Efficacy of a 'Misconceiving' Robot to Improve Computational Thinking in a Collaborative Problem Solving Activity: A Pilot Study

Utku Norman[1], Alexandra Chin[2], Barbara Bruno[1] and Pierre Dillenbourg[1]

***Abstract*— Robot-mediated learning activities are often designed as collaborative exercises where children work together to achieve the activity objectives. Although miscommunications and misunderstandings occur frequently, humans, unlike robots, are very good at overcoming them and converging to a shared solution. With the aim of equipping a robot with these abilities and exploring its effects, in this article we investigate how a humanoid robot can collaborate with a human learner to construct a shared solution to a problem via suggesting actions and (dis)agreeing with each other. Concretely, we designed a learning activity aiming to improve the computational thinking skills of children, in which the robot makes suggestions on what to do, that may be in line with what the human thinks or not. Furthermore, the robot may suggest wrong actions that could essentially prevent them from finding a correct solution. Via a pilot study conducted remotely with 9 school children, we investigate whether the interaction results in positive learning outcomes, how the collaboration evolves, and how these relate to each other. The results show positive learning outcomes for the participants in terms of finding better solutions, suggesting that the collaboration with the robot might have helped trigger the learning mechanisms.**

***Index Terms*— human-robot interaction; mutual understanding; collaborative learning; computational thinking.**

I. Introduction

Social robots as embodied agents have demonstrated a potential for education, where they have been given the role of a tutor, or a peer i.e. a learning companion; typically with the knowledge of an exact solution or the underlying educational concepts [1], [2]. In the educational situation of *collaborative learning*, human participants with no knowledge of the correct solution attempt to learn something together, where their interactions are anticipated to induce learning even though there is no guarantee that these kind of processes will happen [3]. However, the interactions among participants can be 'designed' in order to make it more likely that these productive processes will in fact occur, as they are shaped by the activity and the environment: it is the shared experience and the required *effort* to construct a *mutual understanding* together that may trigger the learning [4], [5].

In this work, we investigate the effects of using a social robot in a collaborative learning activity as a peer for a human.[3] The robot by design has no knowledge of a correct solution. We are interested to see if it can help trigger learning mechanisms by taking actions, and thereby support the effort to build a mutual understanding about the task, even without the knowledge of a correct solution. In particular, we assess the efficacy of the collaboration by considering the learning outcomes based on individual performance in the pre- and post-test (RQ1), the evolution of performance while collaborating with the robot (RQ2), and the link between them (RQ3).

TABLE I

Research Questions and Hypotheses.

No.	Research Question / Hypothesis
RQ1	*How are the learning outcomes after collaborating with the robot?*
H1.1	A participant provides a valid solution more in the post-test than the pre-test.
H1.2	A participant provides a correct solution more in the post-test than the pre-test.
H1.3	A participant provides a better solution (closer to a correct solution) more in the post-test than the pre-test.
RQ2	*How does performance in the task evolve during collaboration with the robot?*
H2.1	A participant submits better solutions (closer to a correct solution) later than earlier.
H2.2	A participant suggests correct actions more later than earlier.
H2.3	A participant (dis)agrees more with (in)correct robot suggestions later than earlier.
RQ3	*How does the evolution of performance in the task link to the learning outcomes?*
H3.1	The more a participant's submissions improve, the better are the learning outcomes.
H3.2	The more a participant's suggestions improve, the better are the learning outcomes.
H3.3	The more a participant's (dis)agreements improve, the better are the learning outcomes.

Table I presents our research questions and hypotheses. In RQ1, we postulate that the collaboration with the robot would have a positive impact in terms of the learning outcomes, as the effort to build a shared solution can help the participant realize his/her misconceptions. Hence, we hypothesize that they perform better in the post-test than in the pre-test, where the tests are designed as individual exercises, with the collaborative activities as counterparts that are done together with the robot. The difference between 'valid' and 'correct' is explained in Sec. III-A. In RQ2, we postulate that while collaborating with the robot, the human would realize his/her misconceptions and improve his/her actions through time. In RQ3 we expect that the more a participant improves

*This project has received funding from the European Union's Horizon 2020 programme under grant agreement No 765955. A. Chin has been supported through a Summer@EPFL fellowship. Ethical approval was granted by the EPFL Human Research Ethics Committee, No. 030-2021/06.04.2021. We are grateful to the schools that made this work possible.

[1]CHILI Lab, EPFL, Lausanne, Switzerland {name.surname}@epfl.ch

[2]Wellesley College, Dept. of Computer Science, Massachusetts, USA ac5@wellesley.edu

[3]The code that represents the activity and governs the interaction with the robot is publicly available online, from the GitHub Repositories https://github.com/utku-norman/justhink_world for the activity, and https://github.com/utku-norman/justhink-ros for the interaction.

during the collaborative activity, the better are the learning outcomes. Previous studies e.g. [6] revealed that the link between the performance in the collaborative activity and the learning outcomes is not trivial: they are not directly, and not necessarily linearly correlated. Success and failure can be productive or not in terms of the learning outcomes, as shown in [7]. Thus, in RQ3 we specifically consider the performance *change* during collaboration, and investigate whether it is reflected in the learning outcomes as a *change* from pre- to post-test.

II. Related Work

Social robotic agents have a potential to become a part of educational environments by undertaking a unique position that extends their functional purpose with personal and social dimensions of interaction [1]. Among agents that provide support with social interaction, physically embodied robots tend to have a higher impact on learning and become more effective for the desired changes in behavior compared to virtual pedagogical agents [1]. To leverage the social dimension, these robots can behave in ways to encourage greater effort towards learning. For instance, within the "learning by teaching" paradigm [8], the robot takes the role of a novice needing help from the human; e.g. to enhance their vocabulary via a robot controlled by Wizard-of-Oz [9], or to improve handwriting via an autonomous robot [10]. Such a role is not very straightforward for a human to take on, as it might not be convincing or even doable, which highlights the unique position of robots. Our work explores another way to benefit from this added social dimension which is by having a physical robot autonomously participate in a collaborative problem solving process together with the child. The need to converge on a shared solution engages the child in a meta-cognitive relationship with the robot that involves trying to understand why the robot acts the way it does. The socio-cognitive conflict [11] likely to arise from this relationship can promote learning [12].

Regarding their role in the interaction, social robots have been used as teachers or tutors to guide children's learning of skills such as a second language [13] (robots *teaching* humans). Conversely (humans *teaching* robots), have been utilized as teachable agents, namely as "surrogate pupils" for children to teach [14], by concentrating on i) the benefits for the robot akin to the "learning from demonstration" paradigm [15]; or ii) for the children's learning as in [10]. Our work focuses on the benefits for the child: by collaborating with a robot to solve a problem and converging to a shared solution, we aim at improving the computational thinking skills of the child. To be effective in achieving the learning goals, we need to investigate how the robot should behave within its role [1]. The robot's behavior is informed by what it knows about the activity and the human. Therefore, our work investigates to what extent the robot needs to model the activity and the human, to produce the desired learning outcomes. For this, we start with the activity, and investigate whether the robot could be effective even while misconceiving the problem.

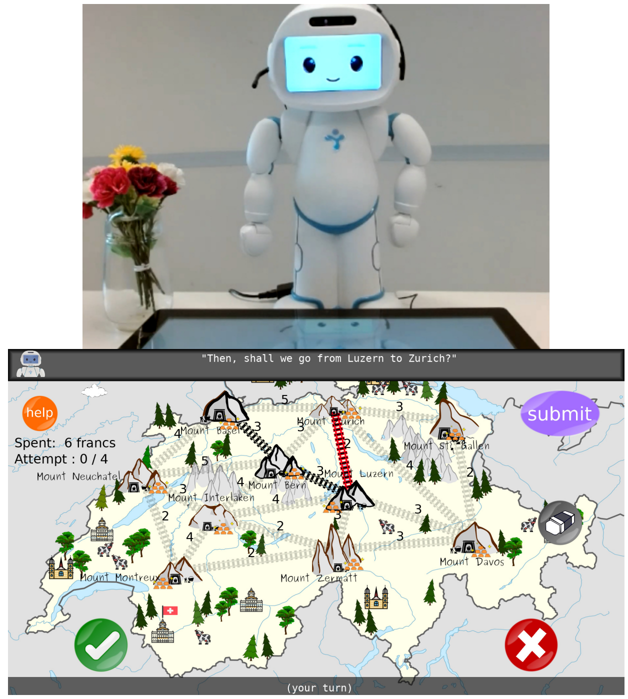

Fig. 1. Robot and first collaborative activity as seen by the participant.

III. Design

A. Instructional Design

We design the learning experiences, that fundamentally shape the participant's interaction with the robot, by the "backward design" approach [16]: we i) identify the desired results i.e. the learning goals, ii) determine acceptable evidence that indicate the desired results are achieved, and iii) design the activities that can lead to learning.

1) Learning goals: The desired results in this study stem from recent research in learning sciences, which emphasizes a need to introduce activities to develop the Computational Thinking (CT) skills of children earlier in schools [17]. Accordingly, we target school children aged 9-12 years old, with the aim to improve their CT skills, by applying abstract and algorithmic reasoning to solve a problem on networks. We choose the minimum-spanning-tree problem on networks as the underlying objective of the task, that children are not expected to be familiar with. Thus, our desired result is that after completing this task, a participant will be able to *correctly choose a subset of connections on a given network*, so that i) all nodes are connected to each other by some path, and ii) the total cost on these connections is minimized. Concretely, given an instance of the minimum-spanning-tree problem, a participant will be able to: (**LG1**) *identify* a valid solution for the given instance, and (**LG2**) *construct* a correct solution to the instance. These goals target "understand" and "create" levels of learning objectives, respectively, in the (revised) Bloom's taxonomy [18].

2) Assessment of learning: As evidence about the desired results, we measure the learning outcomes by comparing the responses of the participant in the pre-test and the post-

TABLE II
THE SEQUENCE OF ACTIVITIES IN THE HUMAN-ROBOT PEDAGOGICAL SCENARIO.

Activity	What is the human supposed to do?	What does the robot do?	Time (min)
Intro.	Listen to the robot, while looking at the screen that displays the context of the game.	Introduce the context and the goal of the game.	2
Tutorial	Learn how to interact with the activity on screen: i) *pick* a connection (i.e. an edge), ii) *clear* all the selections, and iii) *submit* a solution.	Guide the human through the tutorial. Explain the available actions.	3
Pre-test 5 problem instances	In each instance, connect the mines by spending as little as possible for the network. Upon pressing the submit button, confirm the submission or continue modifying the selection. After confirmation, move on to the next instance.	Motivate the human to solve the problem instances in the test, state the (dual) goal of connecting all and spending as little as possible.	$\leq$ 5
Collab. 2 problem instances	Collaborate with the robot to connect the mines by spending as little as possible for a given network, by *suggesting* which connection to pick and *(dis)agreeing* in turns: selecting only if agreed. The human or the robot can submit $\leq$ 4 times per problem instance. Upon submission, receive feedback on whether the selection is correct.	When it is the robot's turn, suggest and (dis)agree according to a sub-optimal strategy based on greedily traversing the network; submit if there are no outgoing connections.	$\leq$ 30
Post-test 5 instances	Solve the problem instances in the post-test, identical to the pre-test.	Motivate the human to solve the problem instances.	$\leq$ 5
Interview	Answer questions from the experimenter.	Robot is not present.	$\leq$ 15

test. We design the tests as a sequence of instances of the minimum-spanning-tree problem, to be solved individually by the participant. We assess the learning goal of identifying a valid solution (LG1) by checking whether the responses (i.e. solutions) of the participant to the tests are *feasible* solutions or not.[4] We evaluate the learning goal of constructing a correct solution (LG2) by checking i) whether the submitted solutions are *optimal* solutions or not,[5] and ii) how far the solutions are from the optimal solutions in terms of their cost.

3) Pedagogical scenario: As an experience that can make the desired results happen, we design a human-robot interaction scenario where the robot orchestrates a sequence of activities as described in Table II. The robot first introduces the scenario in the context of a game, which takes place on a fictional map of Switzerland with rare metal mines located in mountains, see Fig. 1. The goal of the game is to build a railway network to help the miners go from any mine to any other, and spend as little money as possible to build these railways. In each of the *collaborative activities*, the robot and the participant solve a problem instance together. The goal of the game is the same as the *tests*, where the participant individually solves a series of problem instances. At the end of the interaction, an experimenter conducts a semi-structured *interview* with the participant.

a) Collaborative activities: The human and the robot as (same-status) peers collaboratively construct a solution by deciding together which tracks to build, where each track, if built, connects one mine to another. The cost of each track is visible to both the human and the robot. The human and the robot take turns in suggesting to build a specific connection, where the other agrees or disagrees with this suggestion, and then makes a new suggestion. A track will be built only if it is suggested by one and agreed by the other. The human as well as the robot can submit the current solution and receive feedback on whether it is a correct solution or not: if it is a correct solution the activity ends, the solution is cleared otherwise. In each problem instance, a solution can be submitted up to four times. The human and the robot sequentially work on two problem instances of the same complexity, composed of 20 edges and 10 nodes. Fig. 1 shows the first network. For instance, a 'valid' solution for the problem is selecting all the possible connections: this is however not a 'correct' solution as it contains many redundant connections. The second network is obtained by transforming the layout of the first network, renaming the mines, and modifying the costs on the edges.

b) Tests: The pre-test and the post-test consist of five different instances of the minimum-spanning-tree problem, in the same context of the collaborative activities. The tests are identical and no feedback is given on the submitted solutions. The participant is asked to confirm while submitting. The networks of the instances in a test are created from the same underlying network structure, composed of 12 edges and 7 nodes, by transforming (e.g. mirroring, rotating, scaling, and shifting) the layout and modifying the cost (e.g. by doubling or adding a constant). Thus, the problem instances in the tests are of same complexity, and simpler than the instances in the collaborative activities.

c) Interview: In order to get a deeper insight on the impressions of the participants regarding the activity and their interaction with the robot, at the end of the scenario, an experimenter interviews the participants with 10 questions, listed in Table III. The first half of the questions investigate how the participant perceived his/her own strategy and the strategy of the robot: how the participant and the robot made suggestions and how did they agree or disagree with each other. Then, we ask about the participant's perception of the optimality and the autonomy of the robot, and finally if the participant and the robot made guesses about the other while taking actions. About guesses, the robot behavior is designed such that it occasionally asks the human to make a guess, see Sec. III-B. Since we know by design of the pedagogical scenario and the robot behavior the ground truth for some of these questions, the responses give us a chance

[4] A *feasible* solution to an instance of the minimum-spanning-tree problem connects all the nodes in the network of that instance to each other by a path.

[5] An *optimal* solution is a feasible solution with the minimum possible cost.

TABLE III

THE QUESTIONS IN THE INTERVIEW.

No	Question
1	How did you decide on which connections to pick?
2	How did you decide on your suggestions to the robot? Is it following some rule?
3	How does the robot decide on its own suggestions? Does it know the solution?
4	How did you decide whether to agree or disagree with the robot's suggestions?
5	How did the robot agree or disagree with your suggestions?
6	How good were the suggestions of the robot? Were they correct all the time? Incorrect?
7	How independent do you think was the robot? Do you think it is controlled by someone?
8	What do you think was going on inside the robot's mind?
9	Did you select what you think the robot will do, when the robot asks you to make a guess?
10	Do you think the robot made guesses about what you would do?

to reveal discrepancies in expected and observed behavior of the participants.

B. Robot's Role and Behavior

We design the robot behavior such that i) it is completely autonomous to orchestrate the interaction without a need for intervention by an experimenter, and ii) it works together with the human to construct solutions, in such a way that elicits misconceptions and misunderstandings about the other and the activity.

1) Throughout the pedagogical scenario: The robot fully automates the entire interaction by introducing the context and the goal of the game, (un)pausing the game, and moving the displayed activity to the next until the interaction concludes. Its behavior includes verbal explanations and support to motivate the human to try his/her best, facial expressions to convey emotions such as excitement, and gestures like pointing to the participant or looking at the screen.

2) During the collaborative activities: Unbeknownst to the participant, the robot does not know how to solve the problem correctly: it has the wrong conception of feasible and optimal solutions (i.e. no LG1 or LG2). The robot acts in a naive, simplistic and convincible manner, making suggestions greedily from a node it assumes that they are at. This results in a sub-optimal strategy, that will traverse the map in a local-greedy manner, and end up at a visited node: hence selecting a sub-network that contains a cycle, which is a sub-optimal solution. What the robot implicitly and functionally "believes" about the activity and its human counterpart can be summarized as follows: i) "We are at a node and we move as we select edges.", ii) "If we have nowhere new to select from the node we are at, we are done and we should submit.", iii) "My strategy is correct and your strategy is incorrect.", and iv) "If you are persistent, then you are correct." These beliefs are manifested in the robot's actions: when it is the robot's turn, the robot suggests to pick one of the cheapest outgoing edges from a specific node, or submits if there is no edge to select from that node. When it is the human's turn, the robot either i) asks the human to suggest what they should do (with 80% probability), or ii) asks the human to make a guess on what it (the robot) would do (20%). The node from which the robot selects edges is not revealed to the human, and moves as edges get suggested and selected. While responding to the human's suggestions, the robot agrees if i) it is exactly the same edge it would pick or ii) it is the second time the human suggests that edge; it disagrees otherwise: thus the robot is in a sense "convincible", a trait which also prevents the interaction from getting stuck if a human is to insist on selecting a particular connection.

IV. METHODS

1) Measuring learning outcomes: We quantify the learning outcomes separately around each learning goal, by assessing the quality of the solutions of a participant in the pre-test and the post-test. For LG1, we compute the fraction of feasible solutions in the tests: from 0% (none of the solutions is feasible) to 100% (all solutions are feasible and hence valid). For LG2, we calculate the fraction of optimal solutions in the tests: from 0% (none of the solutions are correct) to 100% (all solutions are optimal and hence correct). As a finer assessment of performance in terms of how far a feasible solution is from an optimal solution, we define `error` as the difference between the cost of the solution and the cost of an optimal solution, normalized by the cost of the optimal solution. We quantify the overall performance in a test by the *average error* per problem in that test. We quantify the change in the quality of responses in the post-test, compared to pre-test, on the basis of the `learning gain` of a participant, defined by the relative difference of the average error in pre-test and post-test.

2) Measuring performance in collaborative activities: Each collaborative activity allows submitting several solutions for a problem instance. Thus, we evaluate each solution in a collaborative activity separately, as the pair's best attempt for a correct solution in that activity: this may be any of the solutions, and necessarily the last solution if it is optimal. We use the *lowest error* of the solutions submitted by the human as the measure of overall performance of a participant for each collaborative activity: the lower the error, the higher the performance, with 0% meaning that an optimal solution was found. Note that the error is only computed for the feasible solutions submitted by the human participant. We measure the `performance gain` as the change of performance between the collaborative activities, by computing the relative difference of the lowest error in the first collaborative activity and in the second.

3) Characterizing actions in terms of optimality: When the human and the robot are constructing a solution together, and while the human is individually constructing a solution in the tests, the actions taken can be qualified as *optimal* or not. To assign a quality label for every action, we consider the set of possible optimal solutions that can be constructed from the current solution state, by only adding more connections (or submitting). Thus, a submit action is optimal if and only if the submitted solution is optimal. A clear action is optimal, if there is no optimal solutions that can be constructed

Fig. 2. The study setup. The participant interacts remotely from a school.

from the current state. We label a suggestion as optimal, if the suggested connection is part of at least one possible optimal solution available from that state (and as sub-optimal otherwise). An agreement is optimal if the suggested edge is part of at least one of the possible optimal solutions, and a disagreement is optimal if the suggested edge is not part of any of the possible optimal solutions.

Optimality in this sense is towards finding an optimal solution in minimum number of actions: a sub-optimal agreement would result in selecting a connection that will need to be removed (by e.g. the clear action) in order to obtain an optimal solution. Likewise, a sub-optimal disagreement is not 'wrong', but rather not the quickest way to reach an optimal solution: in a later suggestion, the same or an equivalent connection could be suggested and agreed, and an optimal solution can be constructed (with some more actions taken).

4) Quantifying the trend of change of action quality: For RQ3, we study the trend of improvement of actions of a specific type, such as suggestions (H3.2), agreements, and disagreements (H3.3). To quantify how and how much the quality of actions changes during a collaborative activity, we i) filter for the specific type of actions taken by the human in that activity, then ii) annotate each action with an action quality value by labeling optimal actions as 1 and sub-optimal actions as 0, iii) annotate the action times as a fraction of the overall progress in that activity from 0 i.e. 0% at the beginning to 1.0 i.e. 100% at the end, and iv) compute the slope of the fitted regression line between the action quality values and action times. Thus, a positive r indicates an improvement, with e.g. the human suggesting optimal actions more later than earlier in the activity: the higher the r, the stronger is the trend.

V. User Study

1) Setup: To comply with COVID-19 school safety regulations, we had to resort to conducting the pilot study without bringing any external person or equipment on school premises. To make as meaningful reference to physical in-person interaction as possible, i.e. the way we intended, each child (physically located in a quiet room at school) interacted with the robot (physically located in our lab) remotely through a videoconferencing application, see Fig. 2. One session lasted $\approx$ one hour and followed the outline in Table II. We note that this was not the intended scenario, where the video-view of the robot is probably less engaging than having the robot in the room: however, this was the only possibility.

2) Participants: We collected a dataset of 9 children (3 females and 6 males), aged 10-11 years old, globally accounting for about 9 hours of interaction, of which around 5 hours spent in the collaborative problem-solving activities between the human and the robot (3 hours in the first with $M = 19.6, SD = 7.7$ min, and 2 hours in the second with $M = 13.0, SD = 5.3$ min). The interview took $M = 5.6, SD = 2.4$ min, with a maximum of 12.1 min.

VI. Results and Discussion

1) User Perception: Concerning their own strategy, all participants reported that they tried to select the cheapest connections that were available and that they suggested to the robot the same ones that they would pick. Participants were not sure how the robot decided on its suggestions: for two of the participants it "looks at the most efficient way", while for three it seemed "random" ("Sometimes it went for bigger numbers, sometimes for smaller").

Concerning what the participants did when asked by the robot to make a guess on what it would do, the replies indicate they overall tended to rather follow their own choices. 5 (out of 9) participants thought that their selection would essentially coincide with what the robot thinks, as the robot in that case would agree with it. Two followed their own choices, as they were unsure or wanted to see the robot's reaction ("Because I didn't really understand the pattern it works in"). The rest occasionally ignored the request ("I picked what I thought the robot would do but sometimes I picked my own"). Therefore, in the following analyses we treat the actions taken when the robot asks to make a guess as suggestions being made to the robot.

2) RQ1 on the Learning Outcomes: Fig. 3(a-c) illustrate how the quality of the solutions of the participants changed from the pre-test to the post-test. A Wilcoxon signed-rank test shows that the average error in the post-test is statistically significantly different than in the pre-test ($W(8) = 0.0$, $p = .008$).[6]. The error is lower in the post-test with a large effect size (Cliff's $\delta(8) = -.61$).[7] This supports our hypothesis H1.3: the participants performed better after collaborating with the robot, by submitting on average better solutions with lower error.

The data is inconclusive about a difference between the pre-test and the post-test in terms of submitting more feasible solutions (H1.1: $W(5) = 5.0$, $p = .63$) or more optimal solutions (H1.2: $W(4) = 1.0$, $p = .25$).[8] On H1.1, we observe that all participants except participant 7 submitted feasible solutions in 80-100% of the tests (see Fig. 3(a)): this indicates that they already had a good conception for identifying valid solutions (LG1) prior to the study. For participant 7, only 40% of the solutions were valid in the pre-test, and none in the post-test: we interpret that

[6] Average errors in the post-test are not normally distributed (Shapiro-Wilk's $W(8) = .82$, $p = .044$)

[7] The magnitude of Cliff's Delta (δ) can be interpreted via thresholds $|\delta| < .147$ "negligible", $|\delta| < .33$ "small", $|\delta| < .474$ "medium", and otherwise "large" [19].

[8] For the Wilcoxon tests for H1.1 and H1.2, there are 5 and 4 samples resp. due to the ties (see Fig. 3), because we can not compute exact p-values with ties, and a p-value via normal approximation that allows ties is used for a sample size of typically > 50 [20].

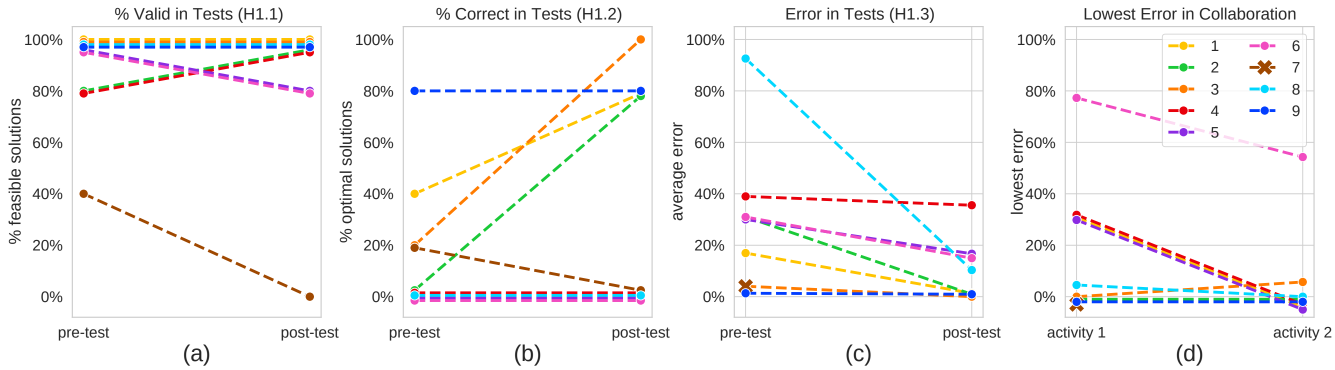

Fig. 3. (a-c) Learning outcomes for each participant, in terms of the quality difference of their responses from the pre-test to the post-test; (d) change in performance through the collaborative activities. The overlapping lines are demarcated by a slight offset. A cross indicates that the error can not be computed. The legend applies to all figures.

7 was confused about the goals of the activity. On H1.2, we observe that one-third of the participants improved by 40-80%, whereas another third did not submit any optimal solutions in either of the tests: while the collaboration with the robot was beneficial for some, it did not seem to help everyone.

Overall, the results show positive learning outcomes for the participants in terms of finding better solutions, suggesting that the collaboration with the robot might have helped trigger the learning mechanisms. However, the interaction was not necessarily beneficial to all the participants, such as participant 7, for whom it is likely that initial misconceptions remained unresolved.

3) RQ2 on the Evolution of Performance in the Task: Fig. 3(d) shows how the performance changes from the first to the second collaborative activity, as measured by the best submission (i.e. `lowest error`). Performance in the second activity is higher than in the first with a medium effect size (Cliff's $\delta(8) = -.38$), but the difference is not statistically significant ($W(6) = 2.0$, $p = .09$).[9] The data is consistent with H2.1: all participants except for participant 3 performed better or the same (by finding an optimal solution in both of the activities); for participant 3 performance only slightly decreased.

From the beginning (0% progress) to the end of the activity (100%), the median optimal suggestion times have the mean at $= 50.7\%$ (SD = 30.5%) and 49.4% (SD = 31.1%) in the first and the second collaborative activity, respectively. In addition, median optimal (dis)agreement times have the mean $= 50.0\%$ ($SD = 31.1\%$) and 59.2% ($SD = 31.4\%$). While we hypothesized in H2.2 and H2.3 that the optimal suggestions and (dis)agreements would come later in the interaction, we observe that they tend to occur throughout the interaction, spread around the middle of an activity.

All in all, participants tended to visibly improve across the collaborative activities, showing better performance in the second compared to the first, even though the robot did not know how to solve the problem. Surprisingly, within the collaborative activities, we do not observe such a trend of improvement in the quality of the actions over time: the notion that optimal actions tend to come later in the interaction is thus probably too simplistic a view for how the interaction evolves.

4) RQ3 on the Link Between Performance and Learning: The change of performance across collaborative activities, i.e. `performance gain`, has a weak correlation coefficient with the learning outcomes as measured by the `learning gain` (Pearson's $r(8) = -.21$, $p = .60$).[10] Thus, the data does not support H3.1. When we look into the specific action types, the trend for suggestions only has a moderate correlation coefficient with the learning outcomes (first activity: $\rho(8) = -.31, p = .46$; second: $\rho(8) = .64, p = .09$): hence, the data is inconclusive on H3.2. The trend for optimal (dis)agreements has a strong negative correlation with the learning outcomes for the first activity ($\rho(8) = .83, p = .01$) and a very weak correlation coefficient for the second ($\rho(8) = .07, p = .87$). To investigate this result in the first activity further, we separately check optimal agreements and disagreements: for these, the trends have very weak to moderate correlation coefficients ($\rho(8) = -.14, p = .74$ and $\rho(8) = -.41, p = .32$, respectively). Thus, there seems to be a non-trivial combined effect of treating them together, without a clear trend for either optimal agreements or disagreements occurring later (median slope for both are at 0). Therefore, we conclude that H3.3 is not supported.

Fig. 4 shows how the performance of the participants evolve through the scenario: from the pre-test instances, to collaborative activities and to post-test instances. We observe in general better performance in collaborative activities compared to tests, even though the robot does not know how to correctly solve the problem by itself. Furthermore, we see that high performance in collaborative activities does not always result in a high performance in the test afterwards.

Overall, the results indicate that the trend of change in the

[9]Lowest errors in the first & second collab. activities are not normally distributed (Shapiro-Wilk's $W(8) = .80$, $p = .030$ and $W(8) = .48$, $p < .00001$, respectively).

[10]The magnitude of Spearman's ρ can be interpreted by: $.00 - .19$ "very weak", $.20 - .39$ "weak", $.40 - .59$ "moderate", $.60 - .79$ "strong", and $\geq .80$ "very strong" [21].

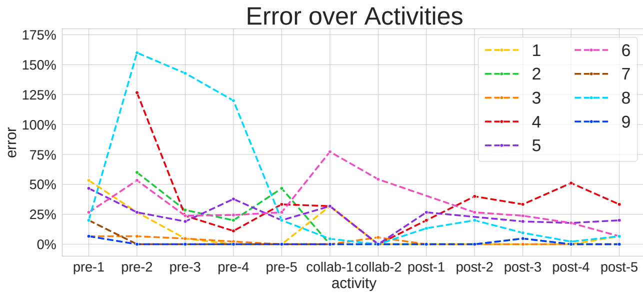

Fig. 4. Evolution of the performance of each participant throughout the scenario.

optimality of actions is not the right level of assessing and predicting the learning outcomes: how the participants think about the activity is not necessarily synchronized with their behavior in the interaction as objectively seen by the actions they take. It is likely that their improvement reflects in the mental representations of the participants about the activity and their collaborator - the robot.

5) Observations on the Dynamics of Interaction: In the collaborative activities, the robot essentially implements a dialogue model. Thus, the goal of this study was, beyond testing the learning activity, to see how much the children 'enter' into the 'dialogic game' more or less as if it was a real dialogue, not just an interaction based on a few rules. For this purpose, we inspect the dynamics of interaction to see patterns of good and bad agreements and disagreements between a child and the robot over time.

Fig. 5 shows the distribution of the optimal and sub-optimal actions taken by each participant. We observe that the children have no problems with disagreeing with the robot. Furthermore, there are some episodes of 'fight', in which the robot's proposed connections are repeatedly rejected. For example, Participant 4 systematically disagreed with the robot's suggestion five times consecutively near the end of the first collaborative activity, and all the disagreements were optimal. Since the robot has an incorrect strategy, optimal disagreements are necessary to achieve the task's goal: i.e. the participant needs to say 'no' to an unnecessary connection, or a costlier connection which if selected would lead to a sub-optimal solution. Thus, it seems that the child asserted his/her belief that the robot was making incorrect choices.

We also observe intervals of possible dis-engagement, in which the child agrees systematically. For example, Participant 8 simply agreed with every suggestion of the robot within the last attempt (seven times, during $t > 20$ min). The second agreement was sub-optimal, and thus the latter agreements and suggestions lead to a sub-optimal solution.

VII. Conclusion

In this study, we observed positive learning outcomes for participants in terms of finding better solutions, after collaboration with the robot to solve a computational problem together. Furthermore, we observed better performance in the collaborative activities compared to the tests, even though the robot did not know how to solve the problem at hand; however, this high *collaborative* performance was not always carried over to the *individual* post-test. We did not observe a correlation between learning outcomes and the evolution of the quality of actions, which indicates a need to delve deeper into the participant's representation of the activity and the robot.

As the robot in our study is not aware of a correct solution, but rather of the process of the interaction, this indicates a shift from the focus on robot knowing and enacting a correct approach to a problem, to rather better modeling the child. We hypothesize that a robot that maintains *beliefs* about the human would be able to work at a more suitable level, with a better proxy to track the learning process: it could guide the interaction using these beliefs and thus bring about learning more effectively. Also, we believe that this type of behavior is easy to port to other activities; as a correct approach need not be included in the behavioral design. Lastly, although the study was conducted with a remote-telepresent robot and was limited to a small number of participants, we believe the results allow further investigation.

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, p. eaat5954, 2018. doi: 10.1126/scirobotics.aat5954

[2] W. Johal, "Research Trends in Social Robots for Learning," *Current Robotics Reports*, 2020. doi: 10.1007/s43154-020-00008-3

[3] P. Dillenbourg, "What do you mean by collaborative learning?" in *Collaborative-learning: Cognitive and Computational Approaches*, P. Dillenbourg, Ed. Oxford: Elsevier, 1999, pp. 1–19.

[4] A. King, "Scripting Collaborative Learning Processes: A Cognitive Perspective," in *Scripting Computer-Supported Collaborative Learning: Cognitive, Computational and Educational Perspectives*, ser. Computer-Supported Collaborative Learning, F. Fischer, I. Kollar, H. Mandl, and J. M. Haake, Eds. Boston, MA: Springer US, 2007, pp. 13–37. doi: 10.1007/978-0-387-36949-5˙2

[5] P. Dillenbourg, S. Järvelä, and F. Fischer, "The Evolution of Research on Computer-Supported Collaborative Learning," in *Technology-Enhanced Learning: Principles and Products*, N. Balacheff, S. Ludvigsen, T. de Jong, A. Lazonder, and S. Barnes, Eds. Dordrecht: Springer, 2009, pp. 3–19. doi: 10.1007/978-1-4020-9827-7˙1

[6] M. Kapur, "Productive Failure," *Cognition and Instruction*, vol. 26, no. 3, pp. 379–424, Jul. 2008. doi: 10.1080/07370000802212669

[7] J. Nasir, B. Bruno, M. Chetouani, and P. Dillenbourg, "What if Social Robots Look for Productive Engagement?" *International Journal of Social Robotics*, Mar. 2021. doi: 10.1007/s12369-021-00766-w

[8] C. C. Chase, D. B. Chin, M. A. Oppezzo, and D. L. Schwartz, "Teachable agents and the protégé effect: Increasing the effort towards learning," *Journal of Science Education and Technology*, vol. 18, no. 4, pp. 334–352, 2009. doi: 10.1007/s10956-009-9180-4

[9] F. Tanaka and S. Matsuzoe, "Children teach a care-receiving robot to promote their learning: field experiments in a classroom for vocabulary learning," *Journal of Human-Robot Interaction*, vol. 1, no. 1, pp. 78–95, 07 2012. doi: 10.5898/JHRI.1.1.Tanaka

[10] 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 Magazine*, vol. 23, no. 2, pp. 56–66, 2016. doi: 10.1109/MRA.2016.2546700

[11] G. Mugny and W. Doise, "Socio-cognitive conflict and structure of individual and collective performances," *European Journal of Social Psychology*, vol. 8, no. 2, pp. 181–192, 1978. doi: 10.1002/ejsp.2420080204

[12] F. Butera, N. Sommet, and C. Darnon, "Sociocognitive conflict regulation: How to make sense of diverging ideas," *Current Directions in Psychological Science*, vol. 28, no. 2, pp. 145–151, 2019. doi: 10.1177/0963721418813986

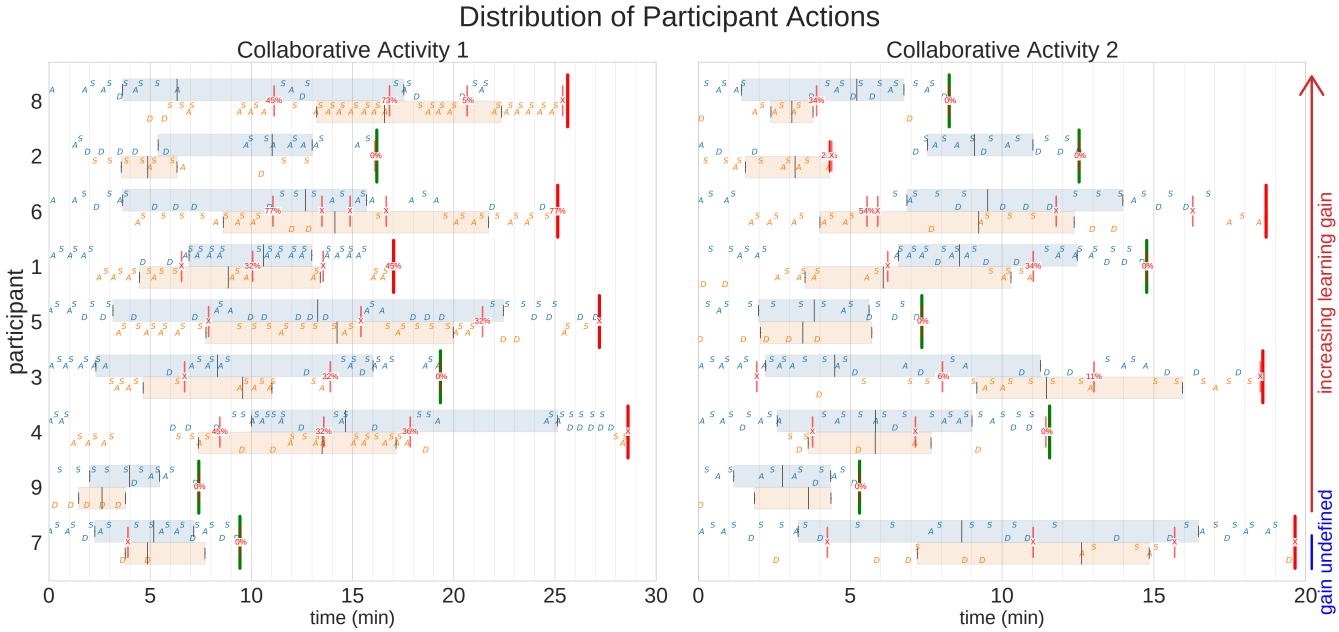

Fig. 5. The participants' optimal and sub-optimal action times in the collaborative activities, shown in blue and orange, respectively, and sorted by decreasing learning gain. Suggestions, agreements and disagreements are designated by 'S', 'A' and 'D', respectively. The box plots show the distribution of optimal and sub-optimal actions through the total duration of interaction. Thick lines indicate the end of the interaction, by submitting a correct solution (in green) or timing out (in red). The thin red lines indicate submission of a solution, with the number showing the error quantifying how far the submitted solution is from an optimal solution in terms of its cost (while 'X' indicates the submission was not feasible and therefore the error can not be computed).

[13] P. Vogt, R. van den Berghe, M. de Haas, L. Hoffman, J. Kanero, E. Mamus, J.-M. Montanier, C. Oranc, O. Oudgenoeg-Paz, D. H. Garcia, F. Papadopoulos, T. Schodde, J. Verhagen, C. D. Wallbridgell, B. Willemsen, J. de Wit, T. Belpaeme, T. Goksun, S. Kopp, E. Krahmer, A. C. Kuntay, P. Leseman, and A. K. Pandey, "Second language tutoring using social robots: A large-scale study," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019. doi: 10.1109/HRI.2019.8673077 pp. 497–505.

[14] J. Werfel, "Embodied teachable agents : Learning by teaching robots," *New Research Frontiers for Intelligent Autonomous Systems*, pp. 1–8, 2014.

[15] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009. doi: 10.1016/j.robot.2008.10.024

[16] G. P. Wiggins and J. McTighe, *Understanding By Design*, 2nd ed. Alexandria, VA: Association for Supervision and Curriculum Development, 2005.

[17] D. Menon, S. Bp, M. Romero, and T. Viéville, "Going beyond digital literacy to develop computational thinking in K-12 education," *Smart Pedagogy of Digital Learning*, 2019. doi: 10.4324/9780429319501-2

[18] L. W. Anderson and D. R. Krathwohl, Eds., *A Taxonomy for Learning, Teaching and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*. New York: Addison Wesley Longman, 2001.

[19] J. Romano, J. D. Kromrey, J. Coraggio, J. Skowronek, and L. Devine, "Exploring methods for evaluating group differences on the NSSE and other surveys: Are the t-test and Cohen's d indices the most appropriate choices," in *annual meeting of the Southern Association for Institutional Research*, 2006, pp. 1–51.

[20] W. J. Conover, *Practical nonparametric statistics*, 3rd ed., ser. Wiley Series in Probability and Statistics. John Wiley & Sons, 1999.

[21] J. D. Evans, *Straightforward statistics for the behavioral sciences*. Thomson Brooks/Cole Publishing Co, 1996.