



HAL
open science

Let's Set Up Some Subgoals: Understanding Human-Pedagogical Agent Collaborations and Their Implications for Learning and Prompt and Feedback Compliance

Jason M Harley, Michelle Taub, Roger Azevedo, François Bouchet

► To cite this version:

Jason M Harley, Michelle Taub, Roger Azevedo, François Bouchet. Let's Set Up Some Subgoals: Understanding Human-Pedagogical Agent Collaborations and Their Implications for Learning and Prompt and Feedback Compliance. *IEEE Transactions on Learning Technologies*, 2018, 11 (1), pp.54-66. 10.1109/tlt.2017.2756629 . hal-02015700

HAL Id: hal-02015700

<https://hal.science/hal-02015700>

Submitted on 26 Aug 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

“Let’s Set Up Some Subgoals”: Understanding Human-Pedagogical Agent Collaborations and Their Implications for Learning and Prompt and Feedback Compliance

Jason M. Harley¹, Michelle Taub², Roger Azevedo² & François Bouchet³

¹ Department of Educational Psychology at the University of Alberta, Edmonton, Alberta,
6-102 Education North, T6G 2G5, Canada

² Department of Psychology at North Carolina State University,
2310 Stinson Drive, Poe Hall 640, Raleigh, NC, 2765-7650, USA

³ Sorbonne Universités, UPMC Univ Paris 06, and CNRS, LIP6 UMR 7606,
4 place Jussieu, BC 169, 75005 Paris, France

Abstract— Research on collaborative learning between humans and virtual pedagogical agents represents a necessary extension to recent research on the conceptual, theoretical, methodological, analytical, and educational issues behind co- and socially-shared regulated learning between humans. This study presents a novel coding framework that was developed and used to describe collaborations between learners and a pedagogical agent (PA) during a subgoal setting activity with MetaTutor, an intelligent tutoring system. Learner-PA interactions were examined across two scaffolding conditions: prompt and feedback (PF), and control. Learners’ compliance to follow the PA’s prompts and feedback in the PF condition were also examined. Results demonstrated that learners followed the PA’s prompts and feedback to help them set more appropriate subgoals for their learning session the majority of the time. Descriptive statistics revealed that when subgoals were set collaboratively between learners and the PA, they generally lead to higher proportional learning gains when compared to less collaboratively set goals. Taken together, the results provide preliminary evidence that learners are both willing to engage in and benefit from collaborative interactions with PAs when immediate, directional feedback and the opportunity to try again are provided. Implications and future directions for extending co- and socially-shared regulated learning theories to include learner-PA interactions are proposed.

Index Terms—intelligent tutoring systems, pedagogical agents, collaborative learning, adaptive hypermedia, self-regulated learning, co-regulated learning, socially-shared regulated learning, externally-regulated learning

1 Introduction

Examining the process of self-regulated learning (SRL), an adaptive learning behavior empirically linked to higher learning outcomes (see [1]), has been a primary focus of research on learning with computer-based learning environments (CBLEs). Researchers have paid relatively little attention, however, to examining the roles that tutors, human [2-3] or virtual (i.e., pedagogical agents; PAs) stand to play in helping learners to develop their SRL skills during learning with CBLEs. Specifically, roles that are examined are constrained by limited theoretical lenses that fail to consider more bi-directional and collaborative dynamics between learner and tutor, especially when the tutor is a PA, which this study focuses on.

Indeed, research on PAs has mostly focused on their ability to externally regulate SRL [4-9], while neglecting to examine collaborative regulatory processes such as co- and socially-shared regulated learning [10-12]. This study contributes to addressing this gap in the literature by expanding the theoretical lens used to understand learner-PA interactions to include these recent advances in educational theory. In doing so, ongoing theoretical and analytical challenges are also addressed, such as clarifying the boundaries between different types of regulatory behaviors (described below), their relationship with learning, and learners’ responses, in terms of compliance, to tutorial prompts and feedback.

Specifically, this study addresses these challenges and limitations through (1) the development of a *theoretically-situated qualitative coding scheme* used to describe different learner-PA interaction patterns, and (2) an investigation into potential *relationships between learner-PA interaction patterns*, including prompt and feedback compliance, and *subsequent learning outcomes*. Prior to discussing the methodology, we present the theoretical framework used to

contextualize learner-PA interactions, followed by an explanation of the learner-PA interactions examined in this study before presenting the research objectives, questions, and hypotheses.

2 Theoretical Framework

There are a number of terms to describe variants of self- and other-regulated learning, including externally regulated learning (ERL) [13-14], socially-shared regulated learning (SSRL) [15-16], and co-regulated learning (CoRL) [17]. This study adopts the theoretical framework and assumptions for self and other-regulated learning advanced by Järvelä, Hadwin, and Miller [10-12], that delineate different types of regulated learning and how they can be applied to group learning. We begin our discussion of regulated learning with self-regulated learning (SRL).

2.1 Self-regulated learning (SRL)

Self-regulated learning is a recursive, loosely sequenced, four-phase process of cognitive and behavioral activity (information processing model) [18-21]. Although the information processing model does not focus on social elements of regulation, it provides a detailed description of the discrete elements involved in SRL, including goal-setting, which is of particular importance to this study.

The four phases of Winne and Hadwin's [18-20] information processing model include: (1) task definition, (2) goal-setting and planning, (3) studying tactics, and (4) adaptations to metacognition. In the first phase, learners form an understanding of the learning task that includes assessing the conditions associated with the task. Conditions are affordances or constraints that can be external (e.g., resources, instructional cues, time, social context) or internal (e.g., beliefs, motivation, domain knowledge, knowledge of task, knowledge of study tactics and strategies) to the learner. In the second phase, learners set goals (e.g., regarding what they want to learn) and outline a means of achieving them. By setting multifaceted goals, learners are establishing standards for their learning. Goal setting is the focus of this study. In the third phase learners enact studying tactics (i.e., operations), such as creating summaries, making inferences, and coordinating informational sources to help them achieve their goals (i.e., meet their standards). The last phase of this model involves learners evaluating their learning products (e.g., learner-generated summaries and models, mental representations [e.g., mental models] of instructional content). Evaluating their learning involves comparing their learning products to their standards. If standards are not met, learners exercise control through modifying operations, conditions, standards, or all three.

These phases are loosely sequenced, such that one can modify plans while enacting studying tactics, for example. Effective self-regulators will identify and make necessary changes (i.e., exert metacognitive control) to their studying tactics throughout their learning session, although research has shown that learners are not particularly adept at doing so [22-23].

It should be noted that while SRL is individual, effective self-regulation in a group or team setting (i.e., self-regulation of collaborative learning [12]) requires individuals to regulate their own cognitive processes, behaviors, and beliefs while collaborating with other group members [11-12, 24]. Although not the only important component of regulation in a group setting, Miller and Hadwin [12] propose that SRL is a necessary precondition to group regulation.

2.2 Externally-regulated Learning (ERL)

Interpersonal regulation is an increasingly recognized component of regulated learning, where pairs or groups (not just individuals) engage in intentional, goal-directed metacognitive activity, and advance, constrain, and advocate strategic control over group members' behaviors, cognitions, and beliefs [10,12,25]. Externally-regulated learning (ERL) differs from SRL in that it involves a human or virtual PA tutor prompting an individual learner to deploy key SRL processes during their learning, which may in turn enhance their SRL [2,14]. The CBLE used in this study, MetaTutor, was designed to model and support SRL using externally-regulated learning (ERL) based on the assumptions, sequencing, and empirical findings of the information processing model [26-28] (see 5.2). It is based on the hypothesis that ERL can help learners become more effective at SRL when they are learning on their own after interacting with MetaTutor. In this study, we examine whether one component of learners' interaction with MetaTutor is more collaborative in nature than traditional ERL by extending theoretical considerations to include the ones below.

2.3 Co-regulated learning (CoRL)

Co- and socially-shared regulated learning are based on extensive theoretical and empirical work by Hadwin, Järvelä, and Miller, and outline different inter-individual, collaborative approaches to regulating learning that differ from the unidirectional regulation of learning that is characteristic of ERL [10-12]. Co-regulated learning involves individuals supporting and influencing one another's regulation of learning, typically in an interdependent and reciprocal manner [10-11]. As such, co-regulation requires learners to be aware of one another's individual goals, behaviors, and progress toward achieving collective goals. Järvelä and Hadwin [11] go on to note that although the distribution of regulatory tasks is made between members, CoRL is characterized by learners mutually supporting each other's regulatory behaviors in order to support the pursuit and completion of collective goals that may have complimentary individual components. Miller and Hadwin [12] note that CoRL can take the form of one or more individuals prompting and scaffolding an individual group member's SRL, or one or more individuals prompting and scaffolding multiple group members'

individual SRL. CoRL, therefore, differs from SRL in that group members take responsibility for being aware of one another's learning and learning progress, and will temporarily intervene to scaffold each other if they detect that a partner or group members are deviating from their task objective, which would retract from the collective group goal. In other words, CoRL is about individuals supporting each other's SRL.

2.4 Socially-shared regulated learning (SSRL)

Socially-shared regulated learning (SSRL) is another important component of effective collaboration; one that involves a heightened level of interdependency when compared to CoRL, which Järvelä and Hadwin [11] purport is not sufficient to foster effective collaboration. SSRL entails groups regulating themselves as a collective unit, using consensus building and negotiation to co-construct and make decisions about group task goals, definitions, beliefs, strategies, and knowledge. SSRL differs from SRL, ERL, and CoRL because decisions, planning, goals, etc. are not made by individuals but rather made by and monitored by the group. In other words, the group regulates itself as a collective unit through negotiation and consensus building rather than group members regulating a single individual or individuals' SRL.

3 Related Work

In addition to having theoretical/conceptual boundaries, as outlined above, empirical evidence also supports the differentiation of CoRL and SSRL, which are harder to differentiate than ERL and SRL and represent an emerging literature. These two types of regulatory activities differed across five studies surveyed in a recent, comprehensive review by Panadero and Järvelä [29], where SSRL involved more collaboration and shared regulation activities, and CoRL entailed less joint work and joint use of strategies. Panadero and Järvelä's [29] review also highlighted a relationship between higher levels of SSRL, and group performance, and learning. Although further research is needed due to a small quantity of empirical studies that examined performance outcomes, the authors' astutely noted that this finding is promising as the increased use of learning strategies and SRL does not guarantee higher performance [5; 30-32]. These results do suggest, however, that SSRL should typically lead to higher performance than CoRL.

These findings should be taken with a grain of salt, however, given the difficulties learners have in engaging in regulatory activities in the first place [11]. Nonetheless, recent proposals for scripting and group awareness tools advanced by Miller and Hadwin [12] may help researchers design increasingly effective and reliable studies of collaborative learning in technology-rich and computer-based settings. In the meantime, a better question might therefore be: under conditions empirically proven to facilitate and encourage effective SRL, CoRL, and SSRL, which is the most effective at improving performance and adaptive learning behaviors? While empirical evidence generally supports collaborative learning, including collaborative regulation, a sufficient quantity of empirical evidence does not yet exist to confidently rank-order different types of regulated learning [11,33-36].

An even more nascent area of research is the evaluation of different collaborative strategies between human learners and pedagogical agents (PAs). While a considerable amount of research in the intelligent tutoring system and artificial intelligence in education communities have examined the effects of different ERL prompts and scaffolds, including dialogues involving multiple PAs and a human learner [37], no study to the authors' awareness has been conducted that examined regulatory behaviors from a collaborative perspective. Indeed, ERL interactions between learners and PAs have been primarily characterized by prompt and feedback-driven ERL, where requests are given for students to initiate SRL behaviors by PAs in multi- or single-agent computer-based learning environments. Examples of prompts and feedback issued by PAs to learners include prompts to take notes and summarize their emerging understanding of a topic; prompts to engage in metacognitive monitoring through quizzing a teachable agent on its emerging understanding and adapting one's knowledge, beliefs, and strategies in turn; and being prompted to use a worksheet as a mental model to collect, review, reflect on and revise evidence in an immersive, inquiry-based environment [4-6, 27; 32, 38-39]. Another limitation of previous ERL work with PAs is that it primarily focuses on the third phase (studying tactics) of the information processing model. By focusing on the goal-setting and planning phase, this study extends empirical research with PAs informed by the information processing model.

4 Current Study

All of the above-mentioned diverse examples of human-PA interactions share a common limitation: none examined how learners can collaboratively regulate their learning with PAs to increase academic achievement. This study addresses this limitation by examining a rich interaction between a PA and a learner to establish subgoals for learning with MetaTutor, an intelligent tutoring system that teaches complex science topics. Depending on the experimental condition (explained in 5.2) the learner was assigned to, they could receive more or less detailed feedback from the PA.

For the purpose of this study, we were interested in examining whether different levels of collaboration could be reliably identified and characterized, and if so, whether they influenced subsequent learning outcomes from the learning session with MetaTutor immediately following the establishment of learning subgoals with one of the PAs. We addressed these objectives with the following research questions and hypotheses:

(RQ1) *Can distinct learner-PA subgoal-setting interactions be reliably identified across MetaTutor conditions? In*

response to RQ1, we hypothesized that we would be able to identify at least two condition-specific interaction patterns because the PA has different rules in the two experimental conditions (described below in 5.2). Prior to examining learner-PA interaction log files, we were unsure, however, whether additional interaction patterns could be identified, or how any condition-specific interaction pattern would be characterized, aside from one being less constrained and more instructionally informative than the other.

(RQ2) *In interactions where the PA provided feedback, did learners comply with its instructional suggestions?* With regard to RQ2, we hypothesized that learners would follow the PA's hints most of the time, based on Karden and Conati's [40] finding that learners followed adaptive hints delivered from an intelligent tutoring system approximately 63% of the time, despite the fact that it was not a human tutor. We further hypothesized, that learners would follow less constrained and more instructionally informative response-adaptive subgoal hints more often than learners in Karden and Conati's [40] study because they were more open-ended in nature. Additionally, learners were afforded more autonomy in fulfilling them, which is a positive motivational and affective task feature [41]. This hypothesis is in-line with the differences in hint compliance observed between Karden and Conati's [40] and Conati, Jaques, & Muir's [42] studies, where the former used hints with more than one behavioral option (e.g., which part of a problem to focus on next), and the latter only one (e.g., to use or not use a learning tool; a magnifying glass). In keeping in-line with Karden and Conati's [40] terminology, we referred to learners' tendency to follow or not follow the PA's prompts and feedback as compliance or non-compliance.

(RQ3) *Does the type of learner-PA subgoal-setting interaction have an influence on learning?* In response to RQ3, we hypothesized that more collaborative interactions with the PA would lead to better learning outcomes, in line with Panadero and Järvelä's [29] review. While PAs are qualitatively different learning partners than humans, we believed that if response-adaptive hints, and especially those that are less constrained and more instructionally informative, are being followed (RQ2), then it means a collaboration of some kind is occurring, which opens up the possibility of learners reaping the benefits observed in research on collaborative learning with humans.

We addressed these three research questions by conducting three sets of analyses. The first was an open-coding qualitative analysis of learner-PA interactions (LPAI) during the subgoal setting phase that yielded different learner-PA interaction pattern descriptions across the learning conditions. These qualitative results are situated relative to the different types of regulated learning (SRL, ERL, CoRL, SSRL) in the discussion section (7.1). The second analysis examined whether or not learners followed the prompts they received from the PA when it provided one by classifying and counting the number of followed vs. un-followed prompts. The third analysis used descriptive and inferential statistics to examine whether differences could be observed between different learner-PA interactions and student learning.

5 Methodology

5.1 Learners

107 undergraduate students (72% female) from two large public universities in North America participated in this study. The following demographic information is based on data collected from 76 learners (71% of sample), as the remaining learners did not provide this information. Learners were between 18 and 37 years old ($M = 20.95$; $SD = 2.88$) and enrolled in various undergraduate programs, including social sciences (65%) and nursing (5.3%; which was the only discipline related to human biology). Learners' self-reported GPA ranged from 1.9 to 3.89 (out of 4.00) ($M = 3.03$; $SD = 0.47$). Less than half of the sample (45%) reported taking at least one biology course at the undergraduate level. All learners were paid up to \$40 (\$5/half-hour; \$35 if they finished 30 min. early) for completion of the two-day, four-hour experiment. Learners were randomly assigned to one of two conditions, which varied in the degree of prompting students received and in the quality of the feedback provided after SRL processes were performed. 50 learners were assigned to the condition with the higher degree of prompting and 57 to the other. No significant differences between conditions were observed for learners' pretest scores (see section 5.3).

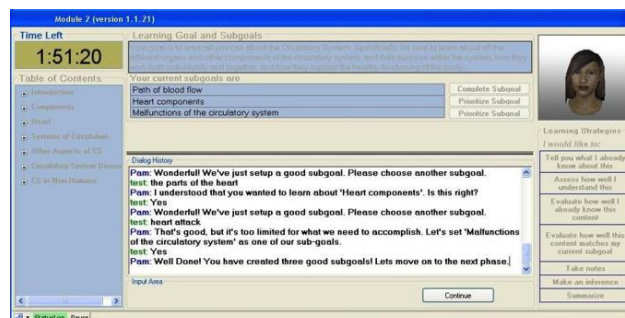


Fig. 1. Screenshot of the MetaTutor interface and goal setting learner-PA interaction example for the feedback scaffolding (PF) condition. Above the dialog history, the three subgoals that have just been set remain visible during the whole learning session. On the top right hand corner the PA currently interacting with the learner is visible (here, Pam the Planner). Non-goal setting interface elements are faded out during this phase.

5.2 MetaTutor and Experimental Procedure

MetaTutor is an intelligent tutoring system (ITS) and hypermedia-learning environment that embeds pages with text and static diagrams organized by a table of contents displayed in the left pane of the environment (see Figure 1). The version of MetaTutor used in this experiment is composed of 41 pages on the human circulatory system taught to learners during their interaction with the four PAs embedded in the system. The experiment took part in 2 different sessions, separated by 3 hours to 3 days.

During Session 1 of the MetaTutor experiment, learners completed a demographics questionnaire, and a pretest on the human circulatory system as well as on their understanding of SRL. Session 1 generally took 20 to 30 minutes for learners to complete. During Session 2, learners interacted with MetaTutor for two hours, after which they immediately completed a counterbalanced 25-item posttest on the human circulatory system.

Prior to the beginning of the learning portion of Session 2, learners were instructed how to complete the think-aloud protocol, and audio and video capture equipment were tested and activated (captured video and audio data is not used in the current study). Several tutorial videos featuring instructions from either a human agent or a combination of a human and PA were shown to the learner at the beginning of the session to explain how to navigate the system and how to use the interface to declare the SRL processes they were performing while learning.

MetaTutor Conditions. Learners were randomly assigned to either an experimental *Prompt and Feedback* condition ($n = 50$) or to a *Control* ($n = 57$) condition. Instructional scaffolding was provided by four PAs and varied depending on the experimental condition to which learners were assigned. In the prompt and feedback condition (PF) condition, learners were prompted by the PAs to use specific self-regulatory processes and were given feedback about their use of those processes during the learning session. In the control condition, learners did not receive prompts or feedback from the PAs and could only perform these self-regulatory processes on their own initiative.

The PAs included Gavin the Guide, Pam the Planner, Mary the Monitor, and Sam the Strategizer. Gavin provided guidance for learners and administered pretest and posttest knowledge assessments and self-report measures. Its interactions with learners did not vary between conditions. Mary prompted and supported learners' monitoring processes (e.g., judgment of learning), and Sam prompted learners to engage in learning strategies and ensured their use (e.g., note-taking, summarizing). In the control condition, Mary and Sam only responded to learners' self-initiated monitoring and self-regulated learning strategies in an instructional manner (e.g., acknowledging completion of a summary), when appropriate, rather than providing feedback on the quality of their responses or recommending that they engage in them. Pam prompted and scaffolded learners' planning processes by encouraging them to activate (i.e., recall) relevant prior knowledge about the topic, and to set three subgoals at the beginning of their learning session, which help them approach the learning task and achieve their overall learning goal: to learn all they can about the human circulatory system. In addition to prompting learners to activate their prior knowledge in the PF condition, Pam provided them with feedback regarding the appropriateness of their proposed subgoals, and offered them the opportunity to try again (when a proposed subgoal was inappropriate). This additional affordance involved learners more in the goal setting process than when they are immediately offered a more suitable alternative (in the control condition). The subgoal setting phase, which is central to this study, is described below in more detail.

The subgoal setting phase in MetaTutor. Before learners began to interact with Pam, who helped them set their subgoals, they had to progress through a few introductory steps, the first of which was the introduction of their overall learning goal and subgoal setting phase by Gavin, who instructed learners that they would have *one hundred and twenty minutes to learn about the circulatory system and that their overall learning goal for the session was to learn all they could about the circulatory system*. Gavin then elaborated that learners should *specifically, be sure to learn about all the different organs and different components of the circulatory system and their purpose within the system, how they work both individually and together and how they support the healthy functioning of the body*. Next, Gavin introduced Pam, the PA who helped students set subgoals, which this study focuses upon: *The first thing you'll be doing is setting some subgoals for your learning session. Pam the Planner will be assisting you in setting these sub goals*. Following Gavin's introduction, Pam introduced herself and requested that learners activate their prior knowledge regarding the circulatory system: *Hi, I'm Pam the Planner, I will help you to set up subgoals. It looks like we're about ready to start learning. First, I want you to tell me everything you know about the circulatory system*. When learners finished entering their prior knowledge regarding the circulatory system into MetaTutor (using the keyboard), they were shown a brief introductory video instructing them on how to set an appropriate subgoal and the PAs role in this process, including an example from another version of MetaTutor about the nervous system.

During the proceeding subgoal setting phase, learners interacted with Pam to successively set three subgoals designed to assist them in learning about the human circulatory system with MetaTutor. Following the subgoal setting example video, Pam proposed that they begin setting subgoals of their own: *Ok, let's set up some subgoals then*. Learners were then required to propose a subgoal using the keyboard. The subgoal setting phase involved several learner-PA dialogue moves, where upon submitting a proposed subgoal, the PA would provide condition-dependent tutorial instruction regarding its appropriateness. Following the PAs tutorial instruction the learner would then have a chance to respond, by, for example accepting or rejecting a new subgoal proposed by the PA, making a selection from subgoal options presented by the PA that it felt the learner may have intended to set, or entering new text following a request to try again.

Ideal subgoals describe, at a level of detail that is not too broad or too specific, something the learner wished to learn about

and understand regarding the circulatory system. The seven appropriate subgoals a learner could set included: (a) path of blood flow, (b) heartbeat, (c) heart components, (d) blood vessels, (e) blood components, (f) purpose(s) of the circulatory system, and (g) malfunctions of the circulatory system. The PA used natural language processing (NLP) to recognize these subgoals as well as goals related to them, but either too specific or too general to be appropriate for the learning session. The PA rejected subgoals, which were unrelated to the circulatory system (e.g., about digestion).

Figure 1 is a screenshot of the MetaTutor system during the subgoal setting phase and displays an interaction between Pam and a mock learner 'test', to whom it is providing feedback regarding their proposed subgoals. Note that when the learner suggests 'heart attack' as one of their subgoals, Pam provides feedback on this choice, specifically that 'heart attack' is too specific a subgoal to guide them in learning all they can about the circulatory system. Furthermore, Pam makes a counter suggestion, which still takes into consideration the topic of heart attack, though it generalizes it to 'malfunctions of the circulatory system', which a heart attack is one kind of.

The previous example can be traced in the flowchart displaying the decision algorithm in Figure 2, where the learner's proposed subgoal, 'heart attack' is: (a) not identified as being too broad; (b) is identified as being related to an ideal subgoal; (c) and one that is not already selected; (d) but one that is not an ideal subgoal; (e) because it is too specific; (f) therefore, the PA proposes the broader, ideal subgoal, which was malfunctions of the circulatory system; (g) which the learner accepted as their subgoal; (h) and is accordingly congratulated; (i) and since this was the third subgoal, they were prompted to move on to the next phase of the learning session with MetaTutor.

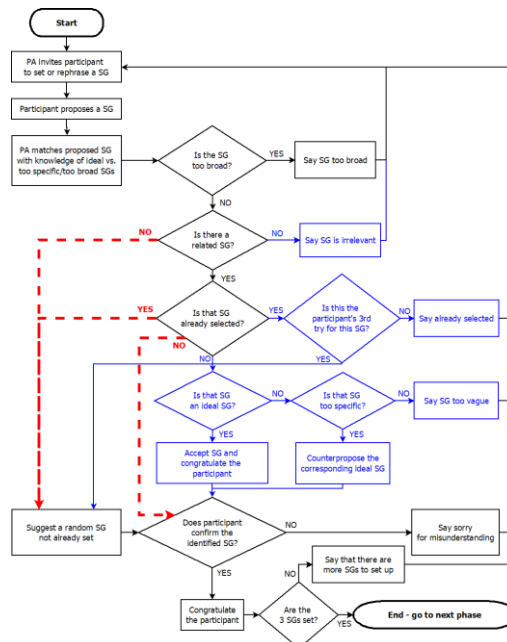


Fig. 2. Flowchart representing the PA's subgoal selection algorithm. Blue boxes and directional arrows are decision nodes and paths available to the PA in the Prompt and Feedback condition, while red paths are those available in the Control condition. Black ones are available in both. SG = subgoal.

5.3 Measures

MetaTutor log files recorded all events in the system, and in particular interactions between the learner and the PA. Appendix A displays an excerpt of such a log corresponding to events occurring during the subgoal setting phase, where column 1 represents the count of the recorded events, columns 2 and 3 are the absolute and relative time (in ms), respectively, since the beginning of the session, and column 4 identifies the type of the event dummy coded by a number. In the case of interactions between the agent and the learner (events of type 3), column 5 represents the initiator of the action (either the PA or the student; "NA" represented the student in Appendix A), column 6 corresponds to the identifier of the script used (by the PA), and column 7 is the associated text either typed as an input by the student or displayed or spoken as an output by the agent.

RQ1: The Learner-PA Interaction (LPAI) coding scheme. The final LPAI coding scheme (qualitative coding process described in 6.1) contained four profiles used by two trained coders to code the PA's responses to learners' proposed subgoals from 247 segments ($\kappa = 0.98$) extracted from the log files (agreement reached 100% following discussion). One coder coded an additional 240 segments. Therefore, we had a total of 487 segments coded from 107 learners; 50 learners from the PF condition and 57 from the control condition. The learner-PA interactions were segmented based on dialogue turns between the PA and the learner where the PA provided feedback regarding the learner's proposed subgoal.

We did not examine data from students' interactions with the PA while they set up their third subgoal in any of our analyses for RQ1 through 3 because only 69.2 % of learners completed it during the learning session. Therefore, its inclusion would have limited our analyses of the relationship between LPAI profiles and subsequent learning. Furthermore, in cases where the PA and

the learner spent more than one dialogue-turn (e.g., segment) setting up an appropriate subgoal, we used the first segment as the profile code for that particular subgoal. The unit of analysis was the PA-learner interaction for each individual subgoal because learners could have different interactions between the first two subgoals we examined for each of them.

RQ2: Examining Prompt and Feedback Compliance (PFC). Three themes emerged from qualitatively examining the PF condition learners' interactions with the PA for instances of compliance and non-compliance. Only learners from the PF condition were coded because learners in the control condition did not have the same opportunities for non-compliance (see 5.2). Frequency counts of the three themes related to compliance and non-compliance were done. Interactions from the 50 learners assigned to the PF condition were coded by the same coders as the previous analyses, where one coder coded all 50 learners' logs, and the second coded 31 logs and 71/110 segments ($\kappa = 0.76$; agreement reached 100% following discussion).

RQ3: Examining the relationship between LPAI profiles and learners' proportional learning gains. Pretest and posttests consisted of two equivalent, counter-balanced sets of 25 multiple-choice items that included questions assessing learners' knowledge of each of the seven subgoals. For example: Why is the heart muscular? (a) because it must pump blood through the system [correct]; (b) because it is connected to arteries and veins; (c) because it sends blood to the muscles; and (d) because all body organs are made of muscle. All 25 items were used to ensure that the learners who were randomly assigned to one of the two conditions, (PF: $M = .70$, $SD = .18$; Control: $M = .68$, $SD = .18$) were equivalent / not significantly different in their prior knowledge scores ($F(1, 105) = .235$, $p > .05$).

To assess differences in learners' pretest to posttest learning, we calculated their proportional learning gain scores using the formula below from [8].

$$\frac{\text{posttest score} - \text{pretest score}}{1 - \text{pretest score}}$$

We made a revision to the original formula because negative values suggest that learners unlearned material, which is unlikely, especially in MetaTutor where it is more probable that they simply guessed successfully on pretest items. The revision involved us replacing negative values with a zero. Therefore, in cases where the pretest score was greater than the posttest score, learners' proportional learning gain was indicated as zero.

Additionally, proportional learning gain scores only drew on items that were related to the subgoals learners set during the initial subgoal setting phase so as to assess the effect of their interaction establishing a subgoal and how they performed on post-test items related to that subgoal. To account for the fact that one learner could be classified as having two different LPAI profiles within the same MetaTutor condition, we drew upon two different proportional learning gain scores per learner: one corresponding to each of the first two (of three) subgoals that they set with Pam, the same ones that were coded. In other words: each subgoal-specific proportional learning gain score (e.g., for subgoal 1) was calculated by excluding items from the pre- and posttest that were not related to the corresponding subgoal (e.g., subgoal 1). Therefore, results from the two subgoals were examined separately.

Subgoal-specific pre- and posttests had different numbers of examined items because there were two different counterbalanced versions of the test (referred to below as "a" and "b") and seven different possible subgoals, which corresponded to varying numbers of pages (and therefore varying numbers of questions in the tests). The number of items assessing knowledge about an individual subgoal ranged between one and 14 where (a) path of blood flow = 6 in version A / 6 in version B of the test, (b) heartbeat = 4 (A) / 1 (B), (c) heart components = 9 (A) / 7 (B), (d) blood vessels = 10 (A) / 14 (B), (e) blood components = 9 (A) / 10 (B), (f) purpose(s) of the circulatory system = 11 (A) / 14 (B), and (g) malfunctions of the circulatory system = 4 (A) / 2 (B) items. The range of relevance is also due to some subgoals being related to more than one of the other seven subgoals. On average, each learner was assessed on 7.97 questions ($SD = 2.40$) relative to each of their first or second subgoal in the pre-test, and on 8.23 questions ($SD = 2.69$) in the post-test.

To statistically examine the effects of learners' LPAI profiles on their proportional learning gains concerning the subgoals they set during the subgoal setting phase, we performed four one-way ANOVAs with profile as the independent variable and subgoal-specific proportional learning gain as the dependent variable (e.g., subgoal 1, subgoal 2). LPAI profiles and proportional learning gains were examined within condition (necessitating four rather than two one-way ANOVAs) because Pam's behavior was not the only PA behavior that varied by condition: the other three PAs' behavior also varied [see 51], as did Pam's post-subgoal setting phase interactions with the learner. Examining the relationship between LPAI profiles and proportional learning gains within condition where post-PA behavior was the same therefore provided a cleaner and more direct evaluation of RQ3. Additional details of the analyses related to the frequency of learners per profile are described in Table 1 of the Results section, following the presentation of the categories derived through the qualitative analyses.

6 Results

The results below are organized to present the multi-stage development of the learner-pedagogical agent interaction (LPAI) coding framework and the accompanying analyses that correspond to research questions one through three.

6.1 RQ1: Can distinct learner-PA subgoal-setting interactions be reliably identified across MetaTutor conditions?

In order to answer the first research question, we engaged in a two-stage qualitative analysis of the subgoal setting interactions between Pam the Planner and the learners that were captured in the log files. We began our analyses using Graesser and colleagues' [44] five-step dialogue framework to code the learner-PA interactions during the subgoal setting

phase of MetaTutor. We adapted the framework through iterative, qualitative analysis to capture the different PA scaffolding and learner behavior themes that emerged. Our analyses led us to develop an initial framework that involved four steps [45]:

1. PA invites student to set up a subgoal
2. Student proposes a subgoal
3. PA gives feedback on proposed subgoal
4. Student responds to feedback.

The next stage of our coding that was used in this manuscript to extend [45], involved increasing our focus on the *nature* of the feedback that the PA provided to students in response to proposed subgoals. Specifically, we coded the PA's feedback based on the *degree* of scaffolding it provided students. This direction emerged from an examination of the learner-PA interactions, which revealed a substantial amount of variance in step 3 (PA feedback) and 4 (learners' response to feedback). The below four categories were the result of coding the PA-learner interactions belonging, primarily, to steps 3 and 4. We drew upon higher education learning and instruction literature that examined the competing goals of students' desires (and expectations) to be given the "right answer", as compared to research illustrating the benefits of students' active involvement and engagement with learning material to label the categories [46-47].

Spoon-feeding Type 1 (SPF1). SPF1 was assigned when the PA told learners what subgoal they should set. In this case, no hints were provided. In other words, the learner was not involved in setting the subgoal, nor were they given any information regarding why their proposed subgoal was inappropriate in relation to their overall learning goal. An example of an interaction that would be coded as SPF1 would be the learner proposing an inappropriate subgoal, such as "brain," and the PA responding "Ok, how about we try to learn about 'Heartbeat?'"

Spoon-feeding Type 2 (SPF2). SPF2 was assigned when the learner was given the answer, but the answer was accompanied with feedback regarding why it was/was not appropriate in relation to their overall learning goal. SPF2 included instances where the PA provided learners with choices regarding which subgoal they would like to set when the proposed subgoal was related to the options it proposed. SPF2 necessitated more involvement and awareness of the subgoal setting task on the part of the learner than SPF1. An example of an interaction that would be coded as SPF2 would be the learner proposing an inappropriate subgoal, such as "blood transport" and the PA responding "that's good, but we need a little more info to pin down what you want to learn about. Did you mean 'Blood components', 'Purposes of the circulatory system' or 'Malfunctions of the circulatory system?'"

Collaboration (Co). Co was assigned when the PA provided the learner with feedback (not a counter-proposed subgoal) regarding why their subgoal was inappropriate and a prompt for them to try again rather than providing the answer. We classified this code as collaboration because it is an example of the PA and learner working together toward a shared objective (to set up three appropriate subgoals). This type of response from the PA afforded the student a greater degree of involvement in setting their own learning objectives for the learning session as well as an opportunity for them to learn from their previous attempts and improve. An example of an interaction that would be coded as Co would be the learner proposing an inappropriate subgoal, such as "heart" and the PA responding "I sense you are onto something here, but I don't have enough details to point you to a good subgoal. Please try again and be more specific this time."

Right Answer (RA). RA was assigned when the PA confirmed its understanding of the learners' proposed subgoal. This involved the PA re-stating what the learner had said using the wording of the ideal subgoal that it believed the learner meant to suggest. Accordingly, the confirmation involved the PA either using closely related (e.g., synonymous), or the exact same wording as the learner to ensure it properly understood the subgoal the learner wanted to set. We classified this type of response as RA because the PA recognized one (or more) appropriate subgoals in the learners' proposal and therefore sought to confirm rather than re-direct their goal setting. An example of an interaction that would be coded as RA would be the learner proposing an appropriate subgoal, such as "I would like to learn about the diseases that could affect the circulatory system" and the PA replying "I understood that you wanted to learn about 'Malfunctions of the circulatory system'. Is this right?"

Table 1 reveals how often each of the aforementioned categories was reflected in learners' interactions with Pam during the subgoal setting phase across conditions. An examination of Table 1 reveals that most of the interactions in the PF condition either had a collaborative opportunity (Co) or involved the learner providing a right answer (RA; appropriate subgoal) and thus not requiring a hint. The majority of the remaining learner interactions were classified as situations where the PA simply provided learners with an appropriate subgoal without explanation of why it was right (SPF1). Few learners across conditions received an explanative account of why their subgoal was not appropriate when another was suggested (SPF2).

6.2 RQ2: In interactions where the PA provided feedback, did learners follow its instructional prompts?

In addition to identifying the type of LPAI, we were also interested in whether learners used the prompts and feedback that the PA provided to them in the PF condition. We referred to learners' use of PA suggestions as learners' prompt and feedback compliance (PFC). PFC was coded when the learner counter-proposed a subgoal that was (1) either more specific or more general than their initially proposed subgoal or (2) a shift toward a new topic if the proposed subgoal was deemed irrelevant or related to a previously established subgoal. A change, rather than the *direction* of the proposed

subgoal change, was used to assign this code because the majority of students (97.1%) were not biology majors and many had low pretest scores. Therefore, we could not be sure that learners knew the hierarchical relationships between the different circulatory system components they proposed, prompting us to focus on whether they recognized that a change in specificity or general topic was being suggested. For example, a learner (PN 23076) who proposed learning about the “heart” as a subgoal and countered a recommendation from Pam to “...please try again and be more specific this time” with “heart parts” would be coded as PFC. Moreover, “heart parts” represents an acceptable subgoal because it is another phrasing for “heart components”, which Pam would subsequently confirm. Similarly, a counter proposal from a learner (PN 33052) to establish “heart parts” as the subgoal following Pam’s feedback that the learner’s proposed subgoal to “...learn about the lungs, and oxygen flow” was related to a subgoal they had previously set (Path of blood flow) would be classified as PFC. This example illustrates a successful counter proposed subgoal following a recommendation to identify a new subgoal topic.

Learners’ interaction with the PA was coded as *no prompt and feedback compliance* (NPFC) when the learner did not adjust their subgoal toward the PA’s recommendation. Accordingly, there were two subordinate codes for this category. The first, *no PF compliance-no shift* (NPFC-NS), was assigned when a prompt for the learner to be more general or specific was followed up by the learner attempting to make only surface changes to or re-proposing their proposed subgoal (e.g., using the same words, a synonym, pluralizing, etc.). A counter-proposal to learn about “blood” (PN 23050) after being informed by Pam that “blood types” was not specific enough is an example of a segment that would be coded as NPFC-NS. Instances where the learner would counter propose a subgoal that was closely related to the topic of their original subgoal after receiving feedback that it was either unrelated to the human circulatory system or too closely related to a previously set subgoal would also be labelled as NPFC-NS. The following example from learner PN 23005 illustrates this second situation where they proposed to establish “digesti[on]” as their subgoal after receiving feedback that their previous subgoal “small intestine” didn’t fit with what they were trying to learn about and the prompt to create a subgoal that was related to the circulatory system.

The second NPFC code, *no PF compliance-abandonment* (NPFC-A) was used if learners abandoned the subgoal they had proposed and received feedback on and counter-proposed a *different* topic regarding the circulatory system when the PA instructed them that it needed to be more general or specific. For example, one learner (PN 23013) proposed to learn about the “lungs” after being told that they were onto something, but to try again and be more specific in response to proposing “heart” as a subgoal. While the students’ lack of expertise with biology prevented us from making inferences about their understanding of how more detailed circulatory system components related to one another, assuming that undergraduate students understood that they were switching topics was a low inference; an assertion supported by the coders’ agreement.

Our coding revealed that learners successfully followed the prompts and feedback that the PA provided the majority (67%) of the time; much more often than they abandoned learning about their originally proposed subgoal to learn about something else (19%) or failed to make any meaningful changes in their counter proposal, clinging to their original proposal (15%). See Figure 3 for details.

TABLE 1
FREQUENCY OF LPAI INTERACTION BY CONDITION AND SUBGOAL

SG	Condition						
	PF (N = 50)				C (N = 57)		
	Co	RA	SPF1	SPF2	RA	SPF1	SPF2
1	24	16	2	8	20	29	8
2	22	19	1	8	17	31	8

LPAI= Learner-PA interaction. PF = Prompt and Feedback condition. SG = subgoal. C = Control condition. Co = collaborative profile. RA = right answer profile. SPF1 = spoon-fed 1 profile. SPF2 = spoon-fed 2 profile. No interactions in the Control condition were classified as CR.

TABLE 2
 PROPORTIONAL LEARNING GAINS BY CONDITION, COLLABORATIVE PROFILE, AND SUBGOAL

		Condition													
		PF (N = 50)						C (N = 57)							
		Co		RA		SPF1		SPF2		RA		SPF1		SPF2	
SG	n	M(SD)	n	M(SD)	n	M(SD)	n	M(SD)	n	M(SD)	n	M(SD)	n	M(SD)	
1	24	40.24 (39.90)	16	36.57 (30.84)	2	17.14 (4.04)	8	29.62 (41.34)	20	51.25 (32.58)	29	27.23 (28.62)	8	47.82 (35.63)	
2	22	41.69 (37.31)	19	32.22 (37.21)	1	58.92 (22.73)	8	46.42 (37.99)	17	32.09 (37.93)	31	38.41 (31.36)	8	36.81 (25.31)	

PF = Prompt and Feedback condition. C = Control condition. SG = subgoal. Co = collaborative profile. RA = right answer profile. SPF1 = spoon-fed 1 profile. SPF2 = spoon-fed 2 profile. Means (M) and standard deviations (SD) reported under each profile correspond to learners' proportional learning gain scores, using items related to the subgoal in question. Italicized frequency and descriptive statistic cells are those with fewer than ten learners. No interactions in the Control condition were classified as Co.

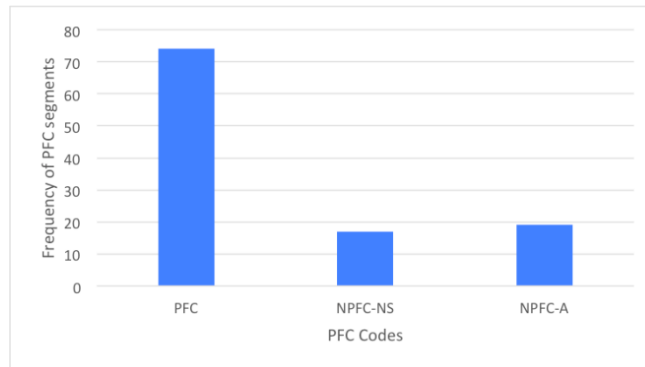


Fig. 3. Frequency of different types of prompt and feedback compliance codes. PFC (prompt and feedback compliance): Follows the PA's recommendation, with either a counter-proposal of a goal more specific or generic or a shift toward a new topic. NPFC (no prompt and feedback compliance). NPFC-NS: Does not follow the PA's recommendation to try something else and stay on the same topic instead. NPFC-A (abandonment): Does not follow the PA's recommendation to reformulate and try another topic instead.

6.3 RQ3: Does the type of learner-PA subgoal-setting interaction have an influence on learning?

Table 2 reveals the frequencies of learners in subgoals 1 and 2 as well as the descriptive statistics for proportional learning gains of related posttest items. Overall, students in the PF condition whom were classified as having a Co interaction profile tended to outperform students with less collaborative interactions with Pam while establishing subgoals, especially when cells with eight learners or fewer are disregarded. Patterns between interaction profiles and proportional learning gains were less apparent in the control condition, though students classified as having a RA interaction tended to outperform students with both spoon-feeding interactions (SPF1 and SPF2) while setting their first subgoal. This pattern was also observed in the PF condition. No Co interactions were identified in the control condition due to the more constrained nature of the PA in this condition.

In order to address the third research question, we also examined the statistical significance of the mean differences within conditions and between LPAI profiles for each of the two subgoals learners set (numbered in the order in which they were set).

Our first set of analyses investigated whether a main effect of LPAI profile on proportional learning gains existed within the PF condition. A one-way ANOVA revealed that there was not a significant main effect of LPAI profile on learners' proportional learning gains for subgoal 1, $F(2, 45) = 0.25, p > .05$. Therefore, post hoc comparisons were not conducted. SPF1 profiles were disregarded in this and the subsequent analyses because of sample sizes of two and one which are unsuitable for mean comparisons (see Table 2). A one-way ANOVA examining the same effect, but for subgoal 2 also failed to reveal a main effect of LPAI on learning gains, $F(2, 46) = 0.71, p > .05$. Therefore, post hoc comparisons were not conducted for subgoal 2 either.

Our second set of analyses investigated whether a main effect of LPAI profile on proportional learning gains existed within

the control condition. A one-way ANOVA revealed that there was a significant main effect of LPAI profile on learners' proportional learning gains for subgoal 1, $F(2, 54) = 3.96, p < .05, \eta^2_p = .13$. Post hoc comparisons with a Bonferroni correction revealed that learners with an RA profile had significantly higher proportional learning gains than those with SPF1 profiles. The second one-way ANOVA failed to reveal a main effect of LPAI on learning gains, $F(2, 53) = 0.19, p > .05$. Post hoc comparisons were therefore not run. Although the sample sizes varied by profile, Levene's test of equality of variances was not violated for any of the aforementioned analyses.

7 Discussion

The purpose of this study was to examine whether different levels of collaboration could be reliably identified and characterized between a learner and PA, and if so, whether they influenced subsequent learning outcomes from the learning session with MetaTutor immediately following the establishment of learners' subgoals. In pursuing these objectives, this study advances the state-of-the-art by extending the examination of learner-PA interactions from ERL to a CoRL and SSRL theoretical perspective by focusing on the *nature* of the collaborative interactions between a human learner and a virtual, artificial intelligence-driven PA. After discussing the empirical results of this study, the collaborative learner-PA interaction profile is situated within the spectrum of CoRL and SSRL theories.

Our qualitative coding related to the first research question (*Can distinct learner-PA subgoal-setting interactions be reliably identified across MetaTutor conditions?*) revealed four distinct learner-PA profiles, in-line with our hypothesis that at least two interaction patterns would be identified because of differences in the decision taken by the algorithm between the PF and control conditions. These condition-related differences were reflected in Table 1, which revealed that the collaborative (Co) interaction profile was only found in the prompt and feedback (PF) condition. Of the four categories, this was the one that held the greatest relevance to the objectives of this study given that collaboration did not truly occur in the other three profiles. The two spoon-feeding profiles (SPF1 and 2) provided students with an appropriate subgoal; the second spoon-feeding (SPF2) profile providing a rationale for why the learner's proposed subgoal was not appropriate in the process. The Right Answer (RA) profile did not offer an opportunity for back and forth communication with the learner either, but in this case it was because the learner proposed a subgoal that was recognized by the PA as being appropriate for the learning session.

As such, the collaborative condition was the most complex from the perspective of numerous decision pathways for the PA to take in response to the learners' replies to its prompts and feedback. While we were unsure at the outset of the study how many collaborative profiles would be identified, the need to examine the extent to which learners followed the PA's prompts and feedback when they were issued and the learner had a choice in how to respond was outlined in the second research question: *In interactions where the PA provided feedback, did learners take it to heart and follow its instructional prompts?* In-line with our second hypothesis and prior research, we found that learners would follow the PA's prompts and feedback the majority of the time. Moreover, also in-line with our initial hypothesis, we found that learners also followed the PA's prompts and feedback more often than those in Karden and Conati's [40] study, potentially, because they were more open-ended in nature, affording learners more autonomy in fulfilling them and thus enhancing task-related motivation.

This finding is interesting from the point of view that the PA is occupying the place of an expert human tutor, and learners are responding to it accordingly. This "buy-in" is an important criteria for collaborative learning to occur between learners and PAs because no scaffolding can take place if learners simply ignore what the PA has to say and/or attempts to game (cheat) the system such as by entering random phrases until the PA provides an answer (bottom-out hints) [42, 48].

In response to our third research question (*does the type of learner-PA subgoal-setting interaction have an influence on learning?*) descriptive statistics revealed that learners who established subgoals with collaborative (Co) learner-PA interaction profiles tended to outperform those with less collaborative interactions with Pam, especially when cells with eight learners or fewer were disregarded. Interestingly, learners classified as having collaborative interactions outperformed learners who were able to establish an appropriate subgoal without the assistance of the PA, though none of the between-profile differences within the PF condition were significant. One explanation for this observation may be that the act of working together with the PA served to heighten these learners' interest and motivation, potentially by underscoring their autonomy (i.e., choice) through requiring them to exert more effort to set the subgoal. While learners who were able to immediately produce an appropriate subgoal had the same degree of choice (more, in fact, by having their first choice accepted), it is possible that they may have taken that choice for granted by having it immediately accepted, confirmed, and then moving on. The motivational underpinnings of the subgoal setting phase represents an area for future research, in particular, knowing that goal orientations have an influence on learner-performance [49]. Another explanation for Co learners outperforming RA learners in the PF condition is the possibility that those students who were able to propose an appropriate subgoal may have felt more confident (perhaps overconfident) in their ability to complete the subgoal and therefore invested less effort in SRL strategies and more time exploring unrelated content out of curiosity. Prior research has documented the negative effect over-confidence can have on learning [50].

Comparisons with SPF1 in the PF condition were not made because of the small sample size of this profile. Similarly, the learners classified as SPF2 (in both conditions) should be interpreted with caution because the small sample size may fail to capture the variation that might be observed in a larger sample.

The significant finding between RA and SPF1 profiles in the Control condition may be explained by learners in the RA

profile having lower initial subgoal 1-specific prior knowledge than those in the SPF1 profile and thus more variance available with which to observe a pre-to-post change in subgoal-related learning. To test this hypothesis we conducted an independent samples t-test comparing the mean subgoal 1-specific pretest means between learners classified as belonging to the RA ($M = .56$; $SD = .28$) or SPF1 ($M = .71$; $SD = .30$) profiles during their first subgoal setting interaction with Pam. While the t-test was not significant, $t(47) = -1.822$, $p = .08$ it was marginally so, suggesting that subgoal 1-specific prior knowledge may have played a limited role and helped contribute toward explaining the aforementioned result. One explanation for this follow-up result is that RA learners may have proposed learning about a subgoal they could think of, but knew little about but, whereas SPF1 learners may have been directed to a subgoal they knew more about, on average.

Taken together, these results provide preliminary evidence that learners are both willing to engage in and benefit from collaborative interactions with PAs when immediate, directional feedback and the opportunity to try again are provided. But how do the collaborative interactions that took place between Pam and learners during the subgoal setting phase of MetaTutor align with the definitions of self-, co-, and socially-shared regulated learning that served as the theoretical framework for this study? In order to answer this question we reviewed Miller and Hadwin's [12] recently published set of guidelines for differentiating between SRL, CoRL, or SSRL and applied them to the learner-PA goal setting interaction described in this study. Notably, the guidelines were intended to differentiate collaborative learning between humans, not between a human and a PA. Although this limited its application, insights can still be drawn.

7.1 Situating LPAI within Theories of Co- and Socially-shared Regulated Learning

In summary, Pam's interactions with the learner during the goal setting phase provided one-way scaffolding for goal setting, which drew upon a pre-existing list of subgoals and a network of related keywords that Pam used in unison to provide feedback on the appropriateness of learners' subgoal proposals and their proximity (too general, too specific) to the ideal subgoal (see Methods for more details). Given that Pam is programmed by a subgoal selection algorithm, however, it cannot be said that the learner truly influenced Pam's regulation per se. Moreover, learners made adaptations to their understanding of what an effective subgoal was based on Pam's prompts and feedback, but Pam's understanding of the task was fixed. As such, several components of the interaction are better classified as externally-regulated learning (ERL).

An important caveat to the goal setting process, however, is that the learner and Pam must *agree* on the subgoals that the learner will use to learn all they can about the human circulatory system. Therefore, the back and forth proposals, feedback, and counter proposals, represent a negotiation between Pam and the learner to mutually establish the subgoals. Although the goal is ultimately the learners' and the regulatory scaffolding is unidirectional, goal negotiation aligns with Miller and Hadwin's [12] characterization of SSRL. Although this makes labeling the interaction challenging, Miller and Hadwin's [12] collaborative regulatory labels are not seen as mutually exclusive, but rather a continuum of elements that shape and characterize a collaborative interaction [12, 29]. Indeed, what can be taken away from reviewing emerging theoretical work on CoRL and SSRL is that learners' subgoal setting interaction with the PA in MetaTutor is more collaborative than traditional conceptualizations of ERL [14]. Moreover, in addition to the Co profile sharing some of the characteristics of SSRL, it also appears to share some of the benefits of this type of regulated learning. Not only did learners attend to the PA's prompts and feedback in negotiating their subgoals, but descriptive statistics revealed that Co profile learners also tended to outperform learners with other LPAI profiles.

7.2 Limitations and Future Directions

One of the limitations with this research is that learners could have been assigned to two different profiles by the coders, one for each of their two first subgoals. As such, relationships between individual learners and LPAI profiles could not be examined. Examination of learners' proportional learning gain scores was also based on different numbers of items (depending on which subgoals were selected), limiting the variance with which to observe differences in learning for some subgoals. It should be noted, however, that the average number of subgoal-specific items used for proportional learning gains afforded reasonable opportunities to observe variance in scores and the mean values for subgoal 1 and 2 were very close (see 5.3, RQ3). A related limitation is the fact that learners' subgoal-specific proportional learning gains drew on different subgoals means that learners were not all assessed on the same content for the purposes of the analyses reported in this article—even though all learners who participated in the study took the same counterbalanced pretest and posttest.

Another limitation of this study is that it did not examine learners' motivational orientations and affective dispositions. This information could have provided additional context to interpret their responses to the PA's prompts and feedback. Moreover, this study did not examine other facets of SRL that have implications for learning, such as metacognitive monitoring during the learning session.

The distribution of learners in the LPAI coding scheme was another limitation, and one that (1) reduced the number of statistical comparisons we were able to properly conduct (e.g., unable to examine SPF1 profiles in PF condition) and limited statistical power for others (e.g., SPF profiles in both conditions). A larger sample size may have yielded enough learners in these conditions to help address these limitations. Finally, it was not appropriate to make direct comparisons of proportional learning gain scores for different LPAI profiles between conditions in this study because of the differences in PA behavior between the PF and control condition. Other studies have, however, found that differences between the PF and control conditions for learning outcomes with the version of MetaTutor examined in this study were not significant [55]. Moreover, conducting mean comparisons between all six LPAI groups in the two conditions would have required a large number of

comparisons (i.e., t-tests), especially when the two separate subgoals were factored in. These analyses would have warranted a severe correction to account for the inflation of Type 1 error (false positive) from running so many analyses.

In light of these limitations, future research should use the coding scheme developed in this study to collect data from a larger number of learners assigned only to the PF condition where the Co profile was identified. Doing so with a large enough sample could allow a study to focus just on the first of the subgoals and therefore draw connections between learners' individual differences, their LPAIs, and learning outcomes from interacting with MetaTutor. Additional information about the learners should also be collected, such as personality traits and goal orientations. Such information could support user modeling where specific individual differences could be connected with learners LPAI-moderated learning gains [51]. Relatedly, examining learners' dynamic psychological states, such as emotional states, could shed light on how learners felt while interacting with the PA and draw relationships between LPAIs, emotional states, and learning [53-55]. Future research using other intelligent tutoring systems should also examine the dynamics between PAs and learners, including prompts, feedback, learner compliance, and learning outcomes to better understand the generalizability of the findings reported in this article, including the conditions (across systems) that facilitate fruitful learner-PA collaboration.

8 Conclusion

In closing, the primary contribution of these results is empirical evidence that collaborative learning interactions can take place between learners and PAs and that they may lead to higher proportional learning gains than more traditional, closed-ended externally-regulated interactions, and even situations where learners complete their objectives without requiring PA assistance. Theoretically, this study was the first to apply CoRL and SSRL perspectives of regulated learning to human-agent tutorial interactions and discuss not only where the PA-learner interaction was best situated within these perspectives, but what affordances and limitations exist when applying them to interactions with an PA rather than a human tutor. Moreover, in doing so, preliminary evidence that learners were willing to afford PAs enough cooperation through PF compliance to foster meaningful and beneficial collaborative interactions was identified and discussed. Much work remains to be done in examining collaborative learner-PA interactions, but this study makes a case for the potential of such interactions, how they can be situated within human-centric educational theory, and outlines important directions for future work.

Acknowledgment

The research presented in this paper has been supported by a doctoral and postdoctoral fellowship from the Fonds Québécois de recherche - Société et culture (FQRSC) and a Joseph-Armand Bombardier Canada Graduate Scholarship for Doctoral research from the Social Science and Humanities Research Council (SSHRC) awarded to the first author. This research has also been supported by funding awarded from the National Science Foundation (DRL 1008282, DRL 1431552, DRL 1660878), the Social Science and Humanities Research Council of Canada, and the Canada Research Chairs program awarded to the third author.

References

- [1] Azevedo, R., & Alevin, V. (Eds.). (2013). *International handbook of metacognition and learning technologies*. Amsterdam, The Netherlands: Springer.
- [2] Azevedo, R., & Hadwin, A. F. (2005). Scaffolding self-regulated learning and metacognition: Implications for the design of computer-based scaffolds. *Instructional Science*, 33, 367-379.
- [3] Graesser, A. C., & McNamara, D. (2010). Self-regulated learning in learning environments with pedagogical agents that interact in natural language. *Educational Psychologist*, 45, 234-244.
- [4] Bouchet, F., Harley, J.M., & Azevedo, R. (2016). Can adaptive pedagogical agents' prompting strategies improve students' learning and self-regulation? In A. Micarelli, J. Stamper, & K. Panourgia (Eds.) *Lecture Notes in Computer Science: Vol. 9684. Intelligent Tutoring Systems (pp. 368-374)*. Switzerland: Springer.
- [5] Bouchet, F., Harley, J.M., Trevors, G., & Azevedo, R. (2013). Clustering and profiling students according to their interactions with an intelligent tutoring system fostering self-regulated learning. *Journal of Educational Data Mining*, 5, 104-146.
- [6] Kinnebrew, J.S., Biswas, G., Sulcer, B., & Taylor, R.S. (2013). Investigating self-regulated learning in teachable agent environments. In R. Azevedo & V. Alevin (Eds.), *International handbook of metacognition and learning technologies* (pp. 451-470). Amsterdam, The Netherlands: Springer.
- [7] Lester, J.C., Bradford, W. M., Robison, J.L., & Rowe, J.P., & Shores, L.R. (2013). Supporting self-regulated science learning in narrative-centered learning environments. In R. Azevedo & V. Alevin (Eds.), *International handbook of metacognition and learning technologies* (pp. 471-483). Amsterdam, The Netherlands: Springer.
- [8] Sabourin, J., Shores, L. R., Mott, B.W., & Lester, J. C. (2012). Predicting student self-regulation strategies in game-based learning environments. In S. Cerri, W. Clancey, G. Papadourakis, & K. Panourgia (Eds.), *Lecture Notes in Computer Science: Vol. 7315. Intelligent Tutoring Systems (pp. 141-150)*. Berlin, Heidelberg: Springer-Verlag.
- [9] Taub, M., Azevedo, R., Bouchet, F., & Khosravifar, B. (2014). Can the use of cognitive and metacognitive self-regulated learning strategies be predicted by learners' levels of prior knowledge in hypermedia-learning environments?. *Computers in Human Behavior*, 39, 356-367.

- [10] Hadwin, A. F., Järvelä, S., & Miller, M. (2011). Self-regulated, co-regulated, and socially shared regulation of learning. In B. J. Zimmerman & D. Schunk, (Eds.), *Handbook of self-regulation of learning and performance* (pp. 65–84). New York, NY: Routledge.
- [11] Järvelä, S., & Hadwin, A. F. (2013). New frontiers: regulating learning in CSCL. *Educational Psychologist*, *48*(1), 25-39
- [12] Miller, M., & Hadwin, A. (2015). Scripting and awareness tools for regulating collaborative learning: changing the landscape of support in CSCL. *Metacognition and Learning*, *52*, 573-588.
- [13] Azevedo, R., Greene, J. A., & Moos, D. C. (2007). The effect of a human agent's external regulation upon college students' hypermedia learning. *Metacognition and Learning*, *2*, 67-87.
- [14] Azevedo, R., Moos, D. C., Greene, J. A., Winters, F. I., & Cromley, J. G. (2008). Why is externally-facilitated regulated learning more effective than self-regulated learning with hypermedia? *Educational Technology Research and Development*, *56*, 45-72.
- [15] Jarvenoja, H., & Järvelä, S. (2009). Emotion control in collaborative learning situations: Do students regulate emotions evoked from social challenges? *British Journal of Educational Psychology*, *79*, 463-481.
- [16] Järvelä, S., Volet, S., & Jarvenoja, H. (2010). Research on motivation in collaborative learning: Moving beyond the cognitive-situative divide and combining individual and social processes. *Educational Psychologist*, *45*, 15-27.
- [17] Hadwin, A. F., Wozney, L., & Pontin, O. (2005). Scaffolding the appropriation of self-regulatory activity: A social-constructivist analysis of changes in student-teacher discourse about a graduate student portfolio. *Instructional Science*, *33*, 413-450.
- [18] Winne, P. H. (2011). A cognitive and metacognitive analysis of self-regulated learning. In B. J. Zimmerman & D. H. Schunk (Eds.), *Handbook of self-regulation of learning and performance* (pp. 15–32). New York: Routledge.
- [19] Winne, P. H., & Hadwin, A. F. (2008). The weave of motivation and self-regulated learning. In D. H. Schunk & B. J. Zimmerman (Eds.), *Motivation and self-regulated learning: Theory, research, and applications* (pp. 297–314). Mahwah, NJ: Erlbaum.
- [20] Winne, P. H., & Hadwin, A. F. (2003). nStudy: Tracing and supporting self-regulated learning in the Internet. In R. Azevedo & V. Aleven (Eds.), *International handbook of metacognition and learning technologies* (pp. 293-308). Amsterdam, The Netherlands: Springer.
- [21] Winne, P., & Azevedo, R. (2014). Metacognition. In R. K. Sawyer (Ed) *Cambridge Handbook of the Learning Sciences* (2nd ed., 63-87). Cambridge, MA: Cambridge University Press.
- [22] Azevedo, R. (2014). Issues in dealing with sequential and temporal characteristics of self- and socially-regulated learning. *Metacognition and Learning*, *9*, 217-228.
- [23] Hadwin, A. F., Webster, E., Helm, S., McCardle, L., & Gendron, A. (2010, April). *Toward the study of intra-individual differences in goal setting and motivation regulation*. Paper presented at the Annual Meeting of the American Educational Research Association, Denver, CO.
- [24] Winne, P. H., Hadwin, A. F., & Perry, N. E. (2013). Metacognition and computer-supported collaborative learning. In C. Hmelo-Silver, A. O'Donnell, C. Chan & C. Chinn (Eds.), *International handbook of collaborative learning* (pp. 462–479). New York: Taylor & Francis.
- [25] Zimmerman, B. J. (2008). Investigating self-regulation and motivation: Historical background, methodological developments, and future prospects. *American Educational Research Journal*, *45*, 166–183.
- [26] Azevedo, R., Feyzi-Behnagh, R., Duffy, M., Harley, J., & Trevors, G. (2012). Metacognition and self-regulated learning in student-centered learning environments. In D. Jonassen & S. Land (Eds.), *Theoretical foundations of student-center learning environments* (2nd ed., pp. 171-197). New York: Routledge.
- [27] Azevedo, R., Martin, S.A., Taub, M., Mudrick, N.V., Millar, G.C., & Grafsgaard, J.F. (2016). Are pedagogical agents' external regulation effective in fostering learning with intelligent tutoring systems? In A. Micarelli, J. Stamper, & K. Panourgia (Eds.) *Lecture Notes in Computer Science: Vol. 9684. Intelligent Tutoring Systems* (pp. 197-207). Switzerland: Springer.
- [28] Azevedo, R., Harley, J., Trevors, G., Feyzi-Behnagh, R., Duffy, M., Bouchet, F., & Landis, R.S. (2013). Using trace data to examine the complex roles of cognitive, metacognitive, and emotional self-regulatory processes during learning with multi-agent systems. In R. Azevedo & V. Aleven (Eds.), *International handbook of metacognition and learning technologies* (pp. 427-449). Amsterdam, The Netherlands: Springer.
- [29] Panadero, E., & Järvelä, S. (2015). Socially shared regulation of learning: A review. *European Psychologist*, *20*, 190-203.
- [30] Panadero, E., Alonso-Tapia, J., & Huertas, J. A. (2012). Rubrics and self-assessment scripts effects on self-regulation, learning and self-efficacy in secondary education. *Learning and Individual Differences*, *22*, 806-813.
- [31] Panadero, E., Alonso-Tapia, J., & Reche, E. (2013). Rubrics vs. self-assessment scripts effect on self-regulation, performance and self-efficacy in pre-service teachers. *Studies in Educational Evaluation*, *39*, 125–132.
- [32] Trevors, G., Duffy, M., & Azevedo, R. (2014). Note-taking within MetaTutor: Interactions between an intelligent tutoring system and prior knowledge on note-taking. *Educational Technology Research & Development*, *62*, 507-528.
- [33] Janssen, J., Erkens, G., Kirschner, P. A., & Kanselaar, G. (2012). Task-related and social regulation during online collaborative learning. *Metacognition and Learning*, *7*, 25-43.
- [34] Lajoie, S. P., & Lu, J. (2012). Supporting collaboration with technology: Does shared cognition lead to co-regulation in medicine? *Metacognition and Learning*, *7*, 45-62.
- [35] Saab, N., van Joolingen, W., & van Hout-Wolters, B. (2012). Support of the collaborative inquiry learning process: Influence of support on task and team regulation. *Metacognition and Learning*, *7*, 7-23.
- [36] Volet, S., Summers, M., & Thurman, J. (2009). High-level co-regulation in collaborative learning: How does it emerge and how is it sustained? *Learning and Instruction*, *19*, 128-143.
- [37] D'Mello, S.K., Lehman, B.A., Person, N. (2010). Monitoring affect states during effortful problem solving activities. *International Journal of Artificial Intelligence in Education*. *20*, 361–389.
- [38] Segedy, J. R., Kinnebrew, J. S., & Biswas, G. (2013). The effect of contextualized conversational feedback in a complex open-ended learning environment. *Educational Technology Research and Development*, *61*, 71-89.
- [39] Kinnebrew, J. S., Segedy, J. R., & Biswas, G. (2014). Analyzing the temporal evolution of students' behaviors in open-ended learning

- environments. *Metacognition and learning*, 9, 187-215.
- [40] Kardan S., Conati C. (2015) Providing adaptive support in an interactive simulation for learning: an experimental evaluation. In B. Bo & K. Jinwoo (Eds.), *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (pp. 3671-3680). NY, NY: ACM Press.
- [41] Pekrun, R., & Perry, R. P. (2014). Control-value theory of achievement emotions. In R. Pekrun & L. Linnenbrink-Garcia (Eds.), *International handbook of emotions in education* (pp. 120–141). New York: Routledge.
- [42] Conati, C., Jaques, N. and Muir M. (2013). Understanding attention to adaptive hints in educational games: an eye-tracking study. *International Journal of Artificial Intelligence in Education*, 23(1-4), 136-16.
- [43] D'Mello, S. K., Lehman, B. A., & Person, N. (2010). Monitoring affect states during effortful problem solving activities. *International Journal of Artificial Intelligence in Education*, 20, 361-389.
- [44] Graesser, A. C., Person, N. K., & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Applied Cognitive Psychology*, 9, 495-522.
- [45] Harley, J. M., Taub, M., Bouchet, F., & Azevedo, R. (2012, June). *A framework to understand the nature of co-regulated learning in human-pedagogical agent interactions*. Paper presented at the Fourth Workshop on Self-Regulated Learning in Educational Technologies (SRL@ET): Modeling, Evaluating, Supporting, and Fostering Metacognition with Computer-Based Learning Environments held in conjunction with the 11th International Conference on Intelligent Tutoring Systems (ITS 2012), Crete, Greece.
- [46] Heikkila, A., & Lonka, K. (2006). Studying in higher education: Students' approaches to learning, self-regulation, and cognitive strategies. *Studies in Higher Education*, 31, 99-117.
- [47] McKay, J., & Kember, D. (1997). Spoon feeding leads to regurgitation: A better diet can result in more digestible learning outcomes. *Higher Education Research & Development*, 16, 55-67.
- [48] Baker, R., Corbett, A. T., Roll, I., & Koedinger, K. R. (2008). Developing a generalizable detector of when students game the system. *User Modeling and User-Adapted Interaction*, 18, 287–314.
- [49] Duffy, M. C., & Azevedo, R. (2015). Motivation matters: Interactions between achievement goals and agent scaffolding for self-regulated learning within an intelligent tutoring system. *Computers in Human Behavior*, 52, 338-348.
- [50] Dunlosky, J., & Rawson, K. A. (2012). Overconfidence produces underachievement: Inaccurate self-evaluations undermine students' learning and retention. *Learning and Instruction*, 22, 271-280.
- [51] Harley, J. M., Carter, C.K., Papaionnou, N., Bouchet, F., Azevedo, R., Landis, R. L., & Karabachian, L. (2016). Examining the predictive relationship between personality and emotion traits and students' agent-directed emotions: Towards emotionally-adaptive agent-based learning environments. *User Modeling and User-Adapted Interaction*, 26, 177-219.
- [52] D'Mello, K. S., Craig, S. D., & Graesser, A. C. (2009) Multimethod assessment of affective experience and expression during deep learning. *International Journal of Learning Technology*, 4, 165-187.
- [53] Harley, J. M., Bouchet, F., Hussain, S., Azevedo, R., & Calvo, R. (2015). A multi-componential analysis of emotions during complex learning with an intelligent multi-agent system. *Computers in Human Behavior*, 48, 615-625.
- [54] Harley, J.M., Lajoie, S.P., Frasson, C. & Hall, N.C. (2017). Developing emotion-aware, advanced learning technologies: A taxonomy of approaches and features. *International Journal of Artificial Intelligence in Education*, 27, 268-297.
- [55] Azevedo, R., Landis, R.S., Feyzi-Behnagh, R., Duffy, M., Trevors, G., Harley, J., Bouchet, F., Burlison, J., Taub, M., Pacampara, N., Yeasin, M., Rahman, A.K.M.M., Tanveer, M.I., & Hossain, G. (2012). The effectiveness of pedagogical agents' prompting and feedback in facilitating co-adapted learning with MetaTutor. In S. A. Cerri, W. J. Clancey, G. Papadourakis, and K. Panourgia (Eds.), *Lecture Notes in Computer Science: Vol: 7315. Intelligent Tutoring Systems* (pp. 212-221). Berlin, Heidelberg: Springer-Verlag.

Jason M. Harley is an Assistant Professor of Educational Psychology and Technology at the University of Alberta. He completed his Ph.D. in Educational Psychology at McGill and served as a Postdoctoral Fellow in the Computer Science and Operations Research Department at the Université de Montréal and at McGill University in the Department of Educational and Counselling Psychology. Dr. Harley's research is funded by the Social Science and Humanities Research Council of Canada (SSHRC) and has also been funded by the Fonds de recherche du Québec – Société et culture (FRQSC). His research examines emotions, emotion regulation, self-regulated learning, and advanced learning technologies, including intelligent tutoring systems, serious games, and mobile augmented reality. Dr. Harley is a Consulting Editor for the international journal, Educational Technology Research and Development.

Michelle Taub is a Postdoctoral scholar in Human Factors and Applied Cognition in the department of Psychology at the University of North Carolina State University. Her research interests include using non-traditional statistical techniques (i.e., multi-level modeling, educational data mining) when analyzing multi-channel data, such as eye tracking, videos of facial expressions of emotions, and log-file data, to examine students' metacognitive monitoring during learning with different types of Advanced Learning Technologies. Specifically, she is interested in investigating what types of metacognitive judgments students deploy, the accuracy in these judgments, and how these judgments can impact students' affective states during learning.

Roger Azevedo is a Professor in the Department of Psychology at North Carolina State University. He received his Ph.D. in Educational Psychology from McGill University and complete his postdoctoral training in cognitive psychology at Carnegie Mellon University. His main research area includes examining the role of cognitive, metacognitive, affective, and motivational regulatory during human-machine interactions during learning with advanced learning technologies. He has published over 200 peer-reviewed papers, chapters, and refereed conference proceedings in the areas of educational, learning, cognitive, and computational sciences. He is the editor of the *Metacognition and Learning* journal and also serves on the editorial board of several top-tiered learning and cognitive sciences journals (e.g., *European Journal of Psychological Assessment*, *International Journal of AI in Education*). His research is funded by the National Science Foundation, the Institute of Education Sciences, and the Social Sciences and Humanities Research Council of Canada. He is a fellow of the American Psychological Association and the recipient of the prestigious Early Faculty Career Award from the National Science Foundation.

François Bouchet is an Assistant Professor at Université Pierre et Marie Curie and Researcher at the Laboratoire d'Informatique de Paris 6 (LIP6). He received his M.Eng from ESIEA, and his M.Sc. and Ph.D. in Computer Science from Université Paris-Sud. The joint research in this paper is a follow-up of his postdoc with Dr. Azevedo at McGill University, where he has been the main architect of MetaTutor and worked on applying data mining techniques across multiple data channels to identify learners' profiles and improve the system adaptation. His latest research interests revolve around MOOC to analyze attrition and improve learners' experience through recommendations and pedagogical agents.

Appendix A

Sample excerpt from a MetaTutor log file during the subgoal setting phase

70	13:55:03	1022720	3	Pam	PamSubgoalFeedbackTooSpecific	That's good, but it's too limited for what we need to accomplish. Let's set [IDEAL-SUBGOAL] as one of our sub-goals.
71	13:55:03	1022736	8	Pam	PamSubgoalFeedbackTooSpecific	Start
						That's good, but it's too limited for what we need to accomplish. Let's set 'Blood components' as one of our sub-goals.
72	13:55:10	1029572	8	Pam	PamSubgoalFeedbackTooSpecific	Stop
73	13:55:11	1030399	3	StudentInput	NA	Yes
74	13:55:11	1030430	3	Pam	PamSubgoalSelected	Wonderful!
						We've just setup a good subgoal. Please choose another subgoal.
75	13:55:11	1030446	8	Pam	PamSubgoalSelected	Start
						Wonderful! We've just setup a good subgoal. Please choose another subgoal.
76	13:55:17	1036564	8	Pam	PamSubgoalSelected	Stop
77	13:55:25	1044290	3	StudentInput	NA	Cardiovascular diseases
78	13:55:25	1044306	3	Pam	PamSubgoalFeedbackAlmostThere	Ok, let's
						set [IDEAL-SUBGOAL] as one of our sub-goals.
79	13:55:25	1044321	8	Pam	PamSubgoalFeedbackAlmostThere	Start
						Ok, let's set 'Malfunctions of the circulatory system' as one of our sub-goals.
80	13:55:31	1050549	8	Pam	PamSubgoalFeedbackAlmostThere	Stop
81	13:55:32	1051391	3	StudentInput	NA	Yes
82	13:55:32	1051407	3	Pam	PamSubgoalFinished	Well done!
						You have created two good subgoals! Please note that the bars at the top center of the interface indicate your progress toward completing each of the subgoals you have set. The top bar will change color to indicate your progress in reading all the pages relevant to your active subgoal. Now, let's move on to the next phase.