Ref. Ares(2019)4486762 - 11/07/2019

DR 2.4: Integration and refinement

Joost Broekens, Rifca Peters, Frank Kaptein, Mark Neerincx

TU Delft, Delft, The Netherlands

‹joost.broekens@gmail.com›

Project, project Id: EU H2020 PAL / PHC-643783

Project start date: March 2015 (48 months)

Due date of deliverable: 29-mrt-2019

Actual submission date: 29-mrt-2019

Lead partner: TU Delft

Revision: final_with_lessons_learned

Dissemination level: PU

In this document we report the integration and refinement of the different systems developed in work package 2 into a consistent whole that is used in an integrated manner in the PAL system. In the PAL project all developments from the work packages have been integrated in early stages in an iterative manner, providing a running system that can be used by the different stakeholders. As such, all outcomes from WP2 were immediately integrated in the PAL system. This means, for example, that the pal authoring tool (PAL control, Task 2.1) was already integrated in the PAL system when it was ready. In this document we will highlight how the final developments of robot style, action explanation, PAL inform and PAL control were integrated in the final experiment in 2018/19.

Executive Summary

Work package 2 has researched and developed 4 main functionalities that resolve around strategic learning goal setting for children, transparency of the system (explanation and style), and progress monitoring of the child. The first is PAL Control (see deliverable D2.1), an interface for children and health care professionals (HCP) to set learning goals for the children. The second is the parental monitoring interface, palInform (see D2.3), which is the software module in the PAL project that enables parents (and health care personnel) to provide insight into how their children use the system, their progress, and what glucose control & dietic values they fill in. The third is a model for robot and virtual agent behavioral style to express behaviors in a more personalized and task-compatible manner (D2.2). The fourth is an explanation method to allow the PAL system (avatar and robot) to explain the actions it proposes to the child (D2.2). In this document we describe how these four functionalities have been integrated with each other and with the PAL system. In general, the integration is guided by the learning goal structure, with at its basis the diabetes learning goal ontology [1]. The ontology describes which goals can be active for the kid, which tasks can be performed to achieve those goals, and what type of goal it is. Learning goals can be set with PAL Control. Learning progress on tasks and goals can be monitored with palInform. The robot's behavioral style (how it expresses itself) can be deduced from the type of learning activity. The explanation of why a robot proposed to the child to perform a particular learning task is generated based on the relation between the task and the goal it helps to achieve. All functions have been implemented and used during the last long term experiment in 2018/19 in Italy and The Netherlands in this integrated manner. The system worked as intended and indeed proposed content and explanations to the children based on individual goal settings. Here we report in more detail on how these functions interact with each other, and how they contribute to the PAL system. For a more detailed overview on how al modules of the PAL system are integrated and deployed, a journal article is in preparation [Annex 5].

Tasks, objectives, results

1.1 Planned work

The work planned for this deliverable relates to Task 2.5 (but also to the other tasks in WP2). It focusses on the integration in the PAL system of the different functions developed in WP2 and on refinement of the different functions developed in WP2.

1.2 Actual work performed

Planned work has been carried out as planned. Integration of the different functions is guided by the learning goal structure, with at its basis the learning goal ontology. The ontology describes which goals can be active for the kid, which tasks can be performed to achieve those goals, and what type of goal it is. Learning goals can be set with PAL Control. Learning progress on tasks and goals can be monitored with palInform. The robot's behavioral style (how it expresses itself) can be deduced from the type of learning activity. The explanation of why a robot proposed to the child to perform a particular learning task is generated based on the relation between the task and the goal it helps to achieve. All functions are operational in the final PAL system that has been running for the 2018/19 final experiment in the Netherlands and Italy. To better understand the integration, we first summarize how each of the 4 different functions was integrated in the PAL system. After that we address how the 4 functions integrate together. Refinement has been done as well over the course of the project. In particular, we have developed more sophisticated ways of expressing style through robot behavior (Annex 1) based on previous work in the PAL project [2]. We have continued to refine our earlier explanation model based on beliefs and goals [3] to also include emotions [4, Annex 2]). We have refined palInform and PAL Control based on feedback form earlier evaluations.

PAL Control Integration

PAL Control (D2.1, M2.1) is an authoring tool for care professionals. This tool enables health care professionals (HCP) to set learning goals together with children during meetings. It further enables the HCP to enter child data including personal data and preferences such as sports and hobbies. This is important for engagement and also contextualized goal setting. Being able to set learning goals appropriately for children contributes to focusing on an achievable set of goals and contributes to structuring the learning process. Three important issues were tackled: how to design the interface (and mechanisms behind) to facilitate easy data entering during an intake conversation, how to support gamification of goal achievement, and, how to formalize the learning goals based on the medical protocols in the system and interface. We have evaluated (with children and healthcare professionals) the PAL Control tool developed (D2.2). The evaluation took place as a qualitative study to find out if the structure and interface was suitable for health care professionals to set goals for children. Based on this evaluation some changes have been made mainly to increase clarity to the users on the goal/task/achievement hierarchy, including a 90 degrees rotation of the goal tree, and, collapsible achievements for overview clarity.. The final result is shown in Figure 1. The goals set by the child and care professional have been integrated as guiding force for the content of the quiz module of the PAL system

(WP4) and enables the system to propose educational quizzes dependent on the context as defined by the learning goals. It is also the driving force behind the explanations generated by the system for why the system proposes to the child a task (see below).

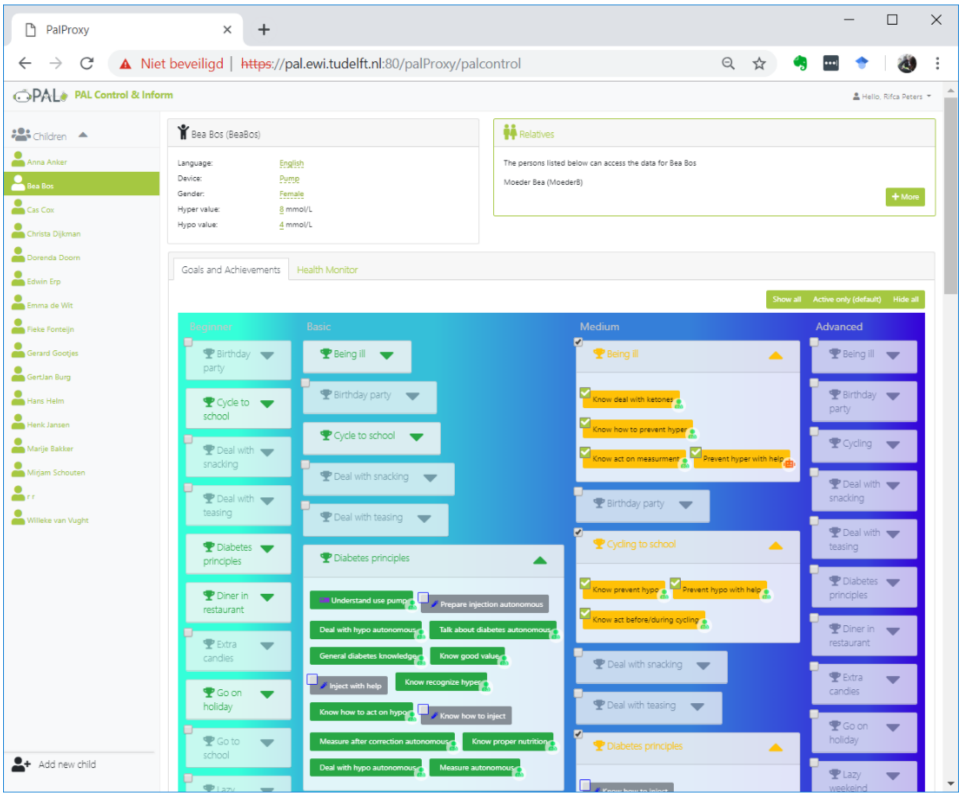

Figure 1. PAL Control goal selection and progress view (green achievements and goals are achieved, achievements are marked with a trophy icon).

palInform prototype Integration

The parental monitoring interface, palInform, is the software module in the PAL project that enables parents (and health care personnel) to get an insight into how their children use the system, their progress, and what dietic values they fill in. It provides a timeline of the most important events, based on system activity and the data the child enters, in an aggregated manner. In this period PAL Inform has been integrated into the pal Control web application. The monitor shows health, nutritional, activity and emotion related data for a child, as entered by the child via the timeline in the MyPAL app. The monitor is available for both healthcare professionals (HCP) and the parent(s). The child and parent can set agreements on what information will or will not be available. Requirements were developed together with the hospitals in the PAL project. palInform is integrated in the PAL system and uses the general goal / task ontology underlying the data in the system (see appendix paper).

In the first evaluation we gathered several recommendations for bettering the tool (see D2.3). The most important ones were integrated in the final version that was used during the final experiment in Italy and The Netherlands. These are shown in Table 1.

Table 1. Changes made after the pilot evaluation with health care professionals reported upon in D2.3.

Desired functionality	1. Glycaemic measurements and value per measurement. 2. Frequency of hypo and hyper vs within target values. 3. Insulin intakes and dosage per intake. 4. Carbohydrate intake (per meal). 5. Activities done. 6. Short-term overview of relations.	All kept as implemented
Not needed	1. Quarterly view. 2. Summary of carbohydrate intake 3. Summary of insulin total and amount	1. Still in but not hindering in interface either. 2. Changed to nr of meals per day. 3. Change to averages
Missing	1. Two week overview 2. Personalized target values for glycaemia 3. Summary of average number of insulin intakes and dosage (units) per day 4. Summary of number of positive and negative emotions 5. Activity duration 6. Goal/achievement description 7. Goal/achievement activation date (and completion time)	1. Implemented 2. Implemented 3. Implemented 4. Implemented 5. Not implemented 6. Implemented 7. Partly implemented (activation date was not).

In the 2018/19 final experiment, palInform has been evaluated by parents in Italy and the Netherlands. The version that was evaluated is shown in Figure 2. The main findings of this evaluation include: parents do not use palInform because they (a) see no need for it, or (b) they look at the myPal kids app and see the goals and progress there and (c) not being able to login. To give an idea of the findings, we report on the Italian data here. Parents perceive palInform to be useful and usable, the usability average score was quite high, 71%, n=14, but 16 of the 18 parents used the app never or only once per month. The high usability score can be explain by the nature of the questions: most questions ask about hypothetical use in case the palInform app would fulfil a need. Apparently parents do not feel the need for a separate overview app. This can be explained in two ways (a) they trust their children, (b) they look at the myPAL app together with the kids and see the progress there. Both reasons were explicitly mentioned by parents in the feedback forms. In general we believe that parents of young children are already overloaded with apps and websites for managing aspects of their and their kid's lives, and unless there is a very clear need, no additional tools should be imposed on them.

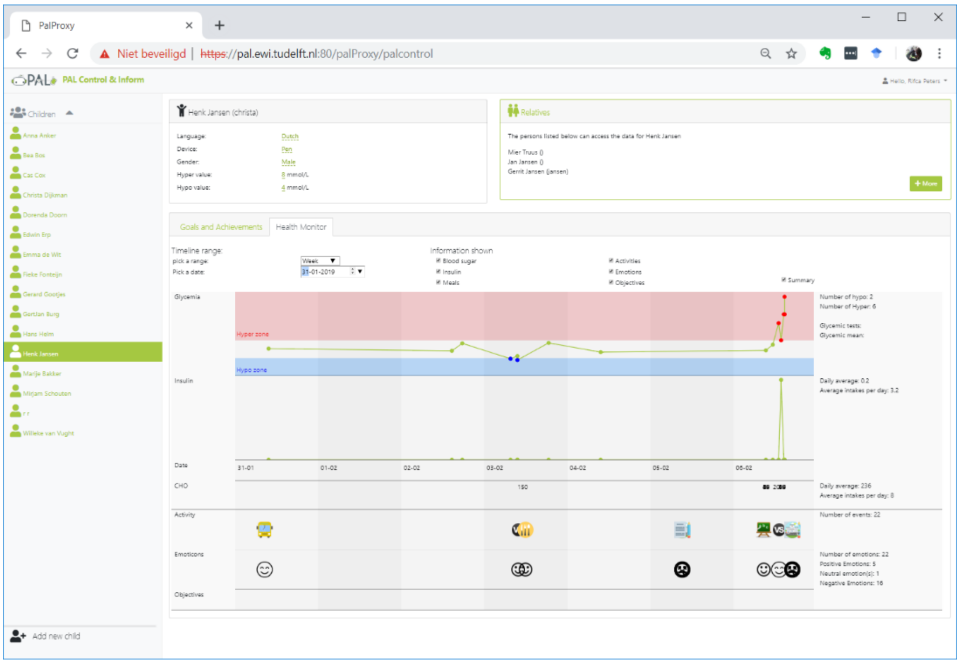

Figure 2. Pal Inform as used in the 2018/19 final experiment.

Behavioral Style for Robots Integration

We investigated if children can perceive differences in robot educational style [2]. This is important for engagement because human teachers and caregivers use different styles to personalize their behavior towards children. This ability to adapt educational style to an individual increases educational efficiency and effectiveness. If we want the PAL system to be able to maximize learning efficiency and effectiveness, personalization of educational style is important. However, for style to have an impact, we first need to know if style is perceived at all. In four experiments, one at two different primary schools, and one at the autumn camp in the Netherlands [2], and two online with MTurk participants [Annex 1, Annex 4], we tested if the perception of style could be manipulated. These experiments showed that this is indeed possible, effects of style were small but significant in the first two experiments, while in the third online experiment perception of style was very present [Annex 1]. In the fourth experiment [Annex 4] we showed that the effect of style (communicated warmth) was significant on the appraisal of a negotiation (the task we selected online to measure effect of style). This showed evidence that style can influence the user.

The results (with the exception of the last due to timing) were integrated in the PAL system for the 2018/19 final experiment in The Netherlands and Italy. When an activity in the myPAL app is started, the style of the robot/avatar changes based on the type of activity. For collaborative and playful activities the style is friendly, while for critical activities such as filling in the timeline the style is direct (see Table 2). Style modulation has been implemented according to Table 3, and based on a review of style in robots in agents [Annex 3]. The style is chosen to support the activity.

During this experiment we also evaluated potential effects of style expression on the usage of the system. Children (within-subject condition to not influence the main PAL research question) were presented with or without style-based expression of the robot during different periods of the experiment. After the first session in the hospital, when children got acquainted with- and played a game with the robot, we collected children’s perception of the robot. These activities have complementary styles, and perception is formed from the experience as a whole, therefore a specific outcome was not predicted for the activity-based condition. Rather we looked if implementing style, and style changes within the session, resulted in a different perception from the neutral condition. Perception seems cultural dependent, with significant differences ($p<0.05$) between the conditions on all factors (perceived affiliation, warmth and competence) except dominance in the Netherlands and no significant differences in Italy. This stresses the importance of contextual validation of expressive robot behaviour. A first look at the usage-data over the full experiment period taught us that overall there does not seem to be a difference in usage time due to behavioral style (no effect of style on the sum of the time a child spend on the tablet).

Table 2. mapping of styles to MyPAL app activities.

Activity	Style
Quiz	Friendly
Games (Memory, Break & Sort, Sorting)	Friendly
Timeline	Direct
Objectives	Direct
Shop	Friendly
Video viewer	Direct
Dance creation	Friendly
Off activity talk	Direct
Introduction	Friendly

Table 3. Motion parameters and values per style.

Style	Style parameter	Gaze	Gesture Size	Gesture Openness	Hand-position Horizontal	Hand-position vertical	Head Tilt	Gesture Speed
Neutral		fixed	medium	centred	sideways	centred	straight	medium
Friendly	(friendliness)	towards	large	centred	peripheral	slight up	up	fast
Direct	(dominance, competence)	towards	medium	open	peripheral	slight up	straight	medium

Explanation Generation Integration

The PAL-system needs to facilitate human-agent collaboration by explaining its users why the Pal-agent behaved in particular ways. This is part of task 2.1, and part of WP2. The PAL-agent autonomously interacts with the children for prolonged periods of time. It helps them to cope with medical health issues. Explainable AI (XAI) facilitates increases in a user's trust and understanding in the system they are working with [5, 6]. Furthermore, it facilitates shared patient-caregiver responsibility for the child's diabetes regimen at preadolescence we developed and tested several techniques to generate robot action explanations. First we have shown in an experiment at a diabetes autumn camp that children and adults prefer robot actions to be explained differently [3]. Although both children and adults prefer goal-based explanations over belief-based explanations (e.g., I propose to play a quiz with you because I want to play a game with you, vs., I propose to play a quiz with you because I know you like to play games). This is important fundamental knowledge enabling us (in the following years) to define how the agent should explain its strategy to the children, parents, and healthcare professionals. Second we investigated the potential role of using emotions in the explanation, or (when simulated by the robot) as source for the explanations [4]. We proposed the use of emotion words in explanations [4] to investigate the effect on system usage: emotion-enhanced explanations or cognitive explanations. This has been implemented in the PAL system (Figure 3) and has been tested during the last experiment in Italy and The Netherlands. Children (within subject condition to not influence the main experiment) were either shown emotionally worded explanations or cognitively worded explanations. The explanation content itself was generated based on the goal-task ontology. For example, if the Nao-robot told Jimmy to take a dextrose when experiencing a hypo it can explain this as follows: I want to teach you to cope with having hypos (which is a goal based cognitive explanation in this case because knowing to take a dextrose when in a hypo contributes to the learning goal of being able to self-manage hypos). A first analysis showed that affective style seems to be associated with a higher probability of requesting a tip-of-the-day by children (Figure 4), i.e., P(click for second tip | style=affective) > P(click for second tip | style=cognitive). However, this effect did not proof to be significant. An interesting effect that is significant is that counter to our expectations, additional explanation (cognitive or affective) is associated with less probability to select the proposed learning task ($p<0.05$). We are still analyzing the meaning of these results.

Figure 3. PAL virtual agent explaining why it proposes the first activity for the child. The robot text balloon shows a goal based explanation for proposing a particular quiz. The text roughly translates to: "I try to achieve that you learn when to ask for help when you encounter problems related to your diabetes".

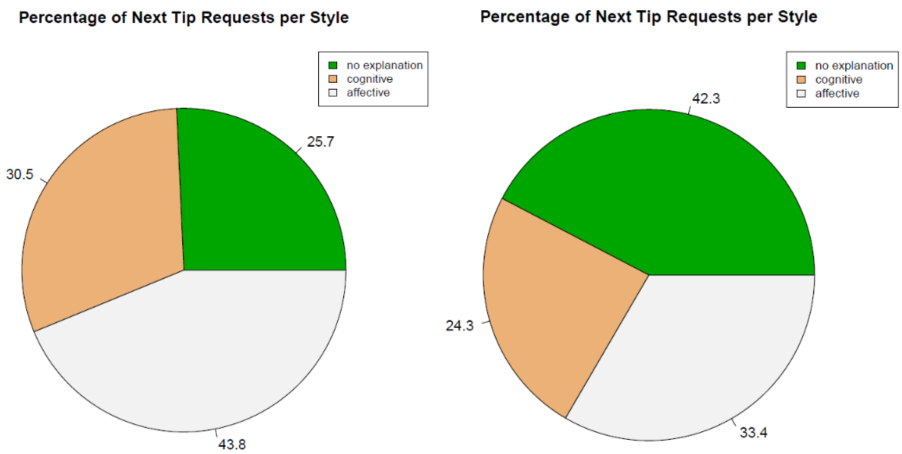

Figure 4. Pie charts (left: NL, right: Italy) show the distribution of second tip-of-the-day requests voluntarily asked for by children after receiving a first tip-of-the-day, i.e., P(click for second tip | style). *Affective style seems to be associated with a higher probability of selecting a tip-of-the-day, however this effect did not proof to be significant.*

Joint Integration

As mentioned above, the integration of the four key components developed in WP2 resolves around the shared ontology of the PAL system and in particular the use of goals, tasks and achievements that are structured therein.

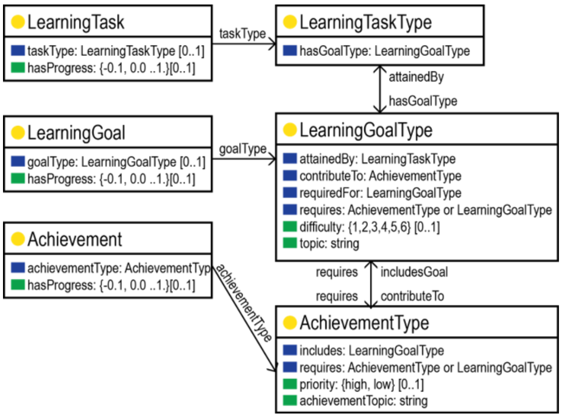

Figure 5. Diabetes learning goal ontology.

In figure 5 a graphical representation is given of the type of elements in the ontology. Learning goals are classified by type and values are given for difficulty and topic. Additionally, restrictions are added for prerequisite goals. (The progress and state properties are specific to a child and values are given at a later time.). Further, achievements (e.g. ‘Cycle to school’ in Figure 1) are added for each topic and difficulty combination, and tasks (e.g., ‘win a quiz on insulin’, not shown in Figure 1) are linked to learning goals (e.g., ‘know how to act on hypos’ in Figure 1). The restrictions and relations allow the system to provide personalized content, and calculate and update goal progress automatically. Goal progress is then visualized in PAL Control and palInform.

Goals are structured at levels of difficulty (Figure 1, left to right). Attaining all learning goals on one topic at one difficulty grants the achievement and unlocks the possibility to advance on this topic (Figure 1, unfolded achievement ‘Diabetes principles’ at basic level). For example, a child who attained all difficulty 1 goals on glucose (i.e., knows why measurement is needed, how to correct a hypo, and understands the measurement value) is granted the achievement ‘Novice Glucose’, and may advance to the next level.

When the child arrives at the hospital for the first time, the caregiver and the child set the learning goals and desired achievements using PAL Control. This customizes the tasks a child needs to perform at home. All children fill in their diary and timeline, but specific quizzes and games are customized by the PAL agent depending on the goals set for the child.

At home, the PAL agent will propose different tasks to perform by the child (WP2,3,4). The tasks depend on the goals. The child can ask for an explanation of why the task is done by clicking the (?) button. This generates an explanation by looking at the goal that the proposed activity (a task) contributes to. This is translated using a standard format to a sentence for the robot/avatar. Further, depending on the type of task (activity) the PAL robot/avatar decides to portray a different style to the child (see Table 2) for its behaviors.

Children sit together with their caregivers and parents to monitor their progress. This is done with palInform. Here, children can see their progress on goals and achievements and an overview of events and activities as well as timeline data. An important element is that PAL Control and palInform, together with the apps developed in WP 5 allow for a gamified setting of goals to leverage user motivation. Goals are set according to difficulty level and children unlock achievements by performing tasks that support learning goals.

All of these systems work together by pushing changes to goals, tasks and activities to the central database used by the PAL system (see D5.3). As such, changes are reflected in real time to all of the modules. For the technical descriptions of the different function and the workings of the other work packages we refer to the other deliverables and relevant papers (see Annexes and refs).

Conclusion, relation to milestones and lessons learned.

Transparency, explanation, and personalization of behavior are of major importance for long-term human robot interaction and contributed directly to the objective of the project and work package 2. We have worked towards these aspects in work package 2 by focusing on personalized goal setting, personalized explanations of the actions of the system, and robot behavioral style that is adapted to the educational task the child has to perform.

We have reached all major outcomes planned. First, we have refined palInform and PAL Control and integrated these tools in the PAL system. palInform has been evaluated in Italy and The Netherlands. Also, the GUI’s of both tools have seen several iterations. As proper goal setting and monitoring is crucial for the children's adherence and motivation as well as acceptance by health care professionals (as they need to work with the goal setting tool), these are important outcomes. PAL Control and palInform are now properly tested tools for educational goal setting and monitoring of child progress (M2.1, M2.3). An important element is that PAL Control and palInform, together with the apps developed in WP5 allow for a gamified setting of goals to leverage user motivation.

Second, we have found consolidating evidence that robots can use communication style in their communication towards children and other adult users. Further, we showed - for the first time - that behavioral style of a virtual agent influence how people think about a human-agent interaction afterwards (negotiation in this case). This is a major breakthrough as this opens up the road towards personalization of robot behaviors towards children and other users related to their learning goals.

Third, we have extended our knowledge on individual differences in preferences towards how a robot explains its actions to users. We have extended this towards the use of emotions in explanations of robot actions by robots. This is an important line of research for both robot transparency as well as personalization of robot interaction.

Finally, all functions have been integrated in the PAL system, and the system has run successfully for several months during the 2018/19 final experiment.

Important lessons learned include the following.

- We have learned how to modulate style for robots and virtual agents and will keep using this knowledge for advanced HRI in future projects.
- We have advanced our understanding of Explainable AI and this has opened the road towards using emotions in generated explanations.
- We have gathered a lot of practical knowledge on doing Child Robot studies, and this work already pays of in e.g. TU Delft related spin off companies including Interactive Robotics. An example is that children and adults seem to have different preferences for explanation.
- We have learned the difficulty and pitfalls of in-the-wild experiments for HRI with children for hypothesis testing and this gives important information to steer future experimental design. Either a large sample size is needed to control for confounding variables, or a more qualitative analysis should be performed to understand the effect of single variables. The approach we took

with lab experiments first, followed by in-the-wild integration did work well and resulted in many publications.

References

1. Neerincx, M., Kaptein, F., van Bekkum, M., Krieger, H. U., Kiefer, B., Peters, R., Broekens, J., Demiris, Y., & Sapelli, M. (2016). Ontologies for social, cognitive and affective agent-based support of child’s diabetes self-management. Artificial Intelligence for Diabetes, 35.
2. Peters, R., Broekens, J., & Neerincx, M. A. (2017). Robots educate in style: The effect of context and non-verbal behaviour on children's perceptions of warmth and competence Robot and Human Interactive Communication (RO-MAN), 2017 26th IEEE International Symposium on (pp. 449-455): IEEE.
3. Kaptein, F., Broekens, J., Hindriks, K., & Neerincx, M. (2017). Personalised self-explanation by robots: The role of goals versus beliefs in robot-action explanation for children and adults Robot and Human Interactive Communication (RO-MAN), 2017 26th IEEE International Symposium on (pp. 676-682): IEEE.
4. Kaptein, F., Broekens, J., Hindriks, K., & Neerincx, M. (2017). The role of emotion in self-explanations by cognitive agents Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW), 2017 Seventh International Conference on (pp. 88-93): IEEE.
5. Lam, D. N., & Barber, K. S. (2005, July). Comprehending agent software. In Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems (pp. 586-593). ACM.
6. Shortliffe, E. H., Davis, R., Axline, S. G., Buchanan, B. G., Green, C. C., & Cohen, S. N. (1975). Computer-based consultations in clinical therapeutics: explanation and rule acquisition capabilities of the MYCIN system. Computers and biomedical research, 8(4), 303-320.

Annex 1

Robots Expressing Dominance: effects of Behaviours and Modulation

Bibliography Rifca Peters, Joost Broekens, Kangqi Li and Mark A. Neerincx (submitted). Robots Expressing Dominance: effects of Behaviours and Modulation. *IVA*

Abstract A mayor challenge in human-robot interaction and collaboration is the synthesis of non-verbal behaviour for the expression of social signals. Appropriate perception and expression of dominance (verticality) in non-verbal behaviour is essential for social interaction. In this paper, we present our work on algorithmic modulation of robot bodily movement to express varying degrees of dominance. We developed a parameter-based model for head tilt and body expansiveness. This model was applied to a variety of behaviours. These behaviours were evaluated by human observers in two different studies with respectively static pictures of key postures (N=772) and real-time gestures (N=31). Overall, specific behaviours proved to communicate different levels of dominance. Further, modulation of body expansiveness and head tilt robustly influenced perceived dominance independent of specific behaviours and observer viewing height and angle. The modulation did not influence perceived valence, but it did correlate with perceived arousal. Our study shows that dominance can be reliably expressed by both selection of specific behaviours and modulation of behaviours.

Relation to WP. Enhances and extends style expression for robots and virtual agents. This is relevant for the perception of the robot by the children and to enhance long-term interaction.

Annex 2

Exploring the Role of Emotions in Human Explanations of Robot Behaviour

Bibliography Frank Kaptein, Joost Broekens, Koen Hindriks, Mark Neerincx (submitted). Exploring the Role of Emotions in Human Explanations of Robot Behaviour. *EXTRAAMA Workshop*

Abstract Robots are slowly but steadily entering our everyday lives. The behaviour and decision making of these robots will become complex. Such robots should be able to self-explain their behaviour to their users. In this paper, we studied whether people use emotions in addition to beliefs and goals when explaining robot behaviour. We found that at least in some cases (4/16) people referred to emotions to explain the robot's behaviour. However, only one of the explanations (1/16) explicitly referred to the emotion of the robot to explain its behaviour. This pilot therefore replicates earlier findings that robot behaviour is indeed explained using folk psychology. It further carefully hints at the possibility that humans use emotions in their explanations of robot behaviour, in line with ideas from the human-human explanation literature

Relation to WP. Enhances and extends our understanding of why and when to use emotions in the explanation of robot behavior. This contributes to better understanding robot transparency.

Annex 3

Validated Expressive Non-verbal Behaviour in Human-Agent Interactions: A Review

Bibliography Rifca Peters, Joost Broekens, Mark Neerincx (in prep). Validated Expressive Non-verbal Behaviour Human-Agent Interactions: A Review

Abstract Attempting to increase agents' abilities to perceive, understand and express social signals, progress has been made in emotion modelling and recognition. Research into interpersonal behaviour is less thorough; to the best of our knowledge no comprehensive, validated model exists of non-verbal behaviour to express stance or style in human-agent interaction. Therefore, we systematically reviewed and analysed existing studies investigating non-verbal behavioural cues to communicate an agent's feelings, thought and stance towards the other or social situation, and perception thereof by a human. This review focus is on experiments manipulating behavioural cues of a artificial agent (either robotic or virtual) and measuring the effect of these manipulations on users perceptions of the agent. Although we are particularly interested in the educational domain, a broad review was done to create a comprehensive overview.

Relation to WP. Enhances and extends our understanding of the effect of behavioral style on users that interact with robots and virtual agents. In particular how to model such style through nonverbal behavior of robots and agents.

Annex 4

The Effect of Virtual Agent Warmth on Human-Agent Negotiation

Bibliography Pooja Prajod, Mohammed Al Owayyed, Tim Rietveld, Jaap-Jan van der Steeg, and Joost Broekens. 2019. The Effect of Virtual Agent Warmth on Human-Agent Negotiation. InProc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,IFAAMAS, 6 pages.

Abstract The perception of warmth and competence in others influences social interaction and decision making. Virtual agents have been used in many domains including serious gaming and training. In this work we study the effect of warmth expressed in the behavior of a virtual agent on a human-agent negotiation. We design and conduct an experiment where participants negotiate with two versions of the same agent displaying varying levels of warmth. The results show that humans are more satisfied with the warm agent, are more willing to renegotiate with it, would recommend the agent more to their friends and had a better interaction experience, even though there is no difference in negotiation outcome (utility, agreement or rounds needed). While studies have shown effects of emotional displays on negotiation and collaboration, this is - to our knowledge - the first time that a clear effect of behavioral style is shown on the post-hoc appraisal of a human-agent collaboration, in our case a negotiation.

Relation to WP. Enhances and extends our understanding of the effect of behavioral style on users that interact with robots and virtual agents.

Annex 5

Developing Cloud-based Social Robots for Long-term Interaction

Bibliography Kaptein et al. Developing Cloud-based Social Robots for Long-term Interaction (in prep).

Abstract In this paper, we show how cloud computing can be used to facilitate social robots in interacting with human users for prolonged periods of time. We present the technical architecture of the PAL system (a Personal Assistant for a healthy Lifestyle). We focus on the architecture and functionality of this system. This system controls a social robot and its avatar and integrates the robot with an expandable set of(mobile) health applications which all connect to a common knowledgebase and reasoning mechanism. The system is capable of autonomously interacting with children "in the wild" for prolonged periods of time without the need for a human-in-the-loop; and, it supports them with and educates them on health-related issues in the domain of diabetes. PAL is an exemplary system that provides stable and diverse long-term human-robot interaction

Relation to WP. Gives and integrated view on the PAL system, directly related to Task 2.5.