Tegos, S., & Demetriadis, S. (2017). Conversational Agents Improve Peer Learning through Building on Prior Knowledge. *Educational Technology & Society, 20* (1), 99–111.

Conversational Agents Improve Peer Learning through Building on Prior Knowledge

Stergios Tegos* and Stavros Demetriadis
School of Informatics, Aristotle University of Thessaloniki, Greece // stegos@csd.auth.gr // sdemetri@csd.auth.gr
*Corresponding author

(Submitted October 29, 2015; Revised March 23, 2016; Accepted April 8, 2016)

ABSTRACT

Research in computer-supported collaborative learning has indicated that conversational agents can be pedagogically beneficial when used to scaffold students' online discussions. In this study, we investigate the impact of an agile conversational agent that triggers student dialogue by making interventions based on the academically productive talk framework. An experimental activity in the context of a university course involved 72 undergraduate students who discussed online in dyads. Two conditions were compared: (a) dyads who received agent interventions while working on a learning task (treatment) and (b) dyads who worked on the same task without any agent interference (control). Utilizing a concept map created by the course instructor, the conversational agent delivered unsolicited interventions that encouraged treatment students to build on their prior knowledge, linking their current contributions to the main domain principles discussed during the course. Study findings indicate that the agent intervention mode substantially improved both individual and group learning outcomes. Evidence suggests that the agent effect on learning performance was mediated by students' explicit reasoning, which was also found to be enhanced in the treatment condition.

Keywords

Conversational agent, Computer-supported collaborative learning, Academically productive talk

Introduction

Conversational agents

In the field of technology-enhanced learning, pedagogical agents have been developed to serve a wide variety of instructional roles, such as expert, motivator, or mentor (Baylor & Kim, 2005). Conversational agents are typically regarded as a subgroup of pedagogical agents involving learners in natural language interactions (Kerly, Ellis, & Bull, 2009). Research has shown using conversational agents to engage learners in one-to-one (student-agent) tutorial dialogues to improve students' comprehension and foster students' engagement and motivation (Veletsianos & Russell, 2014).

During the past decade, researchers also focused on developing conversational agents for collaborative learning support (e.g., Kumar & Rosé, 2011). Despite the established cognitive and social benefits of computer-supported collaborative learning (CSCL), collaborative knowledge construction is not a given but depends on the quality of interactions taking place among learners (Dillenbourg & Tchounikine, 2007; Kreijns, Kirschner, & Jochems, 2002). Under this prism, well-targeted supportive interventions can be used as a method to increase the probability of constructive peer interactions occurring by means of stimulating cognitive processes, such as conflict resolution, mutual regulation or explicit explanation (Tchounikine, Rummel, & McLaren, 2010). Evidence suggests that conversational agents with social interaction capabilities can enhance learning and idea generation productivity by providing dynamic support for learners working together (Kumar & Rose, 2011; Kumar, Beuth, & Rosé, 2011). Chaudhuri et al. (2008) reveal that agents guiding peers through prescribed lines of reasoning on specific topics can improve learning performance. A study by Walker, Rummel, and Koedinger (2011) indicates that an agent displaying reflective prompts in a scripted peer tutoring activity can help students produce conceptually richer statements.

Academically productive talk

Another research direction has recently emerged focusing on an agile form of conversational agent support, which emphasizes the key role of social interaction in student engagement and learning (e.g., Adamson, Dyke, Jang, & Rosé, 2014). This approach draws on the academically productive talk (APT) framework, itself originating from a substantial body of work on useful classroom discussion practices and norms (Michaels, O'Connor, & Resnick, 2008). According to APT, a peer dialogue in class should be accountable to the learning

ISSN 1436-4522 (online) and 1176-3647 (print). This article of the Journal of Educational Technology & Society is available under Creative Commons CC-BY-ND-NC 3.0 license (https://creativecommons.org/licenses/by-nc-nd/3.0/). For further queries, please contact Journal Editors at ets-editors@ifets.info.

community, accurate knowledge and rigorous thinking, irrespective of the subject area (Sohmer, Michaels, O'Connor, & Resnick, 2009). In view of the above, peers should paraphrase and expand on each other's ideas (i.e., being accountable to the learning community), support the validity of their claims making explicit references to a pool of knowledge accessible to the group (i.e., being accountable to accurate knowledge), and logically connect their statements through rigorous argumentation (i.e., being accountable to rigorous thinking).

Unlike other well-known discourse frameworks such as the IRE (Initiation, Response and Evaluation), the APT framework does not entail closing down a conversation after successfully eliciting a correct learner's response; instead, APT aims to promote and scaffold open-ended discussions where learners explicate their reasoning, compare their contributions with their partners' and construct logical arguments based on accurate evidence (Michaels, O'Connor, Hall, & Resnick 2010). Indeed, APT does not expect the teacher to maintain full control over learners' discussions, and prioritizes reasoning over correctness. The importance of the explicit articulation of reasoning is universally acknowledged by researchers, despite the different conceptualization of studies exploring the key features of a productive peer dialogue (for example, "transactivity," "group cognition," and "productive agency") (Stahl & Rosé, 2011). The explicitness of students' reasoning has been shown to enhance conceptual knowledge acquisition and improve collaboration practices (Papadopoulos, Demetriadis, & Weinberger, 2013).

Nevertheless, promoting students' reasoned participation and orchestrating academically productive discussions are often challenging tasks, requiring teachers to provide dynamic support via facilitative conversational moves (Sohmer et al., 2009). Table 1 presents a selection of those APT moves, which can be regarded as conversational interventions (or actions) aiming to model and trigger appropriate forms of peer dialogue (Michaels et al. 2010). Considering that such moves have proven conducive to scaffolding open-ended discussions in class, researchers further proceed to explore their potential contribution to CSCL by tailoring agent supportive mechanisms to display APT-based prompts (Dyke, Adamson, Howley, & Rosé, 2013).

Table 1. A list of academically productive talk moves

APT Move	Example	Accountability dimension
1. Link Contributions A. Agree-Disagree B. Add-On	A. "Do you agree with what your partner said about …?" B. "Would you like to add something to …?"	Learning Community
2. Revoice	"So, are you saying that … Is that correct?"	Learning Community
3. Build on Prior Knowledge	"How does this connect with what we know about …?"	Accurate Knowledge
4. Press for Reasoning	"Why do you think that?"	Rigorous Thinking
5. Expand Reasoning A. Take your Time B. Say More	A. "Please, take your time before answering" B. "That's interesting! Can you elaborate on that?"	Rigorous Thinking

Agent support for academically productive talk

A study involving ninth-grade students as participants indicates that a Revoicing form of support (Table 1, move 2) may have a positive learning effect on novice learners, not familiar with the instructional material (Dyke et al., 2013). Additionally, a study by Adamson, Ashe, Jang, Yaron and Rosé (2013) employing university students as participants suggests that an Agree-Disagree agent intervention mode (Table 1, move 1A) can enhance students' learning and intensify group knowledge exchange. In contrast, a similar study on university students illustrates that experienced learners may not benefit at all from a Revoicing intervention mode since they are already capable of articulating and modeling their own ideas without the need for support (Adamson et al., 2014). Following a similar rationale in higher education settings, a recent study explores the impact of a conversational agent displaying both Agree-Disagree and Add-On interventions (Table 1, moves 1A and 1B) (Tegos, Demetriadis, & Karakostas, 2015). Findings suggest that such agent interventions can improve domain-learning outcomes and increase students' explicit reasoning during their online discussions.

The evidence obtained from the few studies conducted in this area reveals that the efficacy of the APT-based agent support can significantly vary in terms of the APT move used, the difficulty of the instructional domain, and the students' background knowledge. This is consistent with Michaels et al.'s (2008) findings indicating that the effectiveness of APT facilitation strategies may depend on a number of factors, such as the teacher's authority, the students' background, and the education level. Hence, there are a number of issues remaining, far from trivial, concerning the design and use of conversational agents displaying APT-based prompts. For instance, what is the context where each APT move can be used most effectively? Could we design an agent alternating different APT moves in accordance to the learning task or student population? What should the teacher's role be in such a system?

Research objectives

Expanding on previous research in the utilization of agent prompts that foster accountability to the learning community (e.g., Adamson et al., 2013; Tegos et al., 2015), in the following section we present a study exploring the effectiveness of a configurable agent intervention mode promoting accountability to accurate knowledge. More specifically, the study investigates whether a series of agent interventions encouraging students to build on their prior knowledge (Table 1, move 3) can trigger explicit reasoning processes and enhance learning performance in a collaborative activity in higher education. We expect the outcome of the study to inform researchers and designers on the pedagogical value of implementing similar APT-based agent intervention modes to provide collaborative learning support and help students sustain a constructive peer dialogue.

Method

Participants and instructional domain

The study involved 72 undergraduate computer science students (18 females and 54 males). All of them were native Greek speakers and their age ranged from 19 to 23 (*M* = 20.39, *SD* = 1.06). The activity of the study was carried out in the context of a second-year course on Human-Computer Interaction (HCI) as a mandatory course assignment. All students were required to successfully complete the activity in order to pass the course. During the HCI course, students became acquainted with principles of perception and cognition required for effective interaction design (Preece, Sharp, & Rogers, 2015). Additionally, students learned useful practices for designing, prototyping, and evaluating human-centered user interfaces.

MentorChat: A configurable conversational agent system

In the study we used MentorChat, a prototype dialogue-support system (Tegos, Demetriadis, & Karakostas, 2015). As Figure 1 indicates, the interface of MentorChat, resembling an instant messaging application, asks learners to collaborate in order to accomplish one or more learning tasks in an online activity (Figure 1A). Throughout the students' discourse, a conversational agent (Figure 1B) delivers APT-based interventions. The agent uses a text-to-speech (TTS) engine to read its interventions, which are displayed outside the main chat frame (Figure 1C).

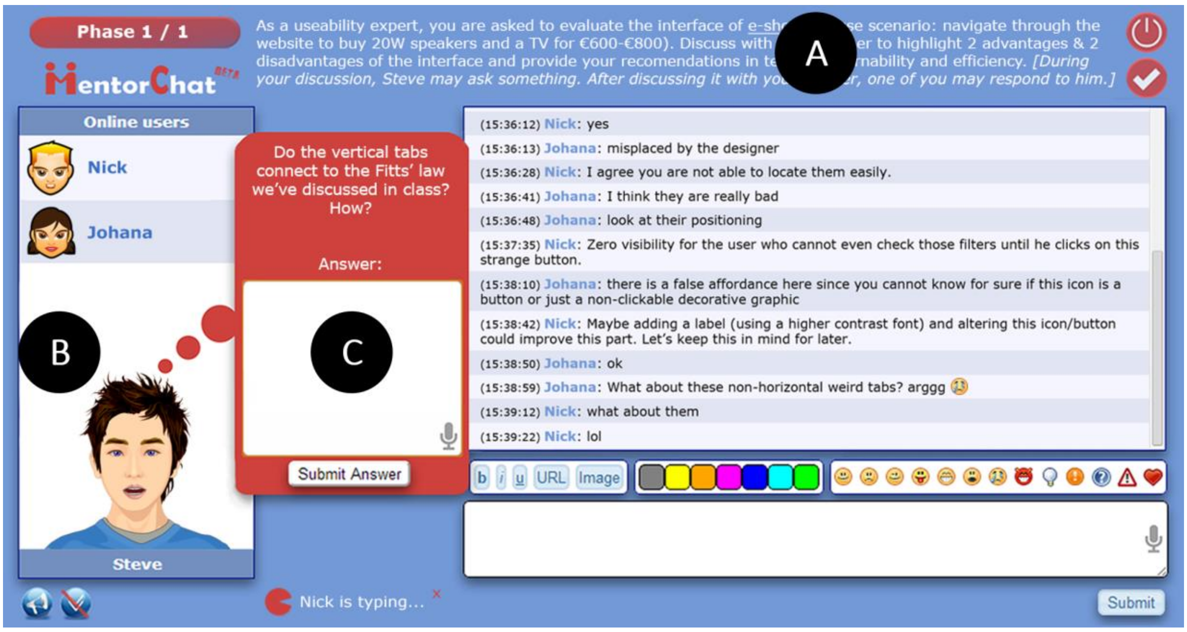

Figure 1. MentorChat learning environment

MentorChat features an administration interface that allows teachers or researchers to register participants and assign them to groups, monitor students' discussions, and create collaborative learning activities. An activity may include multiple dialogue-based tasks (phases), typically asking students in each group to discuss a topic and submit a joint response to an open-ended learning question. The instructor can use the domain configuration panel (Figure 2) to model the agent domain knowledge via a concept mapping interface. A concept map can be created for each task by entering a series of simply structured statements, involving key domain concepts (Figure 2A). Each statement consists of a subject (concept X), an object (concept Y), and a verb or verbal phrase (connection). As illustrated in Figure 2B, these elements are automatically rendered and visualized in a concept map, serving as the domain knowledge representation of the agent. This authoring method facilitates the reusability of teacher-defined domain models by providing access to a library of pre-built concept maps.

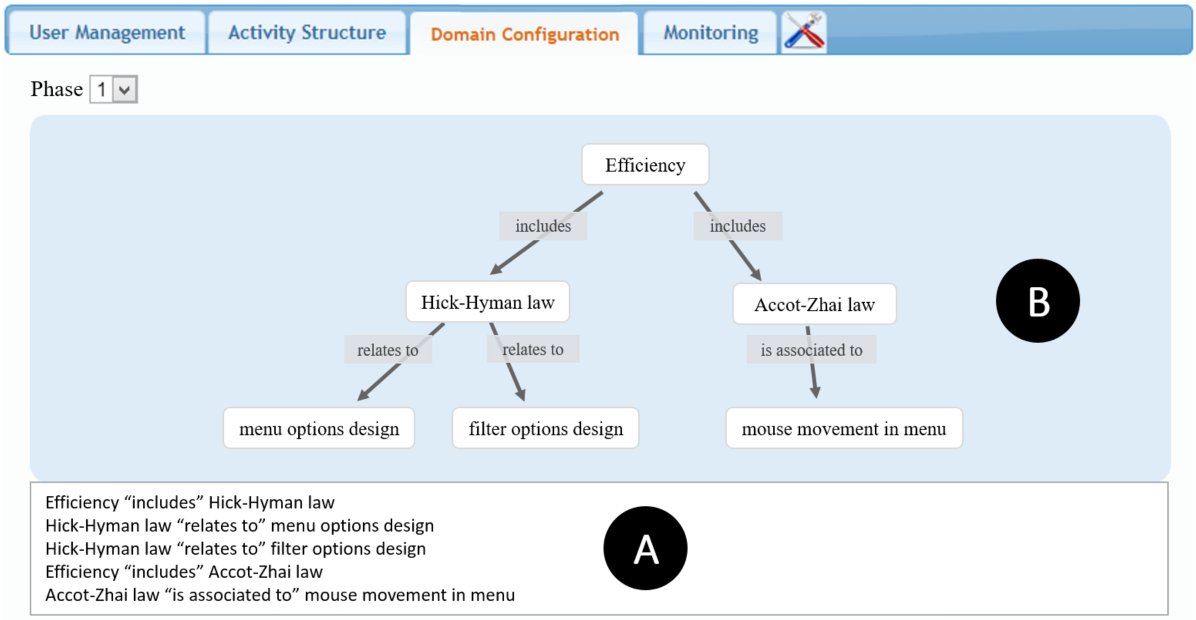

Figure 2. MentorChat domain configuration interface

Since a detailed analysis of all system components is beyond the scope of this paper, we will concisely present the functionality of the conversational agent on the basis of three core models: the peer interaction, the domain, and the intervention models (Figure 3).

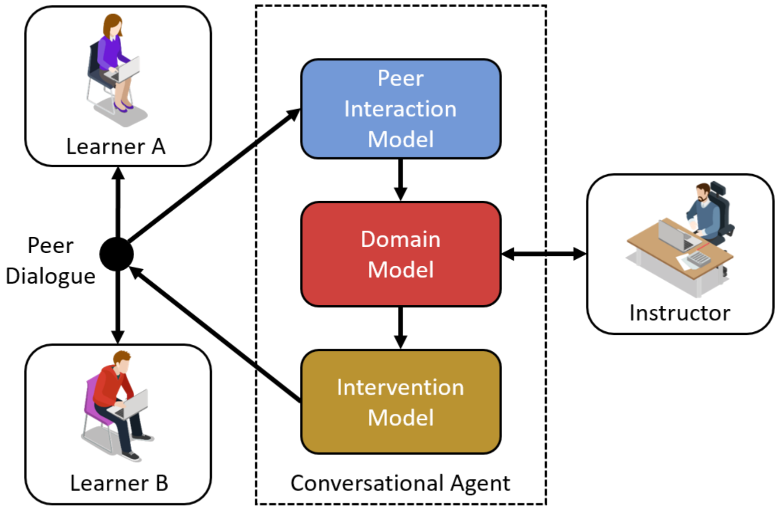

Figure 3. Conversational agent architecture

The peer interaction model processes each dialogue contribution to identify concepts similar to the ones of the agent concept map (Figure 4A). Next, a proximity score is calculated for each concept identified, utilizing a Levenshtein-based string similarity algorithm and a WordNet lexicon for synonyms. If the score exceeds the preset threshold, the associated concept from the agent concept map is added to the concept map of the learner introducing the particular concept in the group discussion (Figure 4B). Thus, each learner's concept map is dynamically updated including new concepts (nodes), as the peer dialogue unfolds. In principle, this allows the agent to direct its interventions to a specific peer (e.g., the least active). In this study, however, the agent was configured to address the group and not any individual student, simultaneously displaying its interventions to both peers.

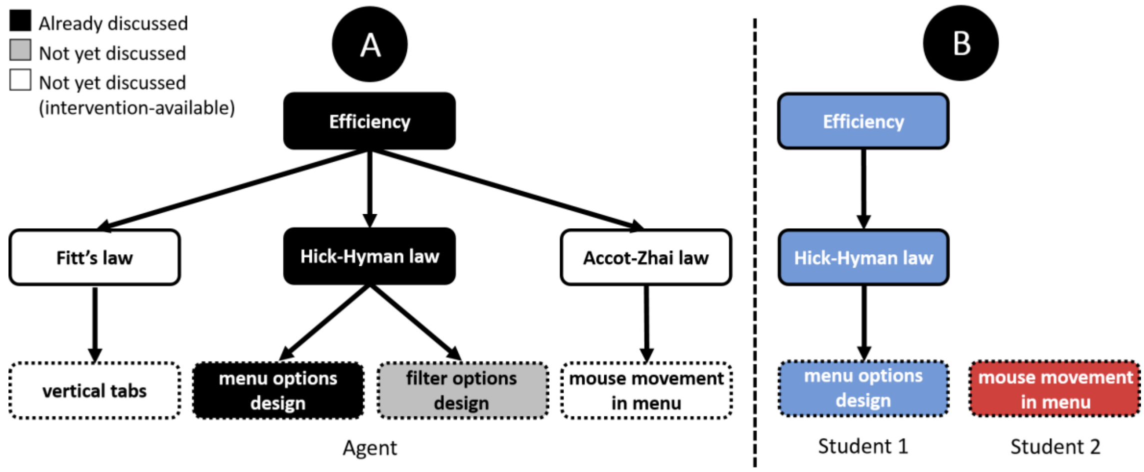

Figure 4. Concept maps formed in MentorChat

The identification of an intervention opportunity is handled by the agent domain model. In the current version of the system, this is accomplished by comparing the nodes of the agent concept map (Figure 4A) to the nodes of the learners' concept maps (Figure 4B). Each time a student makes a contribution, the agent looks for a domain concept that is both marked as "not yet discussed" in the agent map and included in any of the learners' concept maps. If the above criteria are met for a concept Y, the agent proposes an intervention, asking students to logically connect concept Y with an associated higher-level (parent node) domain concept X that is also marked as "not yet discussed" in the agent map. For example, as shown in Table 2, when a peer discusses the low-level domain concept of "mouse movement in menu" for the first time, the agent decides to intervene asking students to link the concept being discussed with a higher-level domain concept introduced by the agent (in this case, the "Accot-Zhai law"). Afterwards, the agent marks both the low- and the higher-level concept nodes as "already discussed" (Figure 4A). This mechanism ensures that no specific agent intervention is activated more than once while the agent is always aware of the domain concepts not yet discussed.

Table 2. A dialogue excerpt including an agent intervention

Name	Message
Mary:	The mouse-over behavior over the menu is fine.
John:	It's good. What else?
Steve (Agent):	Do you think the Accot-Zhai law is somehow associated to mouse movement in menu?
Mary:	John, do you know the answer?
Mary:	Accot's law suggests that the time required for a user to navigate through a multi-level menu depends on the size of its options
John:	True, the arithmetic formula included D/S (D-width and S-height)! Should we reply?
Mary:	*[Submitted Response]* Yes, because the Accot-Zhai law describes the time (a+b*D/S) a user needs to move their pointer through a menu path (tunnel). We believe the specific menu is correctly designed having a sufficient width and height, thus facilitating the user's effort.

Lastly, the agent intervention model manages the synthesis of the intervention text, utilizing the concept map connections defined by the teacher in the domain configuration panel and a pool of key phrases containing system variables. Furthermore, this model examines a series of micro-parameters to approve or cancel the final display of the agent intervention, proposed by the agent domain model. For instance, after examining the number of previously displayed interventions or whether the preset time has passed since the last intervention, the agent may decide to suppress a new intervention to avoid excessive interference from consecutive agent interventions.

Procedure

The instructor of the HCI course set up a MentorChat activity by entering a learning question asking students to collaboratively evaluate the web interface of an e-shop based on the main usability principles (Figure 1A). After their discussion, peers were required to jointly write and submit a report emphasizing two strong and two weak aspects of the interface design, also proposing improvements in terms of efficiency and learnability. Additionally, the instructor set up the agent concept map by entering a series of statements involving key domain concepts, as depicted in Figure 2.

The participants of the study were randomly assigned to dyads and, then, to a control or a treatment group. The control group consisted of 34 students (17 dyads), whereas the treatment group consisted of 38 students (19 dyads). All students were informed that their discussions would be monitored and recorded.

The MentorChat activity took place in the university computer labs, with 40 minutes allotted for all groups. Prior to the activity students filled in a 20-minute pre-test questionnaire. Immediately after the activity, students were required to answer a post-test questionnaire as well as a student opinion questionnaire, lasting 20 and 10 minutes respectively. One week later, the treatment group students participated in a focus group session.

Experimental design

A randomized pre-test post-test design was adopted to assess the impact of the agent intervention mode used in this study. Two conditions were compared: (a) students who received agent interventions while collaborating in dyads to accomplish a learning task (treatment) and (b) students collaborating in dyads to accomplish the same task without any agent intervention (control).

On the basis of the peer interaction, domain, and intervention models described in the above section, the agent employed in the treatment condition delivered dynamic interventions that promoted accountability to accurate knowledge (Table 2, row 3). According to the activity guidelines, peers were expected to respond to the agent in a coordinated way (one of them) using the agent answer box (Figure 1C), which remained visible until an answer had been submitted by a student. The aim of the agent interventions was to encourage students to build on their prior knowledge, linking their current contributions to the main theoretical aspects and usability principles discussed in class.

In contrast to the treatment students, their control counterparts discussed without the presence of the conversational agent and did not receive any APT-based interventions. In both conditions, the system displayed a few static prompts that either supported group awareness (e.g., “John has logged into the activity”) or provided simple instructions on some interface features (e.g., “You can submit your answer by clicking…”).

Data collection and analysis

The pre-test of the study consisted of two parts. The first part included 10 multiple-choice questions requiring students to recall previously learned basic information; thus, it was relevant to Bloom’s taxonomy first level (Krathwohl, 2002). The second part included four open-ended questions asking students to demonstrate their understanding (relevant to Bloom’s taxonomy second level). Answering each of the two questionnaire parts could give students up to 10 points. Students’ answer sheets were mixed and scored by two independent raters with extensive experience in the HCI domain. Holistic rubric scales were used for the assessment of the open-ended questions. The intra-class correlation coefficient indicated a high inter-rater reliability (ICC = .98). Additionally, the internal consistency coefficient of the total pre-test scores was found to be satisfactory (C_α = .85). Since the normality and homogeneity of variance criteria were satisfied, independent samples t-tests were conducted on the pre-test scores of the control and treatment groups to compare students’ knowledge levels prior to the activity.

Individual learning

In order to measure individual learning, we conducted a post-test consisting of six open-ended questions. These asked students to demonstrate their comprehension of the main usability principles (second level of Bloom’s taxonomy). The sum of questionnaire items scores could give a student up to 20 points. Post-test questionnaires were assessed by the same raters as in the pre-test. Their intra-class correlation coefficient was found to be high (ICC = .98). An analysis of covariate (ANCOVA) was performed on post-test scores to explore the difference in individual learning outcomes in the two conditions.

Group learning

In order to measure group learning, the raters assessed the quality of the dyads' answers submitted at the end of the activity to the learning question (Figure 1A). The group answers were scored using a rubric rating scale while a high inter-rater reliability was obtained (*ICC* = .97). An independent samples *t*-test was conducted to compare the group learning outcomes under the control and treatment conditions.

Explicitness

In order to measure the explicitness of students' reasoning, a discourse analysis was performed by the authors. The analysis was based on an adjusted version of the IBIS discussion model widely considered to be an effective framework for analyzing the conversational interactions occurring in remote collaboration sessions (Liu & Tsai, 2008). Drawing on the main categories of the IBIS model (issue, position, and argument), our scheme introduced two additional categories, named explicit position and explicit argument, both focusing on the display of students' reasoning on domain concepts. Table 3 presents the final scheme used for coding students' contributions, each of which could consist of multiple chat posts. Mann Whitney U-tests evaluated the difference in explicitness between the conditions by comparing the explicit position and explicit argument frequencies of the control and treatment dyads.

Table 3. Discourse analysis scheme

Category	Description
Off-task	Utterances playing a purely social function or not relating to the task (e.g., "Greetings")
Repetition	Reiterations of prior contributions
Management	Management-oriented statements used for task coordination (e.g., "Let's submit our answer")
Common Understanding	Short (typically one- or two-word) utterances establishing common understanding on the subject (e.g., "OK," "I see")
Issue	What needs to be done or resolved to proceed with the overall task (e.g., "What other laws are relevant?")
Position	Opinions usually related to the resolution of the issue raised (e.g., "Fitts' law applies here")
Argument	Opinions supporting or objecting to a position (e.g., "You are absolutely right")
Explicit Position	Positions explicitly displaying reasoning on domain concepts (e.g., "According to Hick-Hyman, the reaction time increases logarithmically as the number of options increases")
Explicit Argument	Much as explicit positions, arguments displaying explicit reasoning on domain concepts (e.g., "I disagree, Hick's law cannot be used for randomly ordered lists requiring linear time")

Explicit response ratio

In order to gain a deeper insight into how the intervention mode affected the students' explicitness, an "explicit response ratio" (ERR) was calculated by dividing the number of agent-induced explicit contributions (explicit positions and explicit arguments) by the number of agent interventions. An explicit position or explicit argument was labelled as "agent-induced" only if it was a direct response to the agent or a subsequent comment relating to the agent intervention (Table 2, rows 5-7). Thus, the ERR value of a dyad could range from zero, if the agent interventions had no effect on the generation of explicit contributions, to more than one, when multiple explicit contributions were triggered by each agent intervention.

Following an exploratory content analysis drawing inferences on peers' behavior after an agent intervention, a mediation analysis was implemented. The analysis investigated whether the generation of explicit positions and explicit arguments in students' dialogue could account for a potential association between the conditions and the learning outcomes. In general, a variable is regarded as a mediator if it carries the influence of a given independent variable to a given dependent variable (Rucker, Preacher, Tormala, & Petty, 2011). In our study, the intervention mode (activated in the treatment condition and deactivated in the control condition) was the independent variable, while individual learning (as measured by the post-test) was the dependent variable. Hence, our analysis tested the mediating effect of explicitness (as measured by students' explicit position and explicit argument frequencies) on the relationship between the agent intervention mode and the learning outcomes. Using the PROCESS SPSS macro, the test was performed by applying bootstrapping with biascorrected confidence estimates (Hayes & Preacher, 2014). The 95% confidence interval of the indirect effects was obtained with 5000 bootstrap resamples.

Students' opinions

The student opinion questionnaire included Likert-scale questions ranging from 1 (disagree) to 5 (agree). The first questionnaire part elicited students' opinions regarding their learning experience in MentorChat (C_α = .86). A Mann Whitney U-test was conducted for each questionnaire variable to assess the difference between the scores reported in the two conditions. The second questionnaire part consisted of four pairs of positively and negatively keyed questionnaire items, which were only available for the treatment students and measured their perceptions of the agent interventions. After reverse scoring all negatively keyed items, we calculated a mean score for each question pair (C_α > .70 for all pairs).

The focus group session followed a semi-structured protocol and explored the opinions of the treatment students on the display of the interventions and their interaction with the conversational agent. Students' responses were transcribed verbatim and analyzed in search of common themes following an open coding process.

Results

In respect to the closed-type pre-test questions, the independent-samples t-test yielded no statistically significant differences as the treatment group mean (n = 38, M = 5.89, SD = 1.59) was comparable to that of the control group (n = 34, M = 5.97, SD = 1.96), t(70) = .181, p = .857, d = .043. Likewise, the treatment students (n = 38, M = 4.49, SD = 2.28) performed similarly to control students (n = 34, M = 4.35, SD = 2.16) in the open-ended questions, t(70) = .260, p = .795, d = .062.

Individual learning

The results of the ANCOVA conducted on post-test scores, with total pre-test scores as a covariate, revealed a significant difference in favor of the treatment group. The analysis showed that the students in the treatment condition (n = 38, M = 12.12, SD = 4.41) outperformed their control counterparts (n = 34, M = 9.45, SD = 3.76) on individual learning, F(1,69) = 9.162, p = .003, η_p^2 = .117.

Group learning

The independent-samples t-test evaluating group learning indicated that the dyads in the treatment condition achieved considerably better scores (n = 19, M = 14.03, SD = 3.46) than the dyads in the control condition (n = 17, M = 10.56, SD = 3.93), t(34) = 2.814, p = .008, d = .965.

Explicitness

The discourse analysis identified 2646 students' contributions. From those 1543 came from the treatment group while 1103 came from the control group. The overall results of the coding process are summarized in Table 4.

Table 4. Discourse analysis results

Category	Control (n = 17 dyads)			Treatment (n = 19 dyads)		
	Total	Frequency (%)	Group average	Total	Frequency (%)	Group average
Off-task	85	7.71	5.00	110	7.13	5.79
Repetition	28	2.54	1.65	31	2.01	1.63
Management	187	16.95	11.00	329	21.32	17.32
Common Understanding	160	14.51	9.41	204	13.22	10.74
Issue	93	8.43	5.47	126	8.17	6.63
Position	202	18.31	11.88	189	12.25	9.95
Argument	201	18.22	11.82	249	16.14	13.11

Explicit Position	98	8.88	5.76	192	12.44	10.11
Explicit Argument	49	4.44	2.88	113	7.32	5.95

A Mann Whitney U-test revealed that the treatment groups (*n* = 19, *M* = 13.42, *SD* = 4.55) generated substantially more explicit positions than the control groups (*n* = 17, *M* = 9.22, *SD* = 5.22), *U* = 243.5, *p* = .008, *r* = .433. The treatment groups produced also considerably more explicit arguments (*n* = 19, *M* = 7.48, *SD* = 3.32) compared to the control groups (*n* = 17, *M* = 4.32, *SD* = 2.26), *U* = 258, *p* = .002, *r* = .510.

Explicit response ratio

The average explicit response ratio (ERR) is presented in the last row of Table 5, which depicts the results of the second-level analysis focusing on the agent effect on the students' explicitness.

Table 5. Intervention impact on explicitness

Category	Treatment (*n* = 19 dyads)
Agent Interventions	*M* = 2.74, *SD* = 0.73
Agent-Induced Explicit Positions	*M* = 4.21, *SD* = 1.96
Agent-Induced Explicit Arguments	*M* = 3.05, *SD* = 2.22
Explicit Response Ratio (ERR)	*M* = 2.69, *SD* = 1.19

A preliminary content analysis revealed that in 16 out of the 19 dyads (*F* = 84.21%) the majority of the interventions were addressed by the most active group member, the one with the highest rate of explicit contributions. Although students' responded to the agent using clear statements, a lack of peer coordination was identified following the onscreen display of some interventions (*F* = 15.38%), when one student responded to the agent before communicating with their partner. Additionally, it was noted that only 2 agent interventions (*F* = 3.85%) were ignored by the students. Both interventions were displayed towards the end of the activity, when the students had already begun composing their final group answer.

A regression analysis investigated whether the generation of explicit positions and explicit arguments mediated the effect of the agent intervention mode (treatment condition) on the post-test scores. The results showed that the agent intervention mode was a significant predictor of both students' explicitness (*B* = 3.659, *t*(70) = 3.197, *p* = .002) and individual learning (*B* = 2.668, *t*(70) = 2.744, *p* = .008), while explicitness was a significant predictor of individual learning (*B* = 0.291, *t*(70) = 3.028, *p* = .004). The mediation analysis also resulted in a statistically significant indirect coefficient (*B* = 1.07; *CI* = .280 to 2.420), thus supporting the mediating role of explicitness (Figure 5). In fact, a full mediation was suggested since the agent intervention mode was no longer a significant predictor of individual learning outcomes after controlling for the mediator (*B* = 1.603, *t*(70) = 1.628, *p* = .108), i.e., the frequency of explicit positions and arguments (Rucker et al., 2011).

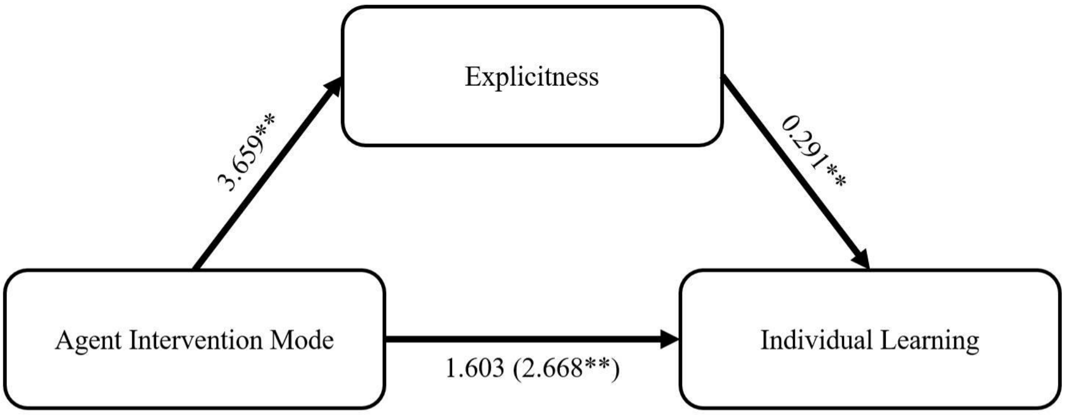

Figure 5. Mediation diagram (*Note.* $^{**}p$ < .01)

Students' opinions

Table 6 depicts the results of two questionnaire items assessing the subjectively perceived benefits of the collaborative activity. Other questions on the interface usability revealed no major issues regarding the ease of use (n = 72, M = 4.07, SD = .70) or the performance (n = 72, M = 4.36, SD = 0.66) of MentorChat in either condition.

Table 6. Student opinion questionnaire results on the collaborative activity

Questionnaire item	Control (n = 34)	Treatment (n = 38)	Mann-Whitney U-test
1. The collaborative activity improved my domain knowledge	M = 3.44, SD = 1.16	M = 3.97, SD = 0.68	U = 825, p = .033*, r = .251
2. The collaborative activity was beneficial for me (regardless of any improvement in my domain knowledge)	M = 4.00, SD = 1.06	M = 4.26, SD = 0.65	U = 711, p = .429, r = .093

Note. $^{*}p < .05$.

In respect to the students' opinions about the conversational agent interventions, Table 7 presents the total scale scores on the four questionnaire measures.

Table 7. Student opinion questionnaire results about the agent intervention mode

Questionnaire item	Treatment (n = 38)
1. The agent questions displayed during the discussion were simple and understandable	M = 4.53, SD =0.60
2. The agent questions helped me recall or retrieve useful domain information for the evaluation of the e-shop interface	M = 4.03, SD = 0.88
3. The agent questions did not disrupt my discussion with my partner	M = 3.95, SD = 0.96
4. The timing and the content of the agent questions were consistent with the on-going discussion	M = 4.08, SD = 0.78

The analysis of the focus group data resulted in the identification of the following themes: (a) the agent interventions appeared to have helped many students to focus their discussion on important usability principles (F = 79%), (b) some students appreciated the fact that the agent did not intervene very frequently and allowed time for their response (F = 39%), and (c) the background color of the agent answer box was perceived as disruptive by a few participants (F = 16%).

Discussion

The assessment of students' post-test answers revealed that the students of the treatment group achieved higher learning outcomes as compared to the students of the control group. Even though the knowledge levels of the students did not differ significantly prior to the activity, the treatment students outperformed their control counterparts in understanding and illustrating conceptual domain knowledge following the activity. Furthermore, when asked whether the collaborative activity improved their domain knowledge, the treatment students were more positive than the control ones (Table 6, item 1). The agent effect on individual learning appears to accord with Adamson et al.'s (2013) results, which showed that the introduction of an agent displaying Agree/Disagree interventions (Table 1, move 1A) could improve learning and intensify the exchange of ideas in higher education.

As for the group learning, it was shown that the final answers of the treatment dyads were superior to those of the control dyads in terms of correctness and comprehensiveness. Indeed, the answers submitted in the treatment condition appeared to be more conceptually complete, including several appropriate references to theoretical principles of usability. The finding is consistent with our previous findings demonstrating that APT-based agent interventions can effectively trigger group discussions to elaborate on more domain concepts, thus leading to more comprehensive group answers (Tegos et al., 2015).

The above evidence suggests that agent interventions had a positive effect on both individual and group learning. This is also supported by students' views expressed in the student opinion questionnaire and the focus group session. Many students commented favorably on the conversational agent interventions, which they believed to have assisted them in recalling and focusing their discussions on central domain principles (e.g., Table 7, item 2). Despite the relatively simple intervention mechanism, we argue that agent interventions effectively scaffolded students' discussions and kept students focused by asking them to support their contributions using fundamental conceptual knowledge, known to be critical to understanding and resolving learning tasks (Streveler, Litzinger, Miller, & Steif, 2008).

The discourse analysis yielded a series of noteworthy results. First, the agent intervention mode was shown to considerably impact on the explicitness of students' reasoning. In particular, the frequency of explicit positions and explicit arguments was found to be significantly higher in the treatment group than the control group (Table 4, rows 8 and 9). The explicit reasoning enhancement seems to be owed to the agent interventions, since the number of explicit contributions identified as agent-induced (Table 5, rows 2 and 3) can justify the difference in the generation of explicit positions and explicit arguments between the conditions (Table 4, rows 8 and 9).

The explicit-response ratio (ERR), measuring the effect of the intervention mode on explicitness, was found to be higher in comparison to the results of our previous studies (e.g., Tegos et al., 2014; Tegos et al., 2015). Although we cannot directly compare the ERR values of different studies due to their dissimilar contexts, the average ERR value reported in this study (Table 5, row 4) suggests that each agent intervention stimulated more than two students' explicit contributions. We argue that agent interventions promoted students' sound reasoning by urging for clear and compelling statements supported by concrete evidence. Students were encouraged to seek out appropriate evidence in order to strengthen their claims with specific domain principles and explicate their reasoning, thus making it available for others (or themselves) to assess, question or challenge (Michaels et al., 2010). This lies in the core of transactivity theory, which highlights the value of explicit reasoning displays, so that the reasoning can be communicated to the partner(s) and then built upon by the latter (Stahl & Rosé, 2011). This kind of explicitness is considered to be a matter of importance in written discussions, where learners may not explicate their thinking to the group, therefore obstructing the negotiation of common ground (Weinberger, Stegmann, & Fischer, 2007).

As illustrated by the path diagram in Figure 5, the influence of the agent interventions on students' explicitness played a significant mediating role. A mediation analysis showed that the impact of the agent intervention mode on individual learning was mediated by the stimulation of explicit reasoning during students' discussions. This partially confirms a previous finding showing that the effectiveness of an APT-based intervention mode depends on its ability to engage students in transactive conversation interactions such as explicit argumentation, where students communicate their reasoning and build on each other's ideas (Tegos et al., 2015). Moreover, there is some converging evidence suggesting that the frequency of students' contributions containing a form of explicit reasoning can act as a robust predictor of students' learning in dialogue-based collaborative learning activities.

As opposed to the directed intervention mode used in our previous work (Tegos et al., 2015), which addressed a specific student of the group, the agent of this study allowed the peers to decide who would respond to its intervention. While the activity guidelines clearly stated that peers should discuss and coordinate their answers to the agent, we observed that the most active group members, those having the highest explicit position and explicit argument frequencies in their dyads, often took the initiative to respond to the agent before communicating with their partners. Based on prior findings on the effectiveness of differently presented intervention techniques (Tegos et al., 2014), we believe this miscommunication between peers to have resulted from the reduced situational constraints, i.e., not requiring students to follow a specific protocol in their interaction with the agent. Although further investigation is needed to draw definite conclusions about how group conversations were affected, we consider that the undirected nature of the agent interventions shifted the overall balance of students' discussions in the treatment condition by allowing the most active peers to regulate the student-agent interaction by themselves.

The relatively low number of interventions displayed in the treatment condition (Table 5, row 1) was commented favorably in the focus group discussion, as many students appreciated the non-frequent intervention mechanism of the agent, which did not appear to have a major interruption effect based on our observations and students' opinions (Table 7, item 3). Additionally, our content analysis revealed that the agent effectively drew students' attention since only two interventions were ignored by the students, who generally perceived the agent interventions as quite comprehensible and consistent with the ongoing discussion topic (Table 7, items 1 and 4).

Yet, in interpreting the above findings one should consider the study limitations as well. First, the participants were aware of their discussions being monitored and recorded, which could have altered the typical conversational behavior of students, as they may have paid more attention to the agent interventions than they would have in a non-controlled environment. Second, due to our limited sample size, study findings need to be confirmed by a larger sample size and other student populations of different backgrounds or ages. Lastly, it is clear that the agent does not possess intelligence capable of engaging into a full-fledged discussion with students. Still, this is in line with our objective of building easily deployed agents that can accomplish substantial learning benefits with the minimum required level of intelligence.

Conclusion

The study provides promising evidence on the utilization of a teacher-configurable agent intervention mode that encourages peers to build on their prior knowledge, drawing on the academically productive talk (APT) discourse framework. Study findings indicate that agent interventions aiming to link students' contributions to previously acquired knowledge can improve both individual and group learning when implemented in the context of a collaborative learning activity in higher education. The agent interventions have also a positive impact on students' conversational behavior by amplifying explicit reasoning. In fact, the level of explicit reasoning appears to mediate the relationship between the agent interventions and individual learning. Overall, we believe the results of this study to contribute to the understanding of how the facilitation of collaborating groups can be effectively automated through agile agent interventions. In the future, we aspire to utilize our configurable conversational agent system to explore the impact of other APT facilitation strategies in various educational domains and levels.

References

Adamson, D., Ashe, C., Jang, H., Yaron, D., & Rosé, C. P, (2013). Intensification of group knowledge exchange with academically productive talk agents. In N. Rummel, M. Kapur, M. Nathan, & S. Puntambekar (Eds.), *To See the World and a Grain of Sand: Learning across Levels of Space, Time, and Scale: CSCL 2013 Conference Proceedings* (Vol. 1, pp. 10-17). Madison, MI: International Society of the Learning Sciences (ISLS).

Adamson, D., Dyke, G., Jang, H., & 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.

Baylor, A. L., & Kim, Y. (2005). Simulating instructional roles through pedagogical agents. *International Journal of Artificial Intelligence in Education, 15*(1), 95-115.

Chaudhuri, S., Kumar, R., Joshi, M., Terrell, E., Higgs, F., Aleven, V., & Rosé, C. P. (2008). It's not easy being green: Supporting collaborative "green design" learning. In *Proceedings of the 9th International Conference on Intelligent Tutoring Systems* (pp. 807-809). Berlin, Germany: Springer.

Dillenbourg, P., & Tchounikine, P. (2007). Flexibility in macro-scripts for computer-supported collaborative learning. *Journal of Computer Assisted Learning, 23*(1), 1-13.

Dyke, G., Adamson, D., Howley, I., & Rosé, C. P. (2013). Enhancing scientific reasoning and explanation skills with conversational agents. *IEEE Transactions on Learning Technologies, 6*(3), 240-247.

Hayes, A. F., & Preacher, K. J. (2014). Statistical mediation analysis with a multicategorical independent variable. *British Journal of Mathematical and Statistical Psychology, 67*(3), 451-470.

Kerly, A., Ellis, R., & Bull, S. (2009). Conversational agents in e-learning. In T. Allen, R. Ellis, & M. Petridis (Eds.), *Applications and Innovations in Intelligent Systems XVI* (pp. 169-182). London, UK: Springer.

Krathwohl, D. R. (2002). A Revision of Bloom's taxonomy: An Overview. *Theory into practice, 41*(4), 212-218.

Kreijns, K., Kirschner, P. A., & Jochems, W. (2002). The Sociability of computer-supported collaborative learning environment. *Journal of Educational Technology & Society, 5*(1), 8-22.

Kumar, R., & Rosé, C. P. (2011). Architecture for building conversational agents that support collaborative learning. *IEEE Transactions on Learning Technologies, 4(*1), 21-34.

Kumar, R., Beuth, J., & Rosé, C. P. (2011). Conversational strategies that support idea generation productivity in groups. In H. Spada, G. Stahl, N. Miyake, & N. Law (Eds.), *Connecting Computer-Supported Collaborative Learning to Policy and Practice: CSCL2011 Conference Proceedings* (Vol. 1, pp. 398-405). Hong Kong, China: International Society of the Learning Sciences (ISLS).

Liu, C. C., & Tsai, C. C. (2008). An Analysis of peer interaction patterns as discoursed by on-line small group problem-solving activity. *Computers & Education, 50*(3), 627-639.

Michaels, S., O'Connor, C., & Resnick, L. B. (2008). Deliberative discourse idealized and realized: Accountable talk in the classroom and in civic life. *Studies in Philosophy and Education, 27*(4), 283-297.

Michaels, S., O'Connor, M. C., Hall, M. W., & Resnick L. B. (2010). *Accountable talk sourcebook: For Classroom that works*. Pittsburgh, PA: University of Pittsburgh Institute for Learning. Retrieved from http://ifl.pitt.edu/index.php/download/index/ats

Papadopoulos, P. M., Demetriadis, S., & Weinberger, A. (2013). "Make it explicit!": Improving collaboration through increase of script coercion. *Journal of Computer Assisted Learning, 29*(4), 383-398.

Preece, J., Sharp, H., & Rogers, Y. (2015). *Interaction design: Beyond Human-Computer Interaction*. West Sussex, UK: John Wiley & Sons.

Rucker, D. D., Preacher, K. J., Tormala, Z. L., & Petty, R. E. (2011). Mediation analysis in social psychology: Current practices and new recommendations. *Social and Personality Psychology Compass, 5*(6), 359-371.

Sohmer, R., Michaels, S., O'Connor, M. C., & Resnick, L. (2009). Guided construction of knowledge in the classroom. In B. Schwarz, T. Dreyfus, & R. Hershkowitz (Eds.), *Transformation of knowledge through classroom interaction* (pp. 105-129). New York, NY: Taylor & Francis.

Stahl, G., & Rosé, C. P. (2011). Group cognition in online teams. In E. Salas, & S. M. Fiore (Eds.), *Theories of team cognition: Cross-disciplinary perspectives*. New York, NY: Taylor & Francis.

Streveler, R. A., Litzinger, T. A., Miller, R. L., & Steif, P. S. (2008). Learning conceptual knowledge in the engineering sciences: Overview and future research directions. *Journal of Engineering Education*, *97*(3), 279-294.

Tchounikine, P., Rummel, N., & McLaren, B. (2010). Computer supported collaborative learning and intelligent tutoring systems. In R. Nkambou, R. Mizoguchi, & J. Bourdeau (Eds.), *Advances in intelligent tutoring systems* (pp. 447-463). Berlin, Germany: Springer.

Tegos, S., Demetriadis, S., & Karakostas, A. (2015). Promoting academically productive talk with conversational agent interventions in collaborative learning settings. *Computers & Education*, *87*, 309-325.

Tegos, S., Demetriadis, S., & Tsiatsos, T. (2014). A Configurable conversational agent to trigger students' productive dialogue: A Pilot study in the CALL domain. *International Journal of Artificial Intelligence in Education*, *24*(1), 62-91.

Veletsianos, G., & Russell, G. S. (2014). Pedagogical agents. In M. Spector, D. Merrill, J. Elen & M. J. Bishop (Eds.), *Handbook of research on educational communications and technology* (pp. 759-769). New York, NY: Springer.

Walker, E., Rummel, N., & Koedinger, K. R. (2011). Designing automated adaptive support to improve student helping behaviors in a peer tutoring activity. *International Journal of Computer-Supported Collaborative Learning, 6*(2), 279-306.

Weinberger, A., Stegmann, K., & Fischer, F. (2007). Knowledge convergence in collaborative learning: Concepts and assessment. *Learning and Instruction, 17*(4), 416-426.