ENCOURAGING COLLABORATION BETWEEN PRIMARY SCHOOL CHILDREN THROUGH A LEARNING ROBOT

K.W. Kaag
s1322273

Faculty of Electrical Engineering, Mathematics and Computer Science
Human Media Interaction (HMI)

Master Thesis Interaction Technology

supervisor: dr. M. Theune
supervisor: prof.dr. T.W.C. Huibers

UNIVERSITY OF TWENTE.

2019-08-23

Contents

Abstract

This research explored how the surface bot, a mobile tablet-based robot, can be used to elicit collaboration between children. Collaboration is seen as a 21st century skill, that children need to learn. A first prototype with the surface bot was developed based on the learning-by-teaching paradigm. The focus was on the "teaching" part, with children acting as tutors of the robot in a story-based activity. The surface bot's tablet is used to display the character and to visualize thoughts about the coming action. Children used a tablet with a slider to give feedback on the robot's actions. The first prototype was controlled by a tele-operator in a Wizard-of-Oz setup. The robot's actions were scripted and it did not learn from the children's feedback. A first study was conducted with 6 pairs of primary school children (age 4-8), aiming to evaluate the activity with the prototype on its effectiveness of encouraging collaboration. In this study, children were engaged and provided consistent feedback over the course of the activity. However, little collaboration was shown during the activity. Children were mainly observed to make individual decisions and to take turns in operating the tablet.

Based upon the outcome of first study, and supported by information found in literature, the prototype was adjusted to encourage more spontaneous collaboration. This was done by introducing more ambiguity to the children's task and making it more challenging for them to track and interact with the robot. The hypothesis was that it would provide more incentive to collaborate, stimulating a division of roles. This improved second version of the prototype made use of Q-learning to learn from the input of children, thereby minimizing the role of the tele-operator during the activity to controlling the robot's movement. In the second study with 9 pairs of primary school children (age 6-10), children were indeed observed to adopt a role division in multiple cases. The level of collaboration was evaluated for each pair of children using a framework of indicators that is adapted from the collaborative problem solving framework by Hesse et al. [15]. The annotation showed higher collaborative scores on average in the second study, compared to a baseline of two pairs of children (age 6-8) from the first study. The pairs of children that participated in an activity with the second prototype, scored slightly better for most indicators of collaboration.

It can be concluded that a concept based on learning-by-teaching can encourage collaboration between primary school children. The reliability of the framework was sufficient for this research, but the validity is inconclusive due to the small sample size. Future work can focus on developing a reliable and valid framework with which different prototypes can be tested and compared on the degree of collaboration they encourage among children. Future research can then focus on longitudinal studies exploring the effect of participating, in activities with the surface bot over a longer period, on the development of collaboration skills of primary school children.

Acknowledgments

I want to thank my supervisors Mariët Theune and Theo Huibers for their guidance through each stage of my thesis. I am very grateful for the time you have taken for regular discussions, giving feedback on my work and sharing expertise. It helped me to sharpen my ideas and shape my research. I also want to thank OBS de Zwaluw in Markelo and BSO de Vlinder in Enschede for their hospitality for conducting user tests at their locations. Furthermore, I want to thank Floris Veldhuizen for his assistance in conducting the user tests of the second study. Finally, I am grateful for my family and friends who supported me during my thesis.

1. Introduction

This research is inspired by the coBOTnity project[1] which aims to explore how hybrid artificial agents can be used in collaborative storytelling to effectively encourage creative thinking and social awareness in children. The coBOTnity project is a project funded by the European Union's Horizon 2020 research and innovation program. Catala et al. [9] mention collaborative or group activities as the preferable structure for storytelling activities with children, and that "embedding storytelling activities in the classroom is time-consuming and not easy." Based on the perspectives of teachers on storytelling, Catala et al. [9] recommended that technology for storytelling should be flexible. It should allow teachers to designate different roles to children and enable them to arrange activities in small working groups. Based on the teachers input, it was recommended to give the children an active role in the creation of stories since it facilitates discussion and children learning from each other. The surface bot was developed by Catala et al. [9] as an affordable, mobile and flexible robot to be used in collaborative storytelling activities.

Collaboration is an important skill for children to learn [2, 19]. It relates to critical thinking, meta-cognition and motivation [19]. Collaboration is referred to as a 21st century skill [1, 6]. It has been shown that beneficial effects regarding learning and development, particularly in the early years or primary education, can occur when children work in small groups or pairs [29]. Furthermore, self-esteem and attitudes towards others are mentioned as beneficial outcomes of collaborative learning in the classroom [5, 25]. But collaboration is not an obvious skill for primary school children, as many young children have difficulties to effectively collaborate [2]. "Children in the age group 5-7 have shown significant changes in the ability to collaborate [21]." But the age group 3 to 7 years is also characterized as being fairly self-centered and doing a lot of parallel play [22]. Literature also points out that children are impulsive and do not yet reason logically. Collaboration is based on communication, cooperation and responsiveness.

Developing collaboration skills takes practice and there might be a long-term education gain when children discover collaboration for themselves [2]. The surface bot can be a tool for activities where children are stimulated to work together in order to contribute to the development of collaborative and social skills of children in the long term.

1.1. Aim and objectives

In this research, I explored how to design an activity with the surface bot to encourage collaboration among small groups or pairs of primary school children. Primary school children in the age

[1] The coBOTnity project: https://www.utwente.nl/en/eemcs/hmi/cobotnity/

of 5-7 were the target group of this research, because their collaboration skills start developing [2]. Children in this age could benefit from settings that encourage social interactions and collaboration. A concept that encourages successful collaboration, makes use of the capabilities of the surface bot and is suitable as a classroom activity. It can be a basis for further collaborative activities with the surface bot that can be integrated into the children's curriculum. The main two questions that were addressed in this research were:

1. How can the capabilities of the surface bot be utilized to create an engaging activity that effectively encourages collaboration between primary school children?

2. How can the extent and manner of collaboration between primary school children be measured in order to evaluate the effectiveness of an activity with the surface bot?

A number of objectives were drafted with which the research questions could be answered. The first objective was to get a background on the surface bot: an overview of the capabilities of the surface bot and the studies in which it has been used in activities with children. The second objective was aimed at gaining an insight into collaboration. I have looked at how children learn collaboration, and how collaboration could be evaluated for pairs of children. The third and last objective is to examine related work regarding implementations of the learning-by-teaching paradigm and studies that describe ways of integrating human input in the learning process of a robot or virtual agent. A concept has been developed based on these three objectives. This concept has been developed into a prototype, which was validated in a first study. A second study was done in which the collaboration between children was assessed on the basis of a framework for evaluating collaboration.

1.2. Overview

Chapter 2 provides a concise description of the surface bot and its capabilities. A selection of related work is described as inspiration for the development of a concept with the surface bot. Chapter 3 addresses collaboration. It deals with the aspects of collaboration, and the conditions that foster collaboration. A brief overview of ways to evaluate collaboration is also provided. Chapter 4 discusses learning robots with the learning-by-teaching paradigm as the basis. Related work on the possibilities of integrating human input into the learning process of an (robotic) agent is described. Chapter 5 motivates and describes a concept based on the learning-by-teaching paradigm. Subsequently, the realization of a first prototype is explained in detail. Chapter 6 describes the first study that aimed to validate the first prototype. The most important results are set out and discussed. An improved prototype is then presented in chapter 7. First, the suggested improvements based on the results of the first study are described. Second, the realization of the second prototype with a reinforcement learning framework is described in detail. Chapter 8 describes the second study that is aimed at exploring the degree of collaboration between pairs of children, and the influence of the robot's action speed on this. Based on the results of both studies, the main research questions are answered and discussed in chapter 9. This chapter also describes the conclusions that were drawn. Finally, a set of recommendations is described for future work in Chapter 10.

2. Introduction to the Surface Bot

The first aim of this chapter is to provide a detailed description of the surface bot. Secondly, the aim is to describe a selection of related work and discuss its relevance to this research.

2.1. What is the surface bot?

The surface bot was developed as an affordable, mobile and flexible robot to be used in collaborative storytelling activities [9]. The surface bot consists of two parts: a tablet and a base with wheels (see Figure 2.1). The tablet and wheelbase make the surface bot capable of movement, sound and visual representations. The tablet is a multi-functional component which is used as a character display [28] and as an interactive interface [7]. Figure 2.2 gives an impression of a surface bot used as character display.

In several studies, the surface bot has been applied in storytelling activities. Catala et al. [8] explored the interaction of children with a surface bot in a storytelling activity. In the test, children (n=22) used an early prototype of the surface bot to tell stories. The screen of the surface bot displayed a character. A special tablet was used to control the movement of the surface bot. The children had a number of small assets, each illustrating a character, location, or object. The children were free to use any asset in their storytelling. During the test, the focus was on four aspects: storytelling, use of assets, character embodiment and movement control. The observations indicated that not all children were able to create coherent stories, and therefore a recommendation was given to have responses or feedback from the robot on the actions of children. With regard to the use of assets, children seemed to expect a response when they tried to give, or show, an asset to the robot. Controlling the movement of the surface bot

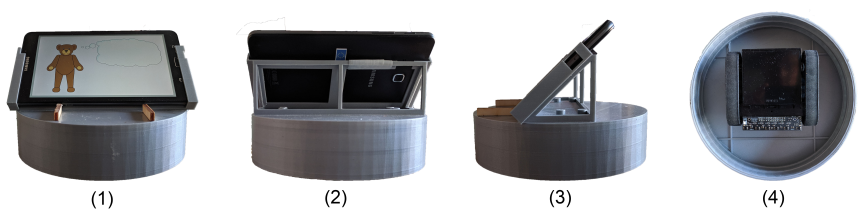

Figure 2.1.: The surface bot. The front view (1) of the surface bot with the tablet. The back view (2) and a side view (3) shows the plastic framework that holds the tablet in position. The image from below (4) shows the wheelbase, with the two small tracks.

Figure 2.2.: An application of the surface bot. [28]

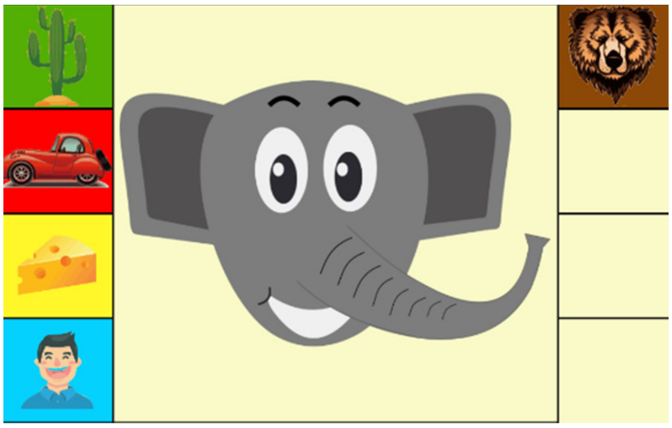

Figure 2.3.: Interface of the surface bot. [28]

was an entertaining experience for the children, but it is suggested that it might take too much from their attention which negatively impacts the storytelling. Although the surface bot had no social behavior, and could not move autonomously, it was seen and treated as an embodied character by the children.

2.2. Teaching the surface bot in a collaborative activity

Verhoeven, Catala and Theune [28] developed an interactive activity with the surface bot as a second-language learner in a story-based activity for children. The aim was to explore how children interacted with the robot and if their French improved during the activity. It was inspired by the learning-by-teaching method, where children acted as teacher of the surface bot and in the process learn themselves. A detailed background on the learning-by-teaching paradigm is provided in Chapter 4, section 4.1. The surface bot was used as a protagonist in a story. The protagonist was described and displayed as an elephant character. The story element was introduced, since it can captivate and motivate children. At the start of the activity, the

Figure 2.4.: Interface of the children's tablet. [28]

surface bot introduced itself as a character located in France trying to learn the language there. Children were asked to assist the robot. They took on the role of teacher and taught French words when the surface bot asked for the translation of a certain object.

The concept was designed as a tabletop activity, making use of the surface bot's movement capabilities. The activity used five different locations that were displayed using tangibles. At each location there were cards with each a unique object on it. Children shared one tablet that they could use to point the robot to a new location. The robot then independently drove towards it. The robot's movement was controlled by a tele-operator according to a Wizard-of-Oz approach. The tablet of the surface bot was used to portray the character and his emotions, see Figure 2.3. In addition, it reflected the words it currently knew. Three emotional expressions were used: happy, sad and neutral. Verhoeven et al. [28] mention the importance of repetition for effective learning, therefore the surface bot would forget the words a couple of times during the activity. The robot would then get sad and ask the children if they could teach the word again. Besides directing the attention of the surface bot towards new locations, the tablet of the children was also used to teach the French words, see Figure 2.4. The tele-operator made use of a corpus of audio fragments to control the robot's speech in order to respond to situations or to initiate interactions from children. This included audio fragments for asking a translation, asking for directions or thanking the children when they taught it something. Figure 2.5 shows the interface of the tele-operator.

Verhoeven et al. [28] evaluated the application in a user test with 22 children at a Dutch primary school. The children were on average 8 years old (min=7, max=9). The French vocabulary of children was tested before and after the session with the prototype. The results suggest a growth in the vocabulary. However, the learning could not strictly and fully be explained by the design

Figure 2.5.: The tele-operator interface. [28]

of the activity. It was argued that children could have learned in the time between the session and the post-test by discussing it in the classroom with other children. Children were observed to have fun during the activity [28]. They communicated about the robot's next location and the usage of the cards displaying objects.

2.3. Conclusion

Verhoeven et al. [28] integrated the learning-by-teaching paradigm into an engaging and fun activity with the surface bot while making use of the robot's main capabilities: movement, speech and and extensive usage of the visual display. The effect of learning-by-teaching was not proven, but has shown to have beneficial educative outcomes in other studies [20]. A concept with the surface bot which aimed to encourage collaboration between children was developed for this research based on the learning-by-teaching paradigm. It is an activity with the surface bot where children act as tutor. The concept was also based on a story, since it can appeal to the imagination of children and can therefore be motivating to participate in an activity with the surface bot. Furthermore, a story-based activity suits the envisioned storytelling, flexible and possibly educative purpose of the surface bot.

3. Defining and evaluating collaboration

This chapter examines what collaboration entails, what the characteristics are and what fosters collaboration among children. First, section 3.1 defines collaboration and discusses the aspects of it. Second, section 3.2 provides an insight in how collaboration can be evaluated. Section 3.3 explored related work for methods and guidelines for encouraging collaboration between children.

3.1. Defining collaboration

Roschelle and Teasley [23] state that collaboration involves a "mutual engagement of participants in a coordinated effort to solve a problem together." First and foremost, a shared goal is needed for collaboration. Secondly, collaboration includes communication, responsiveness and cooperation [15]. Communication is an indispensable requirement for successful collaboration. There should be readiness to exchange knowledge and opinions. Responsiveness involves "active participation and insightful contribution" as described by Hesse et al. [15]. In this research it was seen as an awareness of the perspective of others and providing thoughtful contributions. Cooperation is described as a division of labor. Dillenbourg et al. [13] maintain the same definition of cooperation, however they do not see it as an element of, but rather a state that can arise through collaboration. A division of labor in an activity with children might be a result from a division of roles, in which children each will do something else in order to achieve the shared goal together.

Dillenbourg [12] notes that collaboration is characterized by a symmetrical structure with four factors. First, there should be a symmetry of goals, which implies that people should have a shared goal. Individual goals can give rise to different interests, which may cause conflicts and hinder collaboration. The second factor is a symmetry of actions. This was interpreted as requiring children to have the opportunity to take the same actions. When actions are reserved in advance for certain children in an activity, effective collaboration could be hampered by, for example, jealousy. The third factor is a symmetry of knowledge, which is understood as ensuring that participants have relatively equal knowledge of the activity. It is emphasized, however, that they may differ in perspective. The fourth and last factor is a symmetry of status. This involves "collaboration among peers rather than interactions involving supervisor/subordinate relationships [12]." Another influencing factor on collaboration is interdependence [19]. When children children depend on one another for achieving a shared goal, there is more incentive to collaborate.

3.2. Evaluating collaboration

A method of assessing the level of collaboration is required in order to evaluate the effectiveness of the concept with the surface bot in encouraging collaboration between children. Dillenbourg [12] mentions interactivity and negotiability as aspects that determine the degree of collaboration. Negotiability describes the degree to which individual opinions are imposed on others, when it should be everyone's aim to work towards a common understanding. Interactivity refers to perspective taking and the degree to which people are influenced by the contributions of others. These two aspects did not provide a clear enough distinction to be used as metrics for measuring and quantifying collaboration between children in an activity with the surface bot in my opinion. If one person attempts to perceive and understand another person's point of view, then it can be argued that a high degree of negotiability is already the case, since opinions are not unquestionably adopted at that point.

The framework for assessment of collaborative problem solving described by Care and Griffin [6] consists of more clearly distinguishable factors that determine the collaborative and problem solving skills of individuals. Hesse et al. [15] describe the framework in further detail. They state the framework comprises of cognitive skills and social skills. The social skills relate to the "collaborative" part and the cognitive skills address the "problem solving" part. Each part consists of multiple classes with several indicators. Exploring the cognitive skills of children during an activity with surface bot was outside the scope of this research, therefore this chapter only elaborates on the classes and indicators of the social part of the framework. The social skills category has three classes: participation, perspective taking and social regulation. Participation is about the willingness and readiness of participants to share information or opinions and is described as a "minimum requirement for collaborative interaction [15]." Participation consists of three indicators: action, interaction and task completion. Action is described as the general participation of an individual in a problem solving activity. Interaction refers to interacting and responding to others. An example of participation with high action and low interaction is someone that is highly active, but does not respond or coordinate with others. The third indicator, task completion, refers to perseverance and commitment to the problem or activity.

The second class, perspective taking, refers to "the ability to see a problem through the eyes of a collaborator [15]." The perspective of others must be understood and considered in order to reach a solution or compromise during a discussion, or negotiation. Perspective taking consists of the indicators: adaptive responsiveness and audience awareness. Adaptive responsiveness refers to considering and responding to contributions of others. Audience awareness refers to ensuring that contributions are tailored to the other's perspective, ability or knowledge.

The third class, social regulation, is about coordinating and resolving differences in perspectives. It refers to the strategies used to resolve conflicts and to work together towards solving a problem. It consists of four indicators: negotiation, self-evaluation, transactive memory and responsibility initiative. Conflicts lead to negotiation. Negotiation refers to addressing differences, and working towards a compromise or mutual agreement. Self-evaluation refers to recognizing the strengths and weaknesses of oneself. Transactive memory refers to recognizing the strengths and weaknesses of others. From my perspective, these two indicators relate to the ability to reflect, building a mental model of the knowledge and abilities of oneself and of others. This ability could improve coordination between children in the problem solving activity, as tasks can be tailored to a person's strengths, and weaknesses can be compensated by others. The fourth indicator, responsibility initiative, refers to the collective responsibility in addressing and solving the problem. It relates to whether someone is actively involved or retained in the problem solving process by others.

3.3. Learning collaboration

Besides the factors described in section 3.1, collaboration is also affected by the structure and design of a task [19]. In an activity with the surface bot, the "task" is the responsibility children get and what they are expected to do. It is recommended for tasks to be ambiguous [12] as it tends to foster collaboration. A trivial and obvious task elicits little disagreements between children, and therefore no opportunity arises for negotiation and there is little incentive to engage in a coordinated effort. Disagreements and misunderstandings can cause communication, in the form of explanations and reasons [12]. Communication is an interpersonal skill [21] which will develop when children are provided with the opportunities for social interaction [9]. Benford et al. [2] argue that encouraging collaboration is the right approach and expect positive educational outcomes when children discover the value or pleasure of collaboration themselves.

3.4. Conclusion

I made the decision to integrate a symmetrical structure [12] as well as possible in a concept with the surface bot, with the idea that it would provide the opportunities for children to collaborate. Integrating these factors in the concept, ensures it has a better chance of successfully encouraging collaboration. Therefore, requirements for the concept are that children receive the same introduction, that no division of labor is imposed and that children have the same goals. To guarantee an equality of expertise and skill, a simple and accessible concept was sought. The intent is to let peers participate to ensure the symmetry of status. My expectation was that the mutual relationship of children influences the extent to which they communicate and collaborate. Therefore, in all studies with the prototype, children took part in groups of two classmates. This ensured a symmetry of status, since the children knew each other and were of similar age. As was described in Chapter 2, the aim was to let children act as tutor of the surface bot. The concept's task and activity should therefore be perceived as ambiguous and be designed in a way that social interaction becomes likely. However, the approach was not to force children to collaborate or communicate, but rather create a setting that effectively elicits spontaneous collaboration.

4. Overview of Learning Robots

The goal of the literature study described in this chapter was to explore how existing learning robots are designed, to get inspiration for how the surface bot can be taught by children. The first part, section 4.1, of this chapter provides a theoretical background to the learning-by-teaching paradigm. Then Betty's Brain is explained in section 4.2, a virtual teachable agent developed by Biswas et al. [3]. In section 4.3 the work of Chandra, Dillenbourg and Paiva [10] is described where children assess a robot's handwriting skills. Section 4.4 provides an overview of the Q-learning algorithm. It provides the background on the algorithm used in Sophie's Kitchen; a learning agent developed by Thomaz and Breazeal [26], discussed in section 4.5.

4.1. A background of Learning-by-Teaching

Learning-by-teaching is described as "learning through the act of teaching" [17]. As a pedagogical approach it has shown its effectiveness in terms of learning outcomes and motivational effects [20]. A well-known outcome of the learning-by-teaching approach is the protégé effect [11] where students invest more time and effort to teach others than they do for themselves. Biswas et al. [4] state that students that teach, developed a deeper understanding and were able to express their ideas better, compared to those who were asked to write a summary regarding the same domain. The learning-by-teaching approach can be used between children with one acting as a tutor and the other as a student. This is also referred to as peer-tutoring [10]. Another way is to let children teach a computer agent, otherwise known as teachable agents [3]. Biswas et al. [3] mention that learning-by-teaching includes critical aspects of learning: structuring, taking responsibility and reflecting. Structuring is understood as being aware what can be taught, and what should be taught. It relates to planning, building knowledge and coordinating with each other. Taking responsibility is about the preparation and attention that students put into their role as tutors. My interpretation of reflection was that it concerns monitoring how well ideas and explanations are understood, and that actions are adjusted accordingly in the pursuit of effective teaching.

4.2. Betty's Brain: teaching concepts

Biswas et al. [3] developed the application Betty's Brain, based on the learning-by-teaching paradigm. It is a digital interface with the teachable agent Betty, designed for high school students to teach about river ecosystems. Students could teach Betty by adding and connecting information in a graph structure, referred to as the "concept map". Students could then query Betty about what they had taught her. The answers formulated by Betty were based on the concept map created by the students. Betty did not use machine learning techniques to learn, but reasoned based on the concept map. The interface also displayed a mentor agent that could provide feedback to Betty, or provide hints to the students on how to improve Betty's performance on answering the queries. Three experiments were conducted with each a different role of the mentor agent. In the first experiment, a group of students used Betty's Brain and the mentor agent acted as tutor. It provided feedback directed towards the student in order to improve the concept maps created by them. The second and third experiment used the learning-by-teaching approach. Instead of addressing the students, the mentor agent in the second experiment gave feedback directed towards Betty, based on the answers to queries. This was meant as the baseline group. The third group used a new version of a more responsive Betty's Brain with self-regulated behavior. In this version the mentor agent could provide elaborate explanations and feedback, but only on request of the students by formulating a query. The results showed that students of the three groups had equal performances with regard to memorizing the concept maps they constructed. The group using Betty's Brain with self-regulating behavior "demonstrated better abilities to learn and understand new material."

4.3. Nao: demonstrating handwriting

Chandra et al. [10] conducted a set of experiments with pairs of children of the age 4 to 6, to explore the effectiveness of the peer-learning (PL) and peer-tutoring (PT) method for acquiring handwriting skills. Peer tutoring is another name for learning-by-teaching in which one child is the tutor and the other is a learner. Children get no role assigned in peer learning. A first exploratory study of 20 pairs of children compared the PL against the PT method. Ten pairs were asked to copy letters on a sheet, and give feedback on each others writing. The other 10 pairs were the PT group with one child acting as teacher and the other as learner. Halfway through, their assigned roles were reversed. The "teacher" presented letters one by one to the "learner", who wrote them down. Then, the teacher gave feedback on the learner's handwriting. Children were more excited for the peer-tutoring variant, as they got to act as "teacher". Although the results were too limited to conclude a preference for one of the methods. It was also stated that children of the age 4 to 6 conveyed feedback immaturely, due to their young age. The second study was aimed at exploring the impact of introducing a robot facilitator to see the effect of the PL and PT method on the feedback of children. The focus was on slightly older children, age 6 to 8. Instead of the experimenter, a Nao robot was used as facilitator to provide instructions and accompany the children during the activity. 18 pairs of children participated in the experiment as part of the PL or the PT group. In the PT method, children gave significantly more extended self-disclosure to the robot and significantly more corrective feedback to the learner, compared to the PL method. The improvement of the children's learning gains of the PT method were significant, whereas the PL method showed no significant differences. They concluded that overall the PT method seemed to be more effective. The third study of Chandra et al. [10] used the Nao robot not as a facilitator, but as a peer in a PT activity. The goal was to explore how children perceive, and correct the handwriting of a robot. An experiment was conducted with 24 children of the age 7-8. They participated with the robot under a learning condition or non-learning condition. In the learning condition, the robot's handwriting improved based on the feedback of children. In the non-learning condition, the robot's handwriting did not improve. First, the robot drew a letter on a touch screen. Children were then able to give feedback by changing the shape of the letter using a slider, or they could demonstrate the letter in a specific box on the screen. The results indicated that children were able to notice the robot's learning, as significant higher scores were given by the children on the robot's handwriting performance over time under the learning condition compared to the non-learning condition.

4.4. A background on Q-learning

This section aims to provide a background on Q-learning, a reinforcement learning algorithm, since it is part of the application [26] discussed in the following section. Kaelbling, Littman and Moore [18] describe reinforcement learning as an agent's problem of learning behavior by trial-and-error in an environment. When the problem can be formulated as a Markov decision problem (MDP), then Q-learning can be used to derive the optimal policy on how to act given the environment's circumstances. It is a Markov decision problem when an agent has an accessible, stochastic environment with a known transition model [24]. This means that there is a discrete set of states and a discrete set of actions per state. The transition model describes the state transitions: the state resulting from an action in a given state. In order to acquire a policy, Q-learning requires a reward function which contains the reward received based on a state transition. Rewards can be received in states from where the agent can take no further action - the terminal states - or in any other state. The optimal policy has the sequence of actions that leads to the maximum cumulative reward. There may also be states where the agent receives a negative reward, or penalty. Negative rewards teaches the agent which states to avoid, as they do not contribute to the highest cumulative reward. In the following chapters, the term "reward" is used to describe both the positive and negative rewards.

In order to derive the optimal policy, a Q-function is calculated. The calculation used in Q-learning is based on the Bellman equation, see equation 4.1. In this equation the Q-value $Q(s, a)$ of the last action a and current state s is calculated based on the reward r received for transitioning to the current state and the expected maximum discounted reward, which is the highest Q-value based on the next state s', and a possible action s' of that state. The discount γ is used to determine to what extent future rewards influence the Q-value.

$$Q(s, a) = r + \gamma \max_{a'} Q(s', a') \tag{4.1}$$

In Q-learning, the Q-values are updated based on equation 4.2. In this equation, the Q-values can be iteratively calculated. A new Q-value $Q^{new}(s_t, a_t)$ of last action a and the new state s is based on the previous Q-value $Q(s, a)$. The learning rate α determines the degree to which the Q-value is updated based on the difference between the expected maximum reward and the

Q-value of this state. Similar to the Bellman equation, a discount factor γ is used to determine the importance of future reward. A low discount factor will cause the agent to put more emphasis on the current reward, in contrast to a high discount factor which causes the agent to focus on the long-term reward. The discount factor can be used to navigate through the trade-off between exploration and exploitation. This trade-off is described in further detail in chapter 7.

$$Q^{new}(s_t, a_t) \leftarrow Q(s,a) \ + \ \alpha \ \times \ (r_t \ + \ \gamma \ \times \ \max_a Q(s_{t+1}, a) - Q(s,a)) \tag{4.2}$$

4.5. Sophie's Kitchen: providing feedback and guidance

Thomaz and Breazeal [26] conducted a set of studies that aimed to explore how people want to teach and what they are trying to communicate to a learning agent. The second aim was to use these insights to improve human contribution in guiding a robot's learning behavior. The research meant to contribute to the design and development of robots that can learn more effectively, and are easier to teach by humans. In order to explore how people want to teach and communicate, the application "Sophie's Kitchen" was developed. Sophie is a virtual reinforcement learning agent located in a kitchen environment. She took action independently using a Q-learning algorithm to learn the task of baking a cake. The kitchen environment consisted of three locations; the oven, the shelf and the table. On the shelf were five objects which were necessary to bake a cake. Sophie had a fixed set of actions, which included movement between the locations, picking-up an object and using it. People got the explicit task to teach Sophie using specific feedback channels. These channels communicated a signal that directly influenced the reward for Sophie. So there was no predetermined reward function, but an adaptive reward signal that people controlled. As explained in section 4.4, rewards are used in the calculation of new Q-values. This means people are decisive in Sophie's learning of the optimal policy. The feedback channels were designed for people to provide general feedback on the "whole world state", or provide specific feedback about the state of an object. For general feedback, a participant could click anywhere on the screen. For object specific feedback, the object must be selected. In both cases, a slider appeared to communicate the feedback to Sophie (green=reward, red=punishment). Sophie's exploration lead to a sequence of actions that ended in one of two terminal states: 1. achieving a goal state, which is successful completion of the task (a cake is made) or 2. reaching a disaster state (for example: placing the raw eggs in the oven). Both the goal state and the disaster states resulted in a reset, returning the agent to the initial state, from where the agent could try it again.

In a pilot experiment, 18 participants were asked to give feedback to Sophie. The pilot study resulted in a set of findings that prompted a set of follow-up experiments. Firstly, people used the object-specific feedback to direct the agent's attention as a form of guidance, even though they were told that they could only communicate feedback. Secondly people changed their strategy of teaching, as they began to understand how the agent learned. Thirdly, people showed a rewarding behavior oriented towards positive rewards. Subsections 4.5.1 to 4.5.4 describe the follow-up experiments.

4.5.1. Attention direction

The first follow-up experiment leveraged the tendency of people to direct the attention of Sophie as a form of guidance. A channel of communication was added to distinguish guidance from feedback. With a right mouse click people could select an object to direct the attention of Sophie to it. The Q-learning algorithm was modified to bias the action-selection based on the attention direction signals from a participant. It was expected that the bias of the action-selection process of the agent would result in a faster convergence towards a successful policy. This modification was evaluated with 28 non-expert trainers. The functionality to guide the agent's attention resulted in a significantly faster learning interaction compared to the initial experiment.

4.5.2. Transparency behavior

The second follow-up experiment used gazing of the agent to test if it improved the participant's mental model of Sophie. Gazing is a transparency behavior; a communicative act that reveals the internal state of the agent, as it reveals the next move of the agent. This enabled them to direct the attention to a different object or location. The Q-learning algorithm with guidance and feedback was modified by adding a short delay before the action-selection phase. Preceding the step of taking an action, the agent would gaze at location of the object involved in the next action. During the short delay, the agent waited for a guidance signal from the participant. The duration of the gaze had to show how certain the agent was about its actions. When the duration of the gaze was short, the agent appeared to take resolute action. A longer duration indicates uncertainty, or indecision, which gives people time to provide feedback or guidance. In this experiment, 52 non-expert trainers participated under one of two conditions: with and without the gazing. The results indicated that people without the gaze behavior "overused" the guidance channel, and provided guidance whenever possible. The participants in the gazing behavior condition, provided guidance more when it was required and less when not.

4.5.3. Motivational input

The third follow-up experiment was based on the observation that people tend to provide more positive rewards. It was hypothesized that "people are falling into a natural teaching interaction with the agent, treating it as a social entity that needs motivation and encouragement." In order to test this, a channel for motivational input was created. When Sophie was selected, this was considered to be motivational input. In the third experiment, the ratio of positive and negative feedback was compared for 98 non-expert trainers with or without the motivation channel. The results showed a significant difference between the ratios, with a more balanced ratio in positive and negative feedback for the people that had the motivational channel. How an agent can utilize this motivational signal to improve learning is a question that remained for future work.

4.5.4. Undo behavior

A second hypothesis was formulated based on the positive rewarding bias. It was hypothesized that Sophie would not react as expected when rewarded negatively. In the current application, the agent did not react to any received rewards. It was assumed that it is perceived as if the agent ignores the feedback of the participants. In order to test this hypothesis, an UNDO function is implemented in the Q-learning algorithm that, if possible, reverses the last action. In a fourth follow-up experiment with 97 non-expert trainers, the UNDO function was evaluated. It was concluded that the UNDO behavior significantly improves the learning behavior of the agent. The results showed that the agent's failures during learning decreased by 37%. Failures being the agents transitioning to disaster state. Furthermore, the results indicated a more efficient exploration by the agent.

4.6. Conclusion

As concluded in the previous chapter, the learning-by-teaching paradigm was chosen as the approach for the concept with the surface bot to encourage collaboration. The aspects of learning-by-teaching [3]: structuring, taking responsibility and reflecting, were taken into account during the development of a prototype based on this concept. For structuring, the concept required children to plan and maintain a shared understanding. Taking responsibility relates to their participation which requires the children's engagement. Children should feel responsible for the surface bot and take up the role of tutor with attention to the scenario outlined. Reflection requires children keep track of the progression of the surface bot, and notice their influence on its learning. It also relates to meta-cognition, if a child acknowledges a mistake or is unable to do something and decides to ask the other child for help or information. This research consists of two studies with the prototype to identify the collaborative behavior of children. The suitability of the learning-by-teaching paradigm as a concept with the surface bot was determined by reviewing the children's behavior using the aspects of learning-by-teaching. The study by Chandra et al. [10] showed the proficiency of children as tutor to a robot. They mentioned that children are paying attention to the learning of a robot and are capable of providing corrections using a slider or by demonstration. Furthermore, they stated that children seemed to notice the robot's learning over time and that it had a positive effect on their handwriting skills.

The concept should be based on a problem or task that fits the thinking level and knowledge of children. Betty's Brain relied on students getting an understanding of the concepts and the ability to teach and link concepts in the concept map. I expected that interacting with such a graph structure might be too abstract for children. Reasoning skills for children in the age 3-7 are often not fully developed and it is recommended that products are simple and not too abstract [22]. An interesting aspect of Betty's Brain is the fact that no machine learning techniques were required for "learning", since Betty learns directly from the input from the children.

Thomaz and Breazeal [26] showed successful training of an agent based on human feedback. Therefore, I decided to use their approach and Q-learning implementation as inspiration for the learning capabilities of the surface bot. The concept with the surface bot would require the surface bot to learn a policy based on the guidance and feedback of children. The ability to learn from human feedback, contributes to the intended flexibility of the surface bot, since no explicit transition or reward function needs to be developed for activities. Activities with a learning surface bot only require a model of the possible actions and states, supported by a story that involves children as tutor. The children determine the rewards and ensure that the robot adopts a policy with which it acts logically in an activity. A reinforcement learning framework could make it easier to integrate new learning material into activities with the surface bot. New activities could be based on stories that involve new knowledge, concepts or a new level of complexity. The self-learning aspect of the surface bot could therefore be an answer for the intended flexible use of the robot in the classroom [9].

In this research, the surface bot functions as a teachable agent. However, if children learn from their role as tutor is beyond the scope of this research. The focus is on developing a prototype that forms a basis with which collaboration is encouraged successfully. The first step is to devise a concept that exploits the capacities of the surface bot with a clear role for the children as a tutor. This research continues with exploring how children collaborate in the role of tutor and how this translates into learning by the robot. When the surface bot, as a reinforcement learning agent, effectively encourages collaboration between children and is able to successfully learn from the input of children, the focus can shift towards activities that are informative to the children. The most important measure of the effectiveness of the concept is thus the extent to which collaboration is achieved among the children. The extent to which the aspects of collaboration can be seen during an activity with the surface bot determines whether learning-by-teaching is the right approach to encourage collaboration between children.

5. Prototype 1.0: a proof of concept

This chapter introduces a concept that aims to encourage collaboration between children based on the learning-by-teaching paradigm. The concept has been substantiated on the basis of the related work described in the previous chapters. The concept is explained in section 5.1. Section 5.2 lists a set of requirements that need to be met for successful application of the concept as an activity for primary school children. A detailed description of the realization of a first prototype is given in section 5.3.

5.1. Concept: Ted's Clothing Choice

The concept is based on the characteristics of collaboration, the learning-by-teaching paradigm and reinforcement learning. As described in Chapter 4, learning-by-teaching engages children in an activity where they act as a teacher or tutor of someone. In this concept, children should take responsibility for the learning process of the surface bot. The robot acts as an independent character with a lack of knowledge or skills, and needs the guidance of the children to successfully complete its task. The study by Chandra et al. [10] indicated that children pay attention to the learning of a robot. Inspired by the work of Verhoeven et al. [28] the role of tutor and the robot's task are woven into a story. According to Markopoulos and Bekker [22], children in the age of 3-7 enjoy fantasy. It is meant to rouse the interest of the children and motivate them to participate. The devised story goes as follows:

"Ted is a friendly brown bear. He would like to play outside with his friends. But then he must first get dressed. He looks out the window and sees snow everywhere: winter weather. Ted then realizes that he really doesn't know anything about clothing and he doesn't know what to wear now. Can you help him find the right clothes together?"

The story outlines the situation of the bear, and invites children to guide him towards accomplishing his task: finding the right clothes. The surface bot is the protagonist of the story: the bear. The idea is that children can help the surface bot through guidance and feedback, inspired by the work of Thomaz and Breazeal [26].

The concept uses the physical space around the robot. It should symbolize the bear's house with different locations in it. At these locations the bear can find items of clothing that may be needed for his outfit. Introducing locations utilizes the mobility of the surface bot and makes the activity more dynamic for the children. The robot is always at one of the locations and repeatedly makes the choice between two types of actions: 1. putting on an item of clothing or 2. moving to another location. Children are free to provide input at any given moment. At some point, the robot will decide to go outside: it has the idea that it is wearing the right clothes. When the clothes are inappropriate to the outlined scenario, the story is continued:

"The bear thought it was too cold, and wants to go back in and try again. Can you help him again to find the right clothes?"

The activity is designed with a certain complexity, and randomness in the robot's actions, so that it will not take the right actions in one go. As a result, the activity can be done several times, with the robot taking increasingly targeted actions. The idea is that children exchange their opinions (what clothing should the bear have?) and discuss what feedback (opinion about the bear's action) should be given. An important aspect of the concept is the space that children get for this collaboration. Children can only give input when they know what the robot is doing or wants to do. That is why transparent behavior is needed: a form of communication of the robot to inform the children in advance what action it is about to take.

The concept maintained a symmetrical structure [12], as described in Chapter 3, to encourage collaboration between children. The concept is an activity where children are asked to help the robot (and take on a tutor role). This gives the peers a shared goal. From the start, children have the same possibilities and opportunities to interact with the robot. The hypothesis was that a division of labor is achieved by dividing roles, a result of collaboration. Due to its pace and movement, the robot will make the tutor role difficult in the activity. The hypothesis is that this is an incentive for cooperation. The principle that children teach the robot offers the opportunities to integrate teaching material.

5.2. Concept requirements

The concept requires not only that the robot can learn from human input, but also that children are able to see the robot learn from their provided input. Being able to see the robot learn over a period of time depends, among others, on the speed at which the robot takes actions, the complexity of the activity and the extent to which it receives constructive feedback. A number of requirements have been made that need to be confirmed to ensure that the concept is suitable as an activity with primary school children. It is required that:

1. The surface bot can learn from the input of children.
2. Children are engaged in the role of tutor.
3. Children communicate their feedback correctly.
4. Children perceive the surface bot's learning and relate it to their feedback.
5. The prototype fosters collaboration among children.

In order to meet the first requirement, a surface bot that can be taught by children, a framework is needed that is able to process the feedback of children, and can utilize it to improve its decision making. The second requirement states the necessity of engagement of children in their role as tutor. In order to learn the robot requires feedback, therefore the learning progress solely depends on the children. Children should provide frequent feedback with consistent values. This leads to the third requirement: an interface is needed for rewarding the robot's actions that is simple and clear, so its easily understood by children and used correctly. The fourth requirement assumes that children take their role seriously. In this case, the children must see the robot improve within a period of time that suits their attention span and is fitting for a classroom activity. These four requirements only consider the functioning of the robot and the interaction between child and robot. However, the goal is that it effectively stimulates collaboration between children, hence the fifth requirement. The feedback should come about through forms of collaboration between children.

Figure 5.1.: Impression of prototype 1.0. The surface bot with the character display positioned next to the hallway location. The clothes cards are located on the location card.

5.3. Realization

The first prototype (see Figure 5.1) was meant as a proof of concept. The decision was made to develop the prototype with a focus on the second, third, fourth and fifth requirement. For the first requirement, a reinforcement learning framework was meant to be realized, but given the time of implementation, it was decided to omit this in the first prototype. However, the robot must appear to take actions itself and use the feedback from the children. Furthermore, it must be able to move between the locations. To implement these aspects as autonomous behavior a lot of development is required. It was therefore decided to make use of tele-operation based on a Wizard-of-Oz approach, similar to the work of Verhoeven et al. [28]. The prototype ultimately consisted of three parts, see Figure 5.2 for a schematic overview.

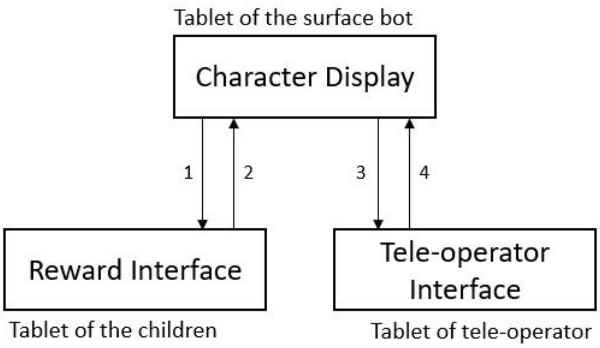

Figure 5.2.: The communication between the three main components of the prototype. The character display is the server, and communicates with the clients: the reward and tele-operator interface. The communication involves (1) status of the activity, (2) value and timing of rewards, (3) the name of the script's next action and (4) controlling the activity (start, stop, next action etc.)

The first part is the surface bot. An application is developed for the surface bot's tablet, further referred to as the character display. The realization of the character display is described in section 5.3.1. The second part is the application developed for the tablet of the children, described in section 5.3.2. This application is the reward interface for communicating feedback to the robot. The third part is the tele-operator interface, a tablet application for controlling the robot's behavior and movement. This application is described in section 5.3.3. Lastly, the objects created for an activity with the prototype are discussed in section 5.3.4.

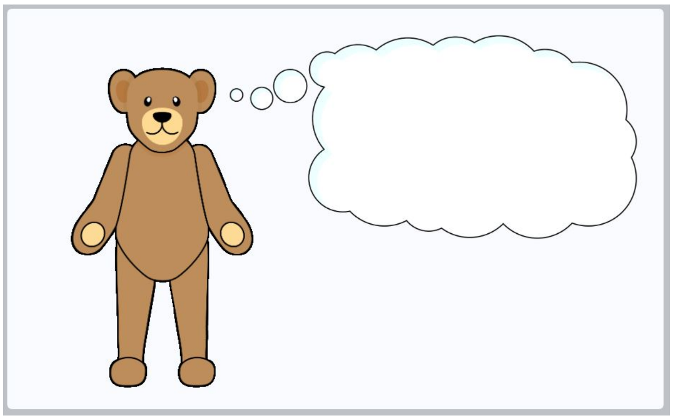

Figure 5.3.: The character display. It shows the character introduced in the story, and the thought cloud. In this figure, it is in the starting state: the bear wears no clothes and has not selected an item yet.

5.3.1. The character display

The character display shows a visualization of the story's character: the bear (see Figure 5.3). The thought cloud was designed as the transparency behavior of the surface bot. Preceding an action, the robot thinks about a piece of clothing for 3 seconds. A visualization of it can then be seen in the thought cloud. There is a repetitive action cycle: the robot thinks of an action and then executes it. The purpose of the thought cloud was to have a visual indication of the character's next action. The expectation is that children recognize the thought visualizations as they correspond with the objects used in the activity, and already start thinking about the appropriateness of the action. If it leads to adequate feedback, then a form of undo behavior would be the next step, since it would likely positively contribute to the surface bot's learning [26]. There are two types of actions that the robot can think of and perform: 1. putting on a item of clothing and 2. moving to a new location.

The character display shows a notification for 2 seconds when it receives feedback. The notification is a thumbs up image, if the value of the feedback was positive, and thumbs down when it was negative. An example of an action cycle would be: the bear thinks of the blue jacket (Figure 5.4). After 5 seconds, the thought disappears and the bear can be seen wearing the blue jacket (Figure 5.5). The children think it is a good decision and send positive feedback, after which a notification appears on the screen (Figure 5.6). The robot sequentially goes through different action cycles till it decides to go outside: the terminal state. It marks the end of an iteration. After this the surface bot started again, without any clothes on, at the starting location. This terminal state was displayed as a winter landscape illustration replacing the white background. Since the story is told that the bear wants to try it again, the intention is to keep the robot's progress in the next iteration. However, the robot did not have any learning capability, therefore a script was developed with sequences of actions. Over three iterations, the robot's actions become increasingly more accurate and ultimately lead to a set of clothes appropriate to the winter weather scenario. The script can be found in Appendix A.3.

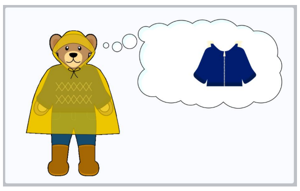

Figure 5.4.: Thought

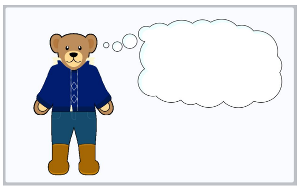

Figure 5.5.: Action

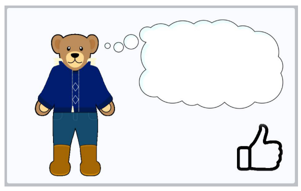

Figure 5.6.: Feedback

The tablet of the surface bot functioned as server that communicated with the clients: the reward and tele-operator interface. The communication was established via a local WiFi network. This structure was based on the work of Verhoeven et al. [28]. Figure 5.2 shows the communication flow between the character display, the reward interface and the tele-operator interface.

5.3.2. The reward interface

The reward interface enabled children to communicate feedback to the robot. The decision was made to use a tablet, since earlier studies indicated that the children had shown proficiency with touch screens in child-robot interactions [28, 8]. The children's tablet interface consisted of two interactive elements: a slider and a send button. Figure 5.7 shows what the application looks like with the slider in neutral position. Dragging the slider to a position in green meant positive feedback, and a position in red was meant as negative feedback. The highest position of the slider would be translated into the most positive feedback value (= 1) and the lowest position into the most negative feedback value (= -1). The idea was that the slider could stimulate negotiation, and could be used to reach a consensus when opinions differ, for example by going for an intermediate value. The send button needed to be pressed to communicate feedback. It was intended as additional confirmation whether the children actually wanted to send feedback to the robot.

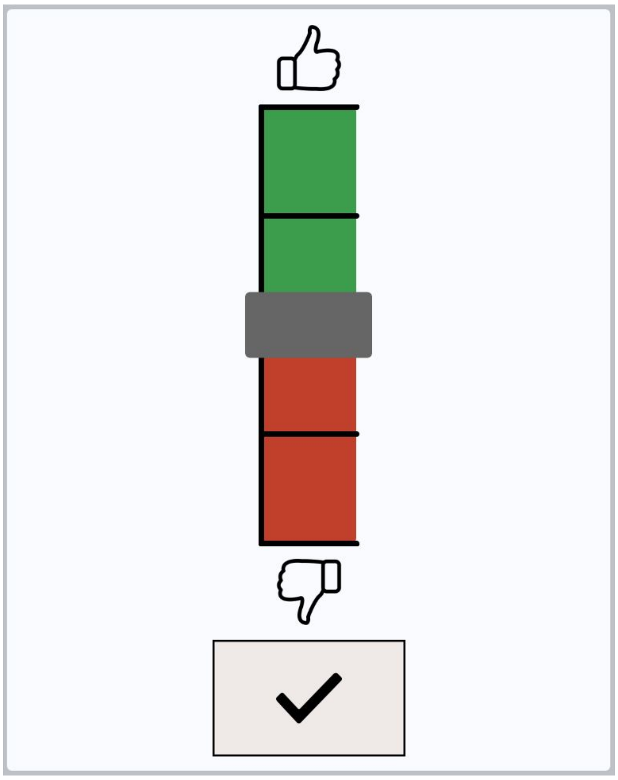

Figure 5.7.: The child's tablet interface. It consisted of a slider and a send button. Children could drag and position the gray rectangle anywhere on the range of the slider. The position determined the value of the feedback. The green color was used to indicate positive feedback, and red color for negative feedback. The send button needed to pressed once to send the feedback to the surface bot.

5.3.3. The tele-operator interface

The tele-operator interface was developed to control the activity and to simulate the autonomous behavior of the robot, see Figure 5.8. The interface had a start/stop toggle to control the activity, and a reset button to start over again at the end of an iteration. The space in the right-top corner displayed notifications that kept track of the actions the robot did, because the surface bot's screen was not always visible from the tele-operator's point of view. The duration of the thought of an action was fixed and automatically transitioned to executing the action. Continuing to the next action of the script was controlled by the tele-operator with the "next action" button. The moment the robot thought of a location, the tele-operator navigated it to the location. With the "next action" button it could be ensured that the robot did not continue until it had arrived at its destination. The movement of the surface bot was established via a Bluetooth connection between the tele-operator tablet and the surface bot's wheelbase. The arrow keys of the tele-operator interface were used for navigation.

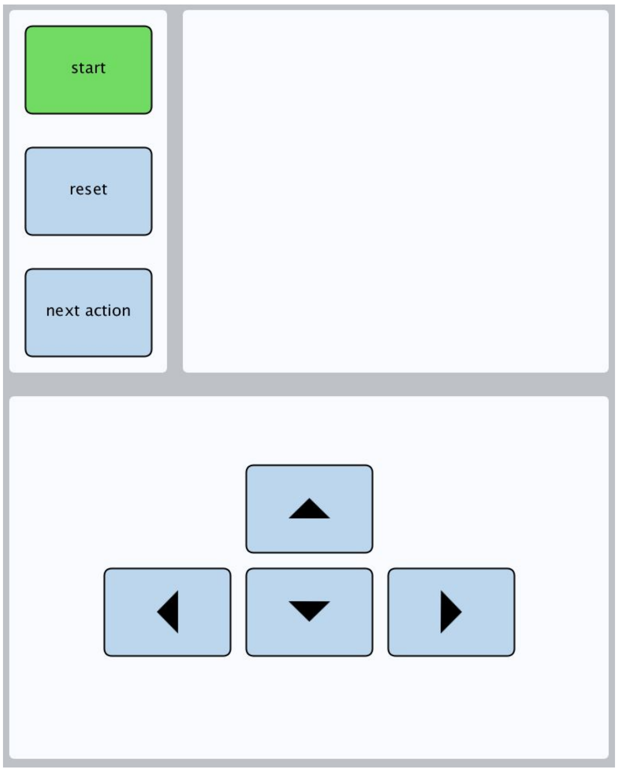

Figure 5.8.: The hidden operator's interface. The "start" button was used to begin or end the activity. The "reset" button returned the robot to the start state and continued to the next iteration of the script. The "next action" button was used to start the subsequent action. The arrow keys were used to move the surface bot another location.

5.3.4. Objects

The activity made use of four locations: hallway, closet, table and a coat rack. The four locations were used to create a square playing field in which the robot could move around. The locations were printed on A4-paper and reinforced with cardboard. Appendix A.1 shows the images that were used as locations. Locations were added to the game to utilize the surface bot's driving capabilities. The advantage of this is that the game becomes more interactive and less easy to follow for the children. Children will have to move along with the robot to see what it is doing or wants to do. The result may be that it encourages them to, for example, divide tasks. Each location had unique and associated items of clothing. These items were designed for one of the four seasons, or for sports. Sixteen 5x5cm cards were made from paper and cardboard. Each card depicted one item of clothing, including the name. An overview of the designed items of clothing, and the names used, are shown in Appendix A.2. The cards were meant as visual support for the children to be able to estimate what the robot wanted to do and could do at each location.

6. First Study: exploring collaboration and validating the concept

The first study explored how children act when they are asked to provide feedback on the behavior of an independently acting surface bot in a story-based activity. The first aim was to validate the concept by examining the reward behavior of children. The second aim was to get an insight into children's engagement, and the way in which they interact and collaborate with one another.

6.1. Goals

The first study intended to test the concept against the requirements established in Chapter 5, section 5.2. The first requirement was left out of consideration. Self-learning as a feature would have taken a lot of time to implement, hence the decision to first test the concept on the interaction between children, and their interaction with the prototype and to include self-learning in a subsequent version of the prototype. The pilot study validated the functioning of the prototype and explored the engagement and collaborative behavior of children during the activity. A number of supporting research questions were prepared:

1. How frequently and consistently do children provide feedback to the surface bot?

The first question aimed to get an insight into the degree of participation and engagement of the children (requirement 2). Feedback consistency relates to the perseverance of children to the given task. Is the feedback given proportionally over the time of the session, or are children actively giving feedback primarily in the beginning? The latter might indicate a loss of attention. The frequency of feedback is an important factor in the envisioned reinforcement learning framework. The surface bot learns every time feedback is communicated. With this, the engagement of children directly influences the progression of the robot.

2. How committed are children to the scenario outlined in the story?

The second question was meant to examine the children's engagement (requirement 2) in greater depth. Given that children provide consistent feedback over time during a session with the prototype, do they adhere to the scenario outlined? The scenario outlined in the story is the bear's search for the right winter weather clothing, and children are asked to teach the robot what clothes are appropriate for this. By observing the children during the sessions, and by placing the timing of feedback given next to the actions of that moment, the feedback values could be used to determine whether children are consistent in the value of their feedback.

Consistent values for the robot's actions indicate that children collaborate and work according to a shared plan. If this is not the case, it might be because children do not reach a consensus, or the children lose sight of the scenario and start focusing on a different set of clothing.

3. How do children tend to reward the surface bot, and what are they trying to communicate?

The third question was aimed to review the children's usage of the reward interface, and what they intended to communicate (requirement 3). The reward interface's slider makes refined feedback possible, and it was expected that it could be a point of negotiation or discussion: a way to encourage collaboration. The distribution of the given feedback values was examined to gain insight into the use of the slider. Furthermore, the observations and timing of the feedback must show the extent to which children assess the two types of actions: putting on an item of clothing and selecting a new location. It must show whether the slider is a good means of communication for the purpose of assessing the robot's actions, or whether an extra or different means of communication is needed.

4. To what extent do children perceive the learning of the robot, and do they understand the effect of their feedback on it?

The fourth question aimed to explore the understanding children have of their role as tutor, and the influence their feedback has on the robot's learning (requirement 4). It is purely exploratory, since the robot works according to a script and does not actually learn from the feedback it receives. However, observing how children react to robot's actions might reveal whether they understand the effect of their input and whether they notice that the robot is converging towards a set of winter clothing.

5. How do children collaborate while acting as tutors of the surface bot?

The fifth and last research question aimed to explore the ways of collaboration between children while they are engaged in the activity with the surface bot (requirement 5). It should provide a baseline for the collaboration between children in follow-up studies with this prototype. The degree of collaboration is assessed on the basis of observations of the children in the sessions with the prototype.

6.2. Participants

The first study was conducted with 12 children (mean age=5.75) at the daycare facility on the campus of the University of Twente. The age of the children varied between 4 and 8 years old. The focus was on children in the age of 5 to 7, but it was difficult to predict how many children of a certain age would be present at the daycare facility on the day of testing. It was therefore decided to let all children participate who wanted to and were allowed to. The children who participated in the test already had a general consent from their guardians for participation and collection of data for user tests of the University of Twente for the year 2018-2019. Each of them had consent for collection of personal data, including recording of the sessions. Table 6.1 shows the age of the participating children.

Group	Age child 1	Age child 2
1	7	8
2	5	6
3	6	7
4	7	5
5	4	5
6	5	4

Table 6.1.: The age of the children that participated in the first study.

6.3. Setup

The study was conducted in a separate room at the daycare. It was a spacious and bright room, next to the room where the children are normally at the daycare. The setup was built on two attached tables instead of on the floor, to prevent children from hindering the robot by walking through the setup. The location cards were placed on the corners of the table. The clothing cards were placed on the intended location. Figure 6.1 gives a top view impression of the setup.

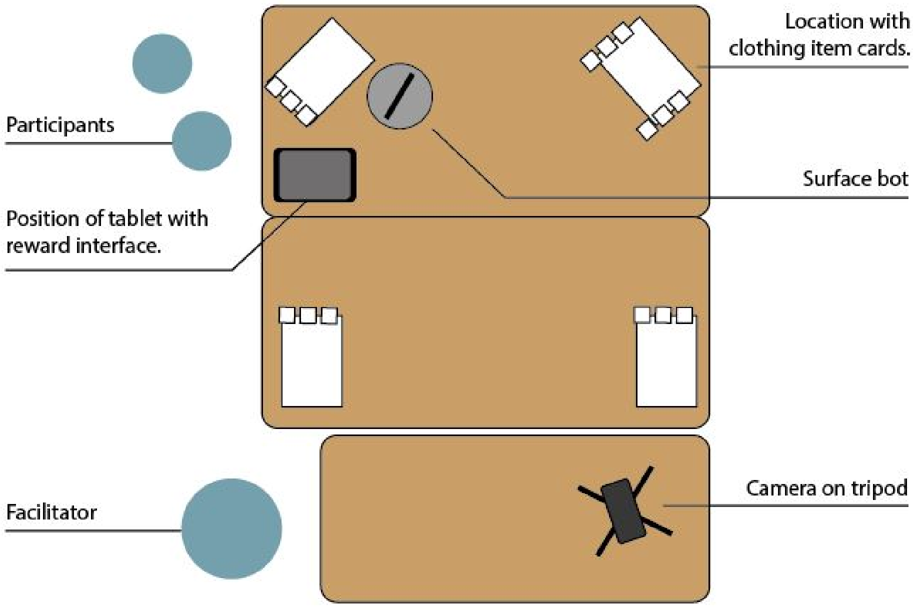

Figure 6.1.: Impression of the setup of the first study. The participants start the activity near the reward interface, next to the "hallway" location. The surface bot would move between the locations at the edges of the table. The camera, on a tripod, was located on a table next to the activity to be able to make good recordings.

6.4. Procedure

A procedure was developed to ensure a consistent approach of the facilitator during the sessions of the first study. The procedure consisted of three parts: the introduction, the test and the closing.

Introduction

The aim of the introduction was to let the participants adapt and become comfortable in the testing environment. Following the recommendation of Hanna, Risden and Alexander [14], I started with some small talk to establish a relationship with children. After the small talk, the story of the activity was told to the children and the surface bot was introduced as the main character. In the story the situation of the robot was explained, and the help of the children was requested. It was also made clear that the robot knew nothing, but wanted to try it for itself. The children must together help the robot in finding the right clothes. The story aimed to introduce the activity and the surface bot. It should also motivate them to participate, as children in the age of 3-7 enjoy fantasy [22].

Then the activity and setup was explained to the children. One of the aims of the explanation was to manage the expectations of the children [14]. Children could start the test with an unrealistic image of a robot that can talk, drive around and interact with them. To prevent them from being disappointed during the activity, it was clearly told what the robot was able to and would do. The explanation ended with a check if children understood the purpose and functioning of the reward interface. The children were asked if they could demonstrate on the tablet what they would do if they thought an action of the robot was very good, very bad or "OK". When the children gave feedback with respectively a positive, negative, and a value in between, the facilitator started the actual test. When children showed that they did not fully understand it, it was demonstrated once by the facilitator. The introduction ended with the researcher giving room to the children for questions related to the setup.

The test

When the facilitator started the test, the surface bot began with the first action cycle. The facilitator stayed relatively close to the setup. It should be noted that this could affect the way children act. It is recommended though for younger children as they need reassurance and encouragement [14]. The decision had a practical reason: the facilitator acted as tele-operator of the surface bot and had to have a clear view of the setup. Hanna et al. [14] do mention that shyer children may find the presence of a facilitator uncomfortable, but they tested only with one child. It was assumed that this did not apply in the same extent to pairs of children. The following actions are being taken by the operator during activity:

- Starting the activity. The operator (as facilitator) asked the children if they were ready. If they were, the activity started. From this moment the robot started taking actions, and children could give feedback via the reward interface.
- Continuing to the next action. The operator would decide when the next action cycle would start. The next action was determined based on a script of actions, see Appendix A.3.
- Controlling the movement of the surface bot. When the next action was a location, the operator was notified on the tele-operator interface. The operator would then use the arrow keys to navigate the robot towards the new location. When the robot was near the destination, the operator indicated on the tablet that the robot had arrived by selecting "next action".
- Continuing to the next iteration. An iteration ended with the surface bot "going outside". When this happened the operator navigated the surface bot to the "hallway" location. The "reset" button was used to display the bear in the starting state (without any clothes) and from here on the next iteration could start.
- Answering questions. Children's questions about the activity were answered during the activity. This concerned uncertainties that children had about their task or how the tablet worked. Questions for confirming whether their feedback was correct, or not, were avoided by emphasizing that it was entirely up to them, that it was their task to assist the robot. The children were not guided or motivated by the facilitator to provide feedback.

Closing

After the test had ended, the children were asked a number of questions. Children tend to confirm and please the researchers, therefore only open and non-confirming questions were asked. Firstly, their opinion was asked regarding the activity to get an insight in their motivation. A second question that was asked: what clothing did you want to give the bear? An A4-sized picture of the bear was placed on the table and the children were asked if they could place the clothing cards on the picture that they had wanted to put on the bear. A third question was asked on their ideas to improve the activity with the surface bot. Older children are known to enjoy "being asked to give ideas about how to make things better [14]." For the purpose of further improving the prototype, this question was therefore asked. Furthermore, children were made aware of the Wizard-of-Oz approach during the activity. It was explained to the children that there was a small deception during the activity, since the facilitator controlled the robot. The time was taken to explain the tele-operator interface to them. They were told that the deception was necessary, because the development of a fully autonomous robot was not possible within the time of this research. The test ended with thanking the children for their participation and helpfulness.

6.5. Measurements

The results are acquired from three sources: 1. observations of the recording of the activity, 2. logging of the interaction of the children with the robot and 3. the questions asked by the facilitator at the end of the research.

To answer the first research question, the consistency in providing feedback was measured based on the logged data of the interactions with the reward interface. This meant the frequency of feedback over time was considered in order to get an insight in children's engagement for each iteration. Whether children adhere to the outlined scenario is determined based on the observations and the feedback values over time.

The tablet of the children kept track of the interactions. Important metrics are the timing of the feedback, the frequency and value (positive/negative, high/low) of the feedback signal. Further analysis relies on observations of the children. If consent is given by the parents, the session was recorded for later analysis. The recordings were observed on how the children collaborate, and interact with the surface bot.

Questions were asked to obtain insights about how children experienced interacting with the prototype. It could also lead to new ideas for the surface bot, the story or the role of tutor.

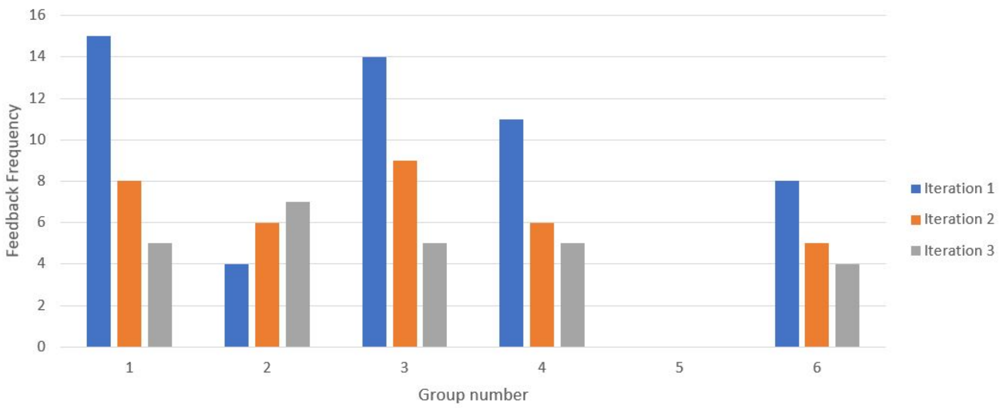

Figure 6.2.: The absolute feedback frequency per iteration for each pair of children, participating in the first study.

6.6. Results

To answer the first research question, the use of the reward interface by the children was examined based on the data obtained from logging the children's interactions with the reward interface. The frequency of feedback was determined per session, as a measure of engagement of each pair. Next, the frequency was determined per iteration to give insight into children's consistency over time in providing feedback. It should be noted that the feedback is only logged once per action cycle. It occurred multiple times that children sent their feedback to the robot one after the other, or that they pressed the feedback button several times in quick succession. In this case, only the last interaction was saved. Figure 6.2 shows the frequency of feedback during each iteration per session. Pair 5, children in the age 4 and 5, did not provide any feedback to the robot. They showed excitement by laughing and running caused by the robot's movement and actions, but showed little understanding of the purpose of the activity and how to operate the tablet. The other five pairs of children gave feedback 22.4 times (on average) to the robot that executed 46 actions over 3 iterations. This is 48.7 percent compared to the total number of actions of the robot. Except for the second pair, the absolute figures show a strong decrease in feedback over the iterations. However, each iteration is shorter and contains fewer actions and thus options for feedback. Figure 6.3 shows the relative frequency of the children's feedback. The relative frequency is a calculated percentage based on the feedback frequency in relation to the number of actions of the robot per iteration. This provides a more accurate picture of the children's consistency in communicating feedback. It shows that the frequency either has a small decrease, or remains fairly stable. Only pair 2 shows an increasing frequency of feedback.

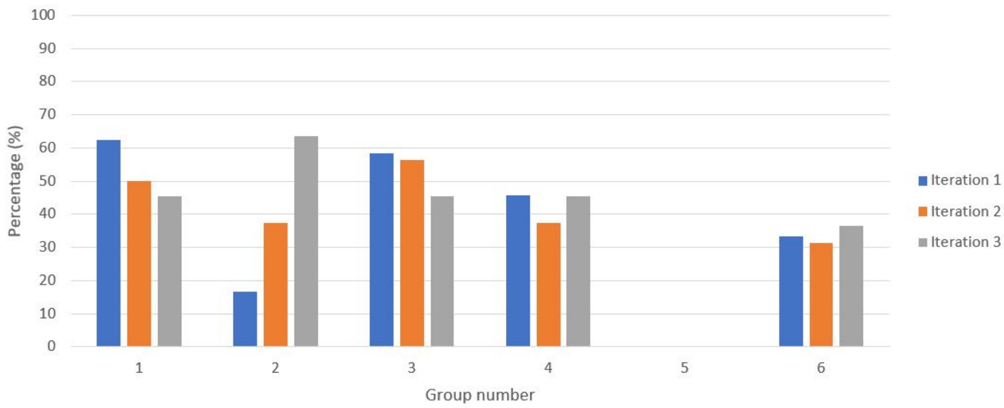

Figure 6.3.: The relative feedback frequency per iteration for each pair of children, participating in the first study. The percentage is based on the absolute feedback frequency relative to the iteration's number of actions.

Answering the second research question required investigating the feedback values given, and how their timing related to the surface bot's actions. This comparison was ultimately not done since it could not be assessed accurately with the data obtained from logging the reward interface interactions. Based on observations, children did seem to take the robot's scenario seriously and adhere to it throughout the session.

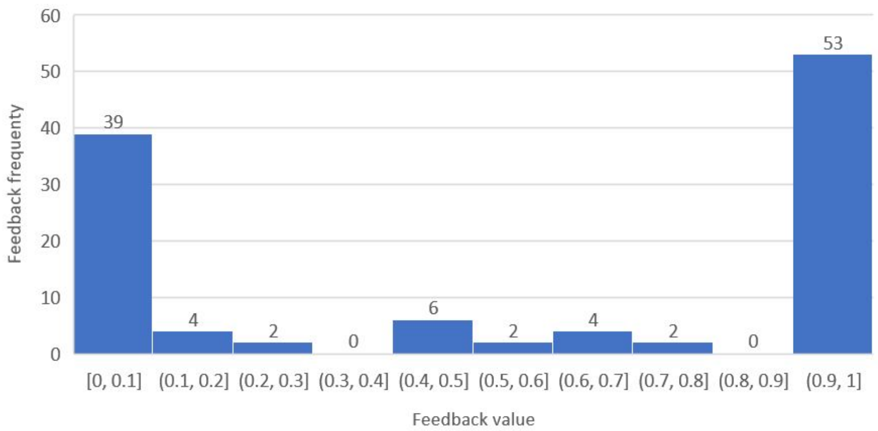

Figure 6.4.: The distribution of the feedback values communicated to the surface bot, based on the feedback of all pairs of children that participated in the first study.

The logged data from the reward interface interactions also gave insight in how children rewarded the surface bot. It showed that children both positively and negatively reward the robot. They thus utilize the range of the slider, but in a very binary way. The distribution of the feedback values was looked at in more detail, in order to provide an accurate answer to the third research question. The feedback provided was expressed in a value between 0 (negative) and 1 (positive), and 82.1% of the feedback had a value between 0 and 0.1 or between 0.9 and 1.0, respectively 34.8% compared to 47.3%. Children generally gave slightly more positive feedback than negative feedback. Furthermore, they were observed to communicate feedback when the action of the robot involved an item of clothing. Children did occasionally voice their opinion about the robot's thought of a next location.

Observations of the sessions enabled answering the fourth research question, whether children understood their role and saw the effect of their input. A child (pair 1) expressed halfway through the session: "Ow! But the more tips we give him..." Although he did not finish his sentence due to an interruption of the other child, in the context it seemed to indicate an understanding of the effect of his feedback. Pair 3 was asked "Did you think the robot was learning?" after the last iteration and one child answered the question with: "Now, it was quite right." This seems to indicate that they saw improvement in the robot's decision making.

The fifth research question meant to explore the collaboration between children during the session with the prototype. Collaboration was observed to be limited between the pairs of children. Children took turns by alternately controlling the tablet. This happened with and without explicit (verbal) coordination. The children of pair 3 agreed at the outset to use the tablet in turn, but repeatedly minor conflicts occurred where one voiced the opinion that the other one was not listening and gave incorrect feedback, or did not want to pass on the tablet. A more extensive form of collaboration, in combination with turn taking, was seen by pair 4. They established a division of roles, where one operated the tablet and the other followed the robot and provided information regarding its action. This pair did not pick up the tablet, it remained at the initial position on the table. It was not made clear to the children whether they could pick up the tablet and walk around with it, but most of the pairs did. Each pair, except pair 4, stayed together most of the time, following the robot from location to location.

6.7. Discussion

With regard to the engagement of the children, they were motivated and enthusiastic during the activity. The results showed that the children provided regular feedback throughout the activity, except for pair 5. Pair 5 were both very young, and seemed shy. This may have contributed to the fact that they communicated little and did not use the tablet. The relative feedback frequency for the three iterations indicated that the other pairs of children maintain engaged by consistently providing feedback. Pair 2 was an exception and showed an increasing frequency of feedback.

The story seemed to ensure children's commitment, and could be further improved in future versions by making the surface bot more interactive and allowing it to introduce and talk to the children. It has been decided not to add this in this research, since there is already a high level of participation by children. This study indicates that the minimal story form already ensures that children give thoughtful feedback over time.

This study also showed that children used the slider in a binary way with mainly extreme values being communicated to the robot. This makes the slider somewhat superfluous. It might be due to a certain unanimity that children had about what was right and wrong. As a result, certain options were consistently assessed as good and others as wrong. It was concluded that the slider was not necessarily unsuitable for further testing, children understood it and made use of it. The focus would rather be on applying changes to the prototype that ensures less unanimity among the children, in the expectation that differences in opinions would lead to more extensive collaboration and more sophisticated feedback. In the first version it was easy to deduce from the objects what the right items of clothing were. By introducing more ambiguity, there may be more need for consultation and collaboration among children [12].

How children started using the reward interface seems to indicate that children understand their role and the effect of their feedback. At least one child seemed to notice the improved performance. Although it should be noted that it was (unintentionally) asked as a closed question, which children tend to answer affirmatively.

The collaboration between children in this study was limited. Collaboration was shown by providing feedback to the robot in turns. There was little conflict about the task, but rather about who operated the tablet. Children decided to walk around with the tablet which created a situation where the activity was doable for one child. It allowed tracking the robot's actions and communicating feedback at the same time. This may have caused little incentive for collaboration. This is confirmed by the pair of children who were observed to divide roles. During this session, the tablet remained on the table. It resulted in a situation where the children relied on each other to share information and opinions. It was decided to change the prototype to further encourage collaboration between children based this observation of pair 4. The second prototype has a setup with the children's tablet in a fixed position. It should stimulate children to divide roles. If a division of roles arises without instruction from the facilitator, it is a result of collaboration where the children are dependent on each other and both actively participate in the activity (as opposed to taking turns, as is now often the case in the first study).

7. Prototype 2.0: a learning surface bot

This chapter describes the second version of the prototype. The prototype was changed based on the insights acquired from the results of the first study. First, the modifications to the prototype are discussed in section 7.1. Second, the learning framework used is explained in section 7.2. Last, section 7.3 describes the realization of the prototype in detail.

7.1. Modifications to the prototype

During one session in the first study a division of roles was observed. This pair might have assumed that it was not allowed to pick up the tablet, or didn't think of it. It resulted in a continuous back and forth walking between the robot and the tablet. At a certain moment, they decided to divide roles, with one of them tracking the robot and telling what it was doing and the other operating the tablet. As discussed in Chapter 6, it created a dependency which resulted in communication regarding the activity. This observation was the inspiration for the first modification of the second prototype: a fixed tablet. Children were not allowed to hold the tablet anymore. The tablet was attached to the table. This makes the effort for both children to track the robot and operate the tablet more time consuming and a less efficient. It was hypothesized that this would stimulate the establishment of role divisions between them.

7.1.1. Environment ambiguity

According to Lai [19], the design of a group task can influence to what extent collaboration is likely to occur. As mentioned in Chapter 3, trivial and unambiguous tasks provide less opportunities for collaboration. The results of the first study indicated that children probably had the same preferences regarding the items of clothing. In order to make the options more ambiguous, alternative items of clothing were introduced. These included items only differing in color and shape, compared to other items. It was hypothesized that the introduction of ambiguity in the options would result in more difference between the preferences of children, which would provide more incentive to discuss and negotiate.

7.1.2. Undo behavior

In prototype 1.0, the robot responded to feedback from the children by displaying a "thumbs up" or "thumbs down" notification. This only lets the children know that the robot received their feedback. The second prototype made use of a learning framework instead of executing a fixed sequence of actions. It was decided to add a direct behavior to the robot that showed the children that their feedback was used and listened to. Inspired by Thomaz and Breazeal [26], a form of undo behavior was added as a response to the negative feedback.

7.2. Learning from feedback: a Q-learning framework

Reinforcement learning was used to enable the surface bot to take actions and to learn from feedback it receives. Reinforcement learning is "the problem faced by an agent that must learn behavior through trial and error interactions within a dynamic environment", as described by Kaelbling et al. [18].

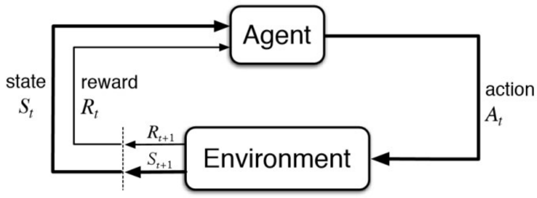

Figure 7.1.: Reinforcement learning

Figure 7.1 shows the basic principle of reinforcement learning. An agent (the surface bot) has an environment (the locations and the bear) in which it can perform actions. An action (putting on the jacket) of the agent changes the state of the environment (the bear wears the jacket). In order to achieve an optimal policy (what are the clothes the bear should have before going outside?), the bear must be rewarded (positive feedback) or punished (negative feedback) for the actions it takes. The reinforcement learning model of this prototype consist of a discrete set of states S' and actions A, and a reward function. Note that there is no explicit reward function, since the input of children determines the rewards on the robot's actions. In this prototype the surface bot's state was based on the clothes it's wearing. It consisted of five independent sub-states: the bear's body parts. Each of the items of clothing belonged to one of the five sub-states. Appendix B.1 gives an overview of the possible items per body part. The bear could only wear one item of clothing per part. The state changed based on the robot's actions. An example of how the state of the robot would progress:

The bear starts without any clothes on, which is expressed as "empty" for each sub-state: [empty, empty, empty, empty, empty]. Then the bear decides to put on the "cap" and the first sub-state (the head) changes from "empty" to "cap". The new state of the bear becomes: [cap, empty, empty, empty, empty]. The bear then continues with the next action, and the state is again changed accordingly.

The actions (A) were the clothes that the surface bot could find and wear in this prototype. There are 17 items of clothing across 4 locations. Based on the work by Thomaz and Breazeal [26] a Q-learning algorithm adaptation was implemented that incorporates human feedback.

As described in Chapter 4 section 4.4, the Q-function was used to determine the Q-values (the utility or quality of being in a certain state) for each state-action pair. The actions of A were possible regardless of the current state. The number of unique states, which are the unique combinations of sub-states, was 1600. If a Q-table (containing the state-action pairs: the states and the possible actions per state) was created based on the states of S' and the 17 possible actions, it would have resulted in a space of 27200 state-action pairs. Given the number of actions the surface bot takes during a session of about 10 minutes, I expected that it would not be possible to see the robot learn with so many states. Therefore, a Q-table was created for each individual sub-state based on the actions that would have resulted in a change of the particular sub-state. The five Q-tables had a combined total of 76 possible state-action pairs. The state-action space was thus considerably reduced. This increased the chance that the robot actually learns to an extent that it can be perceived by children.

The location of the robot was not included as a state of S'. If the robot decided to perform an action at another location, the robot would do an intermediate action which is traveling to that location. The location of the robot was tracked in the software based on the last action of the robot. The software contained a table with the locations and possible actions per location. With the knowledge of the current position of the robot, it was possible to determine whether a chosen action of the robot was at a different location. If this was the case, an intermediate step was taken: a location action cycle. Children saw this as a separate action, where the robot thought of a location and then drove to it. Once arrived, the action cycle of the action began. Children could give feedback during the location action cycle, the robot also showed a notification when receiving feedback. The only difference with the other actions was that the feedback value was not used for the Q-learning algorithm.

The robot had two possibilities for selecting a subsequent action: exploration or exploitation. Exploration is finding out more information regarding the environment, and exploitation is utilizing known information to maximize the reward. If the robot only explores, it selects random actions in order to get information (the reward of being in a certain state) resulting from it. When the robot only exploits, it would always go for the action with the highest Q-value which implies that it takes actions it got previously rewarded for. An action-selection strategy is therefore necessary to address this trade-off, and enable the robot to explore the state-action space while converging to a reasonable policy.

Action selection strategy

The epsilon-greedy approach was used as action selection strategy [27]. An epsilon value ϵ was determined at the start of each iteration, according to equation 7.1, where N is the number of the iteration. The second study consisted of four iterations, with which the epsilon value in the first iteration had a value of 1.0 and in the fourth iteration a value of 0.25. The epsilon value determines the "greed" of the robot: the extent to which the robot exploits its policy to get the highest cumulative reward. An agent with a very greedy approach can exploit a sub-optimal policy that results in a cumulative reward that is lower than the maximum cumulative reward.

In this case, the higher rewards were not found because the agent did not explore the possible actions in every state.

The action-selection strategy starts with generating a random number between 0 and 1. If this value is greater than the predetermined epsilon value the robot will pursue actions based on their Q-values. If the value is smaller than the predetermined epsilon then the next action will be randomly selected. Hence, the agent starts with exploration (epsilon=1.0) and over the iterations it will focus more on its policy. It shifts the focus more and more towards his policy, selecting actions that lead to the highest expected cumulative reward. In the fourth and last iteration, the agent is meant to select the majority of its actions based on the policy. To prevent a situation that the robot keeps choosing between two actions, some randomness was maintained. It is a situation that might occur if children continue to rate two items of clothing positively for the same part of the bear, for example two coats. The bear cannot wear two coats, and every time he wears one coat, the other coat will be selected as the next action since it has the highest Q-value.

$$\epsilon = 1.0/N \tag{7.1}$$

Implementing Undo behavior

UNDO behavior of the surface bot is one of the modifications mentioned in section 7.1. Based on the UNDO behavior [26] described in Chapter 4 section 4.5.4, it is implemented as part of the Q-learning algorithm. Negative feedback is utilized to stop the current action cycle, and select a new action. In this prototype, negative feedback canceled the current thought if the thought was an action concerning an item of clothing. It is further referred to as CANCEL behavior, as UNDO implies the action has already been executed and is reversed. This is not the case, in a visualized thought on the robot's interface an action is proposed and children are able to let the robot "rethink" to a more favorable action by providing negative feedback. When an action is canceled, the Q-function is updated without the robot transitioning to the next state.

7.3. Realization

The second prototype is an improved version of the prototype described in Chapter 5. The realization of the changes, including the modifications described in section 7.1 of this chapter, are described in this section. This prototype had a renewed reward and tele-operator interface. The slider of prototype 1.0 had a distinct difference between green and red, and two horizontal stripes as visual aid to indicate that the slider can be used for more than binary feedback. The first study results showed that the feedback of children consisted mostly out of extreme values. Figure 7.2 shows the new interface where the red and green field and the horizontal lines have been replaced by a green-red gradient symbolizing the transition from completely right to

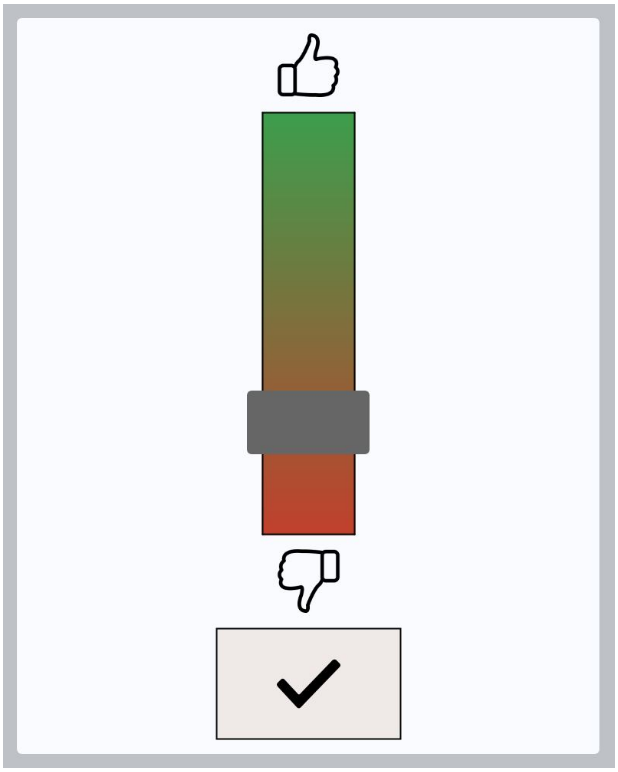

Figure 7.2.: The reward interface for the children. The slider can be used to determine the feedback value. The send button below the slider sends the feedback to the character display of the surface bot.

completely wrong. The transition of color is meant to clarify the utility of the slider, in an attempt to encourage more diverse feedback.

A number of functions on the tele-operator interface were removed and replaced by new ones. The functions “reset” and “new action” were replaced by “continue” and “arrived”. The reset function was used to start the next iteration of the script of actions. The actions of the robot were not predetermined, but a result of the action strategy implemented in the Q-learning algorithm. The “continue” function was used to let the robot proceed in a next iteration starting from the “hallway” location without any clothes on, but with its updated Q-function. The Q-learning framework also made the new action function obsolete, since the facilitator did not control the action nor when the next action started. Only in the case of a location did the facilitator have to send the robot to the correct location. The “arrived” button was used to indicate that the robot had reached the destination and could start the next action. While the robot takes autonomous action, it is still predetermined when it will go outside. A time duration was set for each iteration and as soon as the iteration lasted longer than that duration, the first subsequent action of the robot would be to “go outside”. A session with the prototype had four iterations, where each iteration was shorter than the previous one. The duration of the four iterations, were respectively 5,3,1 and 1 minutes. The idea behind the decreasing duration was that it would fit

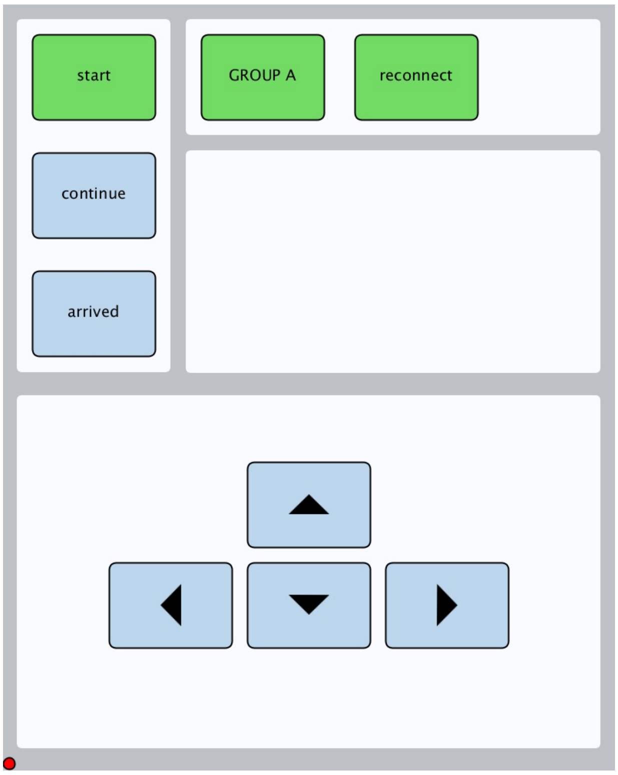

Figure 7.3.: The operator interface. The "start" controls the beginning and ending of the activity. The "group A" button is used to determine the condition (the pace) of a session. If an iteration ended, the tele-operator used the "continue" button to start the subsequent iteration. The "arrived" button was used to indicate that the robot had arrived at a new location, and could continue with the next action. The red dot indicates the connectivity status. Red indicates a disconnection from the character display. The "reconnect" button was used to establish a new connection. The arrow keys were used to navigate the robot.

the robot that is increasingly focused on exploiting based on the current policy. The more you know, the faster you know how to answer. Every iteration should result in an improved policy of the robot, which should lead to the (according to the children) correct items of clothing faster. In addition, with a fixed length of the session, it is guaranteed that children participate long enough and have the opportunity to work together.

Lastly, ambiguity was introduced with new actions for the robot which were choices between similar items of clothing. The objects and visualizations were adjusted accordingly. Furthermore, the names of the objects were made more general to ensure they did not hint to a certain outfit or type of weather. For example, one attribute used in the first prototype had the description: "winter shoes." This was changed to: "shoes." The attributes of the second prototype did not spoil the intended weather, but rather let it open to the interpretation of the children. Appendix B.1 shows the items of clothing used in prototype 2.0.

8. Second Study: measuring collaboration and the influence of pace

This study examined the collaboration that took place between children who participated in a session with prototype 2.0. It aimed to explore if prototype 2.0 is an improvement of prototype 1.0 with regard to encouraging collaboration between children.

8.1. Goals

This study aimed to answer the two main research questions stated in Chapter 1, section 1.1. The first main research question meant to explore how to utilize the surface bot to successfully foster collaboration between children. This study is a step in this, and aimed to answer the following question:

How do children collaborate in the role of tutor of a reinforcement learning surface bot in a story-based activity?

The second main goal of this research was to be able to evaluate the prototype on its effectiveness of encouraging collaboration. Having a method for evaluating collaboration is a prerequisite for answering the main question of this study. In an attempt to answer the second main research question, a framework for the evaluation of collaboration is proposed based on the literature described in Chapter 3. Three sub-questions were formulated to gain a better understanding of the extent to which collaboration is encouraged in an activity with prototype 2.0.

1. How does the collaboration between children with this prototype compare to the collaboration with prototype 1.0?

This study made use of prototype 2.0. As described in chapter 7, the changes made are based on conclusions drawn from the results of the first study. The proposed framework was developed to quantify collaboration to get an insight in the level of collaboration between children. With this, the first study can also be compared with the second study. It should reveal if the changes made to the prototype contribute to more collaboration between children.

2. How does the children's use of the reward interface compare to the reward interface usage of children from first study?

This second research question is formulated to explore this prototype's effect on the reward values given by the children. The first study showed that children tended to use the slider in a binary way communicating mainly extreme values. The assumption was that the ambiguity in the robot's action might result in a more proportionate distribution of feedback values. It was argued that ambiguity should lead to different preferences among children, which should lead to more negotiation. The slider's range of values are expected to be used for compromises as a solution to conflicting opinions or preferences among children.

3. How does the pace of the surface bot's actions influence the collaboration and rewarding behavior among children?

With this third question, it was meant to explore the effect pace has on the collaboration between children. Pace is defined as the speed at which the robot takes actions. In the first study, the pace differed per action cycle as it was regulated by the facilitator. Prototype 2.0 has a reinforcement learning framework that autonomously selects actions and performs them with a fixed pace. It was argued that for a low pace activity it was easier for children to perform the activity individually and therefore collaboration is less likely to occur. The alternative is that a slow paced activity gives children more time for consultation, thereby contributing to collaboration. An activity with a relative higher pace would make it more difficult to track the robot and communicate feedback. The activity becomes more challenging which could be an incentive for children to work together, for example in the form of a division of roles. On the other hand, there is less room for discussion and negotiation. It was therefore decided to look at two conditions of pace and the effect it has on the collaboration of children, and their interaction with the surface bot.

8.2. Participants

The research was conducted at primary school De Zwaluw in Markelo with 9 pairs of children. The age of the children was between 6 and 10 years (mean = 8.00). It was ensured that the pairs consisted out of children from the same class. Table 8.1 provides an overview of the pairs of participants. The demographic data included in this research were the age and gender. The age can provide insight into whether the older children are inclined towards, for example, more or other forms of collaboration. The parents or legal guardians of the children have given permission to participate in this specific study. Hereby they agreed to collect data from recordings and logged user information gathered during the activity. The data was collected only for the purpose of answering the formulated research questions. The data is not used for secondary purposes or shared with any third parties. The data was deleted after the research.

8.3. Setup

The tests were conducted on two different days. The first day, pair 1-6 participated in the study. One week later, pair 7-9 participated. The tests were conducted at the Zwaluw, a primary school in Markelo, but in different classrooms. Both rooms were considered to be a familiar setting to the children. The setup was similar to the that of the first study. The robot's moving space consisted of four tables pushed together (approximately 2x2m). The A4-printed locations with

Group number	Intended Condition	Actual Condition	Age child 1	Age child 2	Gender child 1	Gender child 2
1	A	A	9	9	W	M
2	B	A [1,2]	8	9	M	W
3	A	A	9	9	M	W
4	B	B	8	10	W	W
5	A	A	9	8	M	M
6	B	B	6	7	W	M
7	A	A	9	9	M	W
8	B	A [2]	7	6	W	W
9	A	A	6	6	M	M

[1] The test was temporarily interrupted halfway due to a technical defect. After fifteen minutes, the test was resumed from the second iteration.

[2] Test accidentally performed under the other condition. Discovered afterwards based on the logged data. The cause is a communication error between the operator's tablet and the surface bot tablet.

Table 8.1.: Overview of age, gender and test condition of the participants. The gender is indicated as: W (female) and M (male).

the corresponding item cards (10cmx10cm) were at a fixed angle of the moving surface during the tests. A change that was made, was the fixation of the tablet to the table. The tablet was positioned next to the "hallway" location.

8.4. Pilot

Prototype 2.0 was tested in a pilot with 7 children at BSO de Vlinder at the University of Twente campus with the purpose to identify technical or procedural shortcomings. Most of the children already participated in the first study. It would have been interesting to compare their behavior and degree of collaboration to their first participation. However, one of the technical errors that occurred had to do with the recording equipment, as a result of which these sessions could not be observed and were not further included in this study. Besides this error, a number of technical defects emerged in the communication between the applications. The recording equipment and the observed defects were solved before conducting the actual study.

8.5. Procedure

The procedure of this study was similar to the procedure of the first study, described in Chapter 6, section 6.4. The main difference with the first procedure is the role of the tele-operator and the introduction of a pre-test.

Each session started with the story of Ted which introduced the problem of the surface bot and the task of the children. Subsequently, the children individually made a pre-test that meant to give a picture of the children's preferences. At this point they did not have the opportunity to share or discuss their preferences. The pre-test therefore gives insight of the extent to which they have shared preferences. Appendix C.1 shows the pretest form. The facilitator instructed the children to circle the clothes that they wanted to the bear to wear. The children were not allowed to look at each others form or consult each other during the pre-test. I hypothesized that if the preferences of the children differ, they need to negotiate and would display a higher degree of collaboration. If children are completely in agreement, the task becomes unambiguous and there is little need for negotiation.

During the test, the role of the tele-operator was limited to controlling the robot's movement, since robot's reinforcement learning framework enabled the selection of next actions. The operator still regulated the activity when the game started and under which condition: high pace or low pace. The conditions were chosen alternately, instead of first one condition then the other. If the first groups were all tested with condition A, then condition B would occur at the end of the day (for example). The attention span of the children could be less, which influences the test results. Lastly, the tele-operator controlled the continuation to the next iteration. When the duration of the iteration had expired, the robot decides to "go outside". This marks the end of an iteration, after which the operator navigates the robot back at the start location and, when the children are ready, starts the next iteration. Each session ended with a number of questions about the children's thoughts and ideas regarding the activity.

8.6. Evaluation framework

This section describes the framework used to assess the degree of collaboration between children. The proposed framework of this research is two-fold. The first part is based on the indicators of the social component of the framework for collaborative problem solving [15, 6] discussed in Chapter 3, section 3.2. The second part are possible and expected outcomes of collaboration between children. The first part means to quantify the level of collaboration for each pair of children. The method and metrics used to obtain quantifiable results via this framework, are discussed in section 8.7. The second part is meant to identify the outcomes of the collaborative behavior that children show during the sessions.

8.6.1. Part one: measuring collaboration

The framework of Hesse et al. [15] consisted of indicators that belong the social or cognitive part of the framework. The social component of the framework is related to the collaboration and the cognitive part to the problem solving. The problem solving skills of children fall outside the scope of this research, and only the degree of collaboration is considered. The indicators of the social component provide a good basis for a framework aimed to assess the level of collaboration in pairs of children. However, their indicators and descriptions were aimed at determining the collaborative skills of an individual. The proposed framework for collaboration used the indicators as a starting point, but has redefined them so that the collaboration of pairs of children can be determined in an activity with the surface bot. This section describes the set of indicators and definitions used in this research.

Action

Action refers to an effort made by children within the activity. An action is "participation of an individual, irrespective of whether this action is in any way coordinated with the efforts of other group members [15] ." In this activity, actions were defined as the possible interactions with the robot, i.e. operating the reward interface. This could be adjusting the slider, or sending the feedback. This indicator of participation is of interest, because it makes it possible to compare whether actions go together with the other participation indicator: interaction, and the indicators of the perspective taking and social regulation elements.

Interaction

Interaction refers to "interacting with, prompting and responding to the contributions of others [15]." Contribution can be a communicative initiation or action in the environment. Does a child react to a comment or action (regarding the activity) of another child? It is a useful indicator, as it indicates participation of both children. E.g.:

Child A: communicates positive feedback without consultation.
Child B: "yes, that was right!"

Action and interaction are the two indicators that belong to the "participation" class. The level of participation of children during the experiment can be estimated based on this level of action and interaction.

Adaptive Responsiveness

There is adaptive responsiveness when contributions or prompts of others are adapted and incorporated. It is interpreted in this research as a child accepting or adapting another child's point of view. It differs from interaction, because it deals with the content of the interaction. A reaction regarding the activity is always an interaction, but only counts as adaptive responsiveness when children considered and accepted, or adapted the contribution of the other. An example is:

Child A: "The robot must get item x instead of item y."
Child B: "Yes, let's go for item x."

Audience Awareness

The second indicator of perspective taking is audience awareness, which refers to "awareness of how to adapt behavior to increase suitability for others Hesse et al. [15]." An example given by Hesse et al. [15]:

"Imagine two problem solvers who are placed on different sides of a transparent screen. For a particular object on the left side from a problem solver's point of view, low audience awareness would be exhibited by referring to the object as being 'on the left side'. In contrast, higher audience awareness would be exemplified by referring to the object as being 'on the right side' or even 'on your right side'."

In this study, a situation was identified with audience awareness when a child shared information that was not available to the other. For example, one child can see the surface bot's screen and sees what action it wants to do. The child shares this with the other child, who then uses the reward interface to send feedback. It shows the child's awareness of the perspective of the other, and shares the information accordingly.

Negotiation

Negotiation is understood as an attempt to reach a common understanding, achieving a solution, or reaching a compromise. A dialog that could occur between children participating in an activity with the prototype is:

Child A: "Item x is super wrong", and sets slider to absolute negative.
Child B: "No, it is a bit wrong, but not super wrong."
Child A: Adjusts the feedback slider in accordance to feedback of Child B.

Self-evaluation

Comments on own performance in terms of appropriateness or adequacy are considered as self-evaluation. It is the recognition of own strengths and weaknesses. A dialog that would indicate that a child evaluated their own action, would be:

"I was too late, now the robot wears the wrong jacket."

Transactive Memory

There is transactive memory when one of the children comments on performance of the other in terms of appropriateness or adequacy. It relates to the recognition of strengths and weaknesses of others.An example of a situation where one child has evaluated the behavior of the other child, and then makes that clear:

"Let me operate the tablet, you were not fast enough!"

In the first study, children made clear several times whether the robot had received the feedback. If a “thumbs up” appeared on the character display, children made the remark: ”Yes, it has received the feedback!” In this study, it was only considered as a transactive memory if an opinion was also shared as to whether the feedback provided was correct or incorrect.

Responsibility initiative

Responsibility initiative is the eighth indicator and concerns “assuming responsibility for ensuring parts of task are completed by the group.” An indication is the use of first-person plural in communication regarding the activity, for example:

“We should let the bear know that it is the wrong item!”

In this study, taking responsibility was also considered when one child encouraged the other to take action or share information. For example, a situation that occurred in the second study:

Child A: “Item x is OK, right?”
Child B: “Yes, send the feedback!”

8.6.2. Part two: identifying the manner of collaboration

The first part of the framework consists of the indicators to determine the level of collaboration between children. The second part means to identify what the outcome is, or what manner of collaboration is established, from the collaboration between children in an activity with the prototype. Children can reach agreements, compromises or ideas on how to address the activity at hand through collaboration. Based on the observations of the first study, a number of collaborative outcomes were identified that were expected to occur during the activity.

Presence of roles

A division of roles can arise from collaboration. It happens when children divide responsibilities. It creates a state of dependency on another other for successfully completing the activity. Role division relates to the “resource management” indicator of the task regulation component which is described by Hesse et al. [15] as “managing of resources or people to complete a task.” In this research, maintaining of a role division is a desired outcome of the activity. Since it means, children discovering that relying upon another through cooperation results in a more effective way of engaging in the activity. E.g.: One child operating the tablet, the other child is tracking the robot and provides updates on the robot’s actions.

Shared planning

A shared planning is established when the problem at hand is analyzed and a mutual agreement is reached on (part of) the solution. A case of shared planning in the context of the activity with the prototype would be: when both children agree on a item of clothing the bear should wear. It is an agreement between children about how an action of the robot should be judged. An example is:

Child A: "The bear should get this jacket... and then go the hallway."
Child B: "No, it should get the sweater first then go to the hallway."
Child A: "OK, jacket, sweater and then it should get these shoes at the hallway."

Building shared knowledge

It refers to a child that gets an (1) understanding of another child's opinion or preferences, or (2) an understanding of (a part of the) activity. It applies when a child provides information, or shares his/her preferences or opinion, in response to a question or statement of the other child. An example of shared knowledge is:

Child A: "What items do you think the bear should get?"
Child B: "The jacket and the blue jeans!"

Or when they divided roles:

Child A: "What is the item that the robot displays?"
Child B: "It is the red jacket."

Taking turns

The result of the first study showed that children occasionally took turns in operating the tablet, or even played the game individually. Taking turns can only occur after establishing an agreement and is therefore a result of collaboration. However, it is not necessarily an ideal outcome. A scenario that could occur, is that one child temporarily does not actively participate in the activity and the other does everything. It means they did not discover the benefit of cooperation, but made a kind of compromise to both be in full control of the activity for a while. On the other hand, taking turns can occur while maintaining different roles. For example, one operates the tablet, while the other follows the robot from location to location. After a while, they might decide to switch roles. This would be a coordinated effort that maintains the situation where children depend on each another, and both actively participate.

8.7. Measurements

Results were obtained from the 9 pairs of children from (1) the pretest form, (2) observations of the recordings using the evaluation framework and (3) logged data from interactions with the reward interface.

The pre-test was made individually by each participant before the start of the activity. It provided an overview of the preferred item of each state element. By comparing the two tests made, a percentage of agreement can be calculated which expressed the extent to which the children's preferences overlap. It was meant as support information to be able to explain collaborative behavior observed during the session with the surface bot.

The recordings were used for the annotation using the proposed framework for evaluating collaboration. For the evaluation, the recordings were considered from the point the robot starts taking action until the last iteration, where the robot goes "outside". Inspired by the method of evaluation used by Huskens et al. [16], this period is divided into 30 seconds intervals. Each interval was observed and for every indicator of the framework, it was annotated if it was present (+), or not (-). This resulted in a table of features with binary measurements. From this an average score can be determined for each pair of children for the three classes of collaboration, and collaboration in general. The equations used for calculating the scores of the three classes of collaboration and an overall collaboration score, are described in Appendix C.2. These scores can be used to evaluate the influence of pace, and the prototype's effectiveness of encouraging collaboration compared to the first study's prototype.

To validate the framework and method, one recording was annotated by two researchers. The annotation of both researchers were almost the same except for a few minor differences. The differences were explained by the interpretation of some indicators. The definitions of each indicator was refined and example dialogs were added. This resulted in the described indicators in the previous section.

The logged data from interactions with the reward interface provides information regarding the time of a robot's action and the time and value of the participants feedback (see Table 8.2). This information can be used to obtain the frequency and ratio between negative and positive feedback per pair. This provides a measurement of the engagement of children, and can be used to evaluate the reward behavior of children compared tot the first study.

At the end of the test, children were asked a number of questions regarding their experience of participating in the activity with the surface bot. The questions are meant to provide an insight in their opinion of the activity, their role as tutor and the performance of the surface bot. Furthermore, they were asked for ideas or comments to further improve the surface bot and the activity.

start time action (ms)	end time action (ms)	action	time of feedback (ms)	feedback value
1151905	1158960	jacket	1155589	1.0

Table 8.2.: The format of the logged data with example data.

Group number	Condition	Agreement	Collaboration	Participation	Perspective Taking	Social Regulation
8	A	60	54	95	61	30
9	A	40	47	90	44	27
6	B	40	47	97	53	20
5	A	40	46	95	53	18
3	A	40	43	97	23	27
2	A	0	38	95	21	18
4	B	40	38	96	25	15
1	A	40	25	78	11	6
7	A	60	14	54	0	2

Table 8.3.: The measures (%) of collaboration for the 9 sessions of the second study, in descending order of the "collaboration" score. It also shows the level of agreement (%) and test condition (A = high pace, B = low pace) for each pair of children.

8.8. Results

The results of the second study were examined on the degree of collaboration between children. The results were obtained from annotating 160 intervals from the recordings of the 9 sessions with an average of 17.78 ± 3.55 intervals per session. Appendix C.4 contains the cumulative measurements of all 8 indicators, expressed in absolute and relative values for each pair. Table 8.3 contains the relative values per class, based on cumulative values of the corresponding indicators. Collaboration is also measured on the basis of one average based on the scores of each indicator.

Group number	Collaboration	Participation	Perspective Taking	Social Regulation
1	38	100	25	13
3	35	93	14	16

Table 8.4.: The measures (%) of collaboration for the sessions of 2 pairs of children that participated in the first study.

To answer the first sub-question, two sessions of the first study were annotated based on the recordings. Table 8.4 shows the results. The first study's annotations served as comparative material and a baseline for the degree of collaboration seen in this study with the second version of the prototype. The full results of the first study can be found in Appendix C.3. The average measure was calculated over the relative measures of each pair for a comparison between the first and the second study. The result can be seen in Figure 8.1. It shows that, apart from participation, all the other measures have increased for the second study. Figure 8.2 shows some differences between the indicator scores of the first and second study. The second study scores better on each indicator of the social regulation class, and a notable difference can be seen for "audience awareness." In terms of participation, both the first and the second study scored high, with the first study scoring slightly better on both indicators.

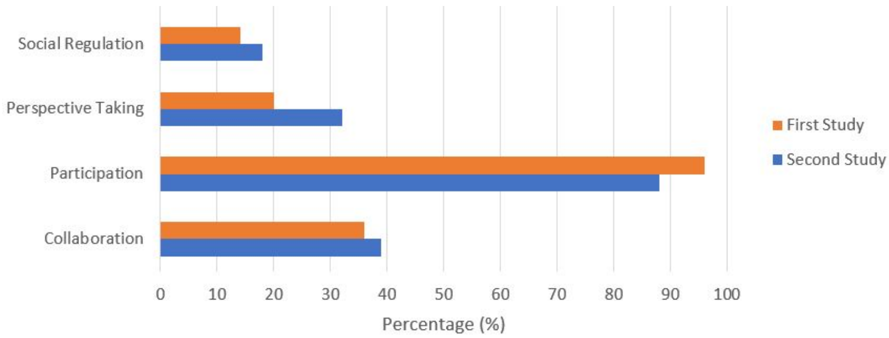

Figure 8.1.: A comparison between the mean scores of collaboration between the first and second study.

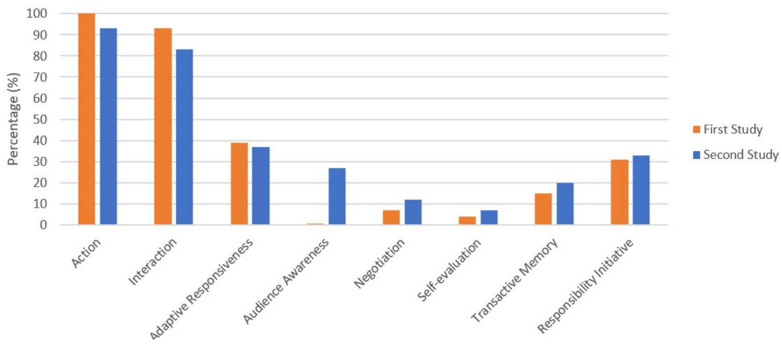

Figure 8.2.: Comparison of the average measure (%) of each indicator for the first and second study.

The recordings have also been evaluated on a set of categories of collaborative behavior. Each category is a possible outcome of a coordinated effort by the children. For each category it was indicated at the end of the recording whether or not it occurred. From this, an average was calculated for the first and the second study based on the cumulative results per category. Figure 8.3 shows how the averages compare per category. The first study's results were based on two sessions which only displayed turn taking. In the second study a division of roles was observed for 4 of the 9 pairs (against 1 out of 5 in the first study). It should be noted that the pair in the first study who established a division of roles, were not recorded (due to failing recording equipment) and therefore not included in this comparison. A few suggestions were made by children in the first study for a division of roles and planning, but the other child did not respond. It is therefore not annotated, since the collaborative behavior was not established. In the second study, the roles were divided in the following way: one child stood by the tablet and gave the feedback, the other child communicated what the robot did. Figure 8.4 gives an impression of the children's positioning at the start of the activity and when roles were divided.

The second sub-question regarding the use of the reward interface was examined based on the logged data. Figure 8.5 shows the frequency of the different feedback values. The possible values were in a range from 0 (extremely negative feedback) to 1 (extremely positive feedback). The

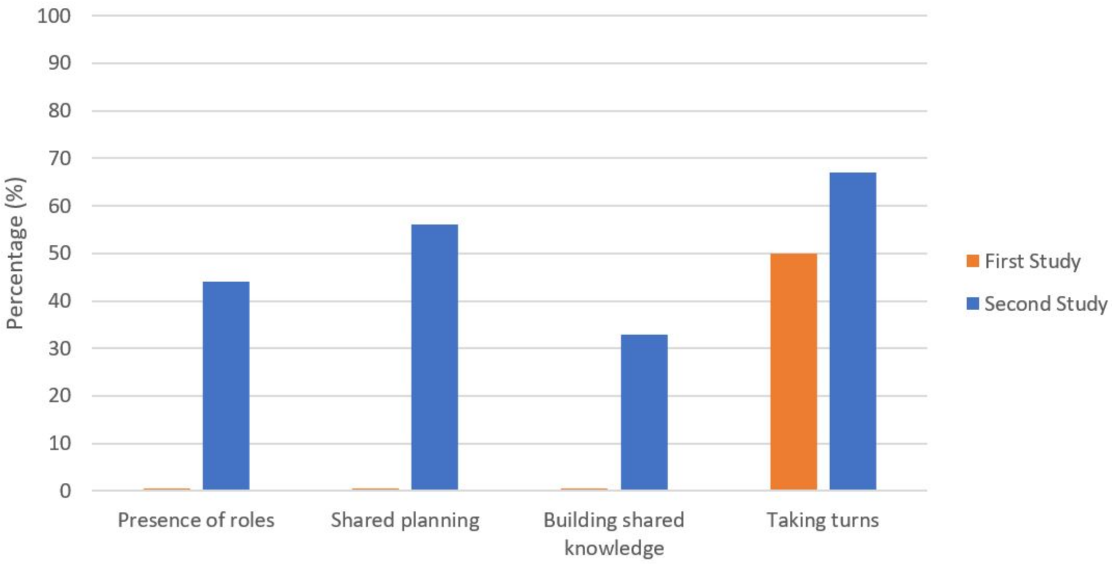

Figure 8.3.: Comparison of the first and second study on the presence of specific collaborative behaviors.

majority of the feedback had a value of the lowest 10% or the highest 10% of the possible values. Of the feedback given, 37.7% had a value between 0 and 0.1, and 41% had a value between 0.9 and 1.0.

The sessions of the second study were conducted with two conditions. For answering the third sub-question, the measure of collaboration and use of the reward interface were evaluated for each condition separately. Table 8.5 shows the average measures of collaboration for each condition. The results seem to indicate slightly better collaboration scores for condition B. However, only two sessions were performed in condition B, compared to 7 in condition A. The distribution was inadvertently caused by a technical defect in communication between the surface bot and the tablet of the tele-operator.

The pretest forms gave insight in the initial agreement among children regarding preferences. The level of agreement is shown in Table 8.3, and describes the percentage to which children's preferences match. Out of their five preferences, the children within each group matched 0 to 3 items of clothing. This shows that each group had several items of clothing that could be discussed. For each group this meant that coordination or negotiation would be necessary, since they differed in opinion or at least had a different preference. Looking at the overall collaboration score, the agreement does not seem to determine the degree of collaboration. There is no indication that children who had little agreement beforehand showed more collaboration, neither in the overall collaboration score, nor in the "negotiation" indicator. Lastly, the questions asked did not yield any substantial results. Children had little input for improvements or what they would like to change.

Figure 8.4.: The first image (left) shows an impression of children at the start of the activity in the second study. The second image (right) shows that they have established a division of roles. The child on the left can see the surface bot's tablet, and informs the other child what action the robot wants to do.

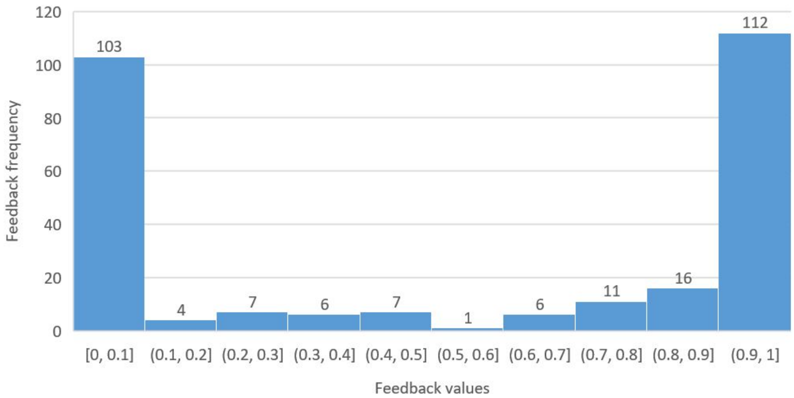

Figure 8.5.: The feedback behavior of the participants.

8.9. Discussion

In order to formulate an answer to the first sub-question, the results were compared in different ways on the level children collaborated. Figure 8.1 and 8.2 indicate a slightly higher degree of collaboration for most indicators in the second study, compared to the first study. A notable difference is the absence of "audience awareness" in the first study, compared to 27% in the second study. Audience awareness was defined in this study as sharing information that is not directly accessible or known to the recipient. A situation that was considered audience awareness was when only one child had a view of the robot's screen (for example, because both are on a different side of the table) and then shared information about the robot's action or what feedback should be communicated. In four sessions of the second study, a division of roles was seen with an average relative measure of 55% for audience awareness compared to 5% for the other five sessions where no division of roles was observed. The measurement

Condition	Collaboration	Participation	Perspective Taking	Social Regulation
A	38	86	30	18
B	42	97	39	17

Table 8.5.: The average measures (%) of collaboration for the pairs of children that participated under condition A (high pace) against the pairs of children that participated under condition B (low pace).

differences cannot be statistically substantiated, due to the small sample size. The results of the first study were based on just two recordings (out of five sessions). This makes it not entirely reliable to compare the first and second version of the prototype. The session of the first study where cooperation and a division of roles was established was not even included. The measurements from the second study were strongly influenced by pair 1 and 7 who scored low on collaboration. They were less fanatic in their role of tutor and communicated less regarding the robot's actions in comparison to the other groups. The explanation for this could be embarrassment or a lesser degree of interest in the game. Remarkable is that all four children were 9 years old. With regard to engagement, the current implementation of the concept seems to be particularly suitable for children aged 5 to 7 years. Figure 8.3 shows that in 4 of 9 sessions, a form of role division was established. This is a clear increase, compared to 1 in the first study. It indicates that prototype 2.0 encourages collaboration more effectively, as it elicits more forms of collaborative behavior. This can be explained by the fact that the activity has become more difficult to tackle efficiently without working together. This is probably due to the fixed tablet as it reduced the efficiency of the behavior shown in the first study: both children following the robot from location to location, while providing feedback. The division of roles that was observed is a form of collaboration where children are dependent on each other, a reinforcing factor of collaboration, as communication becomes necessary. The results of the pre-test (table 8.3) shows that children differed to a certain extent in their preferences, but this does not seem to predict a certain level of collaboration. In retrospect, it would have been better if children could have ticked more options instead of just the best choice per part of the bear. The current pre-test did not give insight into the children's dislikes: what items do they not want bear to wear. The current pre-test did not give insight whether children thought other items of clothing were also fine, besides their preferences. If this was the case, then there is still little need for negotiation.

With regard to the second sub-question, it can be said that the use of the slider remains unchanged. The influence of ambiguity, causing a difference of preferences and an incentive to negotiate, did not lead to a more sophisticated use of the slider, as the distribution of the feedback values shows in Figure 8.5. Children still tend to give extreme values. However, forms of negotiation were observed where children tried to determine the slider value together. In that sense, the slider still has added value. The nature of the task is subjective and can still lead to a set of clothing that is found to be generally good and the rest completely wrong. A recommendation for future work would be to experiment with a task for the robot, where its actions have a certain ambiguity but do not rely on children's preferences. A suggestion for future work is to improve or change the surface bot's environment and actions in a way that some actions are positive, but clearly sub-optimal (short term) and other (sequences of) actions contribute over time to a larger goal (long term). It could challenge children to form a plan and reward the robot in more gradations.

With regard to the third sub-question, no definitive conclusions can be drawn on the influence of the pace on the level of collaboration. An insightful and reliable comparison between the conditions was not possible, because of the small sample size of condition B. Nevertheless, at first sight the results showed only minor differences between the two conditions with regard to the level of collaboration, or collaborative behavior displayed by the children.

It can be concluded that in most cases, children were enthusiastic and actively engaged in their role as tutor, especially children between the ages of 5 and 7. In a number of cases a division of roles arose spontaneously. Negotiation took place only to a limited extent. A few attempts at negotiation did not lead to an agreement or response from the other child. Furthermore the changes made in prototype 2.0 seem to have contributed to a version that can encourage collaboration more effectively, because each category of collaboration occurred (more often) in comparison to the collaboration shown by the pairs of children using prototype 1.0. It is not possible to indicate which features have contributed most to the more extensive collaboration, since prototype 2.0 has made several major changes.

9. Discussion

In this research, two versions of a prototype with the surface bot were developed with the aim of encouraging collaboration between children. Two studies were conducted that aimed to get an insight into the extent collaboration is stimulated among children in an activity with this prototype. The two main questions addressed in this research were:

1. How can the capabilities of the surface bot be utilized to create an engaging activity that effectively encourages collaboration between primary school children?

2. How can the extent and manner of collaboration between primary school children be measured in order to evaluate the effectiveness of an activity with the surface bot?

This chapter discusses how the questions can be answered based on the results of the first and second study.

9.1. Research question 1: the prototype and the level of collaboration between children

A concept inspired by the learning-by-teaching paradigm was the basis for prototype 1.0. A story was used to introduce the activity with the surface bot and was meant to capture the children's imagination and motivate them to participate. The concept is based on a symmetrical structure as described by Dillenbourg [12] which means that the children had a shared goal in the activity: helping the surface bot in finding the appropriate set of clothing. Secondly, all participating children had the same level of knowledge of their role as tutor with the equal opportunities for possible actions in the activity. Prototype 1.0 only brought about a limited form of collaboration in the first study. Children tended to make individual decisions, and there was a strong preference among them for using the tablet. They were observed to take turns with only one child actively contributing to the learning of the robot. This corresponds to the characteristics of the developmental stage of children in the age of 3 to 7 years as described by Markopoulos and Bekker [22]. They state that children in this stage are self-centered and prefer parallel play. The majority of the children was engaged in the activity and feedback was provided throughout the activity. This indicated that they can handle their role as tutor, and remain motivated during the activity. It was concluded that just providing a symmetrical structure in an activity with the surface bot was not enough to encourage children to collaborate. Ensuring the symmetrical structure seems to allow collaboration, but it is not enough to elicit it. For prototype 2.0, changes were made aimed to further encourage spontaneous collaboration.

First, the children's task was made more challenging and ambiguous by adding new actions (and objects) to the activity. Ambiguity is described to be an incentive for collaboration [12], as it diverges the children's preferences and provides more opportunity to discuss and negotiate. Second, the tablet for communicating feedback was fixed to the table. It made it more challenging for one individual to simultaneously observe and interact with the robot. It was expected to be an incentive for collaboration by making agreements and maintaining a division of roles, as seen with one pair of children from the first study. Second, the role of the tele-operator was minimized to the movement of the surface bot, as prototype 2.0 comprised a reinforcement learning framework for autonomous action taking and improving its decision making based on the feedback of children. Prototype 2.0 seems to encourage collaboration more effectively compared to prototype 1.0, however it cannot be significantly proven. The second study with prototype 2.0 showed multiple pairs of children establishing cooperation through a division of roles. The division of roles was established spontaneously among the children with one of them operating the tablet while the other providing information about the robot's actions by mentioning the robot's action and/or voicing his opinion about the robot's action. The slightly higher collaborative score of children interacting with prototype 2.0, compared the score of children interacting with prototype 1.0, can partly be explained by this interplay. The division of roles caused children to share information regarding the activity unknown to the other which was marked as audience awareness, one of the indicators for the measurement of collaboration. The comparison with the first study is only indicative, since this baseline is based on just two recordings, which did not include the best case in terms of collaboration (with a presence of a role division). My recommendation for future work is to improve the baseline, in order to reliably compare new versions of the concept on their effectiveness of encouraging collaboration. In addition, there was a lot of age difference between the pairs of children in both the first and the second study. An activity with a robotic character and a story can appeal to younger children, in the age 5-7, more than older children, because children aged 3 to 7 are in a developmental phase where they enjoy fantasy [22]. The level of collaboration is based on participation of the children, therefore age could be decisive. In order to reliably address the difference or improvement of new versions of the prototype, it would be better to keep the age between the (pairs of) children as small as possible.

Both studies showed that children are engaged in the activity with the surface bot, they understood their role as tutor and seemed to perceive the improved performance of the surface bot and the effect of their feedback. Children took responsibility (first aspect of learning-by-teaching) by providing regular feedback and adhering to the outlined scenario. The studies showed that children are able to establish spontaneous manners of collaboration. Sporadic self-evaluation and transactive memory was seen, which is an indication that children are reflecting (second aspect of learning-by-teaching) on the actions of others, and themselves. With the more ambiguous task and a fixed tablet in prototype 2.0, children made agreements that exceeded taking turns in the use of the tablet. They occasionally established a shared planning. This might indicate that children are capable of structuring information (third aspect of learning-by-teaching) in order to develop a strategy. They are able to evaluate the robot's options and reached agreements on the "best" or "preferred" next move of the surface bot. The aspects of learning-by-teaching can be seen in the evaluation of the collaboration between children in the second study. Reflecting and structuring are limited, but it is an indication that learning-by-teaching is a suitable concept for an activity with the surface bot. Although children showed enthusiasm and established collaboration, it was outside the scope of this research to determine if children learned from participating in the activity, or developed their collaboration skills. The extent to which the prototype fosters children's collaborative skills was not investigated in this research. A recommendation for further research is to study the long term effect of participating in activities with the surface bot (that are aimed to encourage collaboration) on the collaborative skills of primary school children. Do children who regularly participate in an activity with the surface bot show improved collaboration compared to children who did not use the surface bot? And if children improved their collaborative skills through the activities with the surface bot, are the knowledge and skills transfered to other collaborative activities without the surface bot? Longitudinal studies could reveal the educative contribution of activities with the surface bot to primary school children when introduced in the classroom.

9.2. Research question 2: the framework for evaluating collaboration

In the first study, collaboration was evaluated based on observations of the sessions of pairs of children with prototype 1.0. The second study explored the degree of collaboration based on a quantifiable method. A framework was proposed based on the indicators described by Hesse et al. [15] providing a structured way to compare the level of collaboration between the first and second study, and between the two conditions used in the second study. The annotation of the recordings was done by one researcher. To validate the usability of the framework, one recording was first checked by another researcher. Comparing the results showed small differences, and based on that the framework was further refined. Using this framework for annotation requires time, since there were situations that were difficult to assess. Fortunately, the recordings made it possible to watch parts multiple times. The use of the current framework requires practice and preparation, but provided sufficiently reliable results for this study. The validity of the framework remains inconclusive, due to the small sample size.

Collaboration was determined by calculating the average scores for each indicator. Furthermore, one measure of collaboration was used that was calculated as the average of these indicator scores. Each indicator contributed equally to this score. However, it can be argued that some indicators are more indicative or more important for measuring collaboration than others. For example, "action" only says something about the participation of children while "negotiation" is a profound expression of collaboration that requires communication and a certain willingness of both children to listen to each other. In future research, a method could be developed to achieve a more enhanced score of collaboration. The indicators could, for example, be weighted according to their importance or contribution to measuring collaboration. This study has not looked at the significance of the results. For future work, it would be useful to develop a procedure that allows statistical analysis of the results obtained by annotation using the framework. A statistical analysis was not applicable for this research due to the small number of participants.

9.3. Conclusion

Although a statistical analysis could not be done, it can be concluded that the prototype 2.0 encourages collaboration between children to a certain extent. Prototype 2.0 showed it can encourage cooperation, in the form of a division of roles. Both prototypes stimulated communication about the robot's actions, and who controls the tablet. Children were engaged throughout the activity as they kept providing feedback and seemed to adhere, in most cases, to the scenario provided by the preceding story. They took responsibility and occasionally showed reflecting and structuring behavior. These are indications that a concept based on the learning-by-teaching paradigm can be a basis for activities that are able to encourage children to start collaborating. Future work can focus on further improving the prototype to more effectively stimulate collaboration, and explore the effect on the development of collaborative skills as a result of engaging with the prototype over a longer period. This research has led to a prototype with which the technical capacities of the surface bot are utilized. It is a prototype that actively involves the children in a role as tutor, and it provides a basis for activities that encourage collaboration more extensively and effectively. Based on the concept, other activities with the surface bot can be developed, to be integrated into the classroom and effectively encourage collaboration: an important skill that primary school children must learn, and the importance of which they should discover.

10. Future Work

This chapter provides an overview of suggestions for future work. This research brought forth a prototype, an activity with the surface bot, with which collaboration between children can be encouraged. The results from the two studies conducted in this research offer starting points for further research.

10.1. Recommendations for future research

As mentioned in the discussion (Chapter 9), this research did not examine the effect of interacting with the prototype on the collaborative skills of the primary school children. The first recommendation is to conduct a longitudinal study to examine how participating in the activity with the surface bot affects the collaborative skills, and the willingness to collaborate, over time. As future work, it can be investigated whether collaboration becomes more intensive after participating multiple times, and whether they incorporate acquired skills and experiences in the activity in other situations where collaboration is important.

Secondly, the discussion addressed the status of the framework for evaluating collaboration. Collaboration was assessed on average scores based on the framework's indicators. Each indicator was regarded as equally important. However, some indicators can be regarded as outcomes or displays of collaboration, while others can be described as prerequisites that enable collaboration. Further research may focus on determining a single score of collaboration, with the indicators weighted by importance. Furthermore, it is important that there is a procedure for statistical analysis based on this framework, and a detailed method for annotating. This would make longitudinal studies possible and also enables comparisons with other applications. In the assessment of collaboration it would also be of interest to see how children interact and collaborate when the facilitator is not present directly next to the activity. The fact that during each session one or two facilitators were present in the room might have affected the engagement of the children, and how they acted. Jamet et al. [17] point out that when a child is alone with a robot their status of teacher is strengthened. The actions and responses of a child cannot be judged and will therefore be more natural and spontaneous. A one-way-mirror or video equipment could be used to observe the children and control the surface bot in further research.

The concept of this research, and the created prototype, are inspired by the learning-by-teaching paradigm. Children act as tutor, and are responsible for teaching the surface bot. The role of tutor is limited to providing feedback in response to actions of the surface bot. The solution or answers to the problem that the surface bot faced had a subjective nature. Children were expected to formulate their preferences and collaborate to provide coherent feedback: a requirement for the robot's learning. In the current prototype, children "teach" and can practice their social and collaboration skills. The activity's task does not allow "learning". It was outside the scope of this research, but it recommended to explore how children's learning material can be incorporated into the concept while maintaining an activity design that encourages collaboration.

Another direction of interest, is utilizing the audio and display of the surface bot for social and emotional behavior. The current prototype is still an early prototype, as it cannot function completely autonomous during the activity. Children are known to react strongly to social behavior of a robot [30]. It can be used to let the surface bot tell the story which makes the surface bot more credible as a character and its problem less trivial to the children. Furthermore, social behavior can be used to regulate the activity. Experiments can be done, in how this behavior can elicit collaboration between children.

10.2. Suggested improvements of the prototype

The prototype used in this research enabled children to provide feedback via a tablet interface as a reward or punishment in reaction to the actions of the surface bot. Children showed an understanding of the interface and the slider's purpose and effect. However, the first study showed that they mainly provided unilateral feedback with only completely negative or completely positive values. After the first study, it was argued that this was due to their unanimity on what the best "solution" is for the robot's problem. In the second study, the "answers" were made more ambiguous in an attempt to diverge the children's preferences in the expectation that it would be an incentive to negotiate and reach consensuses in the form of intermediate feedback. However, the results in second study were similar to the first study. Although not much negotiation took place, when it did the slider (and the value of the feedback) was a topic of discussion. It could therefore still contribute to collaboration between children, but the activity's design should focus more on utilizing the slider. Therefore, my recommendation is to adjust the task of the robot and the actions that the robot can take in a way that children can develop a short-long term strategy.

A second recommendation is to experiment with different interfaces for children to communicate to the robot, and fulfill their role as tutor. Providing feedback is a reaction to the thoughts or performed actions of the robot. To further develop the surface bot as a teachable agent, children should be enabled to teach and do more than conveying their opinion afterwards. My suggestion is to explore possible interactions to convey knowledge or provide elaborate feedback. These interactions should be assessed on their usage, and influence on collaboration between primary school children. Examples of alternative channels of communication with the surface bot are demonstration, explaining or guidance. Demonstration [10] might be interesting as it directly challenges skills of the children. Communicating explanations could challenge the knowledge and ability to structure information of children [3]. Guidance can be explored for more accurate coordination of the robot's movement [26].

A third recommendation would be improving the reinforcement learning framework. The framework needs to change when experimenting with other channels of communication for the children. The Q-learning can only incorporate feedback (positive and negative values). Guidance would require a change in the action-selection process of the algorithm. The framework could also be used to test other forms of transparency behavior. The current prototype was only transparent about which action it wanted to do (the thought cloud). My suggestion is to experiment with prototypes that are capable of explaining why an action was selected or give insight in the robot's confidence in a selected action. Think of the gaze behavior [26] to indicate which actions are being considered, or speech [3] to give an argument to support the selected action. It could stimulate children to reflect on the most logical moves that the robot could have made, and will have to do.

The fourth recommendation is to further explore the effect of pace on the robot's learning, and collaboration between children. The second study did not provide a clear insight into the effect of the surface bot's pace on the collaboration between children. I still expect that the robot's pace has an effect on the robot's learning and collaboration, since it directly affects the time children have for discussing and negotiating decisions. Further studies could pay attention to the alertness of children on the robot's pace. Do children notice when they provide feedback too late, and do they adjust their interactions accordingly? If children do not notice it, technical or design changes are needed to the prototype in order to enable the robot to learn successfully. For example, distributing the rewards over the most recent actions or by having the robot wait for feedback. A second suggestion is to explore how social behavior of the robot can be used to stimulate children to provide (timely) rewards. In addition to the effect on the learning progress of the robot, attention must also be paid how changes in the prototype influence the interaction of the children with the robot, and the effect it has on collaboration between the children.

Bibliography

[1] J.A. Bellanca. *21st Century Skills: Rethinking How Students Learn.* Leading Edge. Solution Tree Press, 2010. ISBN: 9781935542377. URL: https://books.google.nl/books?id=aGYXBwAAQBAJ.

[2] Steve Benford et al. 'Designing Storytelling Technologies to Encouraging Collaboration Between Young Children'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* CHI '00. The Hague, The Netherlands: ACM, 2000, pp. 556–563. ISBN: 1-58113-216-6. DOI: 10.1145/332040.332502. URL: http://doi.acm.org/10.1145/332040.332502.

[3] Gautam Biswas et al. 'Learning By Teaching: A New Agent Paradigm For Educational Software'. In: *Applied Artificial Intelligence* 19 (2005), pp. 363–392.

[4] G Biswas et al. 'Smart machines in education'. In: *Technology support for complex problem solving: From SAD Environments to AI, K. Forbus and P. Feltowich, Eds. AAAIMIT Press* (2001), pp. 61–77.

[5] Steven T. Bossert. 'Cooperative Activities in the Classroom'. In: *Review of Research in Education* 15 (1988), pp. 225–250. ISSN: 0091732X, 19351038. URL: http://www.jstor.org/stable/1167365.

[6] Esther Care and Patrick Griffin. 'An approach to assessment of collaborative problem solving. Special issue: Assessment in computer supported collaborative learning'. In: *Research and Practice in Technology Enhanced Learning* 9 (Jan. 2014), pp. 367–388.

[7] Alejandro Catala, Mariët Theune and Dirk Heylen. 'My Tablet Is Moving Around, Can I Touch It?' In: *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction.* TEI '17. Yokohama, Japan: ACM, 2017, pp. 495–502. ISBN: 978-1-4503-4676-4. DOI: 10.1145/3024969.3025082. URL: http://doi.acm.org/10.1145/3024969.3025082.

[8] Alejandro Catala et al. 'Exploring Children's Use of a Remotely Controlled Surfacebot Character for Storytelling'. English. In: *Intelligent Technologies for Interactive Entertainment.* Vol. 215. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Germany: Springer, June 2017. ISBN: 978-3-319-73061-5.

[9] Alejandro Catala et al. 'Storytelling as a creative activity in the classroom'. English. In: *C&C '17: Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition 2017.* United States: Association for Computing Machinery (ACM), June 2017, pp. 237–242. ISBN: 978-1-4503-4403-6. DOI: 10.1145/3059454.3078857.

[10] Shruti Chandra, Pierre Dillenbourg and Ana Paiva. 'Classification of Children's Handwriting Errors for the Design of an Educational Co-writer Robotic Peer'. In: *Proceedings of the 2017 Conference on Interaction Design and Children*. IDC '17. Stanford, California, USA: ACM, 2017, pp. 215–225. ISBN: 978-1-4503-4921-5. DOI: 10.1145/3078072.3079750. URL: http://doi.acm.org/10.1145/3078072.3079750.

[11] Catherine C. Chase et al. 'Teachable Agents and the Protégé Effect: Increasing the Effort Towards Learning'. In: *Journal of Science Education and Technology* 18.4 (Aug. 2009), pp. 334–352. ISSN: 1573-1839. DOI: 10.1007/s10956-009-9180-4. URL: https://doi.org/10.1007/s10956-009-9180-4.

[12] Pierre Dillenbourg. 'What do you mean by collaborative learning?' In: *P. Dillenbourg (Ed) Collaborative-learning: Cognitive and Computational Approaches* (1999), pp. 1–19.

[13] Pierre Dillenbourg et al. 'The evolution of research on collaborative learning'. In: *Learning in Humans and Machine: Towards an interdisciplinary learning science*. Ed. by Spada et al. Elsevier, Oxford, 1995, pp. 189–211.

[14] Libby Hanna, Kirsten Risden and Kirsten Alexander. 'Guidelines for Usability Testing with Children'. In: *interactions* 4.5 (Sept. 1997), pp. 9–14. ISSN: 1072-5520. DOI: 10.1145/264044.264045. URL: http://doi.acm.org/10.1145/264044.264045.

[15] Friedrich Hesse et al. 'A Framework for Teachable Collaborative Problem Solving Skills'. In: *Assessment and Teaching of 21st Century Skills: Methods and Approach*. Ed. by Patrick Griffin and Esther Care. Dordrecht: Springer Netherlands, 2015, pp. 37–56. ISBN: 978-94-017-9395-7. DOI: 10.1007/978-94-017-9395-7 2. URL: https://doi.org/10.1007/978-94-017-9395-72.

[16] Bibi Huskens et al. 'Improving Collaborative Play Between Children with Autism Spectrum Disorders and Their Siblings: The Effectiveness of a Robot-Mediated Intervention Based on Lego Therapy'. In: *Journal of autism and developmental disorders* 45 (Nov. 2014). DOI: 10.1007/s10803-014-2326-0.

[17] Frank Jamet et al. 'Learning by Teaching with Humanoid Robot: A New Powerful Experimental Tool to Improve Children's Learning Ability'. In: *Journal of Robotics* 2018 (Nov. 2018). DOI: 10.1155/2018/4578762.

[18] Leslie Pack Kaelbling, Michael L. Littman and Andrew P. Moore. 'Reinforcement Learning: A Survey'. In: *Journal of Artificial Intelligence Research* 4 (1996), pp. 237–285. URL: http://people.csail.mit.edu/lpk/papers/rl-survey.ps.

[19] Emily R Lai. 'Collaboration: A literature review'. In: *Pearson Research Report, 2* (2011).

[20] Markus Lindberg et al. 'Does a Robot Tutee Increase Children's Engagement in a Learning-by-Teaching Situation?' In: *Intelligent Virtual Agents*. Ed. by Jonas Beskow et al. Cham: Springer International Publishing, 2017, pp. 243–246. ISBN: 978-3-319-67401-8.

[21] Sharon A. Lynch and Cynthia G. Simpson. 'Social Skills: Laying the Foundation for Success'. In: *Dimensions of Early Childhood* (2010), 38(2) 3–12.

[22] P. Markopoulos and M. Bekker. 'Interaction design and children'. In: *Interacting with Computers* 15.2 (Apr. 2003), pp. 141–149. ISSN: 0953-5438. DOI: 10.1016/S0953-5438(03)00004-3.

[23] Jeremy Roschelle and Stephanie Teasley. 'The Construction of Shared Knowledge in Collaborative Problem Solving'. In: *Computer Supported Collaborative Learning* (Jan. 1995). DOI: 10.1007/978-3-642-85098-15.

[24] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 3rd. Upper Saddle River, NJ, USA: Prentice Hall Press, 2009. ISBN: 0136042597, 9780136042594.

[25] R.E. Slavin. *Cooperative learning: theory, research, and practice*. Prentice Hall, 1990. ISBN: 9780131725942. URL: https://books.google.nl/books?id=jz0mAQAAIAAJ.

[26] Andrea Lockerd Thomaz and Cynthia Breazeal. 'Teachable robots: Understanding human teaching behavior to build more effective robot learners'. In: *Artif. Intell.* 172 (2008), pp. 716–737.

[27] Michel Tokic and Günther Palm. 'Value-difference based exploration: adaptive control between epsilon-greedy and softmax'. In: *Annual Conference on Artificial Intelligence*. Springer. 2011, pp. 335–346.

[28] Gijs Verhoeven, Alejandro Catala and Mariet Theune. 'Designing a Playful Robot Application for Second Language Learning'. In: *Interactivity, Game Creation, Design, Learning, and Innovation*. Ed. by Anthony L. Brooks, Eva Brooks and Cristina Sylla. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Springer, Jan. 2019, pp. 385–394.

[29] David Wood and Claire O'Malley. 'Collaborative Learning between Peers'. In: *Educational Psychology in Practice* 11.4 (1996), pp. 4–9. DOI: 10.1080/0266736960110402. eprint: https://doi.org/10.1080/0266736960110402. URL: https://doi.org/10.1080/0266736960110402.

[30] Cristina Zaga et al. 'Exploring child-robot engagement in a collaborative task'. In: *Proceedings of the Child-Robot Interaction Workshop: Social Bonding, Learning and Ethics*. eemcs-eprint-25527. Instituto de Engenharia de Sistemas e Computadores, Investigacao e Desenvolvimento em Lisboa (INESC-ID), June 2014.

Appendices

A. Prototype 1.0

This appendix gives an overview of: the locations, the items of clothing and the action sequence used in prototype 1.0.

A.1. Locations

The images used for the locations that were part of prototype 1.0. The same (images of) locations were used in the subsequent prototype.

Hallway

Closet

Coat rack

Table

The last location that the surface bot thinks of and executes is: “going outside”. It marks the end of an iteration, and the surface bot returns to the “hallway” location to start again. The image used:

Going outside

A.2. Items of clothing

This section shows the images designed for the items of clothing used in prototype 1.0. Each item belongs to a certain part of the bear. For each part, the bear can only wear one item at a time. The items are listed per part of the bear.

Head:

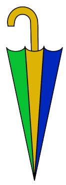

Umbrella

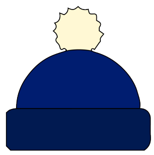

Cap

Sunglasses

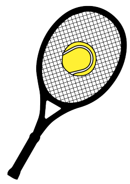

Tennis racket

Outer body:

Winter jacket

Sports jacket

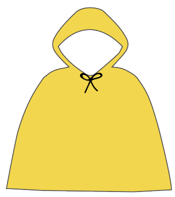

Raincoat

Inner body:

T-shirt

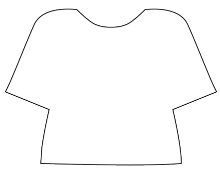

Sport shirt

Sweater

Legs:

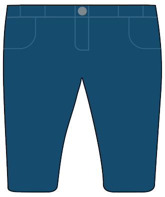

Trousers

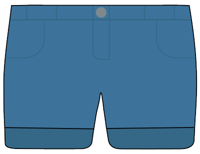

Shorts

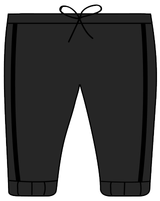

Sport trousers

Shoes:

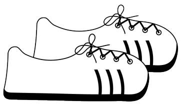

Sport shoes

Flip flops

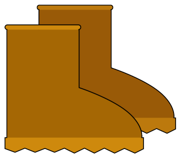

Winter shoes

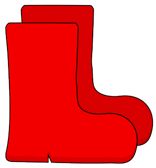

Boots

A.3. The sequence of actions

The detailed scenario for this prototype was about winter weather. A sequence of actions was worked out in which the robot converged in three iterations to the set of winter clothing. The table below shows the sequence of actions per iteration.

Action	Iteration 1	Iteration 2	Iteration 3
1	Hallway	Hallway	Hallway
2	Flip flops	Boots	Winter shoes
3	Sport shoes	Winter shoes	Table
4	Closet	Coat rack	Cap
5	Sport shirt	Rain coat	Closet
6	T-shirt	Winter jacket	Trousers
7	Sweater	Closet	Coat rack
8	Hallway	Trousers	Winter jacket
9	Boots	Coat rack	Closet
10	Winter shoes	Closet	Sweater
11	Closet	T-shirt	Going outside
12	Shorts	Sweater	
13	Sport trousers	Going outside	
14	Trousers		
15	Hallway		
16	Coat rack		
17	Sports jacket		
18	Rain coat		
19	Winter jacket		
20	Closet		
21	Table		
22	Sunglasses		
23	Umbrella		
24	Cap		
25	Going outside		

B. Prototype 2.0

B.1. Items of clothing

The items of clothing used in prototype 2.0. The items are listed per part of the bear. At each moment, the bear could only wear one item per part.

Head:

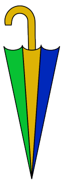

Umbrella

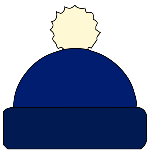

Cap

Sunglasses

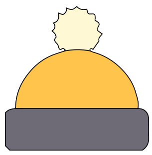

Cap

Outer body:

Winter jacket

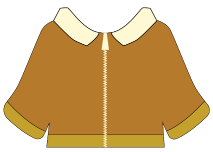

Winter jacket

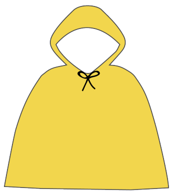

Raincoat

Inner body:

T-shirt

Sweater

Sweater

Legs:

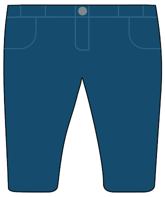

Trousers

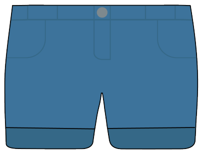

Shorts

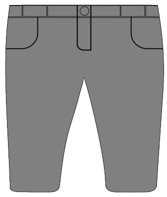

Trousers

Shoes:

Winter shoes

Flip flops

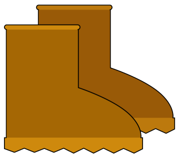

Winter shoes

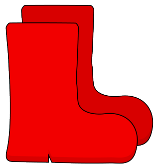

Boots

C. Annotation results of the first and second study

This appendix contains the results obtained from the annotation of the recordings of the first and second study using the framework for evaluation of collaboration, described in Chapter 8, section 8.6.

C.1. Pre-test form

The pre-test form used in the second study. Children completed the pre-test individually. They were given a marker and were asked to tick one item of clothing per line.

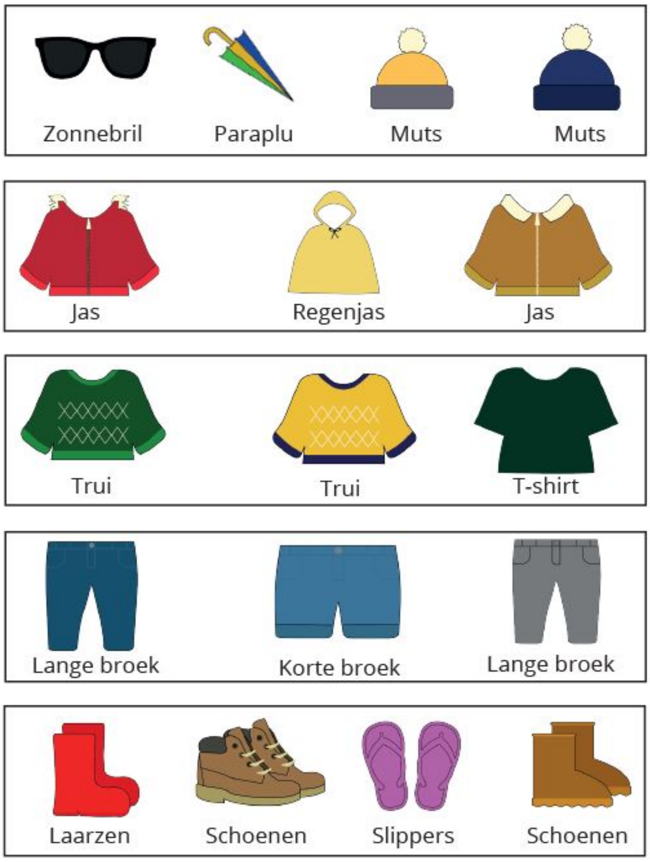

C.2. Equations of the collaboration and class scores

Equation C.1 - C.4 show how the collaboration score and the three class scores are calculated for the 2 recordings of the first study, and 9 recordings of the second study. The scores are based on the relative measures of certain indicators relative to the number of intervals (N) of a recording and the number of indicators.

Participation:

$$score = \frac{action + interaction}{2 * N} \tag{C.1}$$

Perspective Taking:

$$score = \frac{adaptive\ responsiveness + audience\ awareness}{2 * N} \tag{C.2}$$

Social Regulation:

$$score = \frac{negotiation + selfevaluation + transactive\ memory + responsibility\ initiative}{4 * N} \tag{C.3}$$

Collaboration:

$$score = \frac{The\ sum\ of\ the\ eight\ indicator\ scores}{8 * N} \tag{C.4}$$

C.3. First study: the measurements

C.3.1. Relative indicator scores per group

Table C.1 shows the relative indicator scores for each pair of children that participated in the first study. The relative score of an indicator is the percentage of intervals where that indicator occurred, relative to the total number of intervals that were annotated for that recording.

No.	No. of Int.	Action	Inter.	Adp. Resp.	Aud. Awr.	Neg.	Self-eval.	Trans. M.	Resp. In.
1	12	100	100	50	0	0	8	8	33
3	14	100	86	29	0	14	0	21	29

Table C.1.: The calculated relative scores (%) for all indicators of collaboration. The abbreviations used in the table, refer to: No. = Number, No. of Int. = Number of Intervals, Inter. = Interaction, Adp. Resp. = Adaptive Responsiveness, Aud. Awr. = Audience Awareness, Neg. = Negotiation, Self-eval. = Self-evaluation, Trans. M. = Transactive Memory and Resp. In. = Responsibility Initiative.

C.3.2. Mean score of collaboration

Table C.2 shows the mean score indicating the average level of collaboration seen between pairs of children during first study. The standard deviation was not calculated, since it is only based on two samples.

	Collaboration	Participation	Perspective Taking	Social Regulation
Mean score	36	96	20	14

Table C.2.: The mean score (%) over the relative measures of the class scores of the pairs of children that participated in the first study.

C.4. Second study: the measurements

C.4.1. Relative indicator scores per group

Table C.3 shows the relative indicator scores for each pair of children that participated in the second study. The relative score of an indicator is the percentage of intervals where that indicator occurred relative to the total number of intervals that were annotated for that recording.

No.	No. of Int.	Action	Inter.	Adp. Resp.	Aud. Awr.	Neg.	Self-eval.	Trans. M.	Res. In.
1	18	94	61	22	0	6	0	6	11
2	19	100	89	32	11	21	0	26	26
3	15	93	100	33	13	7	20	40	40
4	12	92	100	50	00	17	17	8	17
5	19	89	100	63	42	16	5	11	42
6	19	95	100	26	79	0	5	21	53
7	14	93	14	00	0	0	0	7	0
8	19	100	89	74	47	21	16	32	53
9	25	84	96	36	52	24	0	32	52

Table C.3.: The calculated relative scores (%) for all indicators of collaboration. The abbreviations used in the table, refer to the group number and one of the indicators: No. = Number, No. of Int. = Number of Intervals, Inter. = Interaction, Adp. Resp. = Adaptive Responsiveness, Aud. Awr. = Audience Awareness, Neg. = Negotiation, Self-eval. = Self-evaluation, Trans. M. = Transactive Memory and Resp. In. = Responsibility Initiative.

C.4.2. Mean and standard deviation score of collaboration

Table C.4 shows the mean score, and standard deviation, indicating average level of collaboration seen between pairs of children during second study.

	Collaboration	Participation	Perspective Taking	Social Regulation
Mean	39	88	32	18
Standard deviation	12	14	20	9

Table C.4.: The mean score and standard deviation (%) over the relative measures of the class scores of the pairs of children that participated in the second study.