When Social Adaptive Robots meet School Environments

Completed Research

Daniel C Tozadore
University of São Paulo
tozadore@usp.br

João P H Valentini
University of São Paulo
joao.valentini@usp.br

Victor H Rodrigues
University of São Paulo
victor.henrique.rodrigues@usp.br

Julia Pazzini
University of São Paulo
julia.pazzini@usp.br

Roseli A F Romero
University of São Paulo
rafrance@icmc.usp.br

Abstract

Few robotic adaptive architectures can be programmed by the teachers and used in school environments as a regular tool for teaching, compared to other information systems. This paper describes an experiment in a school environment. The experiment aimed to analyze the benefits of a cognitive adaptive system, which controls a social robot for educational activities, for teaching. The system performance was evaluated based on three criteria: (i) teachers perceptions in using the system, (ii) students perception with regard to the robot and (iii) the system accuracy. Teachers and students perceptions were obtained by means of questionnaires and the system's accuracy was evaluated by manual validation. Overall, the outcomes suggest that the teachers and students identified the benefits of the system for teaching. Furthermore, the system's classifications of verbal answers showed itself efficient when the answers were rightly understood by the speech recognizer.

Keywords

Human-Robot Interaction, Cognitive Robots, Social Educational Robots.

Introduction

Information Systems are commonly used in the educational field to support teachers and students and their advantages are sensible for both sides and since a long time ago (Murray 1999). To the teachers, they can provide tools for creating and approaching theoretical and practical exercises, provide information about the students' performance in these exercises. Whereas the students can take benefit from smart systems that count on difficulty adaptive function and also provide technological novelties.

Regarding social robotics to educational purpose, some disadvantages can be noted compared to other information systems, such as the small number of students that can participate in the interactions with the robot. Another disadvantage is that the students' regular teachers can rarely program the content and have access to the information. One of the main reasons is the lack of preparation of the teachers in using robotics devices in their teaching plan, due to the complexity in programming such systems (Johal et al. 2018). Even after being programmed, the information collected or generated by the student-robot interaction is usually used only by programming experts and the teachers may never have access to them.

Here, a cognitive adaptive system (Tozadore et al. 2018) which can provide adaptation and personalization through open source techniques was developed. It has autonomous analysis and recognition of speech. After creating and running an activity, all sessions' evaluation and information can be accessed for visual analysis, as well as students' preferences throughout the interaction. The system's goal is to intermediate the teachers and the agent, in this case, a NAO robot, to build complete interactions with no programming skill needs. It respects the idea of educational evaluation introduced by Luckesi (Luckesi 2014), which aims a social transformation, not its conservation, considering the learning process of each student.

In this paper, we analyzed how the combination of an Information System with a Cognitive Robotic System can impact both teachers and students in content approaching through pedagogical interactive activities. To validate the system, a set of Brazilian Portuguese grammar exercises was programmed by two teachers in the system Graphical User Interface (GUI) and performed with the students and an interactive robot. The exercises were about the content being addressing in classes by the teachers and approached by the robot in individual sessions.

The robot was controlled by the introduced system that also used rapport building, aiming to ease the children's shyness in interacting with the NAO and possible students demotivation caused by adaptation in the content difficulty. We also analyzed the students' acceptance via survey whether the Human-Robot Personality-Similarity (Robert 2018) is influenced by adapting the robot's behavior in the second meeting in which it said it likes the same things the student said he/she liked in the first meeting.

Related Work

In systems that autonomously interact for long periods with humans, it is essential that its behavior being adapted. Works that adopted this strategy constantly reported a enhance in the user experience. For instance, in a study (Leyzberg et al. 2014) in which the robot was assigned for tutoring a student during tasks that consisted in solve a logic challenge, the result showed that the exercises' resolution times were smaller in the group where there was an adaptive behavior of the robot when compared to the control group.

Adaptive robots can also leads to a more enjoyable experience by itself. Specially to drew attention of the users when tutoring (Vouloutsi et al. 2015). It was found a social adaptive robot can drew more attention from the users than an non-adaptive robot or even a human. Moreover, this method is still more effective for children than adults.

Rapport building and the feeling that the robot is building a relationship with the user is another strategy often used to enhance learning experiences. Some results showed that students tends to put more effort in educational tasks in order to do not disappoint the tutoring robot after create a friendship with it (Brown et al. 2013). This positive impact can also be seeing in group interactions in which teams performed better and were more viable when they were emotionally attached to their robot (You and Robert 2017). Robert (2018) pointed out that humans respond more favorably to extroverted robots, but this relationship is moderated, and humans respond favorably to robots with similar and/or different personalities from them.

However, there is a trade-off needed between content approaching and the robot's intervention in order to do not distract the student. Furthermore, studies have shown that children often do not make the best use of on-demand support and either rely too much on the help function or avoid using help altogether, both resulting in suboptimal learning. In general, the use of robots in educational tasks have shown positives results. Besides potentialize the students learning experience, robots can free up precious time for human teachers, allowing the teacher to focus on what people still do best: providing a comprehensive, empathic, and rewarding educational experience (Belpaeme et al. 2018).

Finally, we identified that the majority of the related studies did not take into account the teachers in their activity planning and evaluations. Thus, we also analyzed their perception of using such systems in this study.

System Scheme

The robotic architecture is a cognitive adaptive system that encapsulates a module based implementation to run over the robot's sensors and actuators. It is also designed for the robot cognition as well. In other words, it is expected the robot also learn from the interaction and use this learning in the next sessions. The agents' interaction flow overview provided by this architecture is shown in Figure 1. The system collects data from the students and uses it both in the current session and in the next ones. These data can be provided voluntarily by the students - such as the verbal answers - or read by the system - such as facial gaze recognition. The majority of them can be used by the adaptive function and other data, such as their personal preferences are used to simulate rapport building in next sessions. The data are stored and the system recovers and uses it after recognizing the user in the next sessions. Moreover, it can ask about some users' preferences and use them in the next conversations. These preferences are manually chosen in the GUI and they may be sport, dance, team, music, toy, hobby, game and food. Every new entry is searched on the internet for their meaning and stored it in the architecture's database.

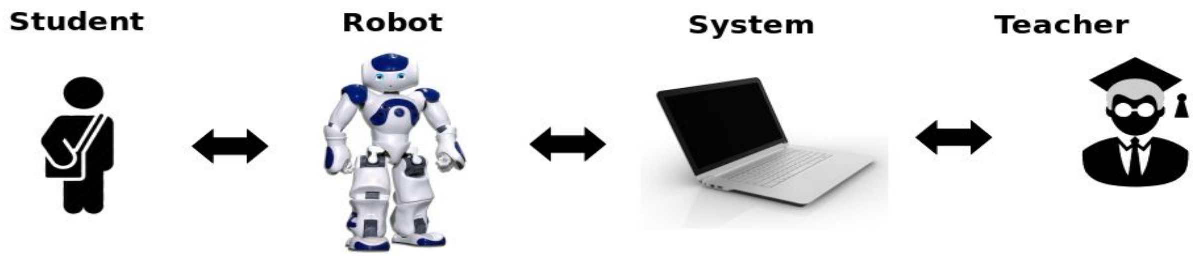

Figure 1: Agents information exchange flow.

The system's dialogue module is the system most important module. It is responsible for the information exchange between the system and the student during the interaction. Three functions compose this module core: read verbal information from the user, processing and understanding the read information and providing verbal information to the user. The user verbal information reading is made by the python SpeechRecognition library in Portuguese-BR idiom. It is one of the best Portuguese speech recognizers. However, it presents a delay to set up the microphone in every requisition and it is noisy sensible sometimes as well.

Sentences interpretation is achieved by a combination of Natural Language Processing (NLP) methods such as the Levenshtein Distance and an embeddings method. It basically compares two strings and returns a distance between them, where 0 means the sentences are the same and 1 means the sentences are totally different. A threshold is configured to accept a user answer as right or wrong compared to a registered expected answer.

The Levenshtein Distance approach is based on the classical implementation of this algorithm and uses the longest alignment between the two sentences, while the embeddings approach is based on the implementation suggested for the ASSIN competition (Hartmann 2016), in which this technique is combined with the TD-IF method. The result is an algorithm that has satisfactory performance in a sparse set of data, such as the user's sentences which the dialogue module is, usually, exposed. The most significant difference between the two approaches resides on the robustness of the methods. When we used the Edit distance the comparison between to sentences stays only into the syntactic field and this presents superficial results, as shown at (Omitted for blinding review), whereas the embeddings plus TD-IF reaches the similarity problem not only syntactically but semantically as well. A relevant weakness noted on this combination is not related directly to itself but to the verbal information caption. When the system does not discern the correct input, because of the pronunciation similarity, the classification tends to be useless.

The dialogue module can alternate among both algorithms by the virtue of some input sentences that can be composed solely by Stopwords. In this situation, the embeddings approach fails because stopwords are removed in the preprocessing stage and the edit distance has shown itself efficient on sentences of less complexity.

The pedagogical model is based on constructivism, as the educational robotic in general (Kafai 2017). The NAO robot plays the main role in the interaction and measure by questions how much the student is rightly constructing its knowledge. The goal is to use this system as practical exercise fixation after regular classes about the topics' concepts. Each meeting between robot and student is called a session. During a session, the robot presents a concept to the student and evaluate if he/she had understood this explanation by asking questions. Every topic aimed to be addressed needs to be registered in the system. The concepts are the topics explanation - and as many questions as the teacher wants to approach. It is mandatory to divide the questions into five levels of difficulty and, at least, one question per level.

The goal of the adaptive function is to make the interactions as attractive as possible to the student, based on the indicators read along the session. The system behavior adapting can take into account three measures read by observing body signals and verbal answers from the user which is interacting with the robot. It can use various reads to adapt the robot's behavior such as face gaze, posture, amount of spoken words per answer, time to answer and answer accuracy. In this case, only the answer correctness was considered to adapt the system and only the questions difficulty was altered. Beyond the importance of respect the following teaching strategies, the levels of difficulty are considered so the students don't have to answer questions that are not in their current level, which could make the student-robot interaction exhausting. In this case, the robot's behavior adapting is considered only the change in the difficulty of the questions the robot is asking.

Experiment

The system was initially tested in laboratory conditions, where the scenarios are considered ideals. It means the interactions were tested in an environment with no noise and performed by users that were already familiarized with the system limitations. So far, its impact with final users - teachers and students - was still unknown. In order to validate the system's data visualization, the programming content and the data visualization by the teachers, an experiment was designed and carried out. This experiment was the first trial with the robot outside the laboratory conditions. Exclusively visual and sound resources were considered in this experiment. The threshold to accept the users' answers as a right answer was set in 0.4

Thus, three independent variables were observed and analyzed in this study:

- Teachers acceptance in using the system's GUI.
- Students' perceptions in interacting with the robot.
- System accuracy in classifying the answers and adapt its behavior (difficulty level).

Further details with regard to the experiment settings are provided in the following sections.

Design

This experiment was divided in 7 steps:

1. *System GUI presentation:* one of the authors presented to 15 teachers of the school how the system should be used/works.
2. *First content programming:* 2 teachers, who participated of the presentation, programmed an activity in the system (Figure 2(a)).
3. *First robot-student interactions:* the robot performed the programmed activities with the students by the first time (Figure 2(a) and 2(b)).
4. *First results analysis and system adjustments:* results of the first interaction were analysed and some adjustments were made in the system.
5. *Second content programming:* evaluation of students performance was shown to the teachers and a second content was programmed by the teachers.

6. *Second robot-student interactions:* the robot performed the second programmed activities with the students by the second time.
7. *Final analysis:* results of the second interaction were analysed and shown to the teachers that answered a questionnaire.

(a) (b) (c)

Figure 2: (a) Teachers programming the content and (b) (c) students performing the activities.

Participants

The participants consisted of 2 female teachers (33 and 54 y.o.) and 32 students (M = 14, F = 18) from the 5th grade with an average age of 9.6 ($\pm$0.49) years old. The selection criteria were all the students that wanted to participate of the sessions and were allowed by their parents.

Interaction Scheme

For the experiment, the teachers programmed two contents in the system, i.e., *Vowel Encounter* and Digraph . The first topic (*Vowel Encounter)* was approached in the two sessions and the second content (Digraph) was only approached in the second session. In general, the questions were about to classify the Vowel Encounter category of a specific word. In the second session, if the words had Vowel or consonant digraph and its category.

The interactions followed a simple scheme. First, the students were called one by one in the classroom where they were developing their regular school activities. The students were head to the school library, where the robot and the researchers were waiting to start the activity. Some instructions were given by one of the researchers, such as how to get the perfect timing to answer the robot, speak louder if the robot did not understand the answer, and so on. Second, the student started the session, the researchers step back and took place behind the table, where the robot was. Each session followed the same scheme in which the robot initiated a dialogue to introduce itself. In the first meet, the robot recognized the student's face and inserted it into the users database. If the student was already registered, all his/her information was recovered to be used in the following conversations. The interaction approach was achieved by the robot asking 3 questions of each topic. The questions were chosen in the difficulty level set by the adaptive function. The level of difficulty of the questions was adjusted based students' answer of a previous question, i.e., from one question to another. In the end, the robot said goodbye to the student and the researches explained possibles system mistakes to the student or why some students' answers might be wrong.

Sessions

The first session was a pilot test for system calibration in the school environment and to the participant students get more used with the robot. In that, the robot asked the student his/her favorite sport, music and food. The system searched the topics on the internet and told them to the student. It also recorded these topics and respective explanations in its database for future use. Next, the robot asked 3 questions, which were randomly chosen according to the level of difficulty in which the student was. The robot was not programmed to repeat any question. It just gave feedback by saying "Congratulations! You are right." or "I'm sorry. You are wrong.". For this session, the robot only considered the content *Vowel Encounter*. The time average of this session was 8.02 ($\pm$3.51) minutes per student.

In the second session, the robot recognized the student by face recognition and recovered his/her information. It also said some things about the student's favorite sport and food to add a personal touch to the interaction. For instance, "I remember you like ice-cream. I tried it since the last time we met. Well... I prefer battery.". Later, the robot asked 3 question regarding *Vowel Encounter*, as in the first session, and 3 questions regarding Digraph, totalizing 6 questions. Differently from the previous session, the robot was able to repeat the questions and confirm the user's answers. Finally, the robot said goodbye to the student. The time average of this session was 9.92 ($\pm$2.83) minutes per student.

Assessment

To evaluate the users (teachers and student) opinion about their experience with the system , a 5-Likert Scale was used. In this scale, each question is an item where the participant gives a score from 1 to 5, in which 1 means absolutely nothing and 5 means totally and 3 is the neutral score.

For the students' part analysis, we applied the Wilcoxon paired test in the result in order to analyze if the users' perceptions in both conditions can be considered the same or if there is a statistical difference between them, it means, if the settings of the interaction may influence the participants. As the number of teachers was not significant - because we used only the teachers that were assigned to the students' classes in this school year - it was analyzed more as a case study than a statistical validation. For the system assessment, it was performed a human validation of the system classifications, taking into account the system's accuracy and tracking the causes of the mistakes.

Results and Discussion

The results of the experiment were grouped according to (i) teachers perceptions in using the system, (ii) students perception with regard to the robot and (iii) the system accuracy analysis.

Teachers perception in using the system

Understanding the teachers' perception of their experience with the system is an important step because the researchers can evaluate f the teachers can easily handle the information through GUI. Two regular teachers participated in the experiment. Both teachers are female and they have 33 and 55 years old. Between step 1 and step 2, they answered a questionnaire together regarding their perception about the system, after attempting the presentation of the system interface (Step 1). After step 7, they answered the same questionnaire again. The questionnaire has the following questions:

1. How much do you think this application can support you in your classes activities?
2. How much do you think it will be easy to create activities in this tool?
3. How much you think this application has potential to become a regular tool in teaching?
4. How much do you think this methodology is efficient for content approaching?
5. How much time do you think it is necessary to get used with this system?
6. How much do you think this methodology can enhance the students' learning?
7. How much do you think that technological methodologies are more efficient compared to the traditional ones?

The teachers answers before using the system (2nd row) and after using the system (3rd row) are presented in Table 1.

Question number	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Before use the system	4	3	3	4	2	3	4
After use the system	4	4	4	4	3	4	4

Table 1: Teachers' scores before and after using the system.

By analyzing the first answers, it is possible to note that the teachers perceived the system potential, once their scores were above the average, except for question number 5 about the time to get used with the system. Although none of the questions got the maximum score, they kept a score above the neutral in the majority of the items (questions of the Likert Scale) after the teachers experienced the system. In open questions, they reported the students' motivation in interacting with the robot and system auto-reports as advantages and system misunderstandings as disadvantages.

Students perception with regard to the robot

We also analyzed the students experience with the system through the interactive robot NAO. A questionnaire was applied aiming to understand if the students were able to notice some robot skills provided by the system. For instance, we investigated if personal information would be a facilitator in rapport building and simulating the robot in becoming more close to the student (item 3). This questionnaire was applied twice to the same students, after the students' first interaction (Step 3) and after the second interaction (Step 6). The questions were:

1. How much do you think you enjoyed this activity?
2. How much do you think you learned with this activity?
3. How much do you think that you and the robot are personal friends?
4. How much do you think that the robot is intelligent?
5. How much do you think difficulty to perform this activity?

The students' average score and Standard Deviation (in parenthesis) after the first interaction (second row) and after the second interaction (third row) are presented in Table 2. We applied the Shapiro-Wilk test to verify if the data follow a normal distribution. Later, we applied the Wilcoxon test aiming to find a significant difference between the samples. In Table 2, the fourth row is the Wilcoxon test comparison between the scores of first and second interaction.

Question Number	**1**	**2**	**3**	**4**	**5**
1st Interaction	4.91 (0.29)	3.57 (1.25)	4.12 (0.89)	4.09 (0.94)	2 (1.25)
2nd Interaction	4.69 (0.58)	4.15 (0.93)	4.32 (0.83)	4.33 (0.92)	2.54 (1.11)
p-value ($\mu = 0.05$)	0.09210	0.05248	0.48979	0.19342	0.02808*

Table 2: Students' average scores after interacting with the robot by the first and the second time.

By observing the fourth row of Table 2 is possible to see that none of the items shown a significant difference but item 5, regarding the activities difficulty, in which the students rated the second activity as harder than the first one.

However, they also rated their learning perception higher in the second section as well. The fact that the robot used the users' stored information in the second meeting did not show a significant difference, despite this item average score was slightly higher in the second evaluation. Similarly, a slightly higher average happened in the second evaluation of item 4, in which they rated the robot as intelligent.

Conversely to other studies (Adam et al. 2016; Charisi et al. 2018; Robert 2018), no influence regarding how they felt about rapport building with the robot and their performance in the activity was found in this experiment, as they rated the second meeting less enjoyable than the first one (4.91 against 4.69) in item 1, although there was no significant difference. This can be associated with the novelty and their expectation of the unknown at the first interaction.

The fact that they rated the second interaction more difficult than the first one but also reported they thought they learned more in the second session leads to believe in the cognitive adaptation potential for learning support. However, the learning gain was not evaluated in this experiment, since it requires more assessment methods to present an accurate analysis.

System accuracy analysis

Only measures of the 2nd interaction (after the adjustments of the pilot test) were considered for a more accurate result. The second interaction, as said before, was composed of 6 questions of 3 questions of each topic. The x-axis points of the graphs in Figures 3(a) and 3(b) represent the questions from the Vowel Encounter (V.E.) and Digraph (D) topics, that happened in this chronological sequence. It means, it was performed 3 questions of Vowel Encounter at first and then 3 questions of Digraph topic. Worthy to remember that the first topic (VE) was already and exclusively approached in the first interaction.

The graph in Figure 3(a) shows the system's classification of the students' answers, in which the green markers are the right ones and the red markers are wrong ones, based on a human validation. A classification is considered right if both the system's classification and human validation about a student answer were the same (right or wrong).

The yellow markers are the answers that were not rightly understood by the speech recognition or that was clearly not a valid answer that some student may say, classified as "Listening problems". These classifications of problems in the system's listening were also validated by one of the researches.

In general, these problems were caused by a change in the outside noise or by a mistake by the student in say the answer while the speech recognition system was readjusting the noisy calibration. The same listening problems were considered for the graph of Figure 3(b) shows the answers read by speech recognizer and classified by a human supervisor. It is the ground truth of the answers evaluations since it is in that way the teachers do to correct the regular tests. The blue markers are the right answers and the red markers are the wrong answers.

The graph in Figure 3(c) shows the 5 levels of difficulty occurrences per question. according to his/her answer in the last question. The system started in medium difficulty level (3). Thus, it is possible to observe the students overall performance during the 6 questions. In other words, the higher this value is, the more the system is evaluating that the student is giving the right answer, according to the expected answer. A very close value to the distance threshold (0.4) - when comparing the student's answer to the expected answer - may not activate the adaptation function, what results in remaining asking questions of the same difficulty level in the next question.

Finally, the pie chart in Figure 3(c) shows the system's accuracy considering all the performed classification. The green area is the system's right classifications and the red area is the wrong ones. The yellow area corresponds to the same listening problems of Figures 3(a) and 3(b). By analyzing the graph in Figure 3(a) is possible to perceive the efficiency of the NPL classification when clearly understanding the users' answers. The same is noticed by analyzing the system overall accuracy graph in Figure 3(d). The system was considered classification had an accuracy of 92.8%, considering only the validated answers.

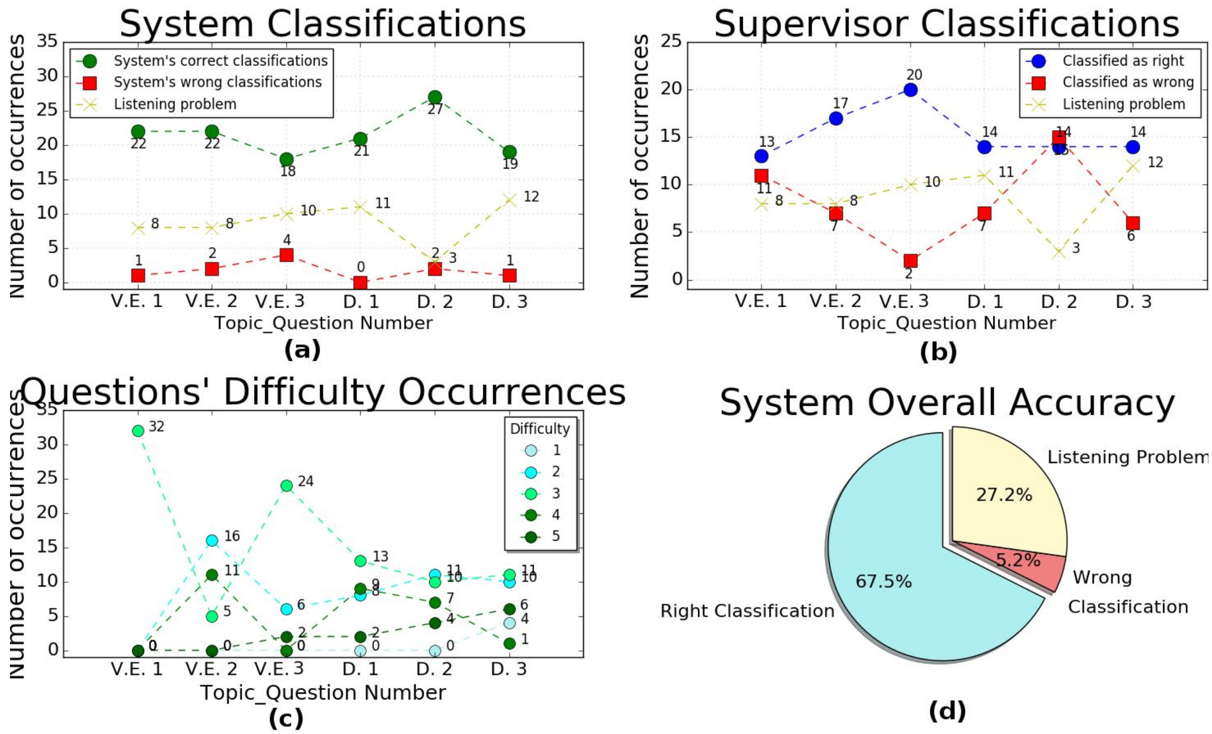

Figure 3: Graphs of (a) System's classifications , (b) Supervisor classifications, (c) Questions' difficulty occurrences and (d) System overall accuracy.

The most critical divergence between the system's classification and the human validation was in the third question of the first topic (V.E. 3) which was the question with the higher number of right answers by the supervisor validation and the lower number of right classifications occurrence by the system. This might has interfered in the adaptation function because it kept the difficulty level for the next topic, but it retook a good response in the second question of the second topic (D. 2). In general, the experiment shown that most demand improvement is the speech recognition treatment, once the system can not validate if the student was really wrong or misunderstanded by the speech recognition.

Conclusion

This paper described an experiment and its results to analyze a cognitive adaptive system by three points of view: by the teachers that used the system, by the students that interacted with the robot and by the researchers that analyzed the system performance. Scores above the neutral state in the teachers' questionnaire showed the teachers understood the potential of this proposal and this finding encourages to keep using this system as a supportive tool for handle the students' profile information along the activities. It seems that the students noticed the potential of the robot being controlled by the system as a tutor in practical exercises, despite it is impossible to analyze the students real learning gain with only correct answers in two interactions with the robot,

Although no statistical evidence was found, a curious point raised with this study was that an improvement in the students' learning perception was observed in the second meeting when the system presented more difficulty adaptation. However, they lower their scores about their interaction enjoyment. Moreover, according to the participants self-statement, their activity enjoyment was not enhanced by the robot's personalized treatment to them.

REFERENCES

Adam, C., Johal, W., Pellier, D., Fiorino, H., & Pesty, S. (2016). Social human-robot interaction: A new cognitive and affective interaction-oriented architecture. In *International conference on social robotics* (pp. 253-263). Springer, Cham.

Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., & Tanaka, F. (2018). Social robots for education: A review. *Science Robotics*, *3*(21).

Brown, L., Kerwin, R., & Howard, A. M. (2013). Applying behavioral strategies for student engagement using a robotic educational agent. In *IEEE International Conference on Systems, Man, and Cybernetics*. 4360-4365. IEEE.

Charisi, V., Alcorn, A. M., Kennedy, J., Johal, W., Baxter, P., & Kynigos, C. (2018, June). The near future of children's robotics. In *Proceedings of the 17th ACM Conference on Interaction Design and Children* (pp. 720-727). ACM.

Hartmann, N. S. (2016). Solo queue at ASSIN: Combinando abordagens tradicionais e emergentes. *Linguamática*, *8*(2), 59-64.

Johal, W., Castellano, G., Tanaka, F. et al. Int J of Soc Robotics (2018) 10: 293.

Lee, M. K., Forlizzi, J., Kiesler, S., Rybski, P., Antanitis, J., & Savetsila, S. (2012, March). Personalization in HRI: A longitudinal field experiment. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 319-326). IEEE.

Leyzberg, D., Spaulding, S., & Scassellati, B. (2014, March). Personalizing robot tutors to individuals' learning differences. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction* (pp. 423-430). ACM.

Luckesi, C. C. (2014). *Avaliação da aprendizagem escolar: estudos e proposições*. Cortez editora.

Murray, T. (1999). Authoring intelligent tutoring systems: An analysis of the state of the art. *International Journal of Artificial Intelligence in Education (IJAIED)*, *10*, 98-129.

Robert, L. P. (2018). Personality in the Human Robot Interaction Literature: A Review and Brief Critique, AMCIS 2018, Aug 16-18, New Orleans, LA. AIS eLibrary, 16-18.

Tozadore, D. C., Valentini, J. P. H., de Souza Rodrigues, V. H., Vendrameto, F. M. L., Zavarizz, R. G., & Romero, R. A. F. (2018). Towards Adaptation and Personalization in Task Based on Human-Robot Interaction. In *2018 Latin American Robotic Symposium* 383-389. IEEE.

Vouloutsi, V., Munoz, M. B., Grechuta, K., Lallee, S., Duff, A., Llobet, J. Y. P., & Verschure, P. F. M. J. (2015). A new biomimetic approach towards educational robotics: the distributed adaptive control of a synthetic tutor assistant. *New Frontiers in Human-Robot Interaction, 22*.

You, S., & Robert, L. (2017). Emotional attachment, performance, and viability in teams collaborating with embodied physical action (EPA) robots. *You, S. and Robert, LP (2018). Emotional Attachment, Performance, and Viability in Teams Collaborating with Embodied Physical Action (EPA) Robots, Journal of the Association for Information Systems*, *19*(5), 377-407.