ISLG 2015
Fourth Workshop on
Intelligent Support for Learning in Groups

held in conjunction with

Seventeenth International Conference on
Artificial Intelligence in Education (AIED 2015)

Monday, June 22, 2015
Madrid, Spain

Workshop Co-Chairs:

Ilya Goldin[1] ,Roberto Martinez-Maldonado[2] ,Erin Walker[3] ,
Rohit Kumar[4] ,Jihie Kim[5]

[1]*Center for Digital Data, Analytics, & Adaptive Learning, USA*
[2]*School of Information Technologies, University of Sydney, Australia*
[3]*Computing, Informatics, and Decision Systems Engineering, Arizona State University, USA*
[4]*Raytheon BBN Technologies, USA*
[5]*Software R&D Center, Samsung Electronics, South Korea*

ilya.goldin@pearson.com, roberto@it.usyd.edu.au,
erin.a.walker@asu.edu, rkumar@bbn.com, jihie.kim@gmail.com

https://sites.google.com/site/aied2015islg

Program Committee

Jessica Andrews, *Educational Testing Service, USA*
Ari Bader-Natal, *Minerva Project, USA*
Tiffany Barnes, *North Carolina State University, USA*
Sherice Clark, *University of Pittsburgh, USA*
Girlie C. Delacruz, *University of California, USA*
Yannis Dimitriadis, *University of Valladolid, Spain*
Toby Dragon, I*thaca College, USA*
Sharon I-Han Hsiao, *Arizona State University, USA*
Seiji Isotani, *Universidade de São Paulo, Brazil*
Charalampos Karagiannidis, *University of Thessaly, Greece*
James Lester, *North Carolina State University, USA*
Alejandra Martínez Monés, *University of Valladolid, Spain*
Bruce McLaren, *Carnegie Mellon University, USA*
Jennifer Olsen, *Carnegie Mellon University, USA*

Table of Contents

Preface

Technological advances in the use of artificial intelligence in education (AIED) over the past two decades have enabled the development of highly effective, deployable learning environments that support learners across a wide range of domains and age groups. Alongside, mass access to and adoption of modern communication technologies have made it possible to bridge learners and educators across spatiotemporal divides. Students can now collaborate using educational technology in ways that were not previously possible.

Intelligent tutoring systems seek to individualize each student's learning experience, but this need not imply a solitary experience. Research on computer-supported collaborative learning (CSCL) has revealed the pedagogical benefits of learning in groups, as well as how to structure the activity to lead to productive interactions. A variety of recent systems have demonstrated ways in which an adaptive learning environment can benefit from the presence of multiple learners. Similarly, students using CSCL systems have been shown to benefit from the introduction of adaptive support. It is of high relevance to the AIED community to explore how AI techniques can be used to support collaborative learning, and how theories of how students learn in groups can inform the design of adaptive educational technologies.

The goal of this series of workshops is to gather the sub-community of AIED researchers interested in intelligent support for learning in groups with learning scientists to share approaches and exchange information about adaptive intelligent collaborative learning support. We invite discussion on how the combination of collaborative and intelligent aspects of a system can benefit the learner by creating a more productive environment. Over the past few years, the AIED research community has started investigating extension of the fundamental techniques (student modeling, model-based tutors, integrated assessment, tutorial dialog, automated scaffolding, data mining, pedagogical agents, and so on) to support collaborative learning. We aim to explore ways that the current state of the art in intelligent support for learning in groups can be informed by learning sciences research on collaborative learning principles.

June, 2015
Ilya Goldin, Roberto Martinez-Maldonado,
Erin Walker, Rohit Kumar, and Jihie Kim

Negotiating Individual Learner Models in Contexts of Peer Assessment and Group Learning

Susan Bull and Lamiya Al-Shanfari

Electronic, Electrical and Systems Engineering, University of Birmingham, UK
s.bull@bham.ac.uk, LSA339@bham.ac.uk

Abstract. This paper introduces learner model negotiation not only as a means to increase the accuracy of the learner model and promote metacognitive activities as in past examples, but also as a way to help learners correct peer assessment entries in their learner model, that they consider inaccurate. While open learner models are not new, and negotiated learner models have been developed before, in today's learning contexts of potentially big data from many sources including other learners, some kind of approach to managing the data as well as helping learners to understand and accept it, or correct it, is needed.

1. Introduction

Benefits of a range of approaches to learning in groups have been argued (e.g. [9]), and there is strong interest in the field of Artificial Intelligence in Education in developing useful support for group learning [18]. Peer assessment and feedback have also been advocated as beneficial to the learning process (e.g. [27],[30]). We introduce a negotiated open learner model (OLM) approach to supporting students in the peer assessment situations that are becoming more common in today's learning contexts.

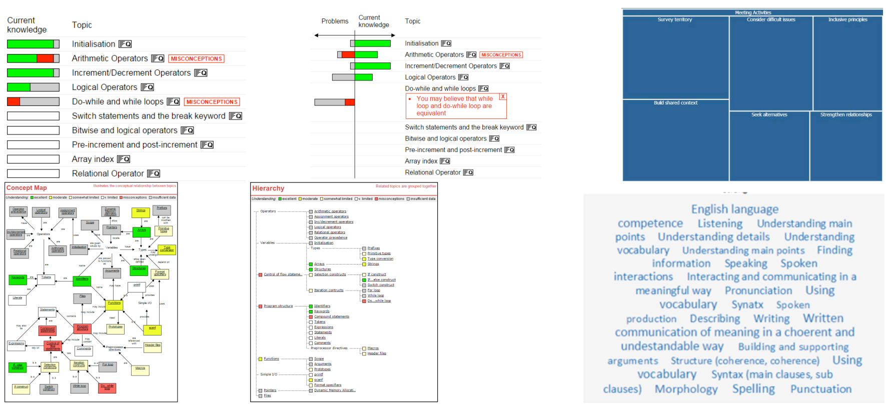

Fig. 1. Examples of open learner model visualisations

OLMs are learner models that are externalised to users in an understandable form, often to support collaboration or metacognitive behaviours [4]. Figure 1 gives OLM visualisation examples of simple skill meters [2], structured concept map and hierarchical tree [20], and newer visualisation approaches of overview-zoom treemap and word cloud [3]. While OLMs to support group learning have been developed (e.g. [1],[2],[6],[28]), the range of activities a student may be engaged in will likely include individual activities. Thus, in this paper we reflect on individual learner models that may be used in a group context. We focus in particular on situations in which peer feedback or assessment contributes to the individual learner model, which may follow the production of an artefact for assessment, or participation in a group activity.

2. OLMs and Peer Assessment in Modern Learning Contexts

Learner modelling has broadened, now being found in contexts with rich collections of digital materials [14]. Recent advances in learner modelling have aimed to address the use of new technologies, e.g.: learner models holding diverse data from different sources [3],[7],[21],[22]); combining e-portfolios and viewable learner models [23]; and OLMs to help learners monitor progress and plan their learning in MOOCs [15].

Peer assessment has become more prevalent in modern learning contexts such as MOOCs [17],[25] and e-portfolios [12],[31]; as well as individual online systems that allow peer assessment and feedback to be given and received [19]. OLMs that include peer assessment and feedback have been proposed [11], and developed (see [3]) in the context of peer assessments (numerical, contributing to the learner modelling algorithm) alongside automated data from a variety of external applications, and feedback (non-interpreted text, to help explain the numerical value of a peer assessment to an assessee). However, although there are many learning benefits for both peer assessors and assessees, there can also be cases of motivation decreasing if a student considers a peer assessment to be unjust [17]; or a learner feels there to be a lack of effort/attention from a peer assessor [10]. Another issue that may cause concern is the outcome of group assessments where there has been unequal contribution from group members [24]. For example, a student who engaged minimally in a group activity or project may receive the same assessment as the other participants. Experiences such as the above can cause strong emotional responses, and a method for learners to either understand learner model representations originating from peer assessors, or to challenge them, would help to relieve this frustration. The solution should allow individuals to understand the reasons for peer assessments and the system's perspective on them, as well as justify why they believe these representations or reasons to be inappropriate. We address these problems in the context of the LEA's Box OLM, where a learner model negotiation mechanism is being developed (based on [5]).

3. Maintaining the Learner Model through Negotiation

Building on the Next-TELL OLM [3], the LEA's Box learner model data may originate from a range of applications. In some cases, activities may be completed away from any tracking software. To address the latter, teacher, self and peer assessments can be entered alongside automated assessments. However, these may themselves differ in quality according to effort, experience and expertise of the assessor. While a learner may accept an automated assessment, or assessment by a teacher, they may be less happy with peer assessments and, indeed, may retain a negative attitude towards peer assessments over teacher assessments [13]. Even though a single peer assessment may ultimately contribute little to the value(s) in their learner model, this negative affective state may remain strong.

Some OLMs have allowed the learner model to be negotiated, where student and system have the same powers and negotiation moves [5],[8],[16]; or to be discussed in some other way, e.g. one partner has greater control over the discussion outcome [26],[29],[32]. Advantages of discussing or negotiating learner models include: the possibility to increase the accuracy of the learner model by allowing the learner to challenge the representations [5]; motivation may be increased by offering an alternative task [26]; significant learning gains may be achieved as a result of the negotiation process [16]. We here add a new benefit resulting from the inclusion of peer-entered data in an individual's learner model, for the increasing number of contexts in which multiple sources of data, including human contributions, are incorporated in the learner model. As well as increasing the accuracy of the learner model, individuals have the opportunity to redress any perceived injustices introduced by peer assessment.

Discussion of learner models typically involves moves such as agree/disagree; requesting information; challenging the other partner (learner or system); stating one's viewpoint; and justifying one's viewpoint by referring to evidence. For example:

- LEARNER: My value for [multiplying matrices] should be [higher].
- SYSTEM: Your last [five] attempts in [OLMlets] showed that you have [multiplied the corresponding items in each matrix]. You are adopting an approach used to [add matrices] when you are trying to [multiply matrices].
- LEARNER: I have since [participated in peer discussion].
- SYSTEM: Please [multiply] the following: ……

This example shows a student-initiated attempt to change a learner model representation (square brackets indicate variables). Having looked at their OLM, the learner notices that they disagree with the strength of the value for multiplying matrices. The learner has the misconception that the process mirrors that of adding matrices. The system has modelled this misconception, and is able to provide the evidence to back its inferred value – that the user's most recent attempts at this type of problem using the OLMlets system [2] illustrate the misconception. On receiving this evidence, the learner may accept the existing value or try to persuade the system further (in this example, stating that they have discussed this problem in their group). In this case the system offers a quick test to see whether it accepts the learner's claim as valid. Negotiation of learner models can take place using a variety of methods – menu-based selection of arguments [5]; buttons to initiate dialogue moves in dialogue games [8]; natural language discussion with a chatbot [16]. In most cases, even where negotiation takes place through a graphical interface, a text version is recorded for reference during the negotiation and afterwards. This record can also be kept as further evidence for the system, in subsequent discussions of the learner model.

As stated above, in LEA's Box the learner model data comes from a variety of activities. These may be simple quizzes, intelligent educational systems, or self or peer assessments. Thus, negotiation may not always be so clearly focussed on specific points such as a misconception as in the above example. However, the evidence used in negotiation can still be meaningful and, indeed, beneficial for both increasing the accuracy of the learner model and facilitating learner reflection. For example:

- LEARNER: My value for [matrices] should be [higher].
- SYSTEM: Your use of [OLMlets] showed [some difficulties].
- LEARNER: In [group discussion 2] I understood [well]. The [peer assessment] is [too low].
- SYSTEM: [Group discussion 2] was [5 days] ago and the [peer assessment] was [4 days ago]. You used [OLMlets] [1 day] ago. The level of [matrices] in [OLMlets] was [easy].

In this example, the system accesses the timestamp of data: in this case data from OLMlets [2]; and a peer assessment following a group discussion. It is able to explain that the first set of data was older, and also that the more recent OLMlets data was from a quite basic task. If the learner did not wish to accept the reasoning, the system could further explain that easier exercises can lead to higher scores, and that the learner was now working on more complex tasks, so old data would be less relevant. Through negotiation, as well as determining the correct representation for the learner model, the learner should come to better recognise their skills as they are required to think about the evidence provided by the OLM as well as in any justifications that they themselves give, supporting their claim. In addition, if the learner has disputed a peer assessment, the interaction will allow them to better understand that assessment, or have the opportunity to persuade the system to correct the disputed value.

Thus, the LEA's Box approach that is currently under development draws on the benefits of OLMs as meaningful visualisations of learning, as well as the benefits of negotiated learner models that can increase the accuracy of the learner model while also promoting learner reflection and other metacognitive behaviours. This is particularly useful when learners may be using disjointed applications, and when the learner model data includes data from other users. For the latter, in addition to the potential to increase the accuracy of the learner model, the process allows learner frustrations and perceived unjust assessments to be handled.

The current method of learner modelling uses a simple weighted algorithm, applying heavier weighting to more recent data, regardless of their origin [3]. However, teachers can adjust the weightings for individual activities according to the relevance of an activity for the learner model. As well as the recency of data as indicated above, the learner model negotiation will take account of these teacher weightings, and include these in its reasoning when 'defending' a representation during negotiation.

4. Summary

We have explained how benefits of negotiating learner models can be applied in today's contexts of multiple applications contributing learner data, as well as other activities which may include group interaction and peer assessment. By giving self, peer and teacher assessments the same status as automated data from various sources, such assessments can offer valuable insights to the learner's current learning state. Including this data also allows a system to better gauge the learner's viewpoint on their understanding (through self assessments), and also take into account learning outcomes from non-computer-based or non-tracked activities (from self, peer and teacher assessment). By negotiating the learner model, users can help maintain their learner model, and through this process they should also benefit from the critical thinking required to justify their viewpoints if they disagree with any representations. In addition, learner model negotiation allows a method to verify peer assessment values, and a means to allow a learner to try to update the learner model in cases of unfairness or perceived unfairness resulting from peer contributions to their model.

Acknowledgement

This work is supported by the European Commission (EC) under Information Society Technology priority FP7 for R&D, contract 619762 LEA's Box. This document does not represent the opinion of the EC and the EC is not responsible for any use that might be made of its contents.

References

1. Bakalov, F., Hsiao, I-H., Brusilovsky, P. & Koenig-Ries B. (2011). Visualizing Student Models for Social Learning with Parallel Introspective Views, Workshop on Visual Interfaces to the Social Semantic Web, ACM IUI 2011, Palo Alto, US.
2. Bull, S. & Britland, M. (2007). Group Interaction Prompted by a Simple Assessed Open Learner Model that can be Optionally Released to Peers, Proceedings of PING Workshop, User Modeling 2007.
3. Bull, S., Johnson, M.D., Demmans Epp, C., Masci, D., Alotaibi, M. & Girard, S. (2014). Formative Assessment and Meaningful Learning Analytics, ICALT 2014.
4. Bull, S. & Kay, J. (2013). Open Learner Models as Drivers for Metacognitive Processes, in R. Azevedo & V. Aleven (eds), International Handbook of Metacognition and Learning Technologies, Springer, New York, 349-365.
5. Bull, S. & Pain, H. 'Did I Say What I Think I Said, And Do You Agree With Me?': Inspecting and Questioning the Student Model, Proceedings of World Conference on Artificial Intelligence and Education, AACE, Charlottesville VA, 1995, 501-508.
6. Chen, Z-H., Chou, C-Y., Deng, Y-C. & Chan, T-W. (2007). Active Open Learner Models as Animal Companions: Motivating Children to Learn through Interacting with My-Pet and Our-Pet, Int. Journal of Artificial Intelligence in Education 17(2), 145-167.
7. Cruces, I., Trella, M., Conejo, R. & Galvez J. (2010). Student Modeling Services for Hybrid Web Applications, Int. Workshop on Architectures and Building Blocks of Web-Based User-Adaptive Systems, http://ceur-ws.org/Vol-609/paper1.pdf.
8. Dimitrova, V. (2003). StyLE-OLM: Interactive Open Learner Modelling, Int. Journal of Artificial Intelligence in Education 13(1), 35-78.
9. Gillies, R.M. & Ashman, A.F. (2003). An Historical Review of the Use of Groups to Promote Socialization and Learning, in R.M. Gillies & A.F. Ashman (eds), Co-operative Learning: the Social and Intellectual Outcomes of Learning in Groups, RoutledgeFalmer, London, 1-18.
10. Hanrahan, S.J. & Isaacs, G. (2001) Assessing Self- and Peer Assessment: The Students' Views, Higher Education Research & Development 20(1), 53-70.

11. Hu, Q., Huang, Y. & Liu, C. (2012). The Design and Implementation of Learner Models in Online Peer Assessment to Support Learning, Int. Conference on Consumer Electronics, Communications and Networks, IEEE, 2961-2963.
12. Hung, S-T.A. (2012). A Washback Study on E-portfolio Assessment in an English as a Foreign Language Teacher Preparation Program, CALL 25(1), 21-36.
13. Kaufmann, J.H. & Schunn, C.D. (2011). Students' Perceptions about Peer Assessment for Writing: their Origin and Impact on Revision Work, Instructional Science 39, 387-406.
14. Kay, J. (2012). AI & Education: Grand Challenges, IEEE Intelligent Systems 27(5), 66-69.
15. Kay, J., Reimann, P., Diebold, E. & Kummerfeld, B. (2013). MOOCs: So Many Learners, So Much Potential, IEEE Intelligent Systems 28(3), 70-77.
16. Kerly, A. & Bull, S. (2008). Children's Interactions with Inspectable and Negotiated Learner Models, Proc. ITS, Springer-Verlag, Berlin Heidelberg, 132-141.
17. Kulkarni, C, Wei, K.P., LE, H., CHIA, D., Papadopoulos, K., Cheng, J., Koller, D., & Klemmer, S.R. (2013). Peer and Self Assessment in Massive Online Classes, ACM Transactions on Computer-Human Interaction 9(4), Article 39.
18. Kumar, R. & Kim, J. (2014). Editorial: Special Issue on Intelligent Support for Learning in Groups, International Journal of Artificial Intelligence in Education 24, 1-7.
19. Lu, J. & Law, N. (2012). Online Peer Assessment: Effects of Cognitive and Affective Feedback, Instructional Science 40, 257-275.
20. Mabbott, A. & Bull, S. (2004). Alternative Views on Knowledge: Presentation of Open Learner Models, Proc. ITS, Springer-Verlag, Berlin Heidelberg, 689-698.
21. Mazzola, L. & Mazza, R. (2010). GVIS: A Facility for Adaptively Mashing Up and Presenting Open Learner Models, Proc. EC-TEL, Springer, Berlin Heidelberg, 554-559.
22. Morales, R., Van Labeke, N., Brna, P. & Chan, M.E. (2009). Open Learner Modelling as the Keystone of the Next generation of Adaptive Learning Environments, in C. Mourlas & P. Germanakos (eds), IUI, Information Science Reference, ICI Global, London, 288-312.
23. Raybourn, E.M. & Regan, D. (2011). Exploring E-Portfolios and Independent Open Learner Models: Toward Army Learning Concept 2015, Interservice/Industry Training, Simulation, and Education Conference Proceedings, Florida USA.
24. Subramanian, R. & Lejk, M. (2013). Enhancing Student Learning, Participation and Accountability in Undergraduate Group Projects through Peer Assessment, South African Journal of Higher Education 27(2), 368-382.
25. Suen, H.K. (2014). Peer Assessment for Massive Open Online Courses (MOOCs), The International Review of Research in Open and Distance Learning 15(3), 312-327.
26. Thomson, D. & Mitrovic, A. (2010). Preliminary Evaluation of a Negotiable Student Model in a Constraint-Based ITS, Research and Practice in Technology Enhanced Learning 5(1), 19-33.
27. Topping, K. (1998). Peer Assessment between Students in Colleges and Universities, Review of Educational Research 68(3), 249-276.
28. Upton, K., & Kay, J. (2009). Narcissus: Group and Individual Models to Support Small Group Work, Proc. User Modeling, Adaptation and Personalization, Springer, 54-65.
29. Van Labeke, N., Brna, P. & Morales, R. (2007). Opening Up the Interpretation Process in an Open Learner Model, Int. Journal of Artifial Intelligence in Education 17(3), 305-338.
30. Van Zundert, M., Sluijsmans, D. & van Merrienboer, J. (2010). Effective Peer Assessment Processes: Research Findings and Future Directions, Learning and Instruction 20(4), 270-279.
31. Welsh, M. (2012). Student Perceptions of Using the PebblePad E-portfolio System to Support Self- and Peer-Based Formative Assessment, Technology, Pedagogy and Education 21(1), 57-83.
32. Zapata-Rivera, J-D. & Greer, J.E. (2004). Interacting with Inspectable Bayesian Student Models, Int. Journal of Artificial Intelligence in Education 14(2), 127-163.

Predicting Student Attrition in MOOCs using Sentiment Analysis and Neural Networks

Devendra Singh Chaplot, Eunhee Rhim, and Jihie Kim

Samsung Electronics Co., Ltd.
Seoul, South Korea
{dev.chaplot,eunhee.rhim,jihie.kim}@samsung.com

Abstract. While there is increase in popularity of massive open online courses in recent years, high rates of drop-out in these courses makes predicting student attrition an important problem to solve. In this paper, we propose an algorithm based on artificial neural network for predicting student attrition in MOOCs using sentiment analysis and show the significance of student sentiments in this task. To the best of our knowledge, use of user sentiments and neural networks for this task is novel and our algorithm beats the state-of-the-art algorithm on this task in terms of Cohen's kappa.

Keywords: Student Attrition, MOOC, Educational Data Mining, Sentiment Analysis, Neural Network

1 Introduction

Massive Open Online Courses (MOOCs) have been gaining lot of interest in academia and industry in last few years. The key reasons in growing popularity of MOOCs include accessibility to every person in the world who has internet, scalability to handle any number of students with wide diversity of needs and expectations, and flexibility they provide to learners to study according to their routine. However, issues such as lack of instructor attention and absence of social learning environment, have led to high rates of attrition in MOOCs. With various unique benefits they offer over traditional classroom setting, online courses have the potential to transform future of education system, which brings out the importance of predicting student attrition in MOOCs.

With scalability, MOOCs also offer huge amounts of data of student activity, which can be utilized to train models for predicting attrition. The absence of physical learning environment makes the forums in MOOCs only medium of interaction with the instructor and peers. In this paper, we analyze the importance of sentiment analysis on these forum posts in predicting student attrition and study the effectiveness of neural network in modeling this problem.

The rest of the paper is divided into the following sections. Section 2 covers related work regarding machine learning techniques used to predict attrition and different kind of features used in them. Our algorithm is described in detail in Section 3. The experiments and results are presented in Section 4. Conclusions and future work are covered in Section 5.

2 Related Work

Recently, there have been many efforts to predict student attrition in MOOCs by extracting a wide variety of features from learner activity data and applying different machine learning approaches. [11] operationalize video lecture clickstream to capture behavioral patterns in student's activity, which is used to construct students' information processing index. [4] use feature such as number of threads viewed, number of forum posts, percentage of lectures watched, etc to predict student attrition. [12] construct a graph to capture sequence of active and passive learner activity, and use graph metrics as features for predicting attrition. [2] use quiz related (attempts and submissions) and activity related (length of action sequences, counts of various activities) features while [7] and [10] extract more than 15 features indicating learner activity and engagement from clickstream log. All these methods use variety of machine learning techniques including Logistic Regression, SVMs, Hidden Markov Models and random forest method.

There has not been much work on use of student sentiments in predicting attrition. [1] conclude that sentiment of students for assignments and course material has positive effects on successful completion of course. [14] also find correlation between sentiment expressed in the course forum posts and student drop out rate while they advice prudence against inconsistencies.

3 Proposed Algorithm

We have used click stream log and forum posts data from Coursera MOOC, 'Introduction to Psychology', which was prepared for MOOC Workshop at EMNLP 2014. The data consists of over 3 million student click logs and over 5000 forum posts. The click stream logs contain clicks made while watching video lectures and requests for viewing forums, threads, quiz, course wiki, etc. with time stamp of each click. More details about the dataset can be found in [7]. The following input features were extracted from the dataset:

- **User ID:** Unique numerical ID of the student.
- **Course Week:** Number of weeks since course has begun.
- **User week:** Number of weeks since student has joined the course.
- **Number of clicks** by the student in the current week.
- **Number of study sessions** by the student in the current week.
- **Number of course pages viewed** by the student in current week which include all pages except the video lectures.
- **Number of forum pages viewed** by the student in current week.
- **Student sentiment** of forum posts in the current week.

All the input features except Student Sentiments were indicated to be most effective by previous works mentioned in Section 2. The output of the algorithm is 1 indicating the user will drop out of the course in next week, and 0 otherwise. Note that we are predicting the exact week when the student is going to drop-out unlike [11] who predict whether student is going to finish the course or not. Our algorithm pinpoints the time when student is predicted to drop-out, which allows the course instructor and his team to take necessary steps to prevent or reduce student attrition during the course.

3.1 Sentiment Analysis

We follow a lexicon-based approach to extract sentiment from forum posts using SentiWordNet 3.0 [3] as the knowledge resource. It assigns a sentiment score to each synset in the WordNet [8]. Given the forum post, we pass the stem of each content word (using MIT JWI [6]) and its POS Tag (using Stanford POS Tagger [13]) to the SentiWordNet which returns a sentiment score. The sentiment score of the forum post is calculated by summing up the sentiment scores of all the words in the post. Fig. 1 shows a block diagram of this process.

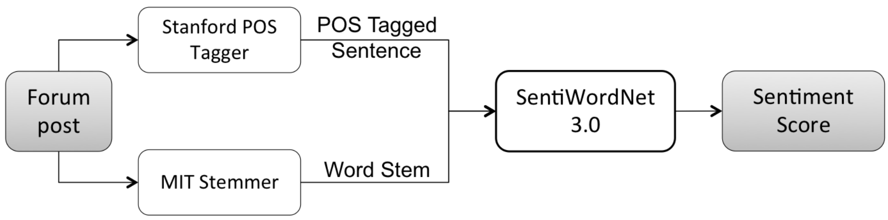

Fig. 1. Block Diagram of lexicon-based sentiment analysis using SentiWordNet 3.0.

3.2 Neural Network

Artificial neural networks are suitable to model the problem of predicting student attrition as there are a large number of inputs, and any mathematical relationship between input and output is unknown. Unlike many other machine learning techniques, neural networks are able to model the output as any arbitrary function of inputs and considered extremely robust if network structure, cost function and learning algorithm are selected appropriately through experiments. Downside of neural networks is inability to interpret the model.

We construct an artificial neural network consisting of 7 nodes in input layer: Course Week, User week, Number of clicks, Number of sessions, Number of page views, Number of forum views and Student sentiment as described above. Output layer consists of single node predicting whether student is going to drop-out in the next week. Each input feature is normalized to take values between 0 and 1. We add a hidden layer of 6 neurons in the neural network between the input and output layer. The number of neurons in the hidden layer were experimentally determined to get best possible results. Fig. 2 shows the structure of the neural network used to predict student attrition. To train the neural network, we use resilient propagation heuristic. It gave best results in our experiments among back propagation, Manhattan propagation and quick propagation.

4 Experiments & Results

In predicting student attrition, our focus is to capture all students who are going to drop-out and thus, minimizing false negative rate is important. False negative rate is the ratio of students who are predicted to stay in the course (predicted negative) in next week but actually drop out in the next week. While minimizing false negative rates, its also necessary to maintain overall accuracy so as to not produce too many false positives for the course instructor to handle.

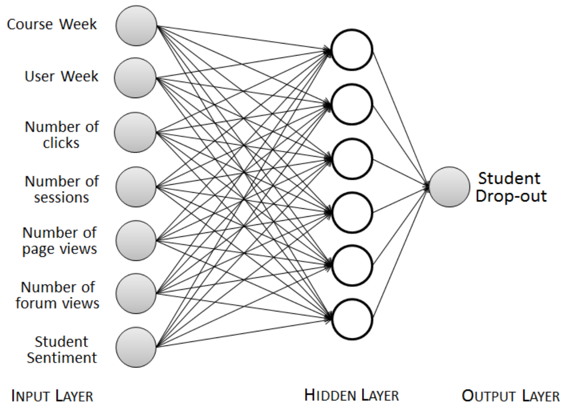

Fig. 2. The structure of neural network used to predict student attrition.

Since we are predicting whether student will drop-out in next week, our data set is highly imbalanced towards negative (will not drop-out) class. This is because a student who joins the course in 1st week, and drops out in 11th week, will have 9 negative class data points (week 1 to 9) and 1 positive class data point (week 10). Since the data set consists of student logs over 19 weeks, it is highly imbalanced with only 22.56% positive data points. Due to high imbalance in data set, we believe comparison of Cohen's kappa [5] is more suitable than comparing total accuracy directly. [9] show that Cohen's Kappa provides a unbiased estimate of performance of a classifier, and is thus much more meaningful than Recall, Precision, Accuracy, and their biased derivatives. It is more robust than total accuracy as it excludes proportion of correct predictions occurring by chance which is important in case of imbalanced data set, as a simple majority classifier would get 77.44% accuracy in this task.

In Table 1 we report our results with and without using student sentiments using 5-fold cross validation and compare them with some other approaches mentioned in Section 2. The proposed algorithm provides the best Cohen's Kappa values as compared to previous algorithms. Fall in accuracy and false negative rate when our algorithm doesn't use student sentiments indicates its importance in predicting attrition. Note that the algorithm which provides the best accuracy [10] also has the highest number of false negatives and the algorithm with best false negative rate has the lowest accuracy (Sinha-14 Baseline + Graph). This is due to imbalance in data which is explained in the following subsection. Note that the proposed algorithm has either better accuracy or better false negative rates than each of the previous algorithms, and this is reason behind better Kappa values. Since the dataset is from a MOOC which had free enrollment, there are many initial lurkers in the first week of the course who just want to browse the contents of the course. Thus, we believe predicting student attrition in first week is not very useful. Substantial improvement in performance of our algorithm without using first week's data is also shown in Table 1.

Algorithm	Accuracy	False Neg.	Kappa
Balakrishnan-13 Stacking [4]	80.5%	0.353	-
Balakrishnan-13 Cross-Product [4]	80.1%	0.442	-
Sharkey-14 [10]	**88.0%**	0.460	-
Sinha-14 Baseline + Graph [12]	62.4%	**0.095**	0.277
Sinha-14 Graph [12]	69.2%	0.157	0.365
Neural Network (NN)	70.7%	0.199	0.365
NN with Sentiment Analysis (SA)	72.1%	0.141	0.403
NN with SA & without Week 1	74.1%	0.132	**0.432**

Table 1. Comparison of accuracy and false negative rates with and without using student sentiments. The best results in each column is marked in **bold**.

4.1 Problem of data imbalance

The high data imbalance leads to biasing of the classifier towards the majority class. The problem of data imbalance in the same task is also addressed by [2] who try to solve it by oversampling the minority class, but were unsuccessful. We counter this problem by setting the boundary for classification to the ratio of drop out data points to total number of data points in the training set. This means that if the value of output neuron is greater than this ratio, then student is predicted to drop out in the next week, and vice-versa otherwise. If complete data set is used as training set, then this boundary would be 0.2256, meaning student is predicted to drop-out if value of output neuron is greater than 0.2256, rather than 0.5 by default. This adjustment to the boundary allows us to train the neural network on highly unbalanced dataset and still achieve very good recall over minority class while maintaining the overall accuracy.

The boundary is essentially a trade-off between accuracy and false negative rate. It can be adjusted to get better accuracy or false negative rates depending upon the application. This boundary can also be calculated using receiver operating characteristic (ROC) Curve.

5 Conclusion & Future Work

We propose an algorithm to predict student attrition using an artificial neural network. Sentiment analysis of forum posts is shown to be an important feature to predict student attrition in MOOCs. We also provide an approach to tackle the problem of data imbalance which can be extended to wide variety of applications in many other domains. This approach allows to find a good middle ground between accuracy and false negative rates and leads our algorithm to beat the previous algorithms in terms of Cohen's Kappa.

Most methods provide analysis of MOOC data which indicate factors responsible for attrition. In contrast, we provide a method to pin-point students who are likely to drop-out during in the following week. Since our algorithm has a very low false negative rate, it can be used in MOOCs to capture most students who are likely to drop-out in near future and take necessary actions specific to the student to prevent them from dropping out. Apart from MOOCs, the proposed algorithm can also used in smart schools using digital methods for learning and interaction, which are becoming increasingly popular in recent years.

References

1. Adamopoulos, P.: What makes a great MOOC? An interdisciplinary analysis of student retention in online courses. In: Proceedings of the International Conference on Information Systems, ICIS 2013, Milano, Italy (2013)
2. Amnueypornsakul, B., Bhat, S., Chinprutthiwong, P.: Predicting Attrition Along the Way: The UIUC Model. In: Proceedings of the EMNLP 2014 Workshop on Analysis of Large Scale Social Interaction in MOOCs. pp. 55–59. Association for Computational Linguistics, Doha, Qatar (October 2014)
3. Baccianella, S., Esuli, A., Sebastiani, F.: Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In: Calzolari, N., Choukri, K., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S., Rosner, M., Tapias, D. (eds.) Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10). European Language Resources Association (ELRA), Valletta, Malta (2010)
4. Balakrishnan, G.: Predicting Student Retention in Massive Open Online Courses using Hidden Markov Models. Master's thesis, EECS Department, University of California, Berkeley (May 2013)
5. Cohen, J.: A Coefficient of Agreement for Nominal Scales. Educational and Psychological Measurement 20(1), 37 (1960)
6. Finlayson, M.: Java libraries for accessing the princeton wordnet: Comparison and evaluation. In: Orav, H., Fellbaum, C., Vossen, P. (eds.) Proceedings of the Seventh Global Wordnet Conference. pp. 78–85. Tartu, Estonia (2014)
7. Kloft, M., Stiehler, F., Zheng, Z., Pinkwart, N.: Predicting MOOC Dropout over Weeks Using Machine Learning Methods. In: Proceedings of the EMNLP 2014 Workshop on Analysis of Large Scale Social Interaction in MOOCs. pp. 60–65. Association for Computational Linguistics, Doha, Qatar (October 2014)
8. Miller, G.A.: Wordnet: A lexical database for english. Commun. ACM 38(11), 39–41 (Nov 1995)
9. Powers, D.M.W.: Evaluation: From Precision, Recall and F-Factor to ROC, Informedness, Markedness & Correlation. Tech. Rep. SIE-07-001, School of Informatics and Engineering, Flinders University, Adelaide, Australia (2007)
10. Sharkey, M., Sanders, R.: A Process for Predicting MOOC Attrition. In: Proceedings of the EMNLP 2014 Workshop on Analysis of Large Scale Social Interaction in MOOCs. pp. 50–54. Association for Computational Linguistics, Doha, Qatar (October 2014)
11. Sinha, T., Jermann, P., Li, N., Dillenbourg, P.: Your click decides your fate: Inferring Information Processing and Attrition Behavior from MOOC Video Clickstream Interactions. ArXiv e-prints (Jul 2014)
12. Sinha, T., Li, N., Jermann, P., Dillenbourg, P.: Capturing "attrition intensifying" structural traits from didactic interaction sequences of MOOC learners. CoRR abs/1409.5887 (2014)
13. Toutanova, K., Klein, D., Manning, C.D., Singer, Y.: Feature-rich part-of-speech tagging with a cyclic dependency network. In: Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1. pp. 173–180. NAACL '03, Association for Computational Linguistics, Stroudsburg, PA, USA (2003)
14. Wen, M., Yang, D., Rosé, C.P.: Sentiment analysis in mooc discussion forums: What does it tell us. In: Proceedings of Educational Data Mining (2014)

Adapting Collaborative Chat for Massive Open Online Courses: Lessons Learned

Oliver Ferschke, Gaurav Tomar, and Carolyn P. Rosé

Carnegie Mellon University, School of Computer Science
Pittsburgh, PA 15213, USA
{ferschke,gtomar,cprose}@cs.cmu.edu

Abstract. In this paper we explore how to import intelligent support for group learning that has been demonstrated as effective in classroom instruction into a Massive Open Online Course (MOOC) context. The Bazaar agent architecture paired with an innovative Lobby tool to enable coordination for synchronous reflection exercises provides a technical foundation for our work. We describe lessons learned, directions for future work, and offer pointers to resources for other researchers interested in computer supported collaborative learning in MOOCs.

1 Introduction

The field of Computer Supported Collaborative Learning (CSCL) has a rich history extending for nearly two decades, covering a broad spectrum of research related to learning in groups, especially in computer mediated environments. In this paper we describe the initial stages of a research program designed to import findings from a history of successful classroom research in the field of CSCL to the challenging environment of MOOCs.

In order to support the growth of student discussion skills, it is necessary to design environments with affordances that encourage transactive behaviors and other valuable learning behaviors. The most popular approach to providing such affordances in the past decade has been that of script-based collaboration [2, 7, 6]. A script is a schema for offering scaffolding for collaboration. Some typical forms of scripts come in the form of instructions that structure a collaborative task into phases, or structured interfaces that reify certain types of contributions to the collaboration. Prior work on dynamic conversational agent based support built on a long history of work using tutorial dialogue agents to support individual learning with technology [11, 10, 5, 12].

The MOOC environment presents a number of challenges that must be addressed in order to introduce synchronous collaboration opportunities into MOOCs. From a research perspective, interesting challenges include exploration of group composition questions with MOOC student populations, which are far more diverse with respect to culture, age, educational level, and goals than typical classroom populations. Another interesting methodological challenge is the lack of control over the context. In a classroom context, a certain amount of time may be set aside for an activity, and students can be expected to be present for the whole activity. In a MOOC, students may come and go as they please, and since they may be logging in from anywhere, any number of events could interfere with the task proceeding as planned. While a collaborative task may have been carefully designed with roles for each student to perform in a serious learning task, those roles may play out differently than intended in cases where the students who take on those roles are actually multiple students, students with a seriously inadequate preparation for the task, or even students with far more expertise than anticipated.

Before any of these issues may even be touched upon, a number of more practical issues must be addressed first to lay a foundation for this research. A major challenge in MOOCs is coordination. Whereas in a face-to-face course and traditional small-scale online courses, students can be expected to be amenable to stipulated meeting times, students in MOOCs typically come from different time zones around the world. The great majority of students make use of resources at their convenience, when they happen to have time to log in, rather than planning ahead and arriving at a scheduled time. The sheer numbers of students make it challenging to coordinate plans for meeting times. Furthermore, not all students adopt the same orientation towards following instructions in general or engaging in a task as presented in particular. Some students may click on an activity in an exploratory or playful fashion rather than with a serious intention of completing the activity. Thus, there is a danger of introducing students into a group in a way that engenders conflict or mismatched expectations.

In the remainder of the paper we first introduce the technical approach we adopted in an initial MOOC deployment. We then summarize our main results and lessons learned. We conclude with directions for continued work and resources to share with the community. Further discussion of the results of our deployment can be found in two separate publications [3, 4].

2 Technical Approach

In order to gain a deeper understanding of the problems that may arise from synchronous collaborative activities in MOOCs, we integrated a chat environment with interactive agent support in a recent 9-week long MOOC on learning analytics (DALMOOC) that was hosted on the edX platform from October to December 2014.

In order to facilitate the formation of ad-hoc study groups for the chat activity, we make use of a simple setup referred to as a Lobby. The Lobby introduces an intermediate layer between the edX platform and the synchronous chat tool. Even though the Lobby allows groups of arbitrary sizes to be formed, we decided that agent-guided discussions in groups of two students are the suited setup in the context of this MOOC. Students enter the Lobby with a simple, clearly labeled button click integrated with the edX platform. In order to increase the likelihood of a critical mass of students being assigned to pairs, we suggested a couple of two hour time slots during each week of the MOOC when students might engage in the collaborative activities. These timeslots were advertised in weekly newsletters. However, the chat button was live at all times so that students were free to attempt the activity at their convenience. Upon entering the lobby, students are asked to enter the name that will be displayed in the chat. Once registered in the lobby, the student waits to be matched with another participant. If they are successfully matched with another learner who arrives at the Lobby within a couple of minutes to interact with, they and their partner are then presented with a link to click on to enter a chat room created for them in real time. Otherwise they are requested to come back later. A visualization is presented to them that illustrates the frequency of student clicks on the button at different times of the day on the various days of the week so that they are able to gauge when would be a convenient time for them to come back when the likelihood of a match would be higher. In the beginning of the course, the graph was based on experiences with past MOOCs while it was later updated with real data from the DALMOOC logs.

When the successfully matched students click on the provided link, they enter a private chat room. This chat setup has been used in earlier classroom research [1]. It provides opportunities for students to interact with one another through chat as well as to share images. The chat environment furthermore has built-in support for conversational agents who appear as regular users in the chat.

In contrast to our earlier work where we support collaborative chat dynamically with conversational agents triggered by real time monitoring of student interactions [1], we employ statically scripted agents in DALMOOC which guide the students with course-related discussion questions (Figure 1). Even though the scripts are linear, the agent prompts are not strictly timed but rather allow the students to interact in their own pace and take as much time as needed to discuss the given topic. Once a team wants to proceed with the discussion, they can move on with the *We're ready*-button. The agent will proceed with the next prompt as soon as both students indicated that they are ready. In case the students never signal their readiness, the agent will inquire after a predefined timeout in order to move forward with the discussion and avoid the students to lose focus.

3 Main Results

Though we encountered many challenges during the DALMOOC deployment, the main results suggest value added by the intervention. In order to begin to assess the added value of integrating reflective chat activities with a MOOC platform, we compared our synchronous collaborative chats with two other communication contexts, namely Twitter and the MOOC discussion forum [3]. What we found is that different subpopulations of learners within DALMOOC tended to gravitate towards different communication contexts. Furthermore, each context was associated with its own unique profile in terms of content focus and the nature of the discussion. The chat conversations showed the highest average of reflective contributions across all the platforms we observed. Furthermore,

Prompt 1	In this collaborative activity, we will reflect on what you have learned about the field of learning analytics. First, take a couple of minutes to introduce yourselves.
Prompt 2	Now that you have viewed the videos, share what you found most interesting about learning analytics.
Prompt 3	Regarding learning analytics tools, did you find the classifications of a) proprietary/open source and b) single functionally/Integrated suites to be useful? How would you improve these classifications to make them more relevant to educators starting with analytics toolsets?
Prompt 4	Reflect on the structure of the dual-layer structure of the course. Describe your experience of coming to understand different course elements.
Prompt 5	Now this activity has come to an end. Thanks for a great chat! Why don't you exchange contact information to stay in touch?

Fig. 1: Agent prompts for the collaborative chat activity in the first week

the one-on-one conversations in Bazaar exhibit a strong constructive character where reflective statements are not merely precompiled by each student and then exchanged, they are rather collaboratively constructed in the course of the conversation. We see ample evidence within contributions across media pertaining to social connection that these MOOC learners crave continuing social engagement with other individuals participating in their MOOC course. The analysis suggests that there is value in providing a diverse set of discussion contexts but that it creates a need for greater efforts towards effective bridging between media and channeling of students to pockets of interaction that are potentially of personal benefit.

We also used a survival analysis to evaluate the impact of participation in collaborative chats on attrition over time in the course [4]. The results suggest a substantial reduction in attrition over time, specifically a reduction by more than a factor of two, when students experience a match for a synchronous collaborative reflection exercise. Nevertheless, these results must be treated with some caution as we experienced significant difficulty in managing the logistics of matches. Even with 20,000 students enrolled in the course, some students had to make as many as 15 attempts to be matched with a partner before a match was made.

4 Lessons Learned

In this first deployment study, we learned valuable lessons that will help to improve our experimental setup in future cycles of our iterative design based research process. In this section, we first describe the main lessons learned and then briefly discuss future directions we are planning to take.

Integrating a synchronous collaborative activity in an inherently asynchronous learning environment used by students in different time zones was one of the greatest organizational challenges to overcome. As mentioned earlier, we attempted to alleviate the problem by introducing dedicated chat hours to increase the likelihood of students getting matched with each other. Nevertheless, the majority of students who entered the lobby could not be matched with a chat partner within 10 minutes. This was a frequent cause of frustration which lead to students abandoning the chat activity in the course of the MOOC.

Since students are matched randomly in pairs for each chat activity, their interaction is naturally limited to a single chat session. Whenever they return to the chat, they will be connected with a different student. From the logs we have seen that especially after longer discussions, students expressed the desire to connect with each other and continue the interaction beyond the chat activity. On several occasions, they exchanged contact information in order to reconnect for further collaboration. However, the intervention did not provide any support for continued social engagement between paired learners.

We are currently developing new strategies for tackling these problems in future deployments of the intervention. First, we will employ a single-chatroom setup that allows students to directly enter at their own volition without the need for explicit matchmaking. The agent in this continuous chatroom will then adapt to the student population in the room at any given time. For instance, a single user in the room would be prompted to reflect on the course material on their own. Once a second user enters, the agent summarized the reflection of the other student and composes a discussion topic for the two users to collaboratively reflect on. The agent keeps track of the topics already discussed by the users currently present in the room to avoid redundant prompts.

Second, we will explore a scheduling system that allows students to sign up for a set of predefined timeslots. This approach has proven effective for multi-party voice chats in MOOCs [8]. Even though the necessity to schedule discussions ahead might negatively affect the engagement of users who merely interact with the MOOC on an ad-hoc basis, the approach could nevertheless help to reduce overall friction by offering an easier transition from the asynchronous nature of the MOOC to the synchronous nature of the chat.

5 Conclusions

This research was motivated by the goal to import best practices and technologies from the field of Computer Supported Collaborative Learning into MOOCs [9]. It is part of a broader effort drawing from two decades of research in Computer Supported Collaborative Learning, where we are working to design an extension of the edX platform to enhance instructionally beneficial discussion opportunities available to students[1]. We are partnering with edX as a satellite collaborative, seeking to involve researchers and developers from multiple universities, foundations, and industrial organizations. Our long term vision is to seek to leverage insights and methodologies from the field of Human-Computer Interaction more broadly and encompassing both synchronous and asynchronous communication very broadly. Our vision includes text, speech, and video based interactions, instrumented with all sorts of intelligent support powered by state-of-the-art

[1] `http://dance.cs.cmu.edu`

analytics and leveraging language technologies and artificial intelligence more broadly in order to offer contextually appropriate support.

Acknowledgments This work was funded in part by NSF grants OMA-0836012 and DATANET 1443068 and funding from Google and Bosch.

References

1. Adamson, D., Dyke, G., Jang, H., Rosé, C.P.: Towards an agile approach to adapting dynamic collaboration support to student needs. International Journal of Artificial Intelligence in Education 24(1), 92–124 (2014)
2. Dillenbourg, P.: Over-scripting CSCL: The risks of blending collaborative learning with instructional design. In: Kirschner, P.A. (ed.) Three worlds of CSCL. Can we support CSCL?, pp. 61–91. Heerlen, Open Universiteit Nederland (2002)
3. Ferschke, O., Howley, I., Tomar, G., Yang, D., Rosé, C.P.: Fostering Discussion across Communication Media in Massive Open Online Courses. In: Proceedings of the 11th International Conference on Computer Supported Collaborative Learning. Gothenburgh, Sweden (2015)
4. Ferschke, O., Yang, D., Tomar, G., Rosé, C.P.: Positive Impact of Collaborative Chat Participation in an edX MOOC. In: Proceedings of the 17th International Conference on Artificial Intelligence in Education. Madrid, Spain (2015)
5. Graesser, A., Vanlehn, K., TRG Group, NLT Group: Why2 report: Evaluation of why/atlas, why/autotutor, and accomplished human tutors on learning gains for qualitative physics problems and explanations. Tech. rep. (2002)
6. Kobbe, L., Weinberger, A., Dillenbourg, P., Harrer, A., Hmlinen, R., Hkkinen, P., Fischer, F.: Specifying computer-supported collaboration scripts. International Journal of Computer-Supported Collaborative Learning 2(2-3), 211–224 (2007)
7. Kollar, I., Fischer, F., Hesse, F.: Collaboration scripts a conceptual analysis. Educational Psychology Review 18(2), 159–185 (2006)
8. Kulkarni, C., Cambre, J., Kotturi, Y., Bernstein, M.S., Klemmer, S.R.: Talkabout: Making distance matter with small groups in massive classes. In: Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW 2015, Vancouver, BC, Canada, March 14 - 18, 2015. pp. 1116–1128 (2015)
9. Rosé, C.P., Goldman, P., Sherer, J.Z., Resnick, L.: Supportive technologies for group discussion in moocs, current issues in emerging elearning. Current Issues in Emerging eLearning. Special issue on MOOCs (2015)
10. Rosé, C.P., Jordan, P., Ringenberg, M., Siler, S., V.K., Weinstein, A.: Interactive Conceptual Tutoring in Atlas-Andes pp. 256–266 (2001)
11. Wiemer-Hastings, P., Graesser, A., Harter, D.: The foundations and architecture of autotutor. In: Goettl, B.P., Halff, H.M., Redfield, C., Shute, V.J. (eds.) Intelligent Tutoring Systems, Lecture Notes in Computer Science, vol. 1452, pp. 334–343. Springer Berlin Heidelberg (1998)
12. Zinn, C., Moore, J.D., Core, M.G.: A 3-tier planning architecture for managing tutorial dialogue. In: Proceedings of the 6th International Conference on Intelligent Tutoring Systems. pp. 574–584 (2002)

Exploring the Effects of Open Social Student Model Beyond Social Comparison

Julio Guerra[1], Yun Huang[2], Roya Hosseini[2], and Peter Brusilovsky[1]

[1] School of Information Sciences, University of Pittsburgh, Pittsburgh, PA, USA
[2] Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA, USA
{jdg60,yuh43,roh38,peterb}@pitt.edu

Abstract. In our journey exploring the effects of Open Student Model (OSM) on students working with programming problems and examples, we have incorporated the idea of social visualizations to extend OSM to Open Social Student Modeling (OSSM). Although comparison features in OSSM, where a student can compare herself to the group or individual peers, have shown to increase students' work, we now shift our attention to other effects. The goal is to explore the OSSM effects beyond comparison, particularly metacognitive support, and we propose a representation of the OSSM towards these lines.

Keywords: Open Student Model, Open Social Student Model, Metacognition, Self-Regulated Learning, Group-Awareness

1 Introduction

Open Student Model (OSM, also called Open Learner Model or OLM) consists of a set of features, usually visual and sometimes interactive, that shows data of progress, mastery of knowledge, or other statistics of the activity of the student to herself [3]. This data comes from the internal user model that the computer-based educational system maintains to bring in adaptive and tutoring functionalities [2]. By showing the user model to the learner, OSM fosters metacognitive processes like self-awareness [4] and can be further used as a navigational tool. In the past we have explored different forms of guidance based on OSM. KnowledgeZoom (KZ) [1] implements a fine-grained student model based on concepts hierarchically organize in an Ontology of Java programming. KZ presents the student model using treemap that shows different levels of details as the student "enters" each of the concepts. This approach allows the student to have an overall view and a detailed view of her progress and knowledge gaps just few clicks away. We have also incorporated the idea of social visualizations and extended the OSM to an Open Social Student Modeling (OSSM) [8, 10]. OSSM seeks for sharing aggregated or individual OSM among the students allowing social comparison and social guidance dynamics. Figure 1 shows a screenshot of the MasteryGrids system, our current OSSM implementation. The first 3 rows represent the progress of the current student, the comparison between the student and the group, and the progress of the rest of the class, respectively. Cells represent topics, ordered as they are covered in the course. Darker colors mean higher progress in the content. The student progress is colored in shades of green, and the group (the average of the class) is represented with a blue color palette. The middle row shows a differential color that turns green when the student is more advanced than the group and blue otherwise. By clicking in a cell, the student has the access to educational material included in the selected topic (in the figure, the cell corresponding to the topic *Loops While* has been clicked.) The second group of cells shows the progress of all individual students in the class, anonymized, ranked by the amount of progress (advanced students at the top) using shades of blue.

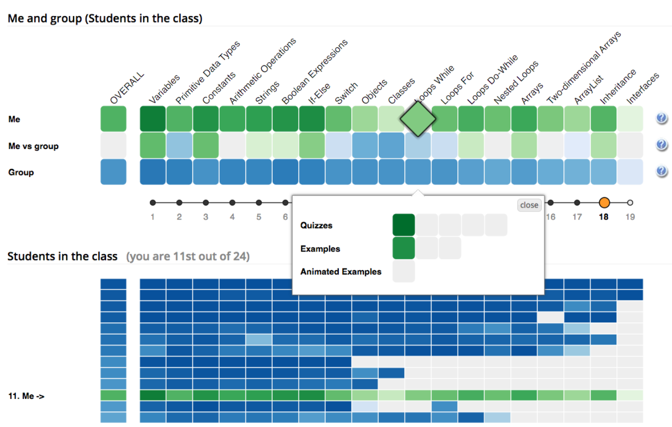

Fig. 1. MasteryGrids OSSM interface.

Overall, our efforts to implement OSSM have been focused on exploiting comparison effects and we have observed in classroom studies that this kind of features make students work more and follow others [8, 10]. Also, the sort of guidance produced by advanced students over non-advanced students is quite conservative, and we further proposed a guidance approach incorporating OSSM and adaptive navigation features (work presented as a poster in AIED 2015 [1]). We now shift our attention to explore the OSSM effects beyond comparison, particularly how OSSM can be applied to support metacognitive processes involved in self-regulated learning activities. The motivation for our vision comes from two areas. On the one hand, the strong ideas behind OSM are related to metacognitive support: OSM increases self-awareness and self-control of the learning process [4]. We believe that the metacognitive support of OSM reaches another level in OSSM. For example, OSSM can give students a sense of common difficulties helping them to make better self-judgments when facing failure. On the other hand, although our approach to OSSM does not incorporate direct interaction and collaborative tools, there is a sort of "indirect interaction" or "soft

[1] Poster title: *Off the Beaten Path: The Impact of Adaptive Content Sequencing on Student Navigation in an Open Social Student Modeling Interface*

collaboration" happening mediated by the cognitive aspects of the group information displayed. This social dimension can be used to enrich OSM features, for example, to guide students using the traces of others. In the next sections we explore related work about open student model, metacognition in self-regulated learning, the measures of metacognitive processes in computer-based learning environments, and social awareness in computer-based collaborative environments, and from these ideas we propose a representation of OSSM.

2 Open Student Model and Metacognition

Open Student Model (or Open Learner Model) discloses the user model that the system maintains to the learner. As a result, OSM is a tool for self-awareness and learning monitoring. In the review of OSM work, Bull and Kay [4] described different systems incorporating OSM features or indicators supporting metacognition, including high level indicators of performance, OSM negotiation features (the learner can negotiate her user model with the system), and fine-grained indicators at finer levels of knowledge components (for example, concepts). Fine-grained conceptual representations of OSM, where the student can discover gaps in her knowledge that are hidden in high level indicators, have been attempted in a number of works [11, 13]. A common approach to fine-grained models involves a detailed domain model that can be represented as a concept-map or concept tree where nodes are concepts in the domain linked by the ontological or semantic relations among them. The learner model is an *overlaid* status of the learner in each of the concepts and it is represented by using, for example, colors [3]. A common problem of fine-grained models is that they can become very complex and hard to understand by the student. Visual techniques has been proposed to deal with this issue, for example, our system KnowledgeZoom uses semantic-zooming [1]. Open Student Model is also acknowledged to be of benefit when shared. For example, the instructor can perform a detailed monitoring of the learners, the learner can find collaborators by inspecting other models, compare with suitable ones, or improve group awareness through open group model [3]. Our vision of OSM incorporates the idea of sharing the OSM (we call it Open Social Student Model) and a fine-grained model that serves the student to make a more precise judgement of her own learning process.

3 Self-Regulated Learning and Metacognition

The research in Self-Regulated Learning (SRL) puts importance to feedback mechanisms in the development of cognitive and metacognitive processes. Feedback is not only related to the learner seeing summaries of her activity traces or providing information to others or the educational system, but also the internal feedback processes the learner develops while reflecting on the activities, for example the update of beliefs about herself and beliefs about the content of study [5]. Moreover, Butler and Winne [5] noted the heterogeneous and adaptive nature of self-regulation (here *adaptive* refers to the behavior of the learner that adapts during the learning experience) and they stressed the need of study it in a finer grain, i.e. continuously, while the learning activity is being performed. They proposed a broader view of self-regulated learning and feedback by describing four stages or elements: knowledge and beliefs, selection of goals, tactics and strategies, and monitoring. We take on this view and see ideas that can be applied in OSSM in each of these 4 stages. For example, for knowledge and beliefs, OSSM might project conflicting information to learner's self-efficacy beliefs (as the learner can compare her performance against others), and this discrepancy could be set to improve self-beliefs when possible. About goals, feedback can help the learner to set her goals and to make a good decision while navigating the content. Establishing a proper strategy to accomplish a goal might be difficult when the task is unfamiliar to the student, and here OSSM can use traces of other students to implement navigational guidance. Monitoring processes need to be supported by feedback information regarding both the current goal and about the progress on the learning activities.

Greene and Azevedo [7] saw the opportunity that Computer-Based Learning Environments (CBLEs) introduce for observing and measuring the learning process in detail, and reported a number of works using different forms of metacognitive measures and interventions in CBLEs. According to them, there are three types of techniques for measuring metacognition: 1) by self-reported instruments usually applied before or/and after the learning activity, 2) by using activity logged by the system or collected by sensors, and 3) inferred from explicit feedback given by the learner as the result of interacting with the system. They emphasized the idea that fine-grained metacognitive measure allows different levels of analyses, including semantic and statistical analyses of the activity, and analysis of sequences of actions in the time, which is in line with what Butler and Winne [5] recognized as necessary to study metacognition: continuous and on-line measurements. From the summary of Greene and Azevedo [7], we consider three broad ideas to incorporate in OSSM. First, it is important that the OSSM system collects all possible information while the learner performs the learning activities. Second, richer analyses and guidance can be achieved by incorporating some sort of dialogue or active interaction in the system that can be used to capture self-reported metacognitive state in real-time (for example asking the student what was the most difficult exercise, or asking the student to verify her model and write down her corrections). Third, the representation of the user model evolution over time (for example the progress in the last week), along with representation of the sequences of actions of the student and of the group or peers, can contribute to a better monitoring and planning tasks.

4 Social Awareness Tools

Janssen and Bodemer [9] summarized ideas of cognitive group awareness and social group awareness in collaborative activities supported by computer. Awareness, a process of inherent metacognitive character, can be of the type *Cognitive Group Awareness*, mainly about the knowledge and expertise of others (content space), or of the type *Social Group Awareness*, mostly about the levels of participation or engagement of peers and the quality of the interaction (relational space). Both broad types of group awareness are not exclusive of each of the spaces. For example, Cognitive Group Awareness also interacts with the relational space. Following this framework, we situate our idea of OSSM as a Cognitive Group Awareness (OSSM shows the knowledge/progress model of peers and the group), and we see the value of incorporating a dimension of Social Group Awareness, for example, by showing indicators of visits, attempts, and other current activity made by peers. Also, as Janssen and Bodemer suggested [9], using Cognitive Group Awareness features will also produce an impact in the relational space and we should not ignore it. For example Glahn, Sprecht and Koper [6] observed that even though a group awareness indicator (showing average of the group performance) produces the longest and strongly positive effect in the amount and the quality of work, it also produces frustration in non-contributing users and in some cases, the belief of vicious competitive behaviors of others. Moreover, the question is how to grasp the benefits of the group awareness features in OSSM on both content and relational spaces. Different group-awareness tools are implemented by Papanikolaou [12] in the system INSPIREus, including indicators of effort, progress, working style, personalization features, and social construction of knowledge (summarizing the type of discussions in forums). Students reported that the indicators allowed them to better understand their weaknesses and helped them to better plan their activities. We take on these ideas to incorporate different types of indicators for reflection, self-monitoring and comparison, specially, by combining indicators of activity with pedagogical information that sets the context of the desired metacognitive processes. Similar to INSPIREus, we maintain a domain model consisting of concepts mapped to the content material and activities, and structured using different semantic relationships, which can be used to provide indicators at different levels, for example high level indicators summarizing a topic.

5 A Concept-Map OSSM

We propose to complement our current OSSM MasteryGrids with a network representation of the concepts as the student progress in the learning activities. Activities are mapped to a set of finer grained concepts and these concepts get connected as the student practices activities containing pairs of concepts. Thus, the network representation or concept map, gets more connected as the student practices the concepts with different other concepts. We recognize that in many domains *mastery* is reached as the student is able to connect different concepts. We hypothesize that this concept map will allow students to have a finer and detailed view of their models, thus engage them in deeper metacognitive processes. On the other hand, the representation grows naturally as the student *connects* concepts, thus giving an idea of the dynamic progress or advance in pursuing learning goals. We plan to incorporate features supporting other metacognitive processes of goal setting and learning strategy. The learner should be able to choose concepts she wants to target, and the OSSM incorporates an indicator of the overall progress of the goal set. Recommendation and navigational clues are giving to signal concepts that should be targeted first and which activities to do to progress on those concepts. Collaborative filtering techniques can be used to grasp the *wisdom of the crowd* in order to power such recommendation mechanisms. For example, once a goal is set, the system can find the traces of other students that set similar goals in the past and use this information to recommend which activities to do. Each concept in the map can show information of the overall activity of the group related to the concept, for example to give an average of the progress on the concept. One important aspect on OSM is letting the learner correct or negotiate the model. Our implementation should allow students to change their knowledge levels through selected assessment items.

References

1. Brusilovsky, P., Baishya, D., Hosseini, R., Guerra, J., Liang, M.: Knowledgezoom for java: A concept-based exam study tool with a zoomable open student model. In: Advanced Learning Technologies (ICALT), 2013 IEEE 13th International Conference on. pp. 275–279. IEEE (2013)
2. Brusilovsky, P., Millán, E.: User models for adaptive hypermedia and adaptive educational systems. In: The adaptive web. pp. 3–53. Springer-Verlag (2007)
3. Bull, S., Kay, J.: Open learner models. In: Advances in intelligent tutoring systems, pp. 301–322. Springer (2010)
4. Bull, S., Kay, J.: Open learner models as drivers for metacognitive processes. In: International Handbook of Metacognition and Learning Technologies, pp. 349–365. Springer (2013)
5. Butler, D.L., Winne, P.H.: Feedback and self-regulated learning: A theoretical synthesis. Review of educational research 65(3), 245–281 (1995)
6. Glahn, C., Specht, M., Koper, R.: Visualisation of interaction footprints for enagement in online communities (2010)
7. Greene, J.A., Azevedo, R.: The measurement of learners self-regulated cognitive and metacognitive processes while using computer-based learning environments. Educational psychologist 45(4), 203–209 (2010)
8. Hsiao, I.H., Bakalov, F., Brusilovsky, P., König-Ries, B.: Progressor: social navigation support through open social student modeling. New Review of Hypermedia and Multimedia 19(2), 112–131 (2013)
9. Janssen, J., Bodemer, D.: Coordinated computer-supported collaborative learning: Awareness and awareness tools. Educational Psychologist 48(1), 40–55 (2013)
10. Loboda, T.D., Guerra, J., Hosseini, R., Brusilovsky, P.: Mastery grids: An open source social educational progress visualization. In: Open Learning and Teaching in Educational Communities, pp. 235–248. Springer (2014)
11. Mabbott, A., Bull, S.: Alternative views on knowledge: Presentation of open learner models. In: Intelligent Tutoring Systems. pp. 689–698. Springer (2004)
12. Papanikolaou, K.: Constructing interpretative views of learners interaction behavior in an open learner model
13. Perez-Marin, D., Alfonseca, E., Rodríguez, P., Pascual-Nieto, I.: Automatic generation of students conceptual models from answers in plain text. In: User Modeling 2007, pp. 329–333. Springer (2007)

Dual Eye Tracking as a Tool to Assess Collaboration

Jennifer K. Olsen[1], Michael Ringenberg[1], Vincent Aleven[1], and Nikol Rummel[12]

[1]Human Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA, USA
{jkolsen, mringenb, aleven}@cs.cmu.edu
[2]Institute of Educational Research, Ruhr-Universität Bochum, Germany
nikol.rummel@rub.de

Abstract. In working towards unraveling the mechanisms of productive collaborative learning, dual eye tracking, a method where two people's eyes are tracked as they collaborate on a task, is a potentially helpful tool to identify moments when students are collaborating effectively. However, we are only beginning to understand how eye gaze relates to effective collaborative learning and how it fits in with other data streams. In this paper, we present three broad areas of analysis where we believe dual eye tracking will promote our understanding of collaborative learning. These areas are: (a) How eye gaze is associated with other communication measures, (b) how eye gaze is associated with task features, and (c) how eye gaze relates to learning outcomes. We present exploratory analyses in each of the three areas using a dataset of 28 4th and 5th grade dyads working on an Intelligent Tutoring System for fractions. Our analyses illustrate how dual eye tracking could be used in conjunction with other data streams to assess collaborative learning.

Keywords: collaborative learning, intelligent tutoring system, dual eye tracking

1 Introduction

Collaboration can be an effective tool for learning; however, it can be difficult to identify the mechanisms of collaboration and how students' actions may lead to learning when working in a group. The communication between partners plays a large role in the success of the group [3], and there are many different processes that happen during a collaborative session that can affect learning such as speech, joint attention, and tutor feedback. By analyzing these different processes separately and together, we may be able to develop a better understanding of the collaborative learning process. In this paper, we specifically focus on dual eye tracking, a method where two people's eyes are tracked as they collaborate on a task, with an Intelligent Tutoring System (ITS) and explore how it could be used with other data streams to analyze students' collaborative interactions. By using multiple data streams that include eye gaze, we may be able to have insights into collaboration that were not otherwise possible.

Research shows eye gaze is tied to communication, making eye tracking a promising method to use for the analysis of collaborative learning [9]. Previous research has shown that there is a link between eye gaze and speech [4], [9]. When people hear a reference through speech, their eye gaze will follow that object [9], and when people are describing a picture, their eye gaze will look at the relevant part of the picture before it is described [4]. These studies show a link between speech and eye gaze that goes in both directions. This same pattern follows when people are working on a task together. There is a coupling of the collaborators' eye gaze around a reference [12]. The eye gaze has a closer coupling when each of the collaborators has the same initial information and when there is a shared selection [7], [12], suggesting that task features influence eye gaze. The coupling of eye gaze between collaborating partners may be an indicator of quality interaction and better comprehension [6], [11]. It also may be associated with better learning because there is more comprehension and understanding from the interactions with a closer coupling of eye gaze. Much of the previous work has focused on the correlation of eye gaze with speech, but it is still an open question of how dual eye tracking can be used to assess the effectiveness of collaboration in terms of learning and how it is associated with other process data, especially within an ITS.

In this paper, we will explore three types of broad questions that can be answered by using dual eye tracking: (a) How is eye gaze associated with other communication measures, (b) how is eye gaze associated with task features, and (c) how is eye gaze associated with learning outcomes. By answering these questions we may have a better understanding of how the interface of an ITS relates to both speech and the learning process while students are collaborating. There are multiple measures that can be gathered through dual eye tracking to understand eye gaze. In this paper, we will focus on one such measure, joint attention, which measures the relative amount of time two students are looking at the same area at the same time and corresponds to a very close coupling of eye gaze. Using a dataset of 4th and 5th grade students working on a fractions ITS, we explore a specific question in each of these three broad areas. These exploratory analyses demonstrate the potential of questions involving dual eye tracking and other data streams to be used to analyze collaborative learning.

2 Methods

2.1 Experimental Design and Procedure

Our data set involves 14 4th and 14 5th grade dyads from a larger study [10]. The dyads were engaged in a problem-solving activity in a collaborative ITS for fractions learning while communicating through audio only using Skype. Each dyad worked with the tutor for 45 minutes in a lab setting at their school. The morning before working with the tutor and the morning after working with the tutor, students were given 25 minutes to complete a pretest or posttest individually on the computer to assess their learning. Through the lab set-up in the school, we were able to collect dual eye tracking data, transcript data, and tutor log data in addition to the pretest and posttest measures for multiple stream of data.

2.2 Tutor Design

The ITS was developed using Cognitive Tutoring Authoring Tools and consisted of two problem sets, targeting procedural and conceptual knowledge. The tutor provides standard ITS support, such as hints and feedback [14], combined with embedded collaboration scripts. Each student had their own view of the collaborative tutor that allowed the students to have a shared problem space and synchronously work while being able to see slightly different information and to take different actions. Three different features supported the student collaboration. On some tutor steps, the students were *assigned roles* where they were either responsible for entering the answer or for asking questions of their partner and providing help with the answer. We supported other problem steps through *individual information* [13]. Here the students were each provided with a different piece of information that they needed to share with their partner. The final feature that was used to support collaboration was *cognitive group awareness* [5]. This feature was implemented in the tutor by providing each student an opportunity to answer a question individually before seeing each other's answers and being asked to provide a consensus answer.

2.3 Data and Dependent Measures

A computer-based test was developed to closely match the target knowledge covered in the tutors. The test comprised of 5 procedural and 6 conceptual test items, based on pilot studies with similar materials. Two isomorphic sets of questions were developed, and there were no differences in performance on the test forms, $t(79) = 0.96$, $p = 0.34$. The presentation of these forms as pretests and posttests was counterbalanced.

In addition, to pretest and posttest measures, we also collected process data including tutor log data, transcript data, and dual eye tracking data. The log data consisted of the transactions that the students took with the ITS. These include attempts at solving each step together with the request of hints and errors.

We coded the dialogue transcript data using a rating scheme with four categories: interactive dialogue, constructive dialogue, constructive monologue, and other. For our analysis, we focused on the interactive dialogue, in which students engage in actions such as co-construction and sequential construction. Interactive dialogue aligns with ICAP's joint dialogue pattern [2]. Our rating scheme was developed to look at groups of utterances associated with subgoals (i.e., a group of steps that all are for the same goal) to account for the interactions between the students. An inter-rater reliability analysis was performed to determine consistency among raters (Kappa= 0.72).

In addition to collecting log data and transcript data, we also collected dual eye tracking data using two SMI Red 250 Hz infrared eye tracking cameras. We calculated a measure of joint attention through gaze recurrence [1], [8]. Gaze recurrence is the proportion of times where the fixations are at the same location for each student. To calculate the joint attention from the gaze data, we used gaze recurrence with a distance threshold of 100 pixels to approximate the percentage of time that students were looking at the same thing at the same time. This distance threshold was chosen to align with prior research [6] and is close to the size of the interface elements.

3 Research Questions and Analysis

The first broad area of analysis is how eye gaze is associated with other communication measures. Within this area, we investigated how joint attention differs between subgoals without talk and subgoals with talk. We also explored whether or not there is an interaction with the subgoals that have errors. Based on previous work, we hypothesize that subgoals with talk will have a higher level of joint attention than subgoals with no talk since talk has been found to be coupled with eye gaze and speech might guide the visual attention [9]. In addition, we hypothesize that subgoals where an error occurred will have a higher level of joint attention compared to subgoals where no error occurred because there will be a visual red mark on the screen for the students to discuss [12]. To investigate the association between talk and joint attention, we used a hierarchical linear model with two nested levels to analyze how the talk during subgoals related to the joint attention. At level 1, we modeled if talk occurred and if one or more errors occurred for the subgoals. At level 2, we accounted for random dyad differences. We found no effect of errors on joint attention, so it was removed from the model. We found greater joint attention for subgoals that had talk ($M = 0.25$, $SD = 0.13$) versus those that did not ($M = 0.22$, $SD = 0.14$), $t(1705)= 12.66$, $p < .001$, showing a coupling between talk and joint attention that extends previous results to younger learners working in an ITS environment.

The second broad area of analysis is how eye gaze is associated with task features. For this area, we investigated how eye gaze is associated with the tutor's three types of collaboration support. Based on previous work, we hypothesize that subgoals supported through individual information would have the lowest joint attention since there is no joint reference for the students on the screen [7]. To investigate the association between collaboration features and joint attention, a hierarchical linear model with two nested levels was used to analyze how collaboration features relate to the joint attention. At level 1, we modeled the type of collaboration support of the subgoals along with the talk type to control for this covariate. At level 2, we accounted for random dyad differences. We found that the joint attention for subgoals that were supported through cognitive group awareness ($M = 0.19$, $SD = 0.11$) was lower than that for subgoals supported through roles ($M = 0.25$, $SD = 0.14$), $t(1705)= -4.19$, $p < .001$, indicating that task type has an impact on joint attention.

The third broad area of analysis is how eye gaze is associated with learning. Within this area, we investigated how joint attention correlates with learning gains for conceptual and procedural knowledge. Based on previous work where we analyzed the first four questions (opposed to the entire session) [1], we hypothesize that joint attention will be correlated with conceptual learning gains, but not procedural learning gains, because a deeper understanding is needed to acquire the conceptual information that can be supported through joint attention [11]. To investigate this question, we computed a linear regression between posttest score and joint attention while controlling for pretest scores. Individual pretest and posttest scores were averaged for each member of the dyad for a single score for each dyad, and the joint attention was calculated for each dyad for the entire 45-minute session. Our results replicate previous findings, where for the conceptual condition, there were no significant results for conceptual or procedural posttest scores. For the procedural condition, there was no significance for procedural posttest scores, but joint attention significantly predicts conceptual posttest scores when controlling for conceptual pretest score, $t(10) = 2.6$, $p = 0.03$, showing joint attention may be more important for gaining conceptual knowledge on procedural problems, whereas students working on the conceptual problems were able to learn the same information with less joint attention.

4 Discussion

Although the correspondence of eye gaze with speech has been studied before, it is still an open question of how dual eye tracking can be used to assess the effectiveness of collaboration in terms of learning and how it is associated with other process data. In this paper, we explore the importance of eye gaze for collaborative learning analysis by presenting three different areas of analysis for using dual eye tracking data. Although the results are preliminary, these questions provide a broad structure and illustrate the potential of dual eye tracking to be used with other data streams.

Can dual eye tracking be used to understand the collaborative learning process? Through our analysis, we found that subgoals where talk occurs have a higher level of joint attention, extending previous results to younger learners and an ITS environment [12]. This result indicates that in an environment where there is step-by-step guidance and steps are revealed one at a time, which may guide eye gaze, there is still a benefit of speech for referencing items on the screen. Although we did not find any impact of errors on joint attention, analyzing the joint attention immediately after an error may provide a better indication of the effect of errors on joint attention. In addition, we found that subgoals supported through cognitive group awareness had a lower level of joint attention compared to those supported through roles showing the importance of the task features on collaboration. The difference between collaborative features on joint attention may be because the students would be looking at different points while answering individually and would then be looking at their partner's answer after it is revealed on cognitive group awareness, which may split the attention of the partners. We also used dual eye tracking to identify moments where collaboration may successfully support learning. We found joint attention as a significant predictor of conceptual posttest scores in the procedural condition, showing collaboration and joint attention may be important for conceptual knowledge specifically when it is not being directly supported. Although the results are preliminary, they show the potential of using dual eye tracking along with other data streams to better understand collaboration. For collaborative learning, dual eye tracking can provide insights into tasks that elicit collaboration as well as providing insights into how joint eye gaze interacts with other communication measures to impact learning.

For future work, we would like to expand the three areas of analysis around dual eye tracking beyond joint attention. There are other measures such as AOIs (areas of interest) analyses and gaze patterns that would be of interest in each of the three areas and can be measured through dual eye tracking. These different measures of eye gaze would not only provide additional ways of comparing collaboration within groups by looking at AOIs and gaze patterns that occur for partners at the same time, but would also allow the comparison to students working individually to see how collaboration affects the learning process. In addition, in our analyses so far, we have analyzed joint attention at the subgoal level and the dyad level, but analysis at additional grain sizes, such as a few seconds around errors and the problem level, would allow us to ask a wider range of questions. This future work will build upon the analysis presented in this paper to further explore the three broad areas of analysis for dual eye tracking to shed light on the mechanisms of collaborative learning.

Acknowledgments. We thank the CTAT team, Daniel Belenky, and Amos Glenn for their help. This work was supported by Graduate Training Grant # R305B090023 and by Award # R305A120734 both from the US Department of Education (IES).

5 References

1. Belenky, D.M., Ringenberg, M., Olsen, J. K., Aleven, V., & Rummel, N.: Using dual eye-tracking to evaluate students' collaboration with an intelligent tutoring system for elementary-level fractions. Paper in the *36th Annual Meeting of the Cog. Sci. Society.* (2014)
2. Chi, M.T.H.: Active-constructive-interactive: a conceptual framework for differentiating learning activities. Topics in Cognitive Science, 1, 73-105 (2009)
3. Chi, M. T., & Wylie, R.: The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes. *Educational Psychologist*, *49*(4), 219-243 (2014)
4. Griffin, Z. M., & Bock, K.: What the eyes say about speaking. Psychological science, 11(4). (2000)
5. Janssen, J., & Bodemer, D.: Coordinated computer-supported collaborative learning: Awareness and awareness tools. Educational Psychologist, 48(1), 40-55 (2013)
6. Jermann, P., Mullins, D., Nüssli, M.-A., and Dillenbourg, P.: Collaborative Gaze Footprints: Correlates of Interaction Quality. In Proceedings of CSCL Conference 2011, 184-191 (2011)
7. Jermann, P. and Nüssli, M. A.: Effects of sharing text selections on gaze recurrence and interaction quality in a pair-programming task. In Proceedings of the 2012 ACM CSCW conference. (2012)
8. Marwan, N., Romano, M. C., Thiel, M., Kurths, J.: Recurrence Plots for the Analysis of Complex Systems, Physics Reports, 438(5-6), 237-329 (2007)
9. Meyer, A. S., Sleiderink, A. M., & Levelt, W. J. M.: Viewing and naming objects: Eye movements during noun phrase production. Cognition, 66(2), (1998)
10. Olsen, J. K., Belenky, D. M., Aleven, V., & Rummel, N.: Using an intelligent tutoring system to support collaborative as well as individual learning. In 12th International Conference on Intelligent Tutoring Systems, 134-143 (2014)
11. Richardson, D. C., & Dale, R.: Looking to understand: The coupling between speakers' and listeners' eye movements and its relationship to discourse comprehension. Cognitive Science, 29, 1045-1060 (2005)
12. Richardson, Daniel C., Dale, R., & Kirkham, N. Z.: The Art of Conversation Is Coordination. Psychological Science, 18(5), (2007). doi:10.1111/j.1467-9280.2007.01914.x
13. Slavin, R. E.: Research on cooperative learning and achievement: What we know, what we need to know. Contemporary educational psychology, 21(1), 43-69 (1996)
14. VanLehn, K.: The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. Educational Pyschologist. 46(4) 197-221 (2011)