Effects of an Adaptive Robot Encouraging Teamwork on Students' Learning

Parastoo Baghaei Ravari[1], Ken Jen Lee[2], Edith Law[2] and Dana Kulić[3]

***Abstract*— In this work, we designed a teachable robot that encourages a pair of students to discuss their thoughts and teaching decisions during the tutoring session. The robot adapts to the students' talking activity and adjusts the frequency and type of encouragement. We hypothesize that the robot's encouragement of group discussion can enhance the social engagement of group members, leading to improved learning and enjoyment. We ran a user study (n = 68), where a pair of participants (dyad) worked together to teach a humanoid robot about rocks and minerals. In the adaptive condition, the robot uses reinforcement learning to maximise interaction between the dyad members. Results show that the adaptive robot was successful in creating more dialogue between dyad members and in increasing task engagement, but did not affect learning or enjoyment. Over time, the adaptive robot was also able to encourage both members to contribute more equally to the conversation.**

I. INTRODUCTION

Robots in education have demonstrated great potential for improving learning, enabling students to improve their self-esteem [1], motivation [2], and engagement [3]. While robots can take a variety of roles [1], teachable robots—robots that take the role of a novice taught by students—leverage learning by teaching to further enhance student learning. The idea is that teaching an agent can induce the so-called Protégé Effect [4], i.e., students learn better by teaching because of the increased sense of responsibility [4], [5].

While pedagogical research highlights the importance of group interaction on students' learning and overall performance in school [6], [7] and the importance of effective facilitation of small groups to increase communication, self-motivation and learning [8], many of the research experiments involving teachable robots are based on one-to-one interactions with the robot [2], [3], [9], [10], [11], [12]. In order to address this gap, we examine social interactions and study the effect of social engagement on the learning experience.

In this work, we developed a teachable agent that dynamically adjusts its behaviour to the communicative activity of a pair of users. The robot encourages social engagement by inviting the users to discuss their thoughts and decisions with their group-mate. The frequency and style of the robot's encouragement is decided by a reinforcement learning (RL) algorithm based on a reward signal that encourages group communication, measured from the users' real-time audio input. During the interaction, a pair of participants teach the robot through a tablet web application called Curiosity Notebook (Figure 1) [13], [14]. The Curiosity Notebook can be used to teach a variety of taxonomy-style topics, in this study we used a geology domain and a rock classification task. The study was initially scheduled to be run in a physical school setting; however, due to Covid-19 restrictions, the study was modified to be conducted remotely with 34 pairs of adult participants. The teaching material was unmodified as the task of rock classification is unfamiliar to most adults. We hypothesized that the adaptive robot will increase users' social engagement, improve their communication, and allow for equal contribution from each member. Our work is one of the few studies analyzing social interactions in a learning context with teachable robots. The robot's objective is to encourage within-group human communication. We performed a user study to assess the effects of adaptive encouragement on increasing group (social) engagement and its effects on learning.

[1]Parastoo Baghaei Ravari is with the Department of Electrical and Computer Engineering at University of Waterloo, Waterloo, Canada `parastoo.baghaei.ravari@uwaterloo.ca`

[2]Ken Jen Lee and Edith Law are with the Cheriton School of Computer Science at University of Waterloo, Waterloo, Canada `kj24lee@uwaterloo.ca`, `edith.law@uwaterloo.ca`

[3] Dana Kulić is with the Faculty of Engineering at Monash University, Melbourne, Australia `dana.kulic@monash.edu`

II. RELATED WORK

Teachable Social Robots in Education: Robots used in education can be categorized based on the robot's role in the learning process, which could be a tutor, a peer, or a novice [1]. Teachable robots are novices that allow the users to take on the teaching role [1] and elicit the protégé effect, which has been shown to increase students' effort [4] and improve learning [15]. Teachable robots can keep students engaged for longer periods [2], [3] and provide increased motivation [2].

Adaptive Teachable Robots: Characteristics of the teachable robot can be designed to be adaptive [16], [1], [17]. Robots that adapt to the educational level and performance of the students led to greater learning gains [1], [12] and task performance [18] in comparison to robots that do not adapt. Robots with adaptive characteristics (e.g. dynamically changing voice, knowledge progression, verbal and non-verbal social behaviour) increase social presence [10], and learning gains [11].

Lubold et. al. [10] performed a study with a peer robot that conveyed emotional information with its manner of speaking, adapting its tone, intensity and speaking rate to that of the students. Participants had the highest social presence (defined as the perceptual illusion of non-mediation) when the robot used both social dialogue and an adaptive voice. However, there was no significant effect on learning gains.

Baxter et. al. [11] compared a personalized and non-personalized robot in two different classrooms. The personalized robot adapted its non-verbal behaviour (gaze and movement), friendliness (e.g. calling the children by their name), and progression (responsiveness) to the students. The students showed significantly increased learning when using the personalized robot. However, despite the robot being in a classroom, the robot only supported one-on-one interaction and the turn-taking was moderated by the teacher.

Group Interaction: Group interactions are rare in studies of teachable robots in education [19], [20], [21], and such studies do not tend to discuss the effects of group interaction on the experiment results. However, outside of educational HRI research, there have been studies attempting to manipulate human team dynamics, performance, and perception of group cohesion during human-robot interaction [22], [23], [24], [25], in which the robot is not an active member of the team, and it is not learning or teaching, rather only providing hints and comments.

Hood et. al. [21] observed group-mates giving each other advice while teaching a humanoid robot handwriting. This collaboration occurred naturally without researchers' or robot suggestions. Similar collaboration was seen in our earlier pilot study [13], where a teachable robot interacted with a team of students in a fixed turn-taking structure.

Strohkorb et. al. [22] focused on the influence of social robots on collaboration in the context of children playing a game in 3 conditions, the robot giving the group task-focused comments, the robot giving relationally focused comments and a control condition (no robot). Children in the task-focused condition performed better than the children in the two other conditions, while the children in the relational condition perceived their performance to be higher than the students in the task-oriented condition.

Short and Mataric [23] used a robot moderator to either help the group score more points with hints related to playing the game (performance - reinforcing) or make sure individual scores are equalized across the group (performance - equalizing). Contrary to expectations, participants scored higher in the performance-equalizing condition while the group cohesion and interaction were higher in the performance-reinforcing condition. Participant helpfulness and group cohesion were positively correlated with the frequency of robot mentions of their name. Participants were also more inclined to take the robot's advice when it was targeted towards improving the performance of the whole group.

Micbot [25], a microphone robot used to shape the group dynamics and team performance in a game context, used back-channelling or encouraging the least active member to join the discussion. The results showed that the robot with encouraging behaviour and matching movements (instead of random movements or no movements) balanced participation, and improved group task performance.

Outside of the context of educational robots, adaptability has been used to increase engagement in group settings. Meng et al. [26] considered a robotic "Living Architecture System" (LAS) automatically adapting to a group's preferences in an experiment conducted within a museum exhibition. The reward for the reinforcement learning (RL) algorithm was user engagement, which was measured with ambient sensors. The results of the experiment showed that the RL-selected robot actions could increase engagement.

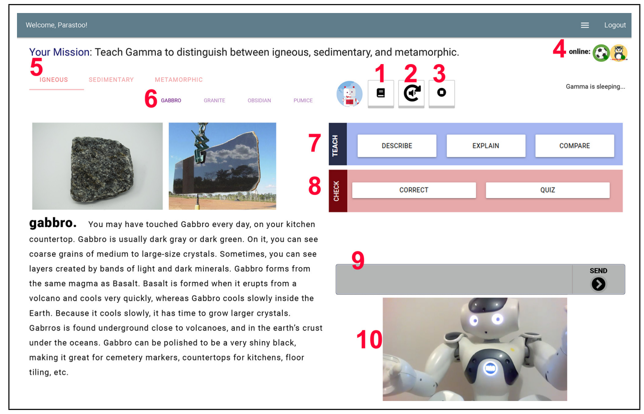

Fig. 1. Curiosity Notebook with Zoom, 1: Robot's Notes, 2: Repeat Request 3: End Teaching, 4: Online Users, 5: Categories, 6: Example Articles, 7: Teach Buttons, 8: Check Buttons, 9: Chat Window, 10: Zoom window of Gamma's video

In contrast to these prior works, our teachable robot is an active member of the team (the learner), influences group dynamics by encouraging and balancing social engagement and is personalized to the dyads' social engagement behaviours during the course of interaction.

III. System Overview

Our proposed system consists of three main components: a web application called Curiosity Notebook [13], [14], which we adapted for this study; an adaptive encouragement algorithm with audio data as the reward signal; and a NAO V6 [1] Humanoid robot, connected to the Curiosity Notebook.

A. Curiosity Notebook

The Curiosity Notebook enables users to read articles on various topics, structured as taxonomies, and teach the robot about them. Figure 1 shows the Curiosity Notebook interface during the task of rock classification. The users start the conversation by clicking the interactive buttons of their choice. There are two categories of buttons to interact with the robot: the *teaching* buttons and the *checking* buttons. Amongst the teaching buttons, users can choose the describe button to teach the agent about an object's features, the explain button to explain the feature and the compare button to discuss similarities or differences between rocks. If the user clicks any of the three teaching buttons, the user is asked to choose a rock, then the robot guides the interaction by asking different types of questions about the features of the chosen rock. If the user clicks the compare button, the robot picks a second rock after the user provided the name of the chosen rock. The conversations initiated by each button are associated with a state machine (an example is shown in Figure 2). No other interactive buttons can be clicked

[1] https://www.softbankrobotics.com/emea/en/nao

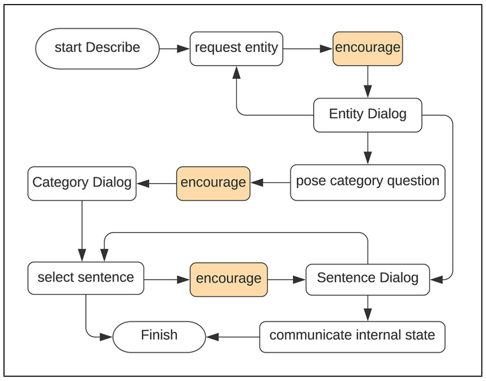

Fig. 2. The state machine that is executed after the Describe Button is clicked. 1) Request Entity: the robot will ask for the user to pick an object, 2) Entity Dialog: reflection on what the robot already knows about the selected object, 3) Pose Category Question: what category does the selected object belong to, 4) Category Dialog: robot reflecting on the new knowledge, 5) Select Sentence: robot asking for a sentence that describes one of the features of the rock, 6) Sentence Dialog: reflect on the sentence, 7) Communicate Internal State: robot communicating excitement about learning.

until the state machine for the current button has reached the termination state. The checking buttons are for testing the robot's learning. The users can choose to either quiz the robot by asking it to categorize an object or correct a previously learned concept in the robot's notes. The users can read and move between the categories and articles of each category at any time.

There are also three special-function buttons. Unlike the interactive buttons, the special-function buttons are all single action with no states and can be clicked at any time. The first button is the *Robot's Notes* button. Clicking on the Robot's Notes button brings up a notebook containing notes of all the knowledge that the robot has learned so far. The *Repeat* button can be used by users in case they misheard or did not hear the robot. The *End Teaching* button allows the users to indicate that they have finished the teaching process. Users can choose to stop teaching at any time.

The Curiosity Notebook supports group or one-on-one interaction. In our work, the group mode was used, in which case the Curiosity Notebook uses automated turn-taking to facilitate group interactions. In the turn-taking mechanism, one user will be the active teacher; only the active teacher can chat with the robot. The second user can click on different articles or use any of the special-function buttons. The turn-taking was shown to be effective in moderating the conversation during our initial pilot [13], otherwise, the user clicking the button faster would always be the active teacher, leaving some users discouraged to contribute at all. Turn-taking gave everyone an opportunity to contribute equally to the learning process, independent of their level of social engagement with their team. The Curiosity Notebook can be used on its own or connected to a physical robot. For our study, we used the Curiosity Notebook with a humanoid robot. The connection to the robot is described in section III-C.

B. User Audio Input

During the remote experiments, the voice activity was coded manually in real-time as we couldn't develop a reliable Zoom audio speaker diarization technique. The WebRTC[2] Voice Activity Detection (VAD) algorithm was used to capture the duration of voice activity for data analysis after the experiment.

C. The Humanoid Robot

The robot connects to the Curiosity Notebook via a Postgres database provisioned on the Heroku platform, such that each robot utterance is sent to the database and then to the robot. In addition to the text of the dialogue, each sentence is coded with an emotion: *happy*, *sad*, *neutral*, *bored* or *curious*. Most of the robot's sentences are neutral. The emotion code is happy if the robot just learned something (e.g., "I love learning about rocks"). Sentences are coded as sad if the robot makes a mistake during the quiz. The emotion is *curious* with a 50% probability if the robot is asking a question, and a sentence is coded *bored* if the notebook was idle for more than 2.5 minutes, in which case the robot asks participants to continue teaching. The emotion coding was also used to select the appropriate movement for the robot. The majority of the motions were predefined in naoqi, the operating system of NAO V6 robot, with a few additional nod motions, developed manually by joint manipulation for the neutral category. The movements are summarized in Table II. The utterances are spoken out by the robot while it is acting as per the corresponding movement.

IV. Adaptive Encouragement

Our goal is to increase group/social engagement between the team members using RL, hence, we modified the robot to deliver encouraging statements during the teaching conversation, specifically, after posing a question to the user and waiting for their answer. The states where encouragement could be provided are highlighted in Figure 2 in yellow. To create an adaptive encouraging system, we use a Q-Learning algorithm [27] to selectively choose if and how to encourage social engagement based on the users' talking activity. The reinforcement learning framework is defined as follows:

1) **States**: At each time step, the interaction can be in one of four states: a) there is not enough conversation, b) user A is dominating the conversation c) user B is dominating the conversation d) both users are equally and fully contributing.
2) **Actions**: There are five possible actions: encourage the active user, encourage the inactive user, encourage both users, pick a non-encouraging sentence, or say nothing.
3) **Rewards**: A weighted sum of the total time spent talking and the ratio of talking time of two users. The total reward is calculated as shown (4).

[2]To read more on the Open Source WebRTC project please visit `https://webrtc.org/`

TABLE I
ENCOURAGING SENTENCES

Encouragement Type	Sentence
encourage active user	"Make sure you let your partner know what you are thinking."
	"Try to explain your answer to your partner before telling me".
encourage other user	"Why don't you ask if your partner agrees with you?"
	"Why don't you ask what your partner thinks?"
encourage-both	"Why don't you two discuss this?"
	"Can you two talk amongst yourselves first?"
	"Let's do this together team!"
	"You can discuss it together first."
	"We can discuss it as a team!"
	"Do you want to discuss if you both agree?"

TABLE II
ROBOT'S MOVEMENTS

Emotion	Sample Movements
sad	scared, frustrated, hurt, sad, crying, getting shy, looking down
happy	laughs, giggles, clapping sounds, excited noises e.g. "Yoohoo"
neutral	hands and head movements while talking, sneezing, eye contact and turning its head to indicate listening
curious	recalling and thinking motions, e.g. scratching its head, putting hand under the chin.

A. Q-Learning Reward Calculation

To calculate the reward of the Q-Learning at each step, the algorithm reads the voice activity duration of both users (spk_1, spk_2) since the last reward calculation. The reward has two parts. The first part rewards more talking between both users (r_{talk}), which is calculated using the ratio of the total duration that both users spoke ($t_{talking}$) and the total time spent interacting with the robot (t_{total}). The second part rewards the ratio of how much the first user talks to how much the second user talks (r_{ratio}). Both parts of the reward (2) and (3) are calculated as their distance from the ideal value of 1 (i.e., the users spoke the full duration of the experiment and they spoke equally). The total is calculated by summing the two partial rewards and combining them as a cost ($c_{overall}$) by negation (4).

$$ratio = \frac{spk_1}{spk_2} \quad \text{and} \quad t_{talking} = spk_1 + spk_2 \tag{1}$$

$$r_{talk} = (1 - \frac{t_{talking}}{t_{total}})^2 \tag{2}$$

$$r_{ratio} = \begin{cases} (1 - ratio)^2, & \text{if } ratio \leq 1 \\ (1 - \frac{1}{ratio})^2, & \text{otherwise} \end{cases} \tag{3}$$

$$c_{total} = -(r_{talk} + r_{ratio}) \tag{4}$$

B. Encouraging Statements

There are three different types of encouraging statements, in addition to baseline non-encouraging statements, summarized in Table I. The active user is selected by the turn-taking mechanism. The robot utterances were created following techniques used in classrooms by teachers to promote user class participation [28] and encourage cooperative learning and group discussions in class [29]. Some utterances were targeted to individual users by calling their name to specifically engage them [30], or to ask them to engage their partner, while others were aimed at the group as a whole [31].

C. Q-Learning Parameters

The learning rate of Q-Learning (λ) controls how fast the algorithm adapts and was set to 0.6 for this experiment based on simulation results. Simulations were done assuming a simple deterministic dyad model as the reward signal for Q-Learning. In this model, every encouragement increases communication by 30 seconds (plus noise). The values of the discount rate (γ) and exploration rate (ϵ) were also selected based on simulation results. The value of the discount rate controls how much current audio input matters versus how much the next audio inputs will matter, set to 0.2. The exploration rate is the probability ϵ by which the algorithm disregards its Q function values and picks a random action, chosen to be 0.5 in our study. While the Q-Learning algorithm was tested in simulation to determine the hyperparameters, the algorithm was not trained before the user study and the only training data used were the data collected during the user study from the participants.

D. Research Questions

Our first hypothesis (**H1**) is that the robot encouraging teamwork will increase users' social engagement during the study, therefore increasing their communication. The second hypothesis (**H2**) is that the adaptive robot will have greater effects on (**H2a**) task engagement, (**H2b**) enjoyment and (**H2c**) learning in comparison to the baseline robot. Our third hypothesis (**H3**) is that adaptively encouraging teamwork will ensure both users contribute to the conversation more equally, without one user dominating the conversation.

V. EXPERIMENT

The goal of the experiment is to study the effect of adaptive encouragement on social engagement, task engagement, enjoyment and learning. Our user studies are done in dyads of participants. The curiosity of the robot, its learning speed, emotions and the questions it asked were kept constant across all the studies. The experiments were carried out in two conditions, baseline or adaptive. In the baseline condition, the robot does not encourage teamwork. In the adaptive condition, the robot encourages teamwork based on the audio input from the users. In the baseline condition, the robot uses neutral statements (e.g. "I don't know yet, I look forward to your answer") instead of the encouraging statements (Table I) to keep the length and the frequency of the robot's dialogues consistent across both conditions.

A. Participants

The user study was reviewed and received ethics clearance through the University of Waterloo Research Ethics Committee (ORE#40392). The participants were recruited through emails from the University of Waterloo and the social networks of the researchers. They were asked to sign up for a time slot, so they were randomly matched with other participants that signed up for the same time slot, except for three pairs who knew each other and chose to sign up for

TABLE III

DYAD-WISE DEMOGRAPHICS AND MEASUREMENTS

	Adaptive (n = 15)	Baseline (n = 13)	p-value
gender combinations	MM=2, FF=5 Mix = 8	MM=2, FF=4 Mix = 7	$\chi^2 \approx 0.14$ p=0.98
age difference (years)	4.13($\pm$3.38)	3.46($\pm$3.13)	t = 0.54, p= 0.6
dyad age sum (years)	51.47($\pm$4.6)	52.23($\pm$5.33)	t = -0.40, p = 0.69
started speaking English	11.8($\pm$9.21)	12.92($\pm$10.2)	t=-0.31 p=0.76
self-declared prior domain knowledge	4.27, SD=1.22	6, SD=2	t=-2.81, p=0.01

the same time slot. After signing up, the participants received instructions to join a Zoom call. A total of 68 participants were recruited for this study, forming 34 dyads. However, two dyads were excluded from analysis due to system lags and errors, whereas two more dyads were excluded due to one of the participants arriving significantly later than the scheduled time. Two more dyads were excluded after analyzing the data for technical issues and missing data. Of the remaining dyads, 15 dyads were assigned to the adaptive-encouraging condition, and 13 dyads were assigned to the baseline (non-encouraging) condition. The participants were all adults between the ages of 20 to 35. Table III summarizes the dyad-wise demographics.

B. Procedure

On the scheduled date and time, Gamma (the humanoid robot) and the researcher were both in the Zoom call (Figure 3) waiting for the participants. The experimental condition was assigned beforehand, and the participants were not aware of the condition assigned. Participants were given their login credentials for the Curiosity Notebook after they joined the call. All usernames were the preferred first name of the participants and it was the name Gamma used to address them. Upon logging in, they had to sign the consent form, after which they proceeded to fill in the pre-study surveys (described in the next section V-C Measures). After the pre-study survey, the participants watched a three-minute instruction video, in which they learned about the Curiosity Notebook buttons and how to teach Gamma. After the video, participants arrived at the teaching interface, they were asked to wait for their partners to also finish all the previous steps. Before they started teaching, they were offered the chance to ask clarification questions from the researcher, as no questions would be answered during the teaching period. The participants' audio was recorded throughout the experiment. The duration of the experiment was up to the participants and varied between 25 minutes and 69 minutes. After deciding to stop teaching Gamma, the participants were taken to the post-study questionnaires.

C. Measures

The pre-study surveys included demographics, familiarity and interest in robots, conversational agents and the topic of rock classification. They also included questions regarding the participant's interest in group work, and whether they knew their co-participant (group familiarity). The group familiarity survey was designed to measure some of the group characteristics that influence group work [32], [33]. In this study, the task was identical for all the participants,

Fig. 3. Zoom Call Interface, with two participants and Gamma

with no roles assigned, and no participants having prior experience with the system, therefore the focus was on group familiarity. We asked about the participants' familiarity with their co-participants in the study, both at an in-class and outside-class interaction level (of the three dyads who have known each other prior to the study, 2 were assigned to the baseline condition, and 1 to the adaptive condition), and their interest in group work, both measured on a 4-point Likert scale. Participants also answered a questionnaire on their feelings towards, and perceptions of Gamma (Godspeed Questionnaire [34]), in addition to their mood (Pick-A-Mood survey [35]). The last pre-study questionnaire was a knowledge test on rocks.

Post-study surveys included a questionnaire on the participants' experience. The participants answered questions on how much they enjoyed their experience, their interest in participating again, learning more about the topic they taught, and how much they enjoyed working with their partner. They were asked if they thought the robot was giving both of the participants a fair chance and encouraging group work. Other post-study surveys include another questionnaire on participants' perceptions of Gamma, another knowledge test, their motivation behind task completion (Intrinsic Motivation Inventory (IMI) [36] and Types of Motivation [37]).

VI. RESULTS

Talking Duration: As the participants decided when to end the experiment themselves, talking time is normalised by the experiment's overall duration. Figure 4 shows the normalised talking time for both conditions. The normalised talking time was significantly larger for dyads in the adaptive condition, confirming H1 ($t(26) = 2.24$, $p = 0.03$). There was no statistically significant difference in overall experiment duration between the two conditions ($M_{adapt} = 2539(s)$, $SD = 511.8$, $M_{base} = 2684.39(s)$, $SD = 622.69$, $t(26) = -0.66$, $p = 0.51$).

The ratio of talking time to total time also depended on the gender (highest when both dyad members identified as female $\beta = 0.13$, $t(17) = 2.46$, $p = 0.03$), and the dyad's interest in conversational agents ($\beta = 0.03$, $t(17) = 2.2$, $p = 0.04$).

Talking Trend: To analyze the effect of encouragement over time, we compared the trends of talking time between two conditions, shown in Figure 5. The adaptive condition results in higher talking time on average for the duration, but

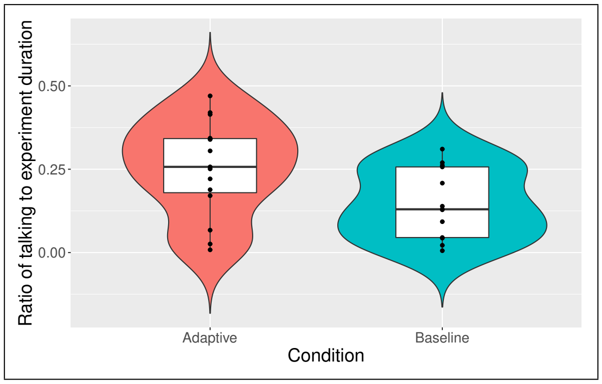

Fig. 4. Average Normalized Overall Talking Time

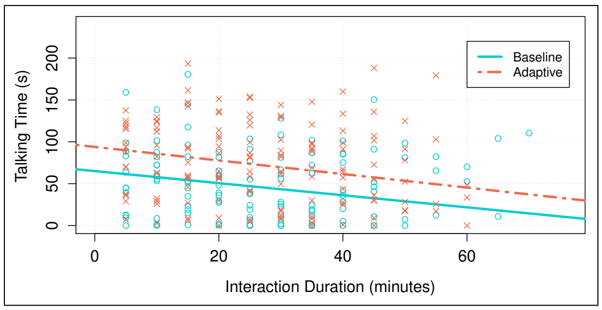

Fig. 5. Slope of Talking Time

the trend is downward for both conditions and the slopes are not significantly different ($M_{adapt} = -0.81$,$SD = 0.75$ and $M_{base} = -0.72$,$SD = 1.27$, $t(26) = -0.23$, $p = 0.82$). Qualitatively, the initial conversations were mostly questions on how the system works or wondering about the functionality of different buttons, regardless of the condition. Those types of conversations faded as the participants learned how to use the system.

Relative Participation: We hypothesized that the adaptive condition would result in a more equal division of speaking between participants (H3). To investigate our hypothesis, we define the speaking percentage of each user at any given period by the ratio of their speaking activity duration to the overall speaking activity in that period. The ideal value of speaking percentage is 50%. Figure 6 illustrates how much the speaking percentage of one participant (in each dyad) deviates from 50% (calculated in (5)), where a lower deviation indicates a more equal contribution. As shown in Figure 6, the average percentage difference is higher in the adaptive condition at the start, but it moves toward 0. In the baseline condition, however, the difference starts at a more desirable value but remains nearly constant, showing that the robot does not influence the balance between the two participants in the baseline condition. In the adaptive condition the slope is $M = -0.13$, $SD = 0.25$ while in the baseline condition slope is $M = -0.002$, $SD = 0.17$. The difference in slope is not statistically significant with $t(26) = -1.55$, $p = 0.13$.

$$||100 \times \frac{(spk_1)}{(spk_1 + spk2)} - 0.5|| \tag{5}$$

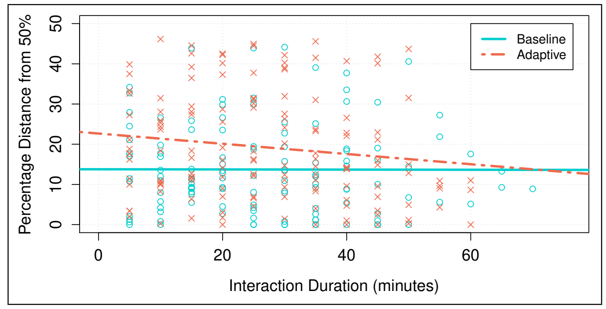

Fig. 6. Distance (5) of Participants' Speaking Percentage from 50%

Other Effects of the Experimental Condition: Task engagement, defined as the number of interactions participants had with the Curiosity Notebook, was higher in the adaptive condition. The number of buttons participants clicked ($F(1, 20) = 4.79$, $p = 0.04$), how many times they taught something ($F(1, 26) = 3.09$, $p = 0.06$), and how many articles they clicked ($F(1, 25) = 8.15$, $p = 0.01$) were all greater in the adaptive condition. The number of articles clicked was also correlated to participants' self-declared desire to teach ($F(1, 25) = 6.29$, $p = 0.02$).

To compare learning gains, we have to consider that despite the random assignment, the participants in the baseline condition self-reported higher knowledge about rocks (refer to Table III), although the pre-study knowledge test scores were not significantly higher for the dyads in the baseline condition. To evaluate if there was any learning gain, we compared the score improvement (post-score - pre-score). The difference in learning gain between conditions was not significant ($M_{adapt} = 1.74$, $SD = 3.47$, $M_{base} = 2.38$, $SD = 2.53$, $t(23) = -0.56$, $p = 0.58$). The improvement in the score was highly and negatively correlated with the pre-study test score($\beta = -0.76$, $t(23) = -6.29$, $p = 0$), which means the participants with less knowledge showed higher learning gains. Secondly, as expected, the improvement was positively correlated with participants' interest in rocks ($\beta = 0.37$, $t(23) = 2.62$, $p = 0.02$). There was a weaker negative correlation between how much the participants thought they knew about rocks and their score improvement ($\beta = -0.4$, $t(23) = -1.71$, $p = 0.1$). Regardless of condition, the average knowledge test score improved after the study by 1 point (out of 6).

Examining our second hypothesis (H2b) on increased dyad enjoyment (how much the participants enjoyed working with their partner), we report no significant difference between the conditions.

Participant perceptions of the robot: From the post-study Godspeed questionnaire, one characteristic of the robot differed significantly between the two conditions. The participants in the adaptive condition found Gamma less pleasant (sum of the dyad's perception on a 5-point scale), and this is affected by both the experimental condition and their interest in rocks ($M_{adapt} = 7.47$, $SD = 1.36$ and $M_{base} = 8.46$, $SD = 0.88$, $p = 0.03$).

From the post-study experience survey (section V-C), most participants in both conditions considered the robot to be fair, in the baseline condition this was due to the turn-taking mechanism, and how the robot gave both participants their turns ($M_{adapt} = 1.93$, $SD = 0.26$ and $M_{base} = 1.58$, $SD = 0.79$, $t(25) = 3.13$, $p = 0.11$). In the adaptive condition, prompts for discussion among dyad members were considered a sign of fairness in 8 instances. In the baseline condition, 5 participants thought the robot was not fair, 2 of which were due to the robot interrupting them. Turn-taking and instance of the robot talking over participants are some of the current challenges of human-robot-collaboration and many of the cues used in face-to-face interaction were not applicable in the virtual setting of our experiment [38].

When it comes to participants' perception of whether the robot was encouraging team work, the participants in the adaptive condition found the robot more encouraging ($M_{adapt} = 1.93$,$SD = 0.28$ and $M_{base} = 1.33$, $SD = 0.65$, $t(24) = 3.13$, $p = 0.004$). Participants in the baseline condition interpreted forced turn-taking as encouragement of teamwork in 6 instances, while others cited unrelated reasons, such as the experiment being fun, or the need to figure out how the system works initially.

The majority (37) of participants in both conditions enjoyed their teaching experience ($M_{adapt} = 2.9$, $SD = 0.9$ and $M_{base} = 2.8$, $SD = 1$, on a 4 point Likert scale). In both conditions, the most frequently mentioned reason for enjoyment was that it was an "exciting experience". Some characteristics of voice and appearance of the humanoid robot (Gamma) were mentioned by the participants, such as "enthusiastic robot", "fun movements", "sense of wonder in Gamma's face", in addition to some non-physical characteristics such as "good notes" and "fast learner". In both conditions, "slow" and "repetitive prompts" were the most common reason behind participants' dissatisfaction. The results from Pick-A-Mood scale [35] show that 20% of participants in the adaptive condition felt tense before the experiment, which dropped to 10% after the experiment. However, in the baseline condition feeling tense increased from 8% before the experiment to 19% afterward. Feeling excited went from 0% to 33% in the adaptive condition and 8 to 23% in the baseline condition. Lastly, more participants in the adaptive condition thought Gamma was excited (83%) in comparison to the baseline (53%).

There is little difference between the two conditions in terms of the participants' perception of their teaching level, or whether they thought of Gamma as a good student. The encouragement of collaboration was mentioned as a characteristic by one participant in the adaptive condition. A participant in the baseline condition called the robot socially unaware and "as if there were no humans in the interaction".

In participants' general feedback, recurring requests were to make the interaction faster, make the Curiosity Notebook easier to interact with, and allow participants more freedom or give them more options. Feedback related to changes in the robot were similar in both conditions, mostly to tone down the excitement, especially laughing. There were also a few reports of Gamma being hard to hear over the call or the voice getting interrupted.

Participant Discussion During the Study: The audio transcriptions of the conversations between the participants recorded during the study were coded using 10 different categories: teaching plans - long term, teaching plans - immediate, questions about interface, discussion about articles, comments on Gamma, comments on teaching, comments on learning, unrelated conversation, unanswered initiation, and conversation about when to end the teaching session. Based on this coding, dyads in the adaptive condition had significantly more discussions compared to those in the baseline condition ($M_{adapt} = 56.64$,$SD = 29.79$ and $M_{base} = 24.92$, $SD = 15.43$, $t(24) = 3.32$, $p = 0.003$). This result differs from the length of the conversation as the discussions were coded irrelevant of the length but based on the change of topic or in case the robot gave a new prompt. The disparity between the number of discussions initiated by each participant was higher in the adaptive condition however this difference was not significant.

Considering the topic of discussion, short-term plans for teaching can happen at every step and included letting the other participant know of the chosen rock or sentence or the immediate next action. Encouragements by the robot were successful in causing the participants in the adaptive condition to discuss their next action with their teammate at a significantly higher level ($M_{adapt} = 26.43$,$SD = 15.95$ and $M_{base} = 4.42$, $SD = 4.03$, $t(24) = 4.64$, $p = 0$). The number of discussions on overall plans was higher in the adaptive dyads as well however this difference was not significantly different. The participants in the adaptive condition were more likely to comment on the teaching of their teammates or proactively offer help and suggestion $M_{adapt} = 3.36$,$SD = 2.53$ and $M_{base} = 1.33$, $SD = 1.87$, $t(24) = 2.82$, $p = 0.03$).

In every other category of discussions, the dyads in the adaptive condition had more discussions, however, the difference was not statistically significant.

VII. Discussion

Our study showed that encouraging group/social engagement increased communication as well as participants' task engagement and exploration. The results also show the potential for an adaptive encouraging robot to create more balanced participation between the dyad members. However, we were not able to show any effect of increased communication on learning outcomes in this short intervention.

Limitations: One of the limitations of this study was that participants' group engagement was only measured through their verbal communication and did not include factors such as emotions and non-verbal behaviours. Additionally, the robot's action repertoire could be expanded by including other forms of encouragement such as eye contact, or other verbal behaviours such as back-channeling [25]. Performing a long-term study will allow for the measurement of learning gains and long-term engagement [1], [3]. The suitability of the system for the tested age-group also needs to be further validated. This system may be more or less suitable for different age groups. furthermore, our study was performed remotely, introducing the possibility of delays and poor audio/video quality that could have affected the engagement negatively. The results might differ for an in-person study. However, the remote study results could be relevant for future studies, given the emergence of online learning platforms.

VIII. Conclusions

This paper proposed adaptive social robots to increase communication within groups in the context of robot teaching. We performed a user study to examine the effect of adaptive encouragement on social engagement, task engagement, enjoyment and learning. Robot encouragement increased team communication, leading to higher social engagement in addition to higher task engagement. Encouraging social/group interactions with an adaptive robot might have the potential for improving learning gains in long-term interactions, which should be investigated in future work.

Further experiments are also needed to investigate the benefit of adaptive over non-adaptive (random) encouragement. Additionally, the participant pool should be carefully selected to prevent the topic knowledge gap observed between participants in the two conditions of our study.

References

[1] T. Belpaeme, J. Kennedy, A. Ramachandran, B. Scassellati, and F. Tanaka, "Social robots for education: A review," *Science Robotics*, vol. 3, no. 21, Aug 2018.

[2] A. Jacq, S. Lemaignan, F. Garcia, P. Dillenbourg, and A. Paiva, "Building successful long child-robot interactions in a learning context," *HRI*, vol. 2016-April, pp. 239–246, 2016.

[3] S. Lemaignan, A. Jacq, D. Hood, F. Garcia, A. Paiva, and P. Dillenbourg, "Learning by Teaching a Robot: The Case of Handwriting," *IEEE Robotics and Automation Mag.*, vol. 23, no. 2, pp. 56–66, 2016.

[4] C. C. Chase, D. B. Chin, M. A. Oppezzo, and D. L. Schwartz, "Teachable agents and the protégé effect: Increasing the effort towards learning," *J. Science Education and Technology*, vol. 18, no. 4, pp. 334–352, Aug 2009.

[5] E. Walker and W. Burleson, "User-centered design of a teachable robot," in *Intelligent Tutoring Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 243–249.

[6] E. Alexopoulou and R. Driver, "Small-group discussion in physics: Peer interaction modes in pairs and fours," *J. Research in Science Teaching*, vol. 33, no. 10, pp. 1099–1114, 1996.

[7] V. Du Preez and V. Barnes, "Encouraging engagement," in *Conf. on Design, Development and Research*. Design, Development & Research Conference- CPUT, 2012, p. 214.

[8] M. Kitchen, "Facilitating small groups: how to encourage student learning," *The Clinical Teacher*, vol. 9, no. 1, pp. 3–8, 2012.

[9] S. Chandra, R. Paradeda, H. Yin, P. Dillenbourg, R. Prada, and A. Paiva, "Affect of robot's competencies on children's perception," in *AAMAS*, 2017, pp. 1490–1492.

[10] N. Lubold, E. Walker, and H. Pon-Barry, "Effects of voice-adaptation and social dialogue on perceptions of a robotic learning companion," in *HRI*, March 2016, pp. 255–262.

[11] P. Baxter, E. Ashurst, R. Read, J. Kennedy, and T. Belpaeme, "Robot education peers in a situated primary school study: Personalisation promotes child learning," *PLOS ONE*, vol. 12, no. 5, pp. 1–23, 2017.

[12] J. Kory and C. Breazeal, "Storytelling with robots: Learning companions for preschool children's language development," in *RO-MAN*, Aug 2014, pp. 643–648.

[13] E. Law, P. Baghaei Ravari, N. Chhibber, D. Kulic, S. Lin, K. D. Pantasdo, J. Ceha, S. Suh, and N. Dillen, "Curiosity notebook: A platform for learning by teaching conversational agents," in *CHI Extended Abstracts*, 2020, p. 1–9.

[14] E. Law, K. J. Lee, and P. Baghaei Ravari, "The Curiosity Notebook: A Platform for SupportingLearning with Teachable Robots," in *HRI 2020 Workshop on Exploring Creative Contents in Social Robots*, 2020.

[15] F. Tanaka and S. Matsuzoe, "Children teach a care-receiving robot to promote their learning: Field experiments in a classroom for vocabulary learning," *J. Hum.-Robot Interact.*, vol. 1, no. 1, pp. 78–95, July 2012.

[16] T. Fong, I. Nourbakhsh, and K. Dautenhahn, "A survey of socially interactive robots," *Robotics and Autonomous Systems*, vol. 42, no. 3, pp. 143 – 166, 2003.

[17] C. Breazeal, *Designing Sociable Robots*. Cambridge, MA, USA: MIT Press, 2002.

[18] R. van den Berghe, J. Verhagen, O. Oudgenoeg-Paz, S. van der Ven, and P. Leseman, "Social robots for language learning: A review," *Review of Educational Research*, vol. 89, no. 2, pp. 259–295, 2019.

[19] C. Zaga, M. Lohse, K. P. Truong, and V. Evers, "The effect of a robot's social character on children's task engagement: Peer versus tutor," in *Social Robotics*, A. Tapus, E. André, J.-C. Martin, F. Ferland, and M. Ammi, Eds. Springer, 2015, pp. 704–713.

[20] F. Tanaka, K. Isshiki, F. Takahashi, M. Uekusa, R. Sei, and K. Hayashi, "Pepper learns together with children: Development of an educational application," *Humanoids*, vol. 2015-Decem, pp. 270–275, 2015.

[21] D. Hood, S. Lemaignan, and P. Dillenbourg, "When children teach a robot to write: An autonomous teachable humanoid which uses simulated handwriting," in *HRI*, 2015, p. 83–90.

[22] S. Strohkorb, E. Fukuto, N. Warren, C. Taylor, B. Berry, and B. Scassellati, "Improving human-human collaboration between children with a social robot," in *RO-MAN*, 2016, pp. 551–556.

[23] E. Short and M. J. Mataric, "Robot moderation of a collaborative game: Towards socially assistive robotics in group interactions," in *RO-MAN*, 2017, pp. 385–390.

[24] S. Strohkorb Sebo, L. L. Dong, N. Chang, and B. Scassellati, "Strategies for the inclusion of human members within human-robot teams," in *HRI*, 2020, p. 309–317.

[25] H. Tennent, S. Shen, and M. Jung, "Micbot: A peripheral robotic object to shape conversational dynamics and team performance," in *HRI*, 2019, pp. 133–142.

[26] L. Meng, D. Lin, A. Francey, R. Gorbet, P. Beesley, and D. Kulic, "Learning to engage with interactive systems: A field study," *CoRR*, vol. abs/1904.06764, 2019.

[27] C. J. C. H. Watkins, "Learning from delayed rewards," Ph.D. dissertation, King's College, Cambridge, UK, May 1989.

[28] K. Sherab, "Strategies for encouraging behavioural and cognitive engagement of pre-service student-teachers in bhutan," *Educational Action Research*, vol. 21, no. 2, pp. 164–184, 2013.

[29] R. M. Gillies, "Teachers' and students' verbal behaviours during cooperative and small-group learning," *British J. Educational Psychology*, vol. 76, no. 2, pp. 271–287, 2006.

[30] B. Maloch, "Scaffolding student talk: One teacher's role in literature discussion groups," *Reading Research Quarterly*, vol. 37, no. 1, pp. 94–112, 2002.

[31] S. Roberts, "Using practitioner research to investigate the role of the teacher in encouraging student interaction within group work," *Nurse Education Today*, vol. 28, no. 1, pp. 85 – 92, 2008.

[32] D. Jung and J. Sosik, "Effects of group characteristics on work group performance: A longitudinal investigation," *Group Dynamics*, vol. 3, no. 4, pp. 279–290, Dec. 1999.

[33] J. Janssen, G. Erkens, P. A. Kirschner, and G. Kanselaar, "Influence of group member familiarity on online collaborative learning," *Computers in Human Behavior*, vol. 25, no. 1, pp. 161 – 170, 2009.

[34] C. Bartneck, E. Croft, and D. Kulic, "Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots," *Int. J. Social Robotics*, vol. 1, no. 1, pp. 71–81, 2009.

[35] P. Desmet, M. Vastenburg, and N. Romero, "Mood measurement with pick-a-mood: Review of current methods and design of a pictorial self-report scale," *J. Design Research*, vol. 14, no. 3, pp. 241–279, 2016.

[36] E. Deci, H. Eghrari, B. Patrick, and D. Leone, "Facilitating internalization: the self-determination theory perspective." *J. personality*, vol. 62 1, pp. 119–42, 1994.

[37] F. Guay, R. Vallerand, and C. Blanchard, "On the assessment of situational intrinsic and extrinsic motivation: The situational motivation scale (sims)," *Motivation and Emotion*, vol. 24, pp. 175–213, 09 2000.

[38] S. Andrist, D. Bohus, B. Mutlu, and D. Schlangen, "Turn-taking and coordination in human-machine interaction," *AI Magazine*, vol. 37, no. 4, pp. 5–6, Jan. 2017.