Second Language Tutoring using Social Robots

Project No. 688014

L2TOR
Second Language Tutoring using Social Robots

Grant Agreement Type: Collaborative Project
Grant Agreement Number: 688014

D5.2 Interaction management for the space domain

Due Date: **31/03/2018**
Submission Date: **13/04/2018**

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 space 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 (13-04-2018)
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.

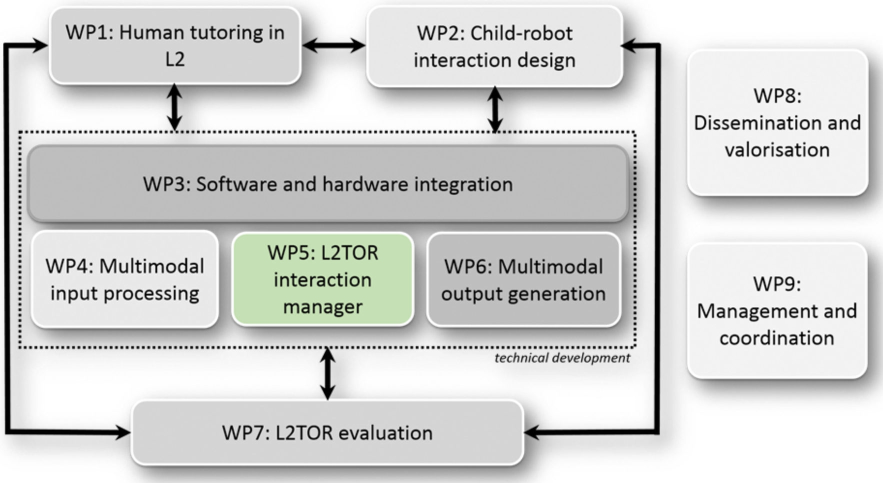

Figure 1: Workpackage overview and how they work together.

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 5.1.

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

As a next step in Task 5.2, an internal "long-term memory" has been implemented. As outlined in Deliverable 5.1, information about the children's answers and the interaction itself should be stored in this part of the system. Therefore, we developed a module plugged into the interaction manager, which is also managing the information flow from and to this memory module. The interaction manager collects the necessary information from the input recognition (WP4) and output-generation modules (WP6) and stores it in the memory, aligned with the current task. Since the knowledge-tracing and affect-tracing modules are not used in the ongoing large-scale evaluation of the L2tor system, this information is not stored in the memory yet. An up-to-date list of all stored information can be found in Appendix A. Because the data stored in the memory will be the basis for all evaluations at the end of the study, we ensured that it is usable even after a system crash. Therefore, we used a json notation instead of a simple binary-file to store the information on the hard-drive. This allows for opening the file by hand and correcting mistakes, if it was not stored correctly during a crash. In cooperation with Softbank Robotics (WP3) we developed a mechanism for storing all the information in different snapshots and copies so that, even if a file is completely corrupted, a snapshot or memory-copy from a previous step in the interaction is still available.

Furthermore, we implemented a so called "checkpoint-system" that uses the already stored information about the current task the child is in, to restart an interaction at, or near the last stored task in the memory. This prevents the child from having to repeat all tasks after a crash occurred, and further ensures that the target words are not presented twice as often as in interactions without a system crash. This is important because a repetition might influence the learning gain measured in the post test. To use the system, the interaction manager first has to run (once) in a specific mode to create checkpoint files for all types of tasks where, e.g., the position of objects or even the whole scene on the tablet will change (to run this mode, a flag has to be set in the start-up configuration of the interaction manager). This means that the interaction manager has to restore the exact status of the tablet game (e.g., objects on the screen, target word) and also "Underworlds" while loading a checkpoint, so a dump of all these information has to be stored at this point in the interaction. Afterwards, if the previous session was not completed yet, the experimenter either has the possibility to start at a specific checkpoint shown in the "ControlPanel" info section after loading a memory (see Fig. 2) or to start the previous session from the beginning again.

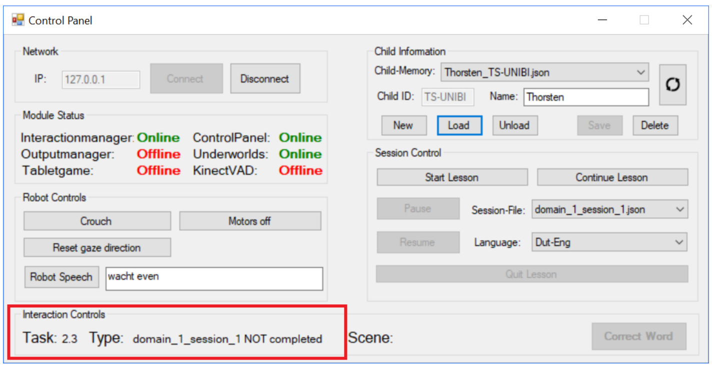

Figure 2: Screenshot of the current version of the experimenters control panel. The red outlined box on the bottom shows the information for the currently loaded memory file. Here it notifies the experimenter, that the previous session 1 in domain 1 wasn't completed and can be continued at task 2.3.

Having all information stored in the memory combined with the cognitive state of the child, it can be used for long-term adaptation of the system. A possible use case for long-term adaptation would be the adaption to a slow or fast learning child. For example, child A needs 5 repetitions on average to memorize a specific word while child B needs only 4. Hence, adaptation to the learning speed of each child is necessary in order to not demand too much or too little from a child. This further means that the knowledge tracing module needs to adapt its update-steps to be a bit smaller (child A) or bigger (child

B). That is, even if this precision error may not influence a short interaction much, it can accumulate in an interaction over several session such that the tracked information are not reliable enough anymore. In addition, storing knowledge about the child will facilitate bonding [1] and establish common ground between robot and child based on prior interactions. Furthermore, the knowledge stored in the memory will be used to make system decisions more transparent to the child by means of commenting and explaining them. It can be assumed that transparency enhances trust in the robot by making its behavior more understandable [2, 3, 4], which in turn should increase the motivation to learn with the robot and thus learning gain (see Section 2.6).

2.3 Basic Interaction Management (T5.3)

The basic interaction management mechanism remained the same (cf. D5.1).

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

As outlined in Deliverable 5.1, another important task for the interaction manager is to use motivational and relational strategies to maintain the engagement of the child [1]. Since not only a bad task performance can influence the motivation of the child, but also tasks that are too repetitive, demanding or boring, it is important to track the affective state of each child during the interaction and to intervene if disengagement occurs. This will be accounted for in the Bayesian Decision Network by a supplementary node that will influence the predictive decision-making process for the course of learning (see Figure 3). That is, if the engagement (E) is too low, a wrong answer (O) to a teaching task (A) is very probable and therefore not beneficial for the learning progress of the skill (S). In this case, it would be more beneficial to use a re-engage strategy (A) first, before teaching the next word, even if this results in higher costs, because of an additional action being performed.

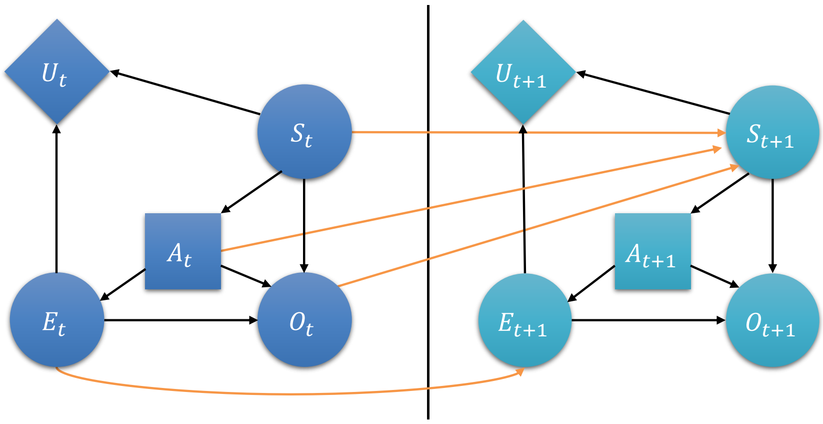

Figure 3: The adaptive decision-making model including the engagement of the child (E), together with skill mastery (S), observations (O), action decision node (A) and and utility node (U) specifying the utility of a specific action for skill S.

To identify the affective states that occur and are important in child-robot tutoring, and how they can be detected based on observations, a qualitative interview study has been conducted (cf. Deliverable 5.1). We used video recordings from a previous study in kindergartens (see [5] for further details) and interviewed five preschool teachers on their perception and interpretation of the children's behavior during the child-robot interactions [6]. Based on the results we were able to identify behavioral cues for engagement or disengagement, upon which most of the asked experts agree. These cues can inform the recognition of the motivational state of a child playing with the robot.

To train an engagement-classifier based on these cues tracked via Microsoft Kinect, we conducted another study in which children's engagement during the interaction had to be rated continuously. Raters were again preschool teachers that watched video recordings of child-robot interactions. By combining the Microsoft Kinect tracking data (recorded during the previous study [5]) with the expert ratings, the parameters for an naive Bayesian classifier can be learned.

2.4.1 Setup & Procedure

For the rating, we used the same study setup as in our earlier expert interviews (see [6] for further details), except that the video was shown in a specially developed rating tool as depicted in figure 4. Ten experts had to rate the child's engagement continuously on a three point scale: 1 = "not attentive at all", 2 = "moderately attentive", 3 = "very attentive" (The labels for the buttons were described in terms of attention to facilitate ratings. However, the raters were informed that we are interested in their perception of the child's engagement, i.e., a state of devotion, attention and willingness to interact with the robot and tablet). The raters also had the option to pause and resume the video by pressing the space-bar and jump back 2 times in 5 second-steps by pressing the backspace key. Videos from four different children (a subset from the expert interviews) were chosen to be rated, two by each expert. That way, we received ratings from different experts for the same video, without overloading the experts with rating all recordings. We further shortened each video-recording, concentrating on the first, mid and last three minutes of the interaction. After watching one part, the tool automatically jumped to the next 3 minutes long part and paused. To continue, the raters had to press the space-bar. When a rater finished the recording of one child, the experimenter asked two questions:

1. What would you do to raise the child's engagement after an engagement drop?
2. What do you think are the robot's possibilities to raise the engagement in the interaction?

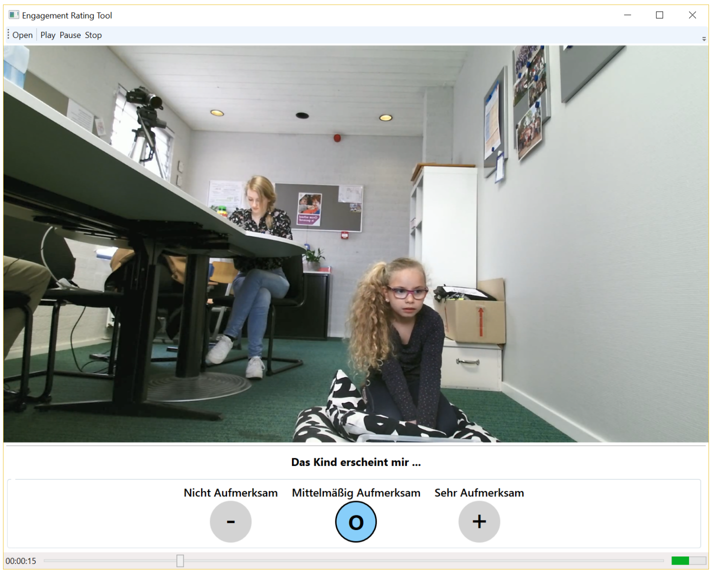

Figure 4: Screenshot of the engagement-rating tool. The buttons below the video represent the current rating (from left: low, mid and high engagement).

2.4.2 Results & Discussion

To validate the quality of the ratings we calculated inter-rater agreement (Fleiss' Kappa) for each video. Despite having high agreement with experts from different kindergartens who have a huge variety in work experience in the first expert interviews, we were not able to achieve sufficient agreement among the experts in the present rating study (see Table 1).

	Video 1	Video 2	Video 3	Video 4
Fleiss' Kappa	0.224	0.338	0.342	0.465

Table 1: Inter-rater agreement values for each of the four different videos presented to the experts.

Several reasons might explain these results. First of all, the amount of ratings assigned to one child varied tremendously (see Table 2). This may be due to the fact that the raters were (deliberately) not forced to give a rating at predefined times. This means, that if the raters did not press a button over a longer period of time, the rating remained the same. Ideally, not changing a previous rating should indicate that the expert does not notice a change in engagement. However, it is possible that the raters overlooked changes in the child's engagement due to distractions, fatigue, or the unfamiliarity with the task. In the future, the ratings should be enforced by the system, e.g., by stopping the video every 10 seconds and asking for an engagement rating.

Rater	Video 1	Video 2	Video 3	Video 4	Mean
Rater 1	14	23			18.5
Rater 2	48	66			57
Rater 3			24	20	22
Rater 4			14	28	21
Rater 5	5	3			4
Rater 6			10	11	10.5
Rater 7			11	14	12.5
Rater 8	6	9			7.5
Rater 9			14	6	10
Rater 10	9		7		8
Mean	16.4	25.25	13.33	15.8	

Table 2: Counts of given (changed) ratings by expert and video.

Furthermore, there can be a temporal dis-alignment between seeing a change in engagement and pushing the button in the interface, caused by different reaction times. In addition, we often saw that the experts started to explain their thought instead of changing the rating and we had to remind them to do so. This behavior might have been favored by the experimental setup and might be resolved by pausing the video in fixed intervals and forcing a rating, before the study can go on (cf. above). Unfortunately, this would result in less labeled data to train a classifier for the same amount of experts. Besides that, the experts might have had different thresholds when switching from low to mid engagement or from mid to high, what might have been caused by their working experience. We showed only foreign children to the experts, but some of them might know children from their kindergarten who show similar behaviors, so that they were able to give more detailed ratings than others for the same child. Furthermore, the experts had to adapt to each child and try to find out the “typical” behavior for the child to classify deviating behaviors as low or high engagement. For instance, if a child moves a lot but is still attentive, it is probably still engaged, but in general more active, while some other children barely move because they are shy and quite. While a lot of movement might be a weak indicator for disengagement of one child, it might be a strong indicator for another. Additionally, this adaptation process might have taken a different amount of time depending on the expert and some of the experts might have adapted better than others.

2.4.3 Other Approaches & Future Work

Other software frameworks which provide a measure of engagement, e.g., AffDex from Affectiva [7] or the approach of Booth et al. [8], use facial features to classify it. However, we saw during our previous experiments, that the children barely show facial expressions while interacting with a robot in absence of a human. Hence, this approach seems unsuitable to detect engagement of children in the age of 4-6 years. Reviewing the literature on emotion classification like boredom and frustration, a lot of approaches use written text [9, 10, 11, 12, 13], the user’s voice [14, 15, 16], physiological cues [17, 18, 19] or even brain imaging [20, 21], which is hardly applicable for child-robot tutoring interactions. On the one hand, cues like spoken sentences or written text are not provided during our interaction and on the other hand, to get the other cues, e.g., ECG data, would make the interaction even more artificial. However, another approach from Kapoor et al. [22] combines facial features with the body language and interaction activity logs, which seems to be most promising with an accuracy of 86%. But they are using a so called “sensing chair” to track the body movements of the user, which is again not applicable in our case. Using the Microsoft Kinect sensor to replace this part of the system would be a possibility, which is exactly what we are aiming for. Although the Microsoft Kinect provides a lot of data, it remains unclear which features should be used for the classification of engagement. Approaches based on deep learning to filter out the important cues would require huge datasets, while we tried to narrow down the feature space with our expert interviews.

Overall, the inter-rater agreement in our labeled data is not promising and there is also no of-the-shelf solution fitting our need. Thus, for the purpose of our interaction management studies, we opt to classify the child’s engagement on the fly by means of a Wizard of Oz interface. That is, a single trained human wizard uses the interface to inform the system about the behavioral cues of disengagement as derived from the expert interviews [6]. Since the reaction to disengagement is critical to ensure a fruitful learning environment, we decided to use this approach instead of an automatic engagement tracking to test different re-engagement strategies in the following.

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)

We currently concentrate on the motivational-relational strategies by testing different possibilities to reach a better long-term motivation of playing and learning with the robot. We try to reach this by implementing preventive as well as repair actions to counter disengagement. An extensive study of the literature on preschool second language learning was carried out to identify possible actions.

The results suggest that communicating underlying decision processes to children can have a positive impact on their language learning [23]. Therefore, we have worked on enabling the robot to explain its inner reasoning (based on the current beliefs represented in the decision-making model) so that the child understands the course taken during vocabulary learning. This strategy should allow children to gather an idea about their own knowledge state concerning each target word, which should scaffold their learning strategies for future vocabulary learning. In addition, we assumed that explanations provided by the robot make the children perceive it as a competent partner with reasonable decisions, which in turn should foster their trust in the robot and hence increase their motivation to interact and learn [4, 2, 3].

To test these strategies we consider in a first study (currently running) the impact of verbalizing system's decisions on the children's liking of the robot, their motivation to continue learning with the robot, and their ability to estimate their own vocabulary knowledge after the interaction. Comparing different verbalization strategies is out of the scope of this exploratory study, but can be considered in WP6 later on. We are currently carrying out a study with two conditions: one with and one without having the robot explicitly verbalize (explain) the decisions made by the adaptive system.

2.6.1 Participants

We are aiming for 40 children aged 5 to 6 years (20 per condition) in German kindergartens.

2.6.2 Study Design

Preventive Strategies: As suggested by our expert interviews [6], preventive strategies such as allowing verbal input by the child and heightening the activity of the robot should be included in child-robot interactions to raise engagement from the beginning. Hence, we modified the animal game in line with the recommendations gathered from experts. Verbal input by the child is encouraged by asking the child to repeat the target words. Therefore, the robot introduces the words and then asks the children to repeat the English word (e.g., "Say ladybug"). Further, the robot's activity has been heightened in three ways: (1) The robot switches its gaze between the tablet and the child (look at the tablet when the child has to choose an animal on the tablet, look at the child after the child had chosen and feedback about the answer is provided). (2) The robot immediately reacts with a feedback concerning the child's answer. If the answer is correct the robot praises the child and displays blinking eyes in rainbow colors. If the answer is wrong, the robot translates the target word one more time and asks the child to repeat it in English. Afterwards the child has to choose the target animal on the tablet again, but this time only out of two options (the target animal and the one that was wrongly chosen before). This procedure will be repeated until the correct answer was provided. (3) If the child does not react (i.e., answer verbally or touch an object on the tablet) the robot addresses the child by his name. For that purpose, the wizard will enter the child's name into the system at the beginning of the interaction to allow the robot to address the child directly during the interaction.

Since the inclusion of gestures has been demonstrated to enhance word learning [5], but the execution of a gesture for each target word presented would lengthen the session too much, we decided to introduce each target word with a gesture in the beginning (further heightening the activity of the robot), but use the words without gestures during the remaining session.

Re-Engaging Strategies: In addition to the preventive strategies, situational actions to re-engage the child were included. The wizard monitors the child's behavior during the interaction (cf. wizard interface) and thus can trigger a re-engagement action by clicking on behavioral cues of disengagement (signs of fatigue, looking away, heightened body movements by the child, and turning away). The adaptive system then randomly chooses one of several possible re-engagement actions according to the cause, e.g., stand up and stretch if a sign of fatigue was remarked.

Decision Verbalization: While we use the preventive strategies for each child and allow for situated reactions to disengagement to test whether these interventions help to keep the children engaged, we divide participants into two groups to test the impact of the robot given explanations. The question of interest is how verbalization of the system's decisions will affect the perception of the robot, the willingness to interact with it, and finally learning gain. Therefore, two conditions are being compared: with and without verbalization. In the verbalization condition the robot explains

1. why an upcoming task is chosen (based on knowledge tracing), and
2. why an action to re-engage the child is started (based on behavioral cues of disengagement monitored by a wizard)

Measurements (dependent variables):

- engagement: how often a re-engagement strategy has to be triggered
- willingness to interact: amount of rounds played
- learning gain: amount of words learned after the interaction (immediately and delayed)
- liking of the robot: how many grapes children share with the robot after the interaction
- estimation of knowledge state: Children's own estimation of how well they already know each target word after the interaction.

Hypotheses: We assume that the children's engagement continuously decreases over the course of a longer learning interaction. We further assume that preventive strategies, as described above, help to slow down this decrease. In addition, evidence from previous work revealed that especially unrelated activities, such as standing up and stretching, can be used as re-engagement strategies as they lower boredom and distraction and increase the willingness to continue learning (cf. [6]). Furthermore, we assume that children are even more willing to continue interacting with the robot, if the robot explains its decisions. We therefore hypothesize:

H1: Children will be more engaged in the interaction if the robot explains its internal decisions.

H2: Children will be more willing to continue interacting with the robot if it explains its decisions.

H3: Children will show a higher learning gain after interacting with a robot that explains its decisions due to a higher engagement and willingness to interact with the robot

In addition, the literature on transparency (e.g., [4] suggests that user's trust a system more and will thus like it more if it uses transparency cues such as explanations. We hence assume that a robot that explains its decisions will be more liked by children.

H4: Children will demonstrate a higher liking of a robot that explains its decisions.

Finally, explaining the knowledge state the robot assumes for the child at each step (e.g.: ”I choose a word of which I am certain you already know it well.”) should foster a better understanding of the child’s own knowledge. Conclusively, the child will be more capable of estimating its own vocabulary knowledge after interacting with a robot that explains its knowledge.

H5: Children will be more capable in estimating their own vocabulary skills after the interaction, if the robot explained its estimation of the child’s knowledge.

2.6.3 Experimental Setup

The experimental scenario is based on our earlier developed “animal experiment” (see [5] for more details), in which children learn animal names in an “I spy with my little eye” game. During the experiment, the child is sitting in front of a tablet while the Nao robot from Softbank Robotics is standing vis-à-vis (see Figure 5, left). During the interaction, the tablet shows the content and accepts touch input, but all instructions are given by the robot, which reacts to the content, shown on the screen. The whole interaction is recorded by two cameras from different angles. One camera is recording the child’s front, i.e., the face and hands, while the second camera records the tablet screen and the robot.

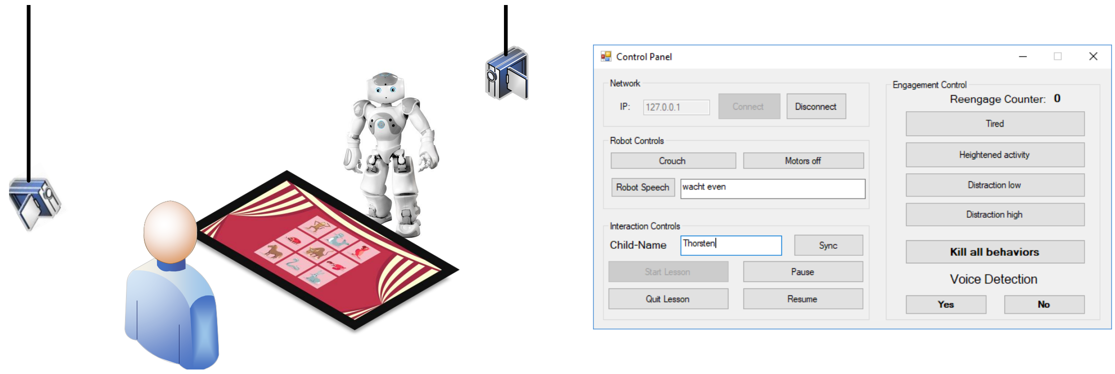

Figure 5: Left: study setup with tablet, robot, child and 2 cameras. Right: WOz GUI used by the experimenter.

The order of the skills being taught as well as the task difficulty for the selected word, is chosen by our adaptive system (see [24] for more details). In contrast to the earlier setting, we embedded the game in a circus environment to make the interaction more appealing to children. We added and changed animals to ensure that the vocabulary was not too easy (e.g., removed “chicken”, because most children in earlier experiments knew the word in advance), and we elaborated the variety of task difficulty by controlling not only for the amount of animals presented simultaneously, but also considering the animals’ colors as a distractor. That is, we assumed that choosing an animal from three animals with the same color is more difficult than choosing out of three animals with distinct colors. Finally, we end up with a set of 9 animals (snake, monkey, lobster, parrot, horse, rabbit, seal, ladybug and bull; see Figure 6) and 4 different difficulties (3 animals with different colors, 3 animals with the same color, 6 animals and 9 animals).

We further adapted the number of rounds to the knowledge and affective state of the child. The maximum number of rounds children play with the robot remained the same (30 rounds), but the system is now able to end the learning interaction earlier if the child’s knowledge is estimated to be good enough (skill mastery above 75 percent on average). If the system decides to end the interaction before scheduled, it will start a short recap of the 3 weakest words for the child before the interaction ends. If the child did not make mistakes on at least 3 skills (weak words), the list of recap skills will be

Figure 6: Tablet screen for the highest task difficulty with all 9 different animals used in the experiment.

filled up by random picks of the taught skills. In addition to ending the game due to knowledge gain, the interaction is terminated if the child is highly disengaged (as spotted by a wizard based on the cues identified in the expert interviews). If the wizard notices a cue of disengagement, s/he pushes a button in the wizard interface (see below) which triggers a re-engagement action executed by the robot. After each attempt to re-engage the child, the robot asks if the child is willing to continue interacting with the robot. If not, the robot thanks the child and introduces the post-test.

Wizard Interface: The wizard can use four different buttons to enter disengagement cues while the child is playing the main learning game (see Figure 5, right): fatigue, high activity, low distraction, and high distraction. Each cue is linked in the adaptive decision-making to different re-engagement strategies extracted from interviews in our previous work (see [6] and Section 2.4 for more details).

2.6.4 Procedure

A few days before the actual learning interaction: The robot is introduced to the children in a group session. This first meeting between children, robot and researchers is designed to help the children accustom to the new situation. The researchers introduce themselves and the robot, who then talks to and dances with the children. The children are also invited to ask questions, which the experimenters will answer. As previous experience has shown, this takes away initial fears the children might experience when immediately interacting with the robot on their own (cf. [25]).

At the actual day of the study: After entering the room, the child sits down at a table together with an experimenter. During a short conversation the child is asked if she/he knows already all German words for the 9 animals selected to be learned during the interaction. Afterwards a recognition pretest is conducted in which the child is asked to point at a specific animal on a piece of paper while only the English name for this animal is given. After finishing the pretest, the child sits down in front of the tablet and gets a short introduction before the interaction with the robot starts. At the beginning, the robot introduces itself and asks for the child's name and age. The child is supposed to answer verbally, while the input is wizarded by one experimenter.

After a short familiarization, the robot introduces all animals separately by showing an image on the tablet, using an iconic gesture and pronouncing the German and English name of the animal. Afterwards, the child is asked to repeat the English name. During the whole interaction all English words used by the robot are phonetically transcribed in the alphabet of Acapela[1], which is the underlying TTS engine for the German voice, so that no voice switch is necessary and no difference between

[1] http://www.acapela-group.com/

German and English words is perceptible.

Subsequent to this introduction, the main learning interaction begins. The robot explains the game and start with a short test-run to make sure the child understands the rules of the game. Afterwards the real game starts. At the end of the learning game, the robot conducts a post-test by means of telling each animal name again and ask the child to feed a grape to said animal. Therefore a bunch of grapes is displayed next to the animals and the children are requested to drag and drop the grapes on the animal referred to by the robot (in English). During the post-test neither positive nor negative feedback is given in order not to affect the children. The acceptance of the answer is noted by a simple neutral comment, e.g., “Okay, next!”. After the post-test the interaction with the robot ends, and one experimenter gives 5 real grapes to the child. The child then is asked if she wants to share these with the robot who the child is told likes grapes (the experimenter counts how many grapes are shared with the robot). Finally, the experimenter conducts a short interview with the child about his/her perception of the robot and thanks the child for participation.

Around two weeks after the main interaction a second post test will be conducted (retention test) to compare the learning gain immediately after the interaction to the gain after some time has passed. This is done because we observed in earlier experiments that the learning gain after a week or two can be great although the gain measured immediately after the interaction was low.

3 Conclusion

In summary, the goal of WP5 is the development of an interaction manager for the L2TOR system, which is responsible for 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 give an update on the changes for the different tasks compared to Deliverable 5.1 and also what still needs to be addressed in the future.

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

- **(2) Interpretation (Tasks 5.2, 5.4, 5.7) and (3) Decision Making (T5.5)**
  A memory module, that is capable of storing all necessary information about the child (e.g., name and birthday) and the interaction (e.g., last task and given answer), has been developed and integrated into the system. In addition, a checkpoint system was implemented, that allows for restarting the interaction at a set point, after a system breakdown. Besides the work on the memory, we put effort into the understanding of and reaction to children’s engagement. A continuous rating of engagement by experts was collected in combination with Microsoft Kinect data, meant to provide input for training classifiers. Unfortunately, the inter-rater reliability was unusually low and ratings were difficult to compare, so that we discarded our attempt to train the classifier with this data and will continue the test of re-engagement strategies as derived from expert interviews, without automatic engagement recognition. Instead, we will have a trained wizard monitoring the behavior of the children and choose out of a given set of cues of disengagement, which the experts from our previous investigation agreed upon. We are currently running an experiment that incorporates re-engagement strategies into child-robot language

tutoring. These are necessary as a decreasing engagement over time can be expected in general. In addition, we are about to test the impact of transparency (i.e., verbalization of the systems decisions) on childrens perception of the robot, learning gain, and estimation of their own skills.

- **Future work**
  We continue to work on extending the predictive decision-making process (based on the Bayesian Decision Network) to allow for the children's engagement to also influence the course of the learning interaction. As part of this, we will integrate the gained knowledge about the effects of explanations and verbalization of system knowledge. In addition, we plan to adapt the decision-making to the child learner on two different timescales, a short timescale during the current interaction (existing mode) and on a longer timescale spanning over several learning sessions. The goal is to develop a hierarchical decision-making process in which the higher-level policy (longer-term, based on the information stored in the memory module developed in T5.2) can adjust the short-term adaptation policy.

References

[1] Tony Belpaeme, Paul E Baxter, Robin Read, Rachel Wood, Heriberto Cuayáhuitl, Bernd Kiefer, Stefania Racioppa, Ivana Kruijff-Korbayová, Georgios Athanasopoulos, Valentin Enescu, et al. Multimodal child-robot interaction: Building social bonds. *Journal of Human-Robot Interaction*, 1(2):33–53, 2012.

[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 *Advances in Human Factors in Robots and Unmanned Systems*, pages 127–136. Springer, 2017.

[3] Joseph B Lyons. Being transparent about transparency. In *AAAI Spring Symposium*, 2013.

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

[5] Jan de Wit, Thorsten Schodde, Kirsten Bergmann, Emiel Krahmer, and Stefan Kopp. The Effect of a Robot's Gestures and Adaptive Tutoring on Children's Acquisition of Second Language Vocabularies. 2018. Submitted to ACM/IEEE HRI 2018.

[6] Thorsten Schodde, Laura Hoffmann, and Stefan Kopp. How to Manage Affective State in Child-Robot Tutoring Interactions? 2017.

[7] Daniel McDuff, Abdelrahman Mahmoud, Mohammad Mavadati, May Amr, Jay Turcot, and Rana el Kaliouby. Affdex sdk: a cross-platform real-time multi-face expression recognition toolkit. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 3723–3726. ACM, 2016.

[8] Brandon M Booth, Asem M Ali, Shrikanth S Narayanan, Ian Bennett, and Aly A Farag. Toward active and unobtrusive engagement assessment of distance learners. In *Affective Computing and Intelligent Interaction (ACII), 2017 Seventh International Conference on*, pages 470–476. IEEE, 2017.

[9] Jeffrey T Hancock, Christopher Landrigan, and Courtney Silver. Expressing emotion in text-based communication. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 929–932. ACM, 2007.

[10] Cleveland G Shields, Ronald M Epstein, Peter Franks, Kevin Fiscella, Paul Duberstein, Susan H McDaniel, and Sean Meldrum. Emotion language in primary care encounters: reliability and validity of an emotion word count coding system. *Patient education and counseling*, 57(2):232–238, 2005.

[11] Jeffrey H Kahn, Renee M Tobin, Audra E Massey, and Jennifer A Anderson. Measuring emotional expression with the linguistic inquiry and word count. *The American journal of psychology*, pages 263–286, 2007.

[12] James W Pennebaker, Matthias R Mehl, and Kate G Niederhoffer. Psychological aspects of natural language use: Our words, our selves. *Annual review of psychology*, 54(1):547–577, 2003.

[13] Michael A Cohn, Matthias R Mehl, and James W Pennebaker. Linguistic markers of psychological change surrounding september 11, 2001. *Psychological science*, 15(10):687–693, 2004.

[14] Chul Min Lee and Shrikanth S Narayanan. Toward detecting emotions in spoken dialogs. *IEEE transactions on speech and audio processing*, 13(2):293–303, 2005.

[15] Laurence Devillers and Laurence Vidrascu. Real-life emotions detection with lexical and paralinguistic cues on human-human call center dialogs. In *Ninth International Conference on Spoken Language Processing*, 2006.

[16] Laurence Devillers, Laurence Vidrascu, and Lori Lamel. Challenges in real-life emotion annotation and machine learning based detection. *Neural Networks*, 18(4):407–422, 2005.

[17] Johannes Wagner, Jonghwa Kim, and Elisabeth André. From physiological signals to emotions: Implementing and comparing selected methods for feature extraction and classification. In *Multimedia and Expo, 2005. ICME 2005. IEEE International Conference on*, pages 940–943. IEEE, 2005.

[18] Rosalind W. Picard, Elias Vyzas, and Jennifer Healey. Toward machine emotional intelligence: Analysis of affective physiological state. *IEEE transactions on pattern analysis and machine intelligence*, 23(10):1175–1191, 2001.

[19] Olivier Villon and Christine Lisetti. A user-modeling approach to build user's psycho-physiological maps of emotions using bio-sensors. In *Robot and Human Interactive Communication, 2006. ROMAN 2006. The 15th IEEE International Symposium on*, pages 269–276. IEEE, 2006.

[20] Mary Helen Immordino-Yang and Antonio Damasio. We feel, therefore we learn: The relevance of affective and social neuroscience to education. *Mind, brain, and education*, 1(1):3–10, 2007.

[21] Jonas K Olofsson, Steven Nordin, Henrique Sequeira, and John Polich. Affective picture processing: an integrative review of erp findings. *Biological psychology*, 77(3):247–265, 2008.

[22] Ashish Kapoor and Rosalind W Picard. Multimodal affect recognition in learning environments. In *Proceedings of the 13th annual ACM international conference on Multimedia*, pages 677–682. ACM, 2005.

[23] Lev S Vygotskij. Denken und sprechen. psychologische untersuchungen. herausgegeben und aus dem russischen übersetzt von j. lompscher und g. rückriem, 2002.

[24] Thorsten Schodde, Kirsten Bergmann, and Stefan Kopp. Adaptive Robot Language Tutoring Based on Bayesian Knowledge Tracing and Predictive Decision-Making. In *Proceedings of ACM/IEEE HRI 2017*, pages 128–136. ACM Press, 2017.

[25] 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, et al. Guidelines for designing social robots as second language tutors. *International Journal of Social Robotics*, 2017.

A Information stored in the memory

General Information	Description
id	An unique id for the current child assigned by the experimenter.
name	The name of the child.
phon_name	The phonetically written name of the child, so that the robot can correctly pronounce it.
birthday	The birthday of the child.
gender	The gender of the child.
lang	The mother-tongue spoken by the child.
school	The name of the school the child goes to and where the study takes place for this specify child.
lang_combi	The language combination the child get thought, e.g., Dutch-English.
condition	The study condition the child is in, e.g., “Robot + Iconic”.

Session Information	Description
domain_id	The id of the domain the tasks information belong to.
session_id	The id of the session the tasks information belong to.
timestamp	An absolute timestamp when the session started.

Task Information	Description
task_id	The id of the task to which the information belong.
is_test	A flag which will help to easily identify the post-test tasks in the different sessions.
timestamp	A timestamp relative to the beginning of the session.
type	The type of the task, e.g., “OBJECT_MOVE_CRITERIUM”.
request_answer	A list containing information about every time the child has been requested to answer including: used_l2_words and timestamp.
give_help	A list containing information about every time the robot provided help including: used_l2_words and timestamp.
answer	The given answer including the info.: correct_answer, given_answer, correctness, abs_timestamp and rel_timestamp (response_time)
feedback	A list of all feedbacks given during the task including the info.: is_positiv, variation_id, used_l2_words and timestamp
gestures	A list of all gestures used during the task including the info.: is_iconic, name and timestamp.