Second Language Tutoring using Social Robots

European Commission

Project No. 688014

L2TOR
Second Language Tutoring using Social Robots

Grant Agreement Type: Collaborative Project
Grant Agreement Number: 688014

D5.3 Interaction management for the story telling domain

Due Date: **31/09/2018**
Submission Date: **15/10/2018**

Start date of project: **01/01/2016** Duration: **36 months**

Organisation name of lead contractor for this deliverable: **Bielefeld University**

Responsible Person: **Stefan Kopp** Revision: **1.0**

Project co-funded by the European Commission within the H2020 Framework Programme		
Dissemination Level		
PU	Public	**PU**
PP	Restricted to other programme participants (including the Commission Service)	
RE	Restricted to a group specified by the consortium (including the Commission Service)	
CO	Confidential, only for members of the consortium (including the Commission Service)	

Contents

Executive Summary

This deliverable describes the advancement of the interaction manager including, but not restricted to the story telling domain. We summarize all efforts that have been done since the last deliverable, with respect to each task defined in the proposal.

Principal Contributors

The main authors of this deliverable are as follows (in alphabetical order):

Laura Hoffmann, Bielefeld University
Sonja Stange, Bielefeld University
Stefan Kopp, Bielefeld University
Thorsten Schodde, Bielefeld University

Revision History

Version 1.0 (15-10-2018)
First Version.

1 Introduction

The aim of work package (WP) 5 is the development of an interaction management component that is responsible for planning and choosing the system's reactions. The interaction manager receives its input from the input recognition and interpretation modules (developed in WP4), which record and interpret the children's behavior during the tutoring interaction. Based on this input, the interaction manager generates high level messages that specify semantic, as well as pragmatic aspects (e.g., interaction management, socio-relational and coaching-related functions). These are then sent to the output manager (developed in WP6), which translates the high level messages into actual verbal and non-verbal behaviors, executed by the robot or tablet (see Figure 1).

As exemplified in Deliverable 5.1, the project proposal specifies the content of each deliverable as a report of developing an interaction manager for each domain separately. However, the general interaction management mechanisms will not change considerably across the three domains, and are hence defined in a general manner. The only aspects that are distinct between the domains, are the settings and target words. Hence, this deliverable focuses on the tasks defined in WP5 from a general view that is suitable for all domains and highlights the work that has been done since the last deliverable, to improve the interaction management in the realm of an overall working L2TOR system.

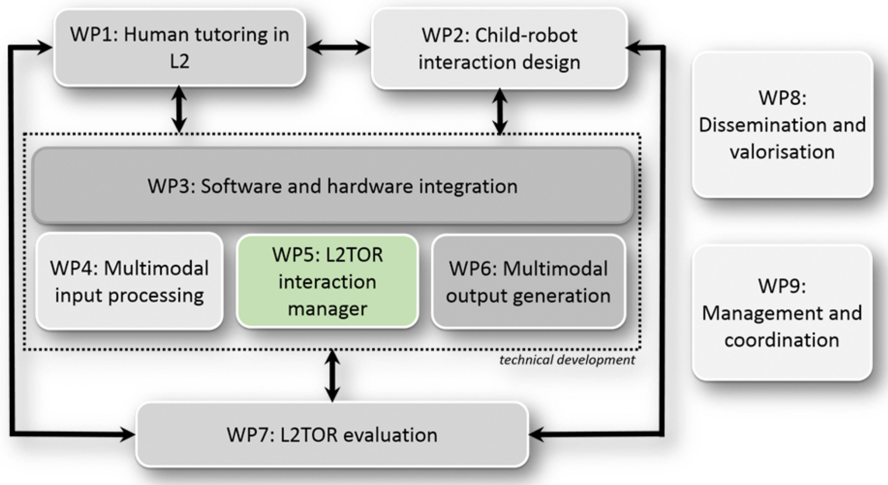

Figure 1: Workpackage overview and how they work together.

2 Tasks of the Interaction Manager

2.1 Input and Output Specification and Representation (T5.1)

Regarding the input and output specifications, nothing changed in comparison to Deliverable 5.2.

2.2 Model of the Child Learner's States and Traits (T5.2)

Regarding the modelling of the learner's states and traits, nothing changed in comparison to Deliverable 5.2.

2.3 Basic Interaction Management (T5.3)

The basic interaction management mechanism remained the same (cf. D5.2).

2.4 Probabilistic State Estimation and Update (T5.4) & Decision-theoretic Dialogue Management (T5.5)

As described in Section 3.6 of the "Revised Objectives"-Document submitted at 03.11.2017 after the first review, we were planning to study the impacts of short-term versus long-term adaptations on a robot–child tutoring interaction. The idea was to allow for the system to trace not only the affective and knowledge states of the children over the short-term but also to store interaction data, which then can be used to adapt the interaction over the long-term. To investigate the resulting effects, we planned to conduct a study to compare a system using short-term adaptation only with a system which also adapts over the long-term (several sessions). In such a system, the long-term interaction data could be used, e.g., for classifying whether a child is a slow or fast learner and to adapt the interaction course accordingly. This is, for slow learner words can be shown more frequently, so that the interaction progresses more slowly, while for fast learners the learning can be sped up to not bore them. The study design comprised the following key points:

- 2 conditions, between-subject design (short-term adaptation vs. short- and long-term adaptation)
- 40 German children in the age of 5-6 years (20 per condition)
- teaching around 20 target words (English)
- 5-6 teaching sessions over 1-2 weeks (20 min per session)
- pre-test to access the knowledge of already known target words
- post-test after the last session to access the resulted learning gain subsequent to the interaction

The study has not yet been conducted due to the time required for preparing the different short- and long-term adaptation strategies. Furthermore, another study (described in the following) was carried out first, addressing many of the other questions we planned to tackle in the remainder of the project. This extensive study was time-consuming and provided many interesting results which needed to be analyzed carefully. Since these results provided substantial insights into how the tutoring interaction and the robot behavior can be adapted to support language learning, we decided to continue from there and to postpone the other, originally envisioned study on long-term adaptation.

2.5 Modeling Interaction Patterns (T5.6)

Besides refining existing interaction patterns no new patterns have been added compared to D5.1.

2.6 Motivational-relational Strategies (T5.7)

As announced in Deliverable 5.2, we developed and implemented strategies to engage children in the robot-child tutoring interaction, while taking their affective and motivational states into account. Therefore, we implemented re-engagement actions that are executable by the robot and further enabled the robot to verbalize its internal decision process to make the interaction flow more transparent to the children. The latter was meant as a scaffolding strategy, which should increase their understanding of

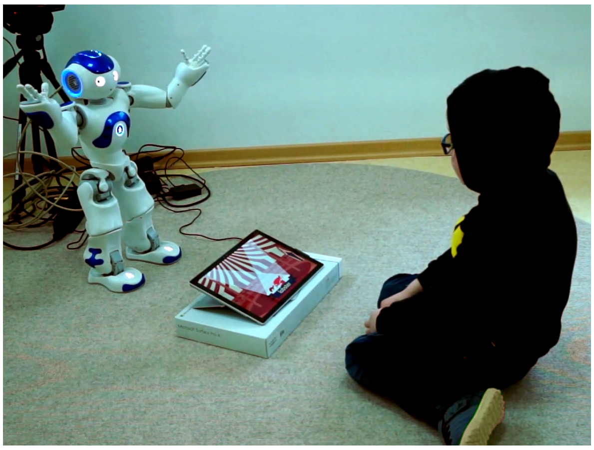

Figure 2: **Example setup of the presented study.** The child learns vocabulary with the tablet app, while the robot supports the language learning by means of demonstrating iconic gesture for the target words, here: lobster.

the interaction flow. A better understanding of the underlying decisions (e.g., to repeat a vocabulary the child is uncertain about) should result in higher motivation to continue learning with the robot and thus better learning gain. The following study combines aspects of two of the studies suggested in the "Revised Objectives"-Document: First, the investigation of effects of re-engagement strategies on the children's motivation and engagement (see Section 3.5 in the Revised Objectives report), and second, the effect of providing explanations to support learning, maintain engagement and raise the long-term motivation to learn (see Section 3.7 in the Revised Objectives report). Since the previously collected data were not usable to train an automated engagement classifier, we decided to use a trained wizard to further investigate the options of re-engaging the children (see Deliverable 5.2). Since we slightly revised the design as reported in Deliverable 5.2, we describe the final design in detail in the following paragraphs along with the study's findings.

The study was conducted with children from three different German preschools (Kindergartens). The children had to solve English learning exercises in the form of an animal-guessing game presented on a tablet, in the presence of a social robot that acts as a knowledgeable peer controlling the tablet (see Figure 2 for the overall setting). We used a one-factorial between-subjects design where the children either interacted with a robot that explained its decisions and beliefs about the child's knowledge before the target word was announced, or not. The children were randomly assigned to one of the conditions, balanced for gender and age. In both conditions the robot performed dedicated re-engagement behaviors, whenever a human Wizard recognized cues of dis-engagement.

Based on findings from work on *open learner models* [1] and *transparency* [2, 3, 4], we hypothesized that verbalizing the robot's belief during the tutoring interaction will make the interaction more meaningful, which will increase the children's motivation to continue interacting (H1), overall learning gain (H2), and perceived learning (H3). For the re-engagement actions, we hypothesized that children who display dis-engagement will re-focus and continue learning after a re-engaging act by the robot (H4a), especially if the initiation of such an action is additionally explained by the robot (H4b).

H1: Children who interact with a robot tutor that explains its belief about the child's knowledge are more motivated to continue learning than children who interact with a robot that does not provide such explanation.

H2: Children who learn with a robot tutor that explains its belief about the child's knowledge show a stronger learning gain than children who interact with a robot that does not provide explanations.

H3: Children who learn with a robot tutor that explains its belief about the child's knowledge gather a better understanding of their perceived learning than children who interact with a robot that does not provide explanations.

H4a: Dis-engaged children will be re-engaged through dedicated actions executed by the robot.

H4b: Dis-engaged children will be more re-engaged when the robot explains its decision to start a re-engagement attempt.

2.6.1 Participants

In total, $n = 49$ preschool children (22 female, 27 male) in the age of 4-7 years from three different preschools in Germany participated in the tutoring interactions. Parents were informed in advance about the purpose and procedure of the study and had to sign consent if their child was willing to participate. The children were randomly assigned to one of the two conditions, balancing age and gender within the conditions. Nine cases had to be removed due to incomplete datasets, resulting from system crashes ($n = 2$), early termination of the interaction due to annoyance or anxiety displayed by the children ($n = 6$), and one case because of impairments in language development that might have affected the overall communication with the robot ($n = 1$).

This leads to a remaining sample of $n = 40$ children (20 female, 20 male), aged 4-7 ($M = 5.43; SD = 0.54$), whereof 22 children participated in the experimental (with explanation) and 18 in the control group (without explanation).

2.6.2 Materials

The experimental setup for the robot–child tutoring interaction included a robot (Softbank Robotics' Nao Version 5), and a Microsoft surface tablet computer. The interaction was monitored by a technician in the same room, however, out of the child's field of view. In addition, two video cameras recorded the interaction from different angles (see Figure 3) for post hoc behavior analyses.

Scenario and chosen vocabulary

To challenge, but not overwhelm the children with the amount of input, we presented them with nine English animal names as target words. The underlying interaction was framed in a child-friendly way, being that the robot and the child are visiting a virtual circus together. A circus is commonly known and liked by children [5] and provides a reasonable frame for different animals showing up in the same place. Furthermore, we took careful decisions in the animal selection based on theoretical assumptions and lessons learned from earlier experiments. For example, we did not include the word "chicken", since we remarked earlier that it is widely known by preschool children. Further decisive factors were a) dissimilarity between German- (first language: L1) and English-word (Second language: L2), b) the possibility to create a comprehensible iconic gesture linked to the animal, e.g., stretching like a monkey (see Figure 4) that is executable by the robot, and c) choice of groups of animals that can be presented in similar colors.

We finally chose: parrot, lobster, and ladybug as red-colored; rabbit, snake, and seal in gray; bull, monkey, and horse for the brown-colored ones as target words for the tutoring interaction (see Figure 5D). Similar to our previous study [6], the task difficulty was varied by means of different

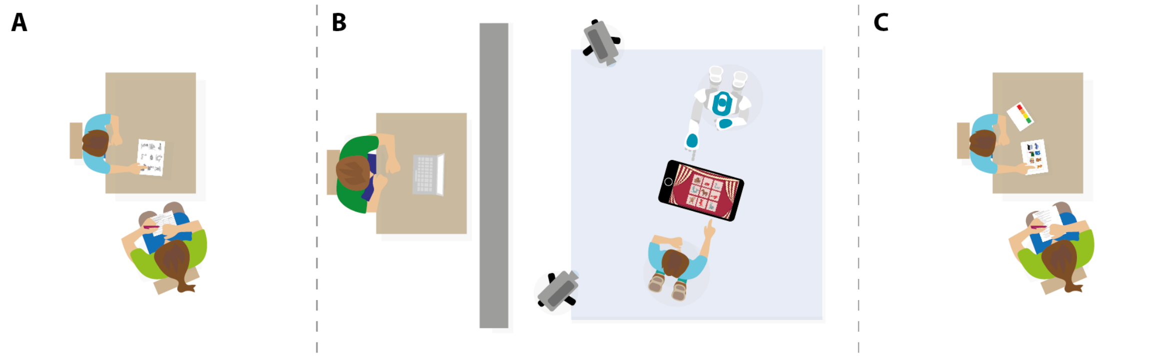

Figure 3: **Experimental setup.** (**A**) Paper and pencil pre-test, (**B**) child–robot interaction supervised by technician and (**C**) post-interaction interview in a separate room.

amounts of distractors, i.e., presentation of three, six or nine different animals on the screen. Additionally, we considered the color as an influencing factor, resulting in four levels of difficulty from easy to hard: 3 animals with 3 different colors, 3 animals with the same color, 6 different animals (2 red, 2 gray, 2 brown), and 9 different animals. Colors were important for our choices as they are a salient feature of objects for children, and might thus result in children remembering the color of the animal instead of the target word [7]. This is why we included different animals with the same color.

Improvements of the tutoring experience

The underlying system used in the present study is based on our previous work with preschool children and teachers [6, 8], and our experiences on how to design a beneficial learning interaction for this particular age group [9]. In particular, we devised a new feedback strategy, new robot behaviors to heighten children's engagement, and verbal explanations by the robot about its current belief about the child's knowledge state.

Feedback strategy: As determined in previous observations [8, 9], the robot gives positive feedback explicitly, e.g., praising the child, when receiving linguistic input or when the right answer is selected. If a wrong answer is given, instead off uttering negative feedback, the target word is repeated by the robot, followed by its translation in the L1. Afterwards, only the wrongly chosen image and the target image are presented and the child is invited to choose once again.

Actions to prevent dis-engagement: In line with suggestions from experts on how to increase children's engagement [10], we implemented small rewarding behaviors in the robot behavior, such as gazing towards the child or nodding when verbal input is received. To signal positive feedback each time the child answered correctly, the robot's eyes were lighted up, blinking in rainbow colors to resemble a smiling face (as has already been used by Fridin [11]). Other than that, iconic gestures for each animal were included to provide multimodal input, which should increase children's interest in the interaction (cf. [6]). However, this time, gestures were not displayed each time an animal name was mentioned by the robot, but only executed once for each animal in the introduction.

Repair actions to recover engagement: In addition to the preventive actions – executed by the robot regardless of the child's state – we designed repair actions that could be carried out spontaneously during the interaction when drops in engagement were observed. We focused on the cues derived from

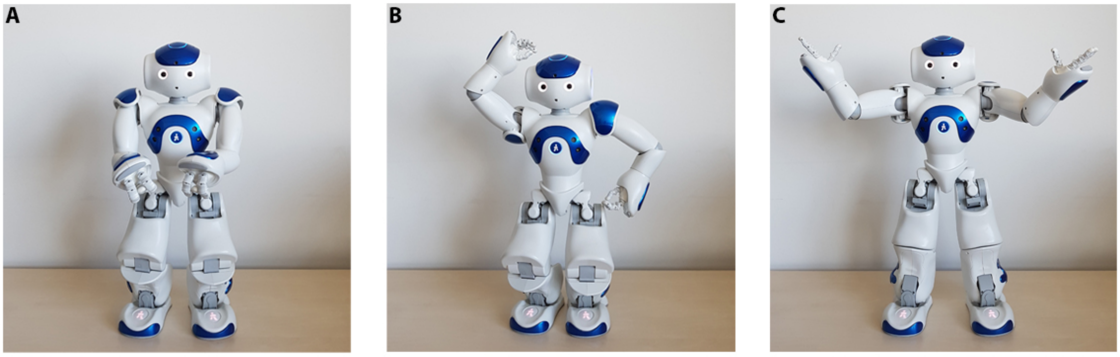

Figure 4: **Iconic gesture examples.** Examples of the robot showing iconic gestures for the animals used in the experiment. (**A**) Imitating riding a horse by grasping imaginary reins and moving the arms up and down; (**B**) Imitating a monkey by scratching armpit and head; (**C**) Imitating a lobster by raising the arms and opening/closing the hand like lobster-claws.

expert interviews to detect dis-engagement [10], summarized in Table 1. As automatic detection is not feasible at the moment, we included a human Wizard of Oz [12] who kept track of the child's behavior. The Wizard was trained beforehand to be aware of the corresponding cues. As soon as the child displayed a behavioral cue of dis-engagement, the Wizard selected a corresponding type of dis-engagement behavior (see Table 1): Tired, heightened activity, low and high distraction. Behavioral cues observable from the child were assigned to these groups, each resulting in a set of repair actions executed by the robot. These actions ranged from simply breathing deliberately and stretching, when the child seemed to be tired, to standing up and moving the arms, or doing squats when the child seemed to have excess energy. All repair actions were primarily executed by the robot, however, most of them invited the child to actively participate, e.g., stand up and squat. Only in cases where the child appeared just slightly distracted, e.g., by gazing away or grimacing, the robot would just wave or whisper to regain attention. When the child would display strong signs of distraction, such as turning completely away or moving away from the interaction, the robot would try to motivate the child to join a dance, as a last resort. These actions were designed to channel the children's attention towards a different aspect and subsequently direct it back towards the vocabulary learning task in order to reconstruct the engagement. By noticing, acknowledging and reacting to a decrease in engagement, the robot should help getting back the child's attention [10].

Explanation Verbalization: The system's adaptive tutoring is based on its continuously updated beliefs about the child's proficiency level (per target vocabulary). Based on these beliefs, the system adjusts the animal order and circumstances of occurrence of target words. The underlying proficiency belief is discrete over four different levels, which are verbally explained by the robot accordingly in the corresponding condition:

1. very low proficiency (skill belief: ∼0-25%), very easy task:
   e.g., *"For practice, I'm selecting an animal which I believe you don't know yet. To keep it simple, only three different animals stand for selection."*
2. low proficiency (skill belief: ∼25-50%), easy task:
   e.g., *"I'm selecting an animal that I believe you are not completely certain about yet. Try and*

Dis-engagement Belief	Behavioral Cues	Repair Actions
Tired	Rub eyes, Yawn, Support head with hands	Joint breathing, Stretching
Heightened activity	Move position, Undirected finger-tapping	Stand up and squat, Stand up and lift arms
Distraction low	Gaze away (from robot & tablet), Grimace	Wave, Whisper
Distraction high	Turn away, Move head from left to right	Chicken dance

Table 1: **Dis-engagement overview for the Wizard.** Overview of dis-engagement beliefs, related cues and repair actions.

find it among the three animals of the same color"

3. medium proficiency (skill belief: ∼50-75%), medium hard task:
e.g., *"Next, we are repeating an animal that I am quite sure you know. To make it a little more challenging, you have to choose from six animals."*

4. high proficiency (skill belief: ∼75-100%), hard task:
e.g., *"I'm now selecting an animal that I am very certain you know. To strengthen your knowledge, it is hidden between eight other animals. I am curious to see whether you find it."*

These sentences occur only in the "with explanation"-condition and are followed by the typical *"I spy with my little eye and see something that is a* ***[target word]****"*, uttered in both conditions. Further, the four wizarded beliefs about the learners' (dis-)engagement are verbally made explicit by the robot:

1. tired: *"I am under the impression that you are a little tired."*
2. activity: *"I think you cannot concentrate anymore."*
3. distraction low: *"I have the feeling that you are a bit distracted."*
4. distraction high: *"I think that you are too distracted to continue learning right now."*

These phrases serve as an introduction to the respective re-engagement behavior, only in the explanation condition, and are followed by the robot inviting the child to join in on a dedicated activity, in both conditions.

The explanations are meant to reveal the system's perception of the children's performance and the resulting consequences for the task difficulty and course of interaction to the children themselves. This should give them a better understanding of their own knowledge and thus provide a scaffold for mastering the task they are faced with.

2.6.3 Measurements

Engagement
Multiple factors are considered to assess children's engagement. First of all, the frequency of dis-engagements: Whenever a child showed signs of dis-engagement, the Wizard of Oz triggered a

re-engagement behavior. It was noted how many children in total showed signs of dis-engagement, and how many attempts to re-engage a dis-engaged child occurred per child. Second, since the children were instructed that termination of the interaction is possible at any given time, the number of rounds they endured are regarded as a further indicator of interaction engagement. Yet, since there were diverse reasons for an early termination of the game, a shorter interaction time does not necessarily account for a lack of engagement, e.g., if the system decided to end early due to a good performance of the child. Therefore, different reasons for the termination of the interaction were also taken into account.

Possible reasons were: 1) The child decided to quit after being asked whether they want to continue after a re-engagement attempt, 2) the WOz triggered the fourth re-engagement behavior, which was set as a threshold, leading to a termination of the interaction, 3) the system decided that the child has learned enough and therefore ended the interaction and 4) the maximum amount of 30 rounds has been reached, and 5) the child quit early due to annoyance without being asked by the robot after a re-engagement action.

Learning Gain

In order to measure the children's learning gain, a pre-, post- and retention-test were conducted. The pre-test was carried out in the following order: First, the child was presented with a paper sheet displaying pictures of all 9 animals in random order. The child was then asked to name the respective animals in its L1, to ensure all used animals are known beforehand. Afterwards, the experimenter verbally presented the animal names in the second language, requesting the child to tap the corresponding picture with their finger. In the meantime, the experimenter noted whether the child answered correctly, but did not provide feedback about the performance to the child.

Subsequent to the interaction, post- and retention-test were used to measure the knowledge increase for the target words. While the post-test was conducted right after the teaching interaction, the retention-test took place at least 1 week later, while both used the same layout and procedure. After the child had received an explanation of the task from the robot and stated to have understood the task, the robot started to present the animal names in the second language while all animals were displayed on the tablet screen. The child was asked by the robot to feed the animal whose name was announced with a grape. For this purpose, the child had to gather a virtual grape from a basket on the bottom right of the screen and move it to the intended animal to feed it (see Figure 5D). Note that during all three tests the same animal pictures were used and the children only got neutral feedback (e.g., *Thanks. The next animal we feed is **[target word]***), so that they did not know if their answer was correct or not.

Perceived Learning

In order to asses how the children estimate their knowledge state regarding the target words, we presented children again with printed pictures of the animals and asked them to sort them into four knowledge categories according to the discretized knowledge levels in the system ($0-25$, $25-50$, $50-75$, $75-100\%$). The categories were discernible via colored squares onto which children were asked to sort the pictures, along with a verbal explanation of each category by the experimenter. The colors were:

- red: *"I don't have a clue what the English name of this animal is."*
- orange: *"I'm uncertain what the English name of this animal is."*
- yellow: *"I'm rather sure what the English name of this animal is."*
- green: *"I definitely know the English name of this animal."*

After having explained the procedure, the experimenter handed over card by card, while each card displayed a printed version of the animals included during tutoring. The child was requested to put each card to one of the four categories, which best fitted their own knowledge estimation, while being advised that they could put as many cards as they wish into one category. For further comparisons of the children's estimated knowledge to their answers in the post-test (which was either coded as right = 1, or wrong = 0), we collapsed the categories into high perceived knowledge (green and yellow) and low perceived knowledge (orange and red).

Re-engagement success
As a measure for the re-engagement success, we considered the continuation of the interaction after an attempt to re-engage a child. Since children were asked whether they want to continue learning vocabulary after each re-engage attempt, we took their decision to continue as an indicator of successful re-engagement. If children quit the interaction early after several attempts of re-engagement, we counted all attempts that resulted in continuation as successful attempts (per child), but the last attempt as unsuccessful.

2.6.4 Procedure

A few weeks before the study, information about the goal, setting and course of the study were handed out to the parents of the children in the respective age group. The parents were informed that for the children to be able to participate a consent form had to be filled out, and encouraged to contact the experimenters in case of any ambiguities.

One week before the actual study took place, the experimenters visited the preschools to introduce the robot and themselves to the children in a group session. This approach was inspired by the work of Vogt et al. [13] and also Fridin [11] who reported positive effects of a first meeting of the robot in a safe and known environment to overcome initial anxiety before interacting with the robot one-on-one. The introductory session was designed as a conversation between the experimenters, the robot and the children. The experimenters began presenting themselves followed by a short background story of the robot which remained initially static. The children were then encouraged to describe the robot's body parts. Here, the similarities (arms, legs, etc.) and dissimilarities (especially face and fingers) between the robot and actual humans were highlighted. The children were then motivated to collectively wake the robot by calling its name. Afterwards, the robot started greeting the children and told a few facts about itself. The robot also suggested to dance the chicken dance to loose its tension. The sound of the chicken dance was then played and robot and children danced together to the sound. Finally the robot went back to a kneeling position in order to rest. The children were then free to touch the robot carefully and ask further questions until the first curiosity was satisfied and the robot said farewell.

On the actual study day, each child entered the room (either with a teacher or alone) and sat down at a small table together with an experimenter. After the child had been informed that they could stop at any time they want, the pre-test was conducted (see Figure 3A). Subsequent to the pre-test, the child was guided to the main interaction spot and asked to sit down in front of the tablet and robot on the ground (see Figure 3B). The experimenter then briefly explained how to use the tablet and that the child was about to learn English vocabulary while visiting a virtual circus together with the robot.

The robot then started with light conversation, asking for the child's name, age and previous knowledge of English vocabulary. Afterwards the circus was introduced as well as all animals living in this circus, one after the other. While the German and English animal names were presented verbally by the robot, the English name and a picture of the selected animal were presented visually on the tablet (see Figure 5A). This input was further enriched by iconic gestures executed by the robot, since

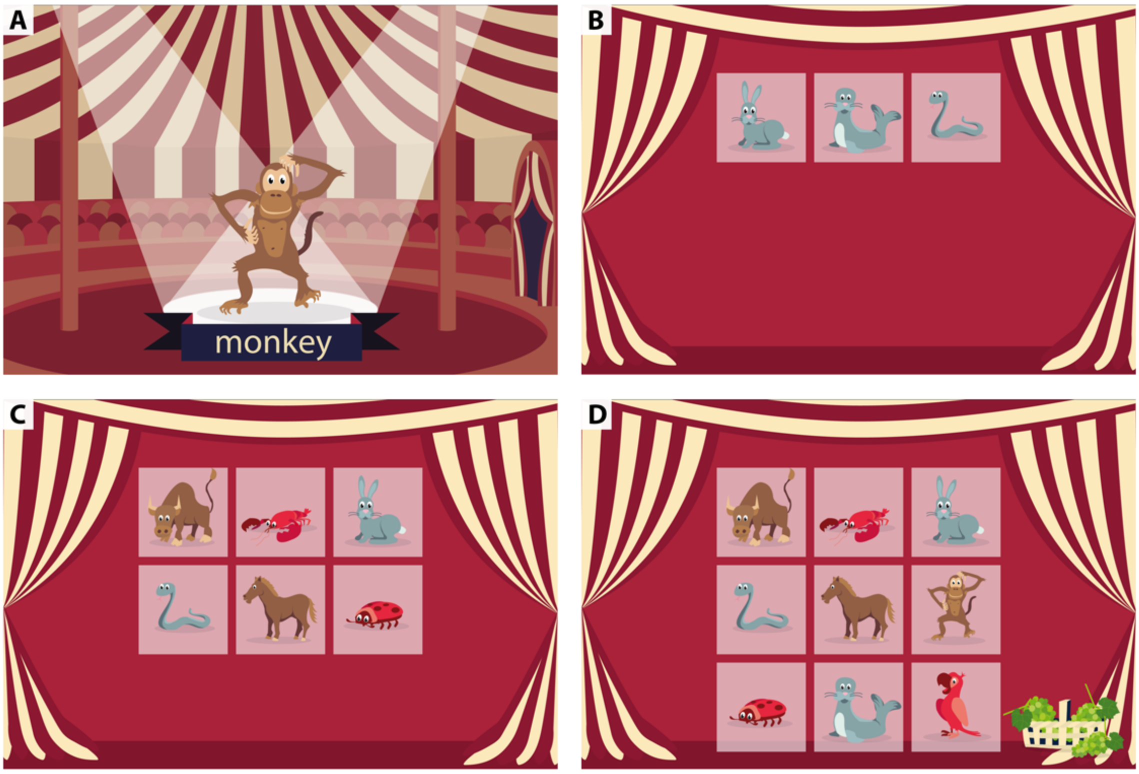

Figure 5: **Scenes from tablet game.** Screen captures from the learning interaction displayed on the tablet: (**A**) introduction of a new animal, (**B**) and (**C**) presentation of a choice of animals of task difficulty 1 and task difficulty 3 during the training, (**D**) post-test including all target words.

it has been shown in an earlier study that word learning is supported by accompanying gestures [6]. The child was further requested to repeat the English name of the animal to strengthen the recall.

After all animals had been introduced, the real learning game, a modified version of the children game "I spy with my little eye..", was explained. To verify if the child understood the game principle, two test rounds were played before the 30 rounds of learning finally started. In each round the adaptive system chose the skill to teach, i.e., an English animal name, in a given task difficulty, expressed through different amounts of animals displayed on the tablet screen (see Figures 5B and 5C). The child then had to tap the animal on the screen which depicted the animal announced by the robot (in English).

After 30 rounds, or as soon as any other reason for the termination of the interaction occurred (see termination reasons), the post-test started (see Figure 5D). After all nine animals had been fed, the robot thanked the child for playing with it and said goodbye. Then the experimenter came back and gave grapes to the child and asked if they want to share them with the robot. Finally, a picture of the child and the robot was taken as a souvenir for the child, before the final interview was initiated by the experimenter, in which children were asked to estimate their perceived learning of the animals (see Figure 3C).

Termination reasons	With Explanation (n = 22)	Without Explanation (n = 18)	Children in total
Played 30 rounds	11 (50%)	11 (61.11%)	22
Finished early due to high knowledge state (system's decision)	7 (31.82%)	5 (27.78%)	12
Finished early after re-engagement attempt (child's decision)	2 (9.1%)	2 (11.1%)	4
Finished early after 4 attempts to re-engage (system's decision)	2 (9.1%)	-	2

Table 2: **Reasons for terminating the interaction.** Amount of children who terminated the interaction due to the listed reasons according to experimental conditions.

2.6.5 Results

In total, data of $n = 40$ children at the age of four to seven years were gathered and analyzed for the impact of explanations on understanding, motivation, learning gain and willingness to continue after dis-engagement.

Children's engagement
As indicators of children's engagement we considered 1) the frequency of behavioral cues of dis-engagement (e.g., yawning or heightened activity; cf. Table 1) during the interaction, 2) the number of rounds children played since they were free to terminate the interaction at any given time, and 3) the reasons for termination.

Based on our earlier experience [10, 6] we expected children to become easily dis-engaged after a couple of rounds due to their short attention spans [14]. However, only twelve children ($n = 7$ in explanation condition, $n = 5$ in control group) displayed one or more behavioral cues of dis-engagement at all. For instance, children played with their clothes or stood up and sat down in a different posture again while not paying attention to the tablet or robot. A comparison of the occurrence of dis-engagement between the conditions revealed that the difference was not significant.

This was somehow unexpected as the tutoring interaction lasted more than 18 minutes in the control condition ($M = 18.16$; $SD = 2.18$) and 23 minutes in the explanation condition ($M = 23.76$ minutes, $SD = 3.97$), which is longer than in previous studies in which children's engagement dropped significantly [6]. Even though the interactions were significantly longer in the explanation condition [$t(38) = 5.36; p < 0.001; d = -1.75$)], we did not observe more dis-engagement or early termination in this condition. The average number of rounds children played with the robot with explanation were $M = 26.55, SD = 5.37$, without explanation $M = 27.67, SD = 3.58$. In principle, this could hint to a stronger engagement in the explanation condition, in which the overall duration of interaction was much longer. This, however, is not true when looking at the numbers of terminations.

Every interaction was either terminated because 1) children decided to quit either by themselves or after being asked by the robot whether to continue after a re-engagement attempt, 2) the system decided to stop after four recognized cues of dis-engagement, 3) the system decided to terminate early due to a high belief that the child's knowledge is high enough (above 75%), or 4) the child reached a maximum of 30 rounds. As visible from Table 2, the majority of children played all 30 rounds or finished early due to a high skill mastery belief, whereas only a few dismissed after attempts to re-engage them. A comparison of the frequency of termination reasons between the conditions yielded no significant differences: Children continued or dismissed similarly often, regardless of the robot's explanation behavior.

In sum, H1 was not confirmed as children did not appear more engaged when the robot explained its belief and decision making, as visible from children's signs of dis-engagement, continuation of playing, and reasons to terminate.

Learning gain

To measure learning gain we administered tests of the to-be-learned words before, immediately after, and 1 to 4 weeks after the tutoring interaction (due to holidays in Kindergartens). Numbers of correct answers are analyzed using a repeated-measures ANOVA, with the timing of the test (pre/post/retention) as a within subjects factor and the experimental condition (with/without explanation) as a between factor. The analysis revealed a significant increase in the amount of words children learned from pre- to post- and retention-test [$F(2, 37) = 44.181; p < 0.001; \eta_p^2 = 0.538$].

Post hoc comparisons revealed that the increase from pre- to post-test, as well as the increase from pre- to retention-test was significant ($p < 0.001$), whereas the difference between post- and retention-test was not. Contrary to our expectations, the increase in the learning gain did not significantly differ between the conditions, although by tendency the mean number of correct answers for the post- and retention-test was consistently higher in the explanation condition (see Figure 6A). Because of the high deviations from the means, we further investigated individual differences between the children. Since the adaptive design of the interaction resulted in a variety of termination points (after 9 to 30 rounds) as well as reasons for termination, we split the data into two groups: *Fast learners* who finished early due to a high skill mastery belief of the system ($n = 12$; average amount of rounds: $M = 24.50; SD = 3.83$) and *slow learners* who interacted for the maximum of 30 rounds ($n = 22$; amount of rounds = 30)[1]. A direct comparison of the test results of fast and slow learners revealed that fast learners performed significantly better in all tests [$F(1, 32) = 26.52; p < 0.001; \eta_p^2 = 0.453$], including the pre-test (see Figure 6B).

Analyzing the effect of explanations on learning gain for fast and slow learners separately, we do find a significant difference: While fast learners' test results were unaffected, slow learners show a big increase in performance when working with a robot that explains its beliefs and decision-making (see Figure 6C). As depicted in Figure 6D, this difference emerges most strongly in the post-test [$t(20) = -2.30; p < 0.05; d = 0.98$]: Slow learners in the explanation condition performed better ($M = 5.09; SD = 1.45$) than slow learners without explanation ($M = 3.55; SD = 1.70$). This difference is not significant anymore in the retention-test after a couple of weeks. In contrast, fast learners all learned similarly, with slightly higher average gains in the explanation condition (pre-to-post: $M = 4.00, SD = 1.73$; pre-to-retention: $M = 3.86, SD = 3.39$) compared to the control (pre-to-post: $M = 3.80, SD = 2.77$; pre-to-retention: $M = 3.80, SD = 3.35$), but with high deviations from the mean.

In sum, the hypothesis that children's learning gain benefits from explanations of the robot's belief and decisions (H2) was not supported for all children. However, while no significant impact on fast learners was found, slow learners showed a significantly bigger learning gain when explanations have been provide as compared to the control group, lending partial support to our hypothesis (H2).

[1] Note that children who finished early due to dis-engagement are excluded from this analysis.

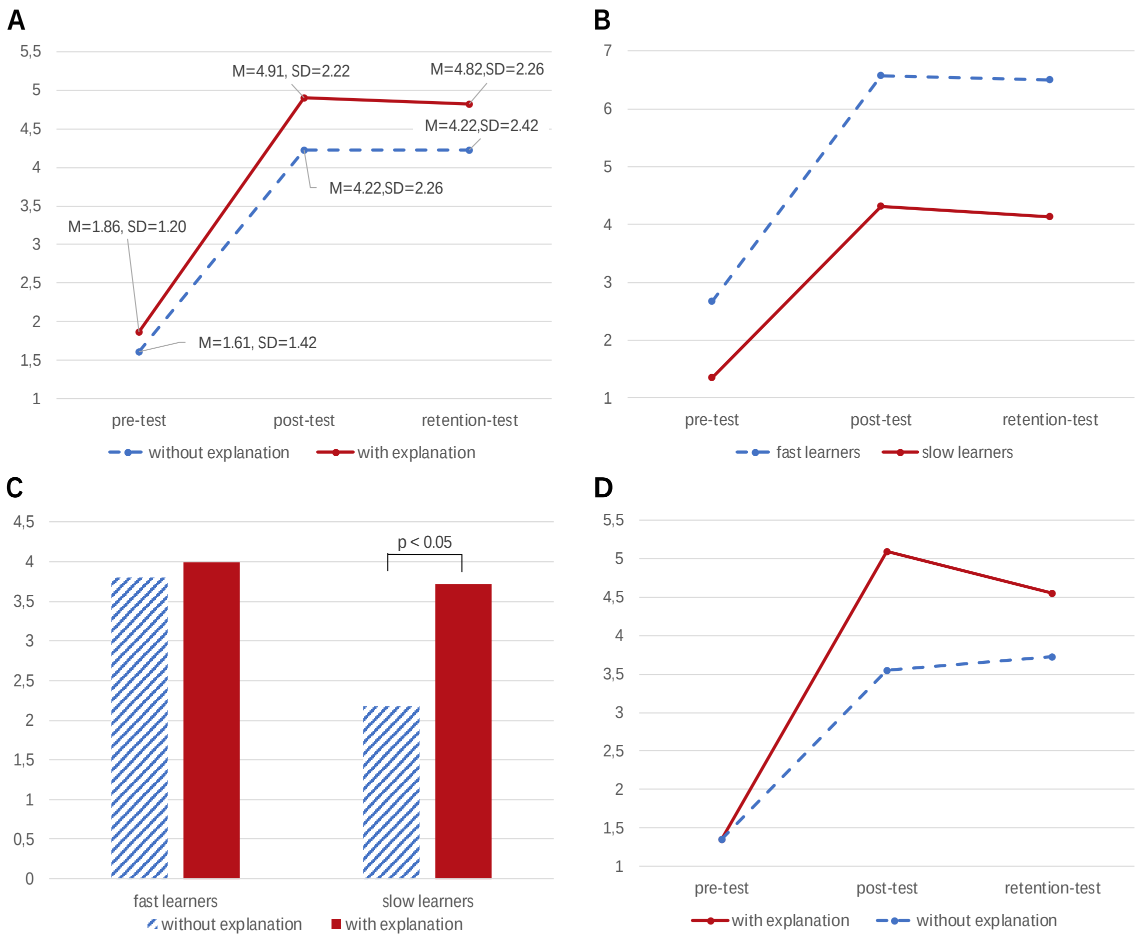

Figure 6: **Visualized learning results.** Pre-, post- and retention-test results with all children included compared between (**A**) study conditions and (**B**) between fast and slow learners. Learning gain (difference pre- to post-test) in the different conditions (**C**) between fast and slow learners and (**D**) for slow learners between the study conditions.

Perceived learning

Children's perceived learning was assessed by asking them to categorize taught target words according to high or low perceived knowledge. To test whether explanations had an impact on the child's estimation of their own knowledge, we first calculated agreement between the children's estimation (high/low perceived knowledge) and their word-specific results in the post-test (right/wrong answer). Agreement was given if the child estimated high knowledge and answered correctly in the post-test, or if they estimated the knowledge as low and gave a false answer.

Calculation of an overall agreement between children's estimates and the post-test results reveal on average equal agreements of 64.44% (with explanation: $SD = 0.21$; without: $SD = 0.15$). We also ran the same analysis only for the slow learners. There, children in the explanation condition showed a higher certainty in estimating their perceived learning (64.33% agreement; $SD = 0.16$) than children in the control group (55.75% agreement; $SD = 0.21$), but the difference was not significant. H3 was thus not confirmed.

Re-engagement after dis-engagement
To analyze whether a social robot could effectively re-engage children in a learning interaction (H4a), we analyzed children's motivation to continue interacting after having shown dis-engagement and being addressed by the robot with re-engaging behaviors. In total, twelve children out of 40 exhibited 23 cues of dis-engagement. Half of the dis-engaged children ($n = 6$) were successfully re-engaged by the robot to continue the interaction until the very end (30 rounds). In two cases, the first attempt to re-engage the child was unsuccessful and the child quit immediately during or after the robot's attempt. In five other cases, dis-engagement re-occurred a second, third and fourth time, despite of several re-engagement attempts, which also resulted in termination of the interaction. In sum, half of the children who got dis-engaged could be motivated by the robot to continue the learning interaction. H4a is thus at least partially supported. An analysis of whether the success of re-engaging attempts is affected by additional explanation for the triggered action (H4b), reveals that 13 out of 17 (76.47%) attempts were successful in the explanation condition, with six out of eight (75%) in the control group. That is, there was no significantly higher success of re-engagement actions when accompanied by explanations.

In summary, 30% of the children received re-engaging actions from the robot, which in turn had an average success rate of 76%. In result, all twelve children who showed signs of dis-engagement continued to work in the learning interaction, with six of them sticking to it for the maximum number of rounds. This seems to be independent of the explanation condition.

2.6.6 Discussion

The overall goal of the present work is to investigate how social robots can be used to scaffold second language learning in preschoolers. We devised a child-friendly learning scenario with a game presented on a tablet computer, augmented with a social robot that guides the child through the interaction as a peer-like tutor. The basis for successful scaffolding is a knowledge-tracing approach that allows for tailoring the interaction to the child's current learning state [8]. Here we studied if we can provide an additional scaffolding for learning by adding a transparency component, similar to an open learner model [1], that enables the robot to give reasons for the learning tasks based on its beliefs on how the child is doing. In addition, we wanted to know if the robot's ability to produce socially supportive behaviors (e.g. nodding and blinking as positive feedback) can be utilized to increase the children's engagement and, even more challenging, to restore engagement when cues of behavioral dis-engagement were visible.

The results of our study with $n = 40$ preschoolers reveal that adaptive, robot-supported language tutoring is overall successful: All children strongly increased their knowledge from pre- to post-test. Retention tests further demonstrated that the knowledge gain persisted several weeks. Regarding the effects of explanations, the results are faceted in interesting ways: Overall, the learning gain for the whole sample did not differ between experimental conditions. A more detailed analysis, however, reveals that explanations do have a strong effect on children that completed all 30 rounds (slow learners, or 'high engagers' for that matter): These slow learners achieved a significantly better post-test result when the robot used explanations, compared to slow learners in the control group. No such difference emerged for fast learners, who overall had a better performance in the pre-test, finished early due to high knowledge, and performed equally well with or without additional explanation by the robot.

One possible explanation is that the verbal explanations mitigated slow learners' uncertainty about the robot's beliefs, their own performance, or the next tasks to be taken – hence reducing cognitive load. In other words, in case of non-optimal performance or increased uncertainty, the robot's explanations can provide a scaffold [15] by reducing felt uncertainty and thus foster learning. However, that although explanations affected the learning gain of slow learners in the expected direction, we did not observe any impact on perceived learning. Fast learners, in contrast, may have a higher task- or learning-proficiency, probably did not work in their “Zone of Proximal Development” [16], and thus do not need additional scaffolding provided by the robot [15]. They already started with higher prior knowledge, performed better and got more positive feedback. Their correspondingly higher affective and cognitive learning may have resulted in a better perceived learning and learning progress. This increase in perceived learning might have further reduced uncertainty, leading to an even smaller effect of the explanations, and resulting in comparable learning gains with and without explanation.

Another pronounced effect was found on overall engagement and motivation of the preschool children, as evidenced by longer interaction times and fewer dis-engagement than in a previous study within the same setting [6]: The use of socially supportive behaviors and re-engaging actions to repair dis-engagement, in addition to the robot’s normal behavior including iconic gestures for the introduction of new words or adapting the tutoring interaction, led to an overall more engaging and thus more concentrated and longer learning interaction. Interestingly, the robot’s re-engagement actions, such as waving or asking the child to stand up and stretch the arms, showed a success rate of 76%, with 50% of the children who started to become dis-engaged continuing the interaction up to the maximum of 30 rounds after a re-engagement attempt by the robot. Note that the robot directly asked the children after each re-engagement action whether they wanted to quit or continue the session, and children had to agree to continue. This indicates that the presence of a social robot can be leveraged to maintain the children’s commitment to the interaction and the task.

Another reason for the children’s willingness to continue could be that the particular re-engagement actions mitigated their negative affective state (e.g., tiredness or boredom). Allowing for bodily action such as standing up or stretching arms could have led to more concentration, as necessary for cognitive learning, and to engagement as necessary for affective learning. In addition, the robot’s reactions to dis-engagement unveil the robot’s active role in keeping track of the learner. In that way, the robot demonstrates a capacity not only to keep track of the child’s knowledge state but also to be attentive to their affective state. Knowing that the robot can notice distraction or inattentiveness might have additionally fostered children’s effort to concentrate on the task. Yet, making this ability even more salient by use of verbal explanation did not increase this effect.

In summary, we can conclude that social robots can indeed be used to scaffold second language learning of preschoolers. Specifically, adapting the course of the interaction and the robot’s behavior, both to the cognitive state and the engagement state of the child, and explaining these adaptations provide effective support to children who struggle with the learning task or start to get dis-engaged. It is this zone where scaffolding is particularly important and where learning occurs.

This also shows that the full potential of an robot tutoring system capable of such strategies can only be exploited when the system is able to respond to the individual differences of children. Future systems must feature a more detailed learner model to profile the children and their needs and to allow for fine-grained adaptation. Such a model should incorporate the child’s knowledge state, engagement level, general attributes such as learning speed, as well as indicators of the child’s responsiveness to specific strategies (e.g., entries about the success of a given strategy such as gestures or explanations). While such demands have been posited before in the realm of intelligent tutoring systems [17], we have provided here insights into how social robots that can add a new element to the learning environment, can support the different dimensions of learning and what they need to adapt to in this. Further research will need to test other scaffolding strategies (e.g. reducing supportive behavior) and investigate their influence on learning as well as side effects of strategy-combinations. For example, previous findings suggest that the combination of adapting the course of the interaction with iconic gesturing by the robot resulted in a higher learning outcome compared to using each strategy alone [6]. Combining an extended learner profile with systematic knowledge about the effects of different supportive strategies (and combinations thereof) should thus be a major goal for the future development of educational social robots that need to do respond to the individual needs of each learner in an optimal manner.

3 Conclusion

In summary, the goal of WP5 has been the development of an interaction manager for the L2TOR system. This component is in charge of planning the course of a second language learning interaction by choosing appropriate actions based on its internal knowledge, affect and memory models, as well as pre-designed interaction patterns. Therefore, this component has to (1) receive/send multi-modal input/output, (2) interpret input for the current state of the interaction and the child, and (3) decide which action should be performed as a reaction to this state. Since the underlying mechanisms incorporated in the interaction manager remain the same for all domains, although the scenes on the tablet and target words are different, we only briefly gave an update on the changes for the different tasks compared to Deliverable 5.2 and also what still needs to be addressed in the future. Summarized, the following work in relation to the different tasks has been done:

- **(1) Input/Output (Task 5.1)**
  The input and output specifications (T5.1) have have not been modified compared to Deliverable 5.2.

- **(2) Interpretation (Tasks 5.2, 5.4, 5.7) and (3) Decision Making (T5.5)**
  We conducted a study as proposed in the Deliverable 5.2 to further investigate the impact of engagement and re-engagement strategies on the learning interaction. Furthermore, we included an additional robot behavior, namely verbalization of the robot's beliefs about the child and the learning interaction, to study whether this influences the learner's performances as well as long-term motivation. Results demonstrated the effectiveness of re-engagement strategies delivered through the robot. Around 76% of the re-engage attempts were successful such that the children went on playing the learning game, often all the way through the full set of learning tasks. Additionally, our study showed that providing explanations through the robot as a further scaffolding strategy help slow-learners to improve significantly and allowed them to catch up with the fast learners. This further highlights the importance of an adaptive interaction management, including the individualized application of scaffolding and learning strategies, so that every learner can optimally benefit without lengthening the interaction too much.

- **Future work**
  The evaluation of an integrated full system, which adapts the interaction by means of planning the interaction flow based on the children's affective and knowledge state and then chooses the right teaching and scaffolding strategies, will not be possible in the remaining project duration (see reasons in Section 2.4). Still, to show how such system could be constructed and to provide a basis for evaluating of the potential benefits it offers, we will implement a demonstrator system in the remaining time. Including and evaluating a setting that combines short-term and long-term adaptation approach will be left for following-up work.

References

[1] Susan Bull and Judy Kay. Open learner models. In Roger Nkambou, Jacqueline Bourdeau, and Riichiro Mizoguchi, editors, *Advances in Intelligent Tutoring Systems*, pages 301–322. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.

[2] Joseph B. Lyons, Garrett G. Sadler, Kolina Koltai, Henri Battiste, Nhut T. Ho, Lauren C. Hoffmann, David Smith, Walter Johnson, and Robert Shively. Shaping trust through transparent design: Theoretical and experimental guidelines. In Pamela Savage-Knepshield and Jessie Chen, editors, *Advances in Human Factors in Robots and Unmanned Systems*, volume 499, pages 127–136. Springer International Publishing, Basel, Switzerland, 2017.

[3] Joseph B Lyons. Being transparent about transparency: A model for human-robot interaction. In *Proceedings of the 2013 AAAI Spring Symposium Series*, 25 to 27 March 2013, Stanford, CA, USA, 2013.

[4] Joseph E. Mercado, Michael A. Rupp, Jessie Y. C. Chen, Michael J. Barnes, Daniel Barber, and Katelyn Procci. Intelligent agent transparency in humanagent teaming for multi-uxv management. *Human Factors*, 58(3):401–415, 2016.

[5] Karin Jampert, Kerstin Leuckenfeld, Anne Zehnbauer, and Petra Best. *Sprachliche Förderung der Kita : Wie viel Sprache steckt in Musik, Bewegung, Naturwissenschaften und Medien?* verlag das netz, Kiliansroda, Germany, 2006.

[6] Jan de Wit, Thorsten Schodde, Bram Willemsen, Kirsten Bergmann, Mirjam de Haas, Stefan Kopp, Emiel Krahmer, and Paul Vogt. The effect of a robot's gestures and adaptive tutoring on children's acquisition of second language vocabularies. In *Proceedings of the 13th Annual ACM/IEEE International Conference on Human-Robot Interaction*, 5 to 8 March 2018, Chicago, IL, USA, pages 50–58, 2018.

[7] Irene M Barrow, Donald Holbert, and Michael P Rastatter. Effect of color on developmental picture-vocabulary naming of 4-, 6-, and 8-year-old children. *Speech-Language Pathology*, 9(4):310–318, 2000.

[8] Thorsten Schodde, Kirsten Bergmann, and Stefan Kopp. Adaptive robot language tutoring based on bayesian knowledge tracing and predictive decision-making. In *Proceedings of the 12th Annual ACM/IEEE International Conference on Human-Robot Interaction*, 6 to 9 March, 2017,Vienna, Austria, pages 128–136, 2017.

[9] Tony Belpaeme, Paul Vogt, Rianne van den Berghe, Kirsten Bergmann, Tilbe Göksun, Mirjam de Haas, Junko Kanero, James Kennedy, Aylin C. Küntay, Ora Oudgenoeg-Paz, Fotios Papadopoulos, Thorsten Schodde, Josje Verhagen, Christopher D. Wallbridge, Bram Willemsen, Jan de Wit, Vasfiye Geçkin, Laura Hoffmann, Stefan Kopp, Emiel Krahmer, Ezgi Mamus, Jean-Marc Montanier, Cansu Oranç, and Amit Kumar Pandey. Guidelines for designing social robots as second language tutors. *Social Robotics*, 10(3):325–341, Jun 2018.

[10] T. Schodde, L. Hoffmann, and S. Kopp. How to manage affective state in child-robot tutoring interactions? In *Proceedings of the 2nd International Conference on Companion Technology*, 11 to 13 Sept. 2017, Ulm, Germany, pages 1–6, Sept 2017.

[11] Marina Fridin. Kindergarten social assistive robot: First meeting and ethical issues. *Computers in Human Behavior*, 30:262–272, January 2014.

[12] Nils Dahlbäck, Arne Jönsson, and Lars Ahrenberg. Wizard of oz studies: Why and how. In *Proceedings of the 1st International Conference on Intelligent User Interfaces*, 4 to 7 Jan. 1993 ,Orlando, Florida, USA, pages 193–200, 1993.

[13] Paul Vogt, Mirjam De Haas, Chiara De Jong, Peta Baxter, and Emiel Krahmer. Child-robot interactions for second language tutoring to preschool children. *Frontiers in Human Neuroscience*, 11(March):1–7, 2017.

[14] Dieter Krowatschek, Caroline Reid, and Gita Krowatschek. *Marburger Konzentrationstraining (MKT) für Schulkinder*. verlag modernes lernen Borgmann, Dortmund, Germany, 2011.

[15] Pauline Gibbons. *Scaffolding language, scaffolding learning*. Portsmouth, NH: Heinemann, Berkeley, CA, USA, 2002.

[16] Lev Vygotsky. *Mind in society: The development of higher psychological processes*. Harvard University Press, Cambridge, MA, 1978.

[17] Lorella Giannandrea and Marilena Sansoni. A literature review on intelligent tutoring systems and on student profiling. *Learning & Teaching with Media & Technology*, 287:287–294, 2013.