at. It is the intelligent system's task to detect this undesired kind of behaviour and to counteract if necessary. As "gaming the system" can potentially distort the learner model that becomes the basis for adaptive system interventions, it is important for our future work to consider the findings of, e.g., [12], where the authors discuss ways to adapt to this kind of behaviour.

In relation to the projected future extensions of our work towards the analysis of collaborative learning behaviour, one can identify similarities between the use of DMMs to represent activity sequences, and "behavioural transition diagrams" (see, e.g., [30]), which also encapsulate the concepts of activities as "nodes", and the utilization of transition probabilities to derive a "significance" metric that is used as the weight of edges connecting the nodes. Yet, such diagrams have been traditionally used for the visual exploration of activity sequences rather than the computational derivation of patterns. Also of relevance is the work reported in [25] where it is shown that patterns indicative of dysfunctions in an online knowledge-sharing community can be algorithmically detected, and potentially acted upon.

Acknowledgements

The work reported in this paper is funded by the "Adaptive Support for Collaborative E-Learning" (ASCOLLA) project, supported by the Austrian Science Fund (FWF; project number P20260-N15). Data used in this research was provided by the Pittsburgh Science of Learning Center DataShop, funded by the National Science Foundation award No. SBE-0354420.

References

[1] A. Soller and A. Lesgold, "Modeling the Process of Collaborative Learning ," in *The Role of Technology in CSCL*, 2007.

[2] A. R. Anaya and J. G. Boticario, "Clustering Learners according to their Collaboration," in *Proceedings of the 2009 13th International Conference on Computer Supported Cooperative Work in Design*, 2009.

[3] C. Li and J. Yoo, "Modeling Student Online Learning Using Clustering," in *ACM-SE 44: Proc. of the 44th Annual Southeast Regional Conference*. ACM, 2006, pp. 186–191.

[4] C. Romero and S. Ventura, Eds., *Data Mining in E-Learning*, ser. Advances in Management Information. WITPress.com, 2006, vol. 4.

[5] V. Aleven, B. McLaren, I. Roll, and K. Koedinger, "Toward Meta-Cognitive Tutoring: A Model of Help Seeking With a Cognitive Tutor," *International Journal of Artificial Intelligence and Education*, vol. 16, pp. 101–128, 2006.

[6] M. Köck and A. Paramythis, "Activity Sequence Modeling and Dynamic Multi-Level Clustering for Personalized E-Learning," In review (available upon request from the authors).

[7] K. VanLehn, C. Lynch, K. Schulze, J. A. Shapiro, R. H. Shelby, L. Taylor, D. Treacy, A. Weinstein, and M. Wintersgill, "The Andes Physics Tutoring System: Lessons Learned," *International Journal of Artificial Intelligence in Education*, vol. 15, no. 3, 2005.

[8] K. Koedinger, R. Baker, K. Cunningham, A. Skogsholm, B. Leber, and J. Stamper, "A Data Repository for the EDM community: The PSLC DataShop," in *Handbook of Educational Data Mining*, 2010.

[9] M. Jarvis, *The Psychology of Effective Learning and Teaching*. Nelson Thornes, 2005.

[10] S. Nelson-Le Gall, "Help Seeking Behaviour in Learning," *Review of Research in Education*, vol. 12, pp. 55–90, 1985.

[11] V. Aleven, E. Stahl, S. Schworm, F. Fischer, and R. Wallace, "Help Seeking and Help Design in Interactive Learning Environments," *Review of Educational Research*, vol. 73, no. 3, pp. 277–320, 2003.

[12] R. S. Baker, A. T. Corbett, K. R. Koedinger, S. Evenson, I. Roll, A. Z. Wagner, M. Naim, J. Raspat, D. J. Baker, and J. E. Beck, *Intelligent Tutoring Systems*, ser. Lecture Notes in Computer Science, 2006, vol. 4953/2006, ch. Adapting to When Students Game an Intelligent Tutoring System, pp. 392–401.

[13] A. M. Martínez and A. C. Kak, "PCA versus LDA," *IEEE Transactions on Knowledge and Data Engineering*, vol. 23, no. 2, 2001.

[14] K. R. Koedinger and V. Aleven, "Exploring the Assistance Dilemma in Experiments with Cognitive Tutors," *Educational Psychology Review*, vol. 19, no. 3, 2007.

[15] A. Soller, A. Martínez, P. Jermann, and M. Muehlenbrock, "From Mirroring to Guiding: A Review of State of the Art Technology for Supporting Collaborative Learning," *International Journal of Artificial Intelligence in Education*, pp. 261–290, 2005.

[16] E. Walker, N. Rummel, and K. R. Koedinger, "CTRL: A Research Framework for Providing Adaptive Collaborative Learning Support," *User Modeling and User-Adapted Interaction*, vol. 19, no. 5, pp. 387–431, 2009.

[17] A. Paramythis, "Adaptive Support for Collaborative Learning with IMS Learning Design: Are We There Yet?" in *Proceedings of the Workshop on Adaptive Collaboration Support, in conjunction with the 5th International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH'08)*, 2008, pp. 17–29.

[18] R. M. Carro, A. Ortigosa, and J. Schlichter, *Web Engineering*, ser. Lecture Notes in Computer Science, 2003, vol. 2722/2003, ch. Adaptive Collaborative Web-Based Courses.

[19] R. M. Carro, A. Ortigosa, E. Martín, and J. Schlichter, *Groupware: Design, Implementation, and Use*, ser. Lecture Notes in Computer Science, 2003, vol. 2806/2003, ch. Dynamic Generation of Adaptive Web-Based Collaborative Courses.

[20] E. Alfonseca, R. M. Carro, E. Martín, A. Ortigosa, and P. Paredes, "The Impact of Learning Styles on Student Grouping for Collaborative Learning: A Case Study," *User Modeling and User-Adapted Interaction*, vol. 16, 2006.

[21] S. Liu, M. Joy, and N. Griffiths, "Incorporating Learning Styles in a Computer-Supported Collaborative Learning Model," in *Proceedings of the International Workshop on Cognitive Aspects in Intelligent Adaptive Web-Based Education Systems (CIAWES 2008)*, 2008, pp. 3–10.

[22] R. M. Felder and L. K. Silverman, "Learning and Teaching Styles in Engineering Education," *Journal of Engineering Education*, vol. 78, no. 7, pp. 674–681, 1988.

[23] R. M. Felder and R. Brent, "Understanding Student Differences," *Journal of Engineering Education*, vol. 94, pp. 57–72, 2005.

[24] H.-J. Suh and S.-W. Lee, "Collaborative learning agent for promoting group interact," *ETRI journal*, vol. 28, no. 4, pp. 461–474, 2006.

[25] S. Kleanthous and V. Dimitrova, "Analyzing community knowledge sharing behavior," in *User Modeling, Adaptation, and Personalization*. Springer, Heidelberg, 2010, pp. 231–242.

[26] I. Roll, V. Aleven, B. McLaren, E. Ryu, R. S. Baker, and K. R. Koedinger, *Intelligent Tutoring Systems*, ser. Lecture Notes in Computer Science, 2006, vol. 4053/2006, ch. The Help Tutor: Does Metacognitive Feedback Improve Students' Help-Seeking Actions, Skills and Learning?

[27] I. Roll, V. Aleven, B. M. McLaren, and K. R. Koedinger, "Can Help Seeking be Tutored? Searching for the Secret Sauce of Metacognitive Tutoring," in *Proceedings of the 13th International Conference on Artificial Intelligence in Education*, 2007, pp. 203–210.

[28] H. J. Pol, E. G. Harskamp, C. J. Suhre, and M. J. Goedhart, "The Effect of Hints and Model Answers in a Student-Controlled Problem-Solving Program for Secondary Physics Education," *Journal of Science Education and Technology*, vol. 17, no. 4, pp. 410–425, 2008.

[29] R. Baker, J. Walonoski, N. Heffernan, I. Roll, A. Corbett, and K. Koedinger, "Why Students Engage in "Gaming the System" Behaviour in Interactive Learning Environments," *Journal of Interactive Learning Research*, vol. 19, no. 2, pp. 185–224, 2008.

[30] H.-T. Hou, "A Framework for Dynamic Sequential Behavioral Pattern Detecting and Automatic Feedback/Guidance Designing for Online Discussion Learning Environments," in *Advanced Learning*. In-Teh, Vukovar, Croatia, 2009, pp. 275–286.