Intelligent Conversational Agents as Facilitators and Coordinators for Group Work in Distributed Learning Environments (MOOCs)

Gaurav Singh Tomar, Sreecharan Sankaranarayanan, and **Carolyn Penstein Rosé**
{gtomar, sreechas, cprose}@cs.cmu.edu
Language Technologies Institute
Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

Artificially intelligent conversational agents have been demonstrated to positively impact team based learning in classrooms and hold even greater potential for impact in the now widespread Massive Open Online Courses (MOOCs) if certain challenges can be overcome. These challenges include team formation, coordination and management of group processes in teams working together while distributed both in time and space. Our work begins with an architecture for orchestrating conversational agent based support for group learning called Bazaar, which has facilitated numerous successful studies of learning in the past including some early investigations in MOOC contexts. In this paper, we briefly describe our experience in designing, developing and deploying agent supported collaborative learning activities in 3 different MOOCs in three iterations. Findings from this iterative design process provide an empirical foundation for a reusable framework for facilitating similar activities in future MOOCs.

Introduction

In this paper we present resources grounded in an empirical foundation for facilitating group learning in online communities using intelligent conversational agents. In particular, we focus on the now widespread MOOCs. The motivation for our work is drawn from literature in Computer-Supported Collaborative Learning (CSCL).

Despite the inclusion of discussion forums in virtually all MOOCs, the typical experience of participation in a MOOC is solitary. In particular, MOOC contexts allow for asynchronous interaction and possibly even collaboration. However, the asynchronous nature leads to some unsavory experiences. For example, sometimes participants spend time posting a thoughtful post but get only a cursory reply. Participants sometimes have to wait days or even weeks to get a response to a question. This lack of immediacy implies that information can be out of date by the time someone views it, or worse, that the student needing help or feedback gave up and dropped out before the response was posted. These types of issues can hinder motivation.

Traditionally, MOOC platforms, including team based MOOC platforms like NovoEd, have not offered synchronous interaction opportunities or even instantaneous forms of social awareness. These limitations at the interface level stem from limitations at the architecture level due to challenges in scaling immediate update protocols. Through integration protocols such as the Learning Tools Interoperability protocol (LTI) chat tools and other forms of synchronous interaction have found their way into a few MOOCs. However, just the ability to integrate a synchronous chatroom into a MOOC does not solve the problem. Challenges remain regarding coordinating the times of the discussions as well as supporting the functioning of an ongoing discussion. We outline a vision for work towards a solution to both of these challenges in this position paper.

The proposed solution to this problem is to support synchronous interaction within the MOOC context. Students may prefer synchronous activities because it offers them a greater experience of active involvement and social connection. It is our hope that this social stimulation will lead to higher commitment to the MOOC as well as greater intensity of participation and further facilitate interactions that lead to learning that includes social learning experiences involving more than one student, with an opportunity for communicating and collaborating. The value of these learning experiences lies in the opportunity for students to see the learning material through the distinct lenses of other students' viewpoints, facilitating self-reflection and ultimately a deeper understanding of their own knowledge.

In the remainder of the paper, we first introduce our theoretical framework explaining the rationale for supporting synchronous collaboration and the types of interactions we hope to foster through this collaboration. This is followed by an outline of an extension to the Bazaar framework (Adamson and Rosé 2012) - a modular framework for designing multi-party collaborative agents, which has been used to develop a chat tool for supporting synchronous group work for MOOC studies mentioned in this paper. Next we describe the experience during our iterative design based research process relating to deployment of chat tools in three different MOOCs spanning a variety of group formation and coordination types. Our description follows a "lessons learned" paradigm in which the description also highlights the major problems that were identified in each cycle and our proposed design decisions made to solve these problems for subsequent cycles and how we incorporated these decisions at a lower level, both in the behavior and dialogues of the agent.

Theoretical Foundation

The fact that interaction is a site of (collaborative) learning has been underscored by many learning theories ranging from cognitive to socio-cultural perspectives (Webb 2013). Interaction is conceived as providing potential opportunities for learning, which does not necessarily mean that all actual interactions lead to learning, but instead that meaningful interactions may facilitate learning. In fact, it is well known that without support, many instances of collaborative learning fail. In light of this fact, the field of collaborative learning has produced a wide variety of forms of scaffolding often referred to as scripts. Scripts may operate at the macro level, providing task structuring and role assignment. Or it may operate at the micro level, structuring the nature of the flow of contributions to the discourse.

With scaffolding, students are able to contribute to their collaborative endeavors in a more skilled manner than they would be able to without. Thus, it is essential for designers of such scaffolding to be aware of what behaviors are valuable for learning, and to design scaffolding that increases the propensity for participants to engage in those behaviors. As will be shown throughout the paper, our chat activity and agent designs seek to build upon this fact by combining macro and micro scripts in the kinds of interactions (learner-agent and learner-learner) elicited.

Our goal is to find the best ways to provide learners with a space to interactively hone their understanding of concepts related to a specific domain so that they have the chance to display their reasoning, experience how others display their own reasoning, challenge and be challenged by others, among other interactions.

In order for the interactions to provide a meaningful learning experience for students, it is essential that they be well integrated into the instructional design, and not treated as an afterthought or an appendage. This leads to more authentic, and ecologically valid learning experiences for students (Tudor 2001; Van Lier 2004). In addition to this embedding, interactions are mainly meaningful for learning when they are structured and scaffolded in an appropriate manner.

In our setup, scaffolding is provided through the facilitation moves of an intelligent conversational agent taking a teacher role. The goal is for the agent to aid the students in integrating their respective understandings of the concepts they previously encountered individually in (the video lectures that precede the chat activity). On the other hand, it also provides an opportunity to develop their collaborative skills (ODonnell 1999).

Bazaar

The chat tool system deployed in MOOC studies mentioned in the paper employed tutorial dialogue agent technology of the Bazaar framework to provide interactive support within a synchronous collaborative chat environment. The design and development of this tool system and the underlying agent followed cycles of an iterative design based research process. During these cycles, we faced challenges of engaging students in more intensive discussion-based interactions in MOOCs. Figure 1 depicts agent architecture networks

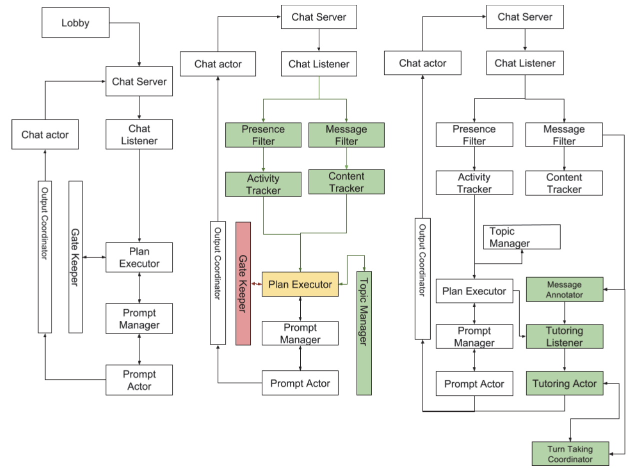

Figure 1: Iterative design process : agent architecture networks for cycle 1, 2 & 3. Green components were added, red ones removed and yellow ones modified as part of design decisions made to solve the problems in previous cycles.

across these cycles and many of the components in these networks will be referred throughout the paper in italics. The independent components in each of these agent architecture networks receive and respond to user, environment, and system generated actions and present the orchestrated output of these components to the user which is decided by the *Output Coordinator* component and is pushed back to the environment by the *Chat Actor* component.

Events come in from a collaborative environment (like text chat rooms and shared whiteboards) via *Chat Server* component, traverse a network of other components generating new events at every step, then a final set of events are pushed back to the environment via the *Chat Actor* component. Among components in this network, listeners/trackers monitor stimuli from the environment and translate them into events internal to the agent, actors perform actions which may be directly observable by other participants in the environment and filters process information that events carry and propagate further based on their programmed conditions.

The system and agent design employed in the studies discussed in next section could be used to very rapidly conduct experiments for investigating a wide range of important questions within the design space of group learning. The system deployment effort will be focused entirely on authoring domain dependent dialogs, while a rich set of domain-independent conversational behaviors are transparently generated by the dialog agent.

Iterative Design Based Research Process

This section describes three cycles of a design based research process in which we made synchronous chat collaboration opportunities available to students in form of weekly reflective chat exercises to be done after completing the course content for that week in three MOOC courses mentioned in table 1. The third one is still running in its third

Course	abbrev.	duration
Data, Analytics and Learning	DALMOOC	9 weeks
Big Data in Education	BDEMOOC	8 weeks
Medicinal Chemistry	MCMOOC	8 weeks

Table 1: Three MOOC courses used for deployment studies

week. These interactions are potentially of personal benefit to the students but involve costs in the form of group formation and coordination issues discussed in this section. We seek effective practices for incorporating synchronous collaboration in MOOC contexts in light of these tradeoffs.

Cycle 1 - DALMOOC deployment study

Team formation: In order to facilitate the formation of adhoc study groups for the chat activity, we made use of a simple setup referred to as a Lobby. The *Lobby* component in first network of Figure 1 introduced an intermediate layer between the edX platform and the synchronous chat tool. Students entered the Lobby with a simple, clearly labeled button integrated with the edX platform. Upon entering the lobby, students were asked to enter a username that would be displayed in the chat. Once registered in the lobby, the student waited to be matched with another participant. If they were successfully matched with another learner who arrived at the Lobby within a couple of minutes to interact with, they and their partner were then presented with a link to click on to enter a chat room created for them in real time. Otherwise they were requested to come back later.

Team management: 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 (Adamson et al. 2014). 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 (Adamson et al. 2014), in DALMOOC we employed statically and macro scripted agents which guide the students with course-related discussion questions. The *PlanExecutor* component delivered a timed sequence of steps in this macro static script, starting when a Launch Event is received which is when a student enters the chat room. The prompts sequenced and phrased by the *PlanExecutor* component are delivered by the *PromptActor* component. The *PromptManager* component takes care of built-in delay related to the length of the prompt, to allow time for reading the prompt.

Team Coordination: Even though the scripts were linear, the agent prompts were not strictly timed but rather allowed the students to interact at their own pace and take as much time as needed to discuss the given topic. Once a pair wanted to proceed with the discussion, they could move on by pressing the "Were ready" button. The agent would proceed with the next prompt as soon as both students indicated that they were ready otherwise it informed the student pressing it to wait for his peer to do the same. In case the students never signaled their readiness, the agent would inquire after a predefined timeout in order to move forward with the discussion and and help to keep the students from losing focus. This coordination is handled by the *Gatekeeper* component in the agent architecture network. ***Lessons learned*** (Ferschke, Tomar, and Rosé 2015): In this first deployment study, we learned valuable lessons that helped to improve the experimental setup in future cycles of our iterative design based research process. The main results suggest educational value was added by the intervention (Ferschke et al. 2015a) and other results based on survival model analysis suggested that there is a substantial reduction in attrition over time when students experience a match for a synchronous collaborative reflection exercise (Ferschke et al. 2015b). But 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. This was due to the lack of a critical mass available during these attempts. Understandably, the analysis showed a negative impact of not getting matched with a partner to do the reflective chat exercises. The next deployment study tried to address this issue.

Cycle 2 - BDEMOOC deployment study

Team formation: Our analysis of the previous deployment study revealed that integrating a synchronous collaborative activity in an inherently asynchronous learning environment used by students in different time zones and conflicting schedules poses inherent organizational challenges. In order to address this, we removed the intermediate Lobby interface as can be seen in the second cycle of Figure 1. Instead we exposed a single, continuous chat room that multiple students could join at anytime without being matched with a particular peer student. The conversation continues as long as there is at least one person in the room. The tool is able to facilitate a chat with only one student present and keeps the conversation going for any number of participants. That way, students can enter the chat on their own schedule and join other students who are engaged in the activity. The room automatically resets once all participants left the chat.

Team Management: The agent kept track of the students in the room and prompts that had been given to the students. For students, it tracked the events of their joining and leaving the room using the *Presence Filter* component. For topic prompts, it tracked the times they were given and the students who were present in the chat room at that time using the *Topic Manager* component. This tracking was done to make sure each student could engage with all question prompts of the reflection exercise and no student saw the same question prompt twice unless the activity script restarted and students decided to repeat it.

During a synchronous chat collaborative activity, the students may not be motivated to stay on topic and may engage in goalless interactions or no interactions at all. It might also be because students don't understand the question properly. To address this issue, the agent kept track of any student and group inactivity that went on for more than two minutes, using the *Activity Tracker* component. This component also keeps track of how many utterances each student has contributed recently so that they can be graded if desired. The *Message Filter* component listened for students' messages and the *Content Tracker* component calculated a similarity score between student messages and the topic prompt to know if the students were on topic. This was also done over a two minute window. If the group was dormant or does not say anything contentful in that two minute window then the agent generated a poke message which rephrased the question again by providing timely hints so that the student could discuss on it more easily. A contentful group discussion on a question prompt could last for 10 minutes.

Team Coordination: Although a single, continuous chat room provides the capacity to solve the problems with synchronous collaboration in a MOOC, it may entail additional coordination challenges. For instance, a student may not be motivated to join in the middle of the discussion or they might feel lost due to a lack of a frame of reference for where to start. To address this issue, we introduced a summarization behavior to be elicited by the agent to ramp up new participants and encourage current participants to voice their understanding of discussed concepts. The agent-facilitated summary can be given by the agent itself by mentioning list of topics already discussed and a summary of current topic (maintained by the *Topic Manager* component) or be elicited by the agent from other participants in the conversation whenever a new participant joins the room. To ensure the agent does not interrupt the discussion with repeated summarization requests, these requests were issued only if at least two topic prompts had been discussed since the last request so that the agent does not interrupt people in the room too frequently. This summarization facilitation move was supported by the *Presence Filter*, *Plan Executor* and *Topic Manager* components.

Lessons learned: The idea of a single continuous chat room seemed to work as matching into pairs was not required, which was especially beneficial during periods in these studies when an odd number of students were present in the chat room (see table 3). A pairing approach like that used in the first deployment study could potentially have left a person waiting for a peer forever. Table 3 shows that more than 85% of the match attempts would have failed as group sizes of odd numbers would not have been possible. By converting these potential failed match attempts into successful opportunities to do reflection exercises, we have taken a step towards alleviating the negative impact found by the first deployment study. Results showed that when summarization moves were successful, they did help new entrants to alleviate some of their confusion and stay in the room.

We observed that many students who were chatting with the agent in the absence of peers felt that agent was not a good listener, felt bored and left room in the middle of the discussion. They expressed a desire to be able to discuss the questions with the agent instead of just giving their reflections on the question prompt being asked. Another reason people left the room turned out to be the strict timing of at least four minutes to be spent on each question prompt, which made the student impatient as he was unable to move to the next question prompt immediately. A single student usually does not need to review each topic for four minutes.The final deployment study sought to address these issues.

Cycle 3 - MCMOOC deployment study

Team formation: There were no changes in team formation techniques compared to the previous deployment study but chat exercises were graded. This resulted in greater critical mass and thereby bigger groups could be formed.

Team management & coordination: To alleviate the boredom felt by a single student in the chat room, the system now supports discussions between a student and the agent in addition to student-student discussions. Specifically, Knowledge Construction Dialogues (KCDs) have been introduced, which can help a student feel more involved even without peers by discussing questions with the agent while providing his reflections on question prompts. If a single student is present in the room, the *Plan Executor* launches an event which is listened by the *Tutoring Listener* component which directs the *Tutoring Actor* component to start a KCD. KCDs ask a series of questions with the goal of leading a student to construct a correct explanation for the question being discussed. This is composed of a sequence of steps in an Initiation-Response-Feedback interaction pattern. The dialog engine executes these steps by presenting the Initiation question, matching the student response, and presenting appropriate feedback before moving on to the next step. The script formalism also provides the opportunity to introduce another intervening sequence of remedial steps as feedback to incorrect responses. Student replies received via the *Message Filter* and the *Message Annotator* components are collected by the *Turn Taking Coordinator* component when the agent is expecting the students to respond to its question. The findings from previous study showed that these rigid timings drove away single students from the chat rooms as they became impatient waiting to be moved to the next prompt. So we incorporated adaptive timings and delays for the prompts based on the number of students in the chat room so that a single student or small groups need not spend same amount of time as bigger groups on a question prompt.

Conclusion

Bazaar is a publicly available architecture inviting further work from a broad and creative community of researchers working on intelligent support for group learning. Our group has conducted numerous experiments using it in both classroom and MOOC learning studies. Its flexibility allows for rapid and iterative development of platforms for investigating a wide range of important questions within the design space of group learning. We presented here a list of our studies using it in a MOOC context highlighting the rapid and iterative nature of its development process. As part of these studies, we investigated and aimed to overcome some important challenges related to group formation, coordination and management in distributed learning environments. As we continue to do so, we expect to discover ways in which the Bazaar architecture can be extended and refined. The findings show that synchronous group collaboration in online learning environments has benefits just like other learning environments as long as the cost (in form of group formation and coordination) for entering and participating into such group work does not appear formidable to learners.

References

Adamson, D., and Rosé, C. P. 2012. Coordinating multi-dimensional support in collaborative conversational agents. In *Intelligent Tutoring Systems*, 346–351. Springer.

Adamson, D.; Dyke, G.; Jang, H.; and Rosé, C. P. 2014. Towards an agile approach to adapting dynamic collaboration support to student needs. *International Journal of Artificial Intelligence in Education* 24(1):92–124.

Ferschke, O.; Howley, I.; Tomar, G.; and Yang, D. 2015a. Fostering discussion across communication media in massive open online courses.

Ferschke, O.; Yang, D.; Tomar, G.; and Rosé, C. P. 2015b. Positive impact of collaborative chat participation in an edx mooc. In *Artificial Intelligence in Education*, 115–124. Springer.

Ferschke, O.; Tomar, G.; and Rosé, C. P. 2015. Adapting collaborative chat for massive open online courses: Lessons learned. In *ISLG 2015 Fourth Workshop on Intelligent Support for Learning in Groups*, 13.

ODonnell, A. M. 1999. Structuring dyadic interaction through scripted cooperation. *Cognitive perspectives on peer learning* 179–196.

Tudor, I. 2001. *The dynamics of the language classroom*. Ernst Klett Sprachen.

Van Lier, L. 2004. *The ecology and semiotics of language learning: A sociocultural perspective*, volume 3. Springer Science & Business Media.

Webb, N. 2013. Information processing approaches to collaborative learning. *The international handbook of collaborative learning* 19–40.

Appendix A

	graded?	#weeks	#rooms	#unique users
DALMOOC	No	9	217	371
BDEMOOC	No	8	598	647
MCMOOC	Yes	2.5	389	487

Table 2: Information about three Bazaar based deployment studies

	1	2	3	4	5	>=6	Total
DALMOOC		217					217
BDEMOOC	486	68	20	12	4	8	598
MCMOOC	332	37	8	4	1	7	332

Table 3: Group sizes in three Bazaar based deployment studies