Designing a Social Robot to Support Children's Inquiry Learning: A Contextual Analysis of Children Working Together at School

Daniel P. Davison[1] · Frances M. Wijnen[2] · Jan van der Meij[2] · Dennis Reidsma[1] · Vanessa Evers[1]

Abstract

Designers of educational interventions are always looking for methods to improve the learning experience of children. More and more, designers look towards robots and other social agents as viable educational tools. To gain inspiration for the design of meaningful behaviours for such educational social robots we conducted a contextual analysis. We observed a total of 22 primary school children working in pairs on a collaborative inquiry learning assignment in a real world situation at school. During content analysis we identified a rich repertoire of social interactions and behaviours, which we aligned along three types of interaction: (1) Educational, (2) Collaborational, and (3) Relational. From the results of our contextual analysis we derived four generic high-level recommendations and fourteen concrete design guidelines for *when* and *how* a social robot may have a meaningful contribution to the learning process. Finally, we present four variants of our Computer Aided Learning system in which we translated our design guidelines into concrete robot behaviours.

1 Introduction

Social robots are increasingly used as educational tools in topics such as STEM and language learning [5,52]. In the

✉ Daniel P. Davison
  d.p.davison@utwente.nl

  Frances M. Wijnen
  f.m.wijnen@utwente.nl

  Jan van der Meij
  j.vandermeij@utwente.nl

  Dennis Reidsma
  d.reidsma@utwente.nl

  Vanessa Evers
  v.evers@utwente.nl

1 Human Media Interaction, Faculty of EEMCS, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands

2 ELAN, Faculty of BMS, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands

EU-FP7 project EASEL[1] we look specifically at learning interventions that use robots in a collaborative inquiry learning assignment, described later in more detail. To develop such interventions, we must design robot behaviour and human-robot interaction patterns that are appropriate for the learner and the educational context. Generalised pedagogical and educational psychology theories inform us about what we might generally expect by way of interactions that occur between learners, but do not specify in what way such social interactions carry across to child-robot learning situations. To ground the design of actual behaviour and interaction of the robot we also need real examples of behaviours exhibited by children in the specific environment and collaborative task in which the robot should operate. From such examples we can derive concrete opportunities and guidelines for *when* and *in what form exactly* robots may contribute to a child's learning.

In this paper we report on a contextual analysis that we carried out to derive guidelines for the design of social robots to support primary school children in a collaborative inquiry learning assignment. This work was guided by the following research questions:

---

1 EASEL—Expressive Agents for Symbiotic Education and Learning (http://easel.specs-lab.com/).

RQ1 Which typical behaviours and social interactions do we identify in learners working with our assignment and how do these behaviours align with theory?

RQ2 How do children progress through the various phases of our learning assignment and do the types of learning behaviour they exhibit vary with these phases?

RQ3 What are opportunities for social robot interventions in regard to interaction patterns, specific behaviours and roles, to contribute to the child's learning in our assignment?

The paper has the following structure: Sect. 2 presents our view on social learning through the lens of the literature. Section 3 describes the approach of our contextual analysis, in which we annotated video recordings of pairs of children collaborating on a learning assignment in their school. Section 4 presents the annotation schemes we used to analyse these video recordings, the results of which are presented in Sect. 5. Based on these results, we formulated four generic recommendations and fourteen concrete design guidelines for robot behaviour, presented in Sect. 6. Finally, Sect. 7 describes how we applied these guidelines to build four variants of child-robot learning scenarios. Here, we also reflect anecdotally on the social learning behaviour exhibited by children interacting with these variants. Limitations and future work are discussed in Sect. 8, after which we conclude the paper in Sect. 9.

2 Background

In this section we briefly discuss the literature and theories from educational psychology and pedagogy related to the type of learning assignments and social interactions that occur in our collaborative inquiry learning context. Insights presented here are partly based on an earlier paper by Charisi et al. [15].

2.1 In Learning, Students Construct Knowledge

The well-known developmental theory of Piaget [55] describes learning as a dynamic process comprised of successive stages of adaption to reality. According to this constructivist framework, children are active learners who construct their knowledge by creating and testing their own theories and beliefs. During play, children explore the real world, organise their thoughts, and perform logical operations [54]. This occurs mainly in relation to concrete objects rather than abstract ideas [59].

Inquiry learning is one of the tools to support a more structured exploration of abstract ideas, often by using physical objects. During inquiry learning tasks learners adopt the scientific method of exploration to discover underlying rules and principles that govern phenomena in the world around them. Knowledge is assimilated by conducting experiments during which learners construct hypotheses and argumentations, and reflect upon applied methods and gathered results [39,40,46, 71]. A typical inquiry cycle consists of five consecutive steps: (1) Preparation; (2) Hypothesis generation; (3) Experimentation; (4) Observation; and (5) Conclusion. Learners often require additional support when they encounter difficulties while working on inquiry tasks [46].

Vygotsky introduced the notion of scaffolding to describe the supportive structures offered to a learner that lead to the construction of higher mental processes [74]. The amount of information and guidance offered to the learner influences the complexity of the task. Bell et al. [3] and Banchi et al. [1] describe a continuum of several levels of inquiry. Learners start with highly-structured "confirmation inquiry tasks" in which questions, hypotheses, methods, and results are explicitly provided. They then gradually move to more "open inquiry tasks" in which they formulate their own questions, hypotheses, and methods, after which they construct arguments to explain experiment results [1]. By offering appropriate scaffolds (e.g. predefined task structures or explicit instructions) learners are able to progress from closed to more open forms of inquiry.

2.2 Learning is Inherently Social

In solitary learning situations, scaffolding is traditionally offered by non-social (technological) tools. For example, by providing templates or background information about the topic, or by constraining the learner's interaction with the learning environment [46,71,72]. However, in most classroom situations learners learn together with "someone else", which can be beneficial to the learning process.

Social educational "partners" help each other to optimally develop their abilities and knowledge. According to Vygotsky [74] cognitive development occurs and is promoted by an individual's interaction with the world and by collaboration with others. He distinguished between two developmental levels: the level of actual development and that of potential development. The actual development is the level that the learner has already reached and where they can comfortably solve problems independently. The potential level of development is the level that the child should be able to reach, if he/she interacts with more knowledgeable others. The space between the actual level of development and the potential level of development is referred to as the Zone of Proximal Development (ZPD) [74]. It is through guidance or collaboration with others that learners bridge the ZPD and manage to grow from their actual level of development to their potential level of development [15,20].

There are many definitions of collaborative learning. For instance, Dillenbourg defines collaborative learning as "a situation in which two or more people learn or attempt to learn something together" [24, p. 1], focusing on participation of multiple parties in the learning process. Similarly, Rogoff's [60] definition of collaboration describes mutual involvements, engagement, and participation in shared endeavours, which may or may not serve to promote cognitive development. Finally, the process of peer learning has been defined by Topping as "the acquisition of knowledge and skill through active helping and supporting among status equals or matched companions" [68, p. 1].

An important shared characteristic of such descriptions of collaboration, is the requirement for a certain interactivity between learners, which is defined by the extent to which learning partners influence each other's cognitive processes [24].

The role of the learner and their social partner shape how they collaborate with each other. Firstly, the social partner can be a "more knowledgeable other", e.g. a teacher/tutor or more advanced student [76]. Secondly, the social partner could be a "similarly or differently knowledgeable other", e.g. a fellow student of a similar level working on the same task [24,68]. Finally, the learner could be working with a "less knowledgeable other", e.g. a less advanced student in need of support of their own. Tutoring other students can contribute to the learning of the more advanced student as well [56].

2.3 ECR: A Three-Level Model to Organise Social Learning Interactions

Related literature and theories on social learning describe several forms of social interactions that take place between collaborating learning partners. In this section we summarise and organise this theoretical background to derive a pragmatic model that can be used as a lens through which to look at social learning. We propose the "**ECR**" model, as illustrated in Fig. 1, which describes three levels of social interactions. Firstly, there are Educational interactions (e.g. asking for help and explaining). This requires "being in sync" with each other, which is pursued through Collaborational interactions (e.g. discussing joint goals and strategies, and performing joint actions). To achieve this, co-learners need to be "working towards the same thing" and they need to "know the value of their partner", which can be calibrated through Relational interactions (e.g. bonding, shared experiences/emotions, team role assignments).

2.3.1 Social Learning Requires Educational Interactions

In collaborative inquiry learning, students working together provide each other with scaffolding. They do this by sharing
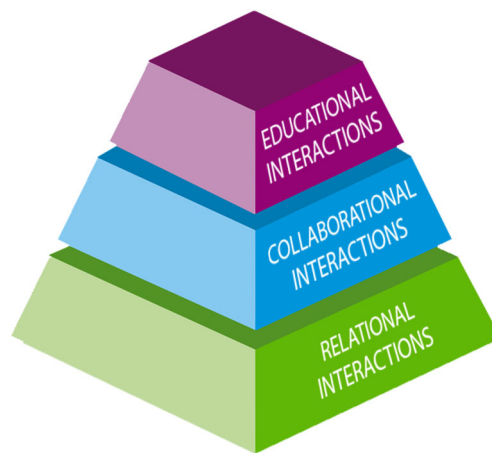

Fig. 1 The **ECR** model illustrates the three levels of social interaction that occur during collaborative inquiry learning: Educational interactions, Collaborational interactions, and Relational interactions

and explaining their knowledge, ideas, and techniques [74]. This indicates that a less knowledgeable learner can benefit from support offered by a more- or differently knowledgeable other. This is done through what we call educational interactions. Through such interactions learners help each other to develop skills on both a cognitive level, related to learning content (how to understand a problem), and on the metacognitive level, related to the learning process (how to approach a problem).

Scaffolds and explanations offered by child A to child B not only help child B, but also contribute to the learning of child A. Offering explanations to a partner triggers certain cognitive processes in the learner. These cognitive processes are considered helpful for learners to gain a deeper conceptual understanding of the subject matter [56]. In inquiry learning situations, this process of explanation often occurs when making predictions or drawing conclusions.

Literature typically describes two forms of explanation: (1) self-explanation, which refers to explanation of the subject of interest to oneself [17]; and (2) interactive explanation, which refers to explanation to another person [56]. In the latter case, the presence of a social partner facilitates children's verbalisation of their thinking (e.g. by asking probing questions). However, depending on the type of the social partner, children may exhibit different behaviours.

Tutors and teachers are typically more knowledgeable than the learner and have a more holistic view of the learner's progress. Therefore, they can offer additional explanations and scaffolding in the various stages of the inquiry learning process. Furthermore, tutors can help learners to appraise their own performance and efficacy. Finally, tutors can use their knowledge about the skill level and personality of children to tailor the explanations and scaffolding.

2.3.2 Social Learning Requires Collaborational Interactions

When working on a joint activity, such as a learning task, collaborators rely on an intricate coordination and communication process to reach a level of alignment [18,41]. Klein and Feltovich [41] describe such joint activities in terms of criteria, requirements, and choreography. As criteria for joint activities they state that there should be an intention to work together, and the work must be interdependent. Regarding requirements for joint activities; there should be a degree of interpredictability between partners, there must exist a sufficient degree of common ground, and each partners' actions must be directable by the other. Throughout the interaction, as they progress through the phases of the activity, partners use signalling to facilitate the choreography of collaborating on joint activities.

Such principles of joint activities apply similarly to the context of collaborative learning. In this context collaboration implies that both co-learners are more or less on the same status level, can perform the same actions, and often have a common objective [24]. Furthermore, offering meaningful scaffolding to one another requires that learners align their goals, ideas, plans, and reasoning in an explicit and understandable way [45,67]. Establishing and maintaining this alignment takes place through certain types of interactions: we call these collaborational interactions.

Learners may seek alignment regarding the overall task objective, the strategy for attaining this objective, and the actions necessary for executing this strategy.

In inquiry learning, the task objective is often imposed by the learning task. Co-learners can align their interpretation of the task objective by agreeing upon the meaning of the task and by discussing possible completion criteria. The phases of the inquiry cycle outline the strategy that learners may take to reach the overall objective. Learners can establish alignment regarding their strategy by discussing the inquiry phases and may maintain this alignment by signalling their progression through the phases. Each phase of the inquiry cycle requires a specific sequence of actions. Learners should discuss such actions and agree upon a division of labour to prevent disrupting each other.

2.3.3 Social Learning Requires Relational Interactions

Effective collaborative inquiry learning requires a healthy working relationship. Firstly, for partners to accurately judge each other's interests, abilities, skills, and knowledge, they need to interact and understand one another on a personal level. Secondly, in order to agree upon shared objectives, strategies, and actions, partners need to be "allies in learning"; meaning that they maintain mutually benevolent attitudes towards each other.

Through their past experiences and interactions, partners shape and maintain common ground [41]. Doing so, they are able to more accurately envision their partner's strengths and weaknesses. This enables them to create a fair and appropriate role and task division.

While working on a task, it is important for partners to remain mutually invested in the learning process. The importance of positive feelings during the learning process has been stressed before [43]. They promote the individual's openness to new experiences and resilience against possible negative situations [26]. Furthermore, dynamic behaviours might involve reciprocal influences between emotion and cognition [25]. For instance, emotions affect the ways in which individuals perceive reality, pay attention, and remember previous experiences, as well as the skills that are required for an individual to make decisions. Shared success and failures are ideal moments in the learning process where partners may offer each other encouragement and motivation.

2.4 Related Work from HRI

We envision that a social robot is able to navigate the social space of all three of the ECR levels by taking on a variety of roles in the collaborative learning process. As a starting point, we look at examples from related work, which help us to highlight and illustrate possible contributions of such a robot.

Robots in educational settings are often presented to the learner as a "peer" or "tutor" [53,64,78]. Within these roles the robot may have a different impact on the child's learning and may use a different repertoire of behaviours and techniques to offer scaffolding.

Robots presented as facilitator, teaching assistant, or tutor, often take on a role of "more knowledgeable other" to guide one or more learners in the learning task (e.g. [13,14,36,48,62,77]). Operating in this role, a robot may contribute to the learning process through traditional tutoring methods; such as direct instruction, praise, encouragement, or explicit feedback [76]. For example, Kanda et al. [36] used a social robot as a facilitator to offer direct instruction and explanations to children working on a collaborative task. Their results suggest that the social nature of the robot contributed to increased motivation in early phases of the learning task. Similarly, Chandra et al. [13] found that a robot who acted as facilitator to two collaborating children, improved their feelings of responsibility. Additionally, Saerbeck et al. [62] have shown that a robotic tutor with social supportive behaviours can have a positive effect on the performance of the learner. However, in certain tutoring situations prone to distractions, *too much* socially appropriate behaviours seem to have a detrimental effect to learning [37].

Alternatively, robots can be presented to the child as co-learners or peers. As such, the robot may portray a "differently (or equally) knowledgeable other" and may seek

collaboration with the child within the learning task. In this role the robot may offer proactive anticipatory helping behaviours to improve the fluency of their collaboration [2,28]. Such peer-like robots often engage with children over multiple interaction sessions, seeking to build a social relationship [34,65]. A common approach for promoting learning with a peer-like robot is during playful game-like interactions. For example, the robot might play either an "ally" or an "opponent" of the child, challenging them to achieve a higher performance (e.g. [4,12,31]). To sustain such interactions over extended periods of time, the robot and child can engage in a diverse variety of such activities [19]. Other approaches focus on interactive storytelling with a social robot [44]. For example, Gordon et al. [27] used such a robot to improve children's second language skills.

Finally, robots portraying a "less knowledgeable other" tend to evoke caregiving behaviour in the child. Children are naturally inclined to help a "robot in need" when it displays distressed behaviour. However, the extent of helping depends on how the child was introduced to the robot [6,7]. Tanaka et al. [66] use such a care-receiving robot to promote vocabulary learning. In such cases, the child themselves may act as a tutor, correcting the robot's mistakes (e.g. [29,47]) or explaining concepts to the robot (e.g. [23]). Through this *learning by teaching* paradigm, the child gains a deeper understanding of the learning content [16,61].

In conclusion, here are some of the various (social) roles the robot could have: explaining to the learner, be explained to by the learner, encouraging the learner, praising the learner in various ways, giving good or bad examples as a fellow learner, and many more. The effectiveness of these actions depends to a certain extent on the relation between the learner and the other, and how the other behaves and presents itself to the learner. To gain further inspiration regarding the design of concrete robot roles and behaviours we conducted a contextual analysis.

3 Contextual Analysis

The type of robot that we investigated would have to navigate the social interactions with learners in the context of an inquiry learning assignment. From theory and related work we found that the robot may navigate several roles, as outlined in the previous section: more-, differently-, and less knowledgeable other. Within such roles, we expect three different types of interactions to take place: educational, collaborational, and relational. As input to our design guidelines we carried out a contextual analysis to determine concrete opportunities for *when* and *how* robots may contribute to a child's learning process.

3.1 Method

The field of Human-Computer Interaction often emphasises the importance of considering the context in which interactive products are to be used (e.g. [8,9,38,49]). Over the years different terms have been used for contextual analysis, such as "Context-of-use analysis", "Usability Context Analysis", or "Contextual Design". Assessing and accounting for contextual factors is considered as an integral part of a user-centred design process aimed at optimising usability of products. By mapping the challenges and opportunities that exist in a certain context, the outcomes of a contextual analysis will typically have a drastic influence on the design of an interactive product. Therefore, such a contextual analysis should take place at an early stage—before having a fully developed prototype or product.

The method of contextual analysis aims to take into account as much as possible, all factors that may influence the use of a product. Such factors include: characteristics of the user group; tasks that users aim to carry out with the product; and the technical, organizational, and physical environment where the product is to be used. A contextual analysis typically consists of approaches like stakeholder analysis, surveys, and focus groups [49], as well as ethnographic strategies like case studies and field studies [50] (or combinations thereof).

In our contextual analysis we used such ethnographic methods to investigate the context in which children learn with our specific inquiry assignment. We conducted a field study where we recorded pairs of children working with our assignment. We then used a grounded theory approach to design an annotation scheme with which we analysed the observed behaviours [51,63].

3.2 Venue and Participants

Since we are interested in gaining insights in typical behaviours of children as they occur in practice, we aimed to maintain a high degree of ecological validity throughout the study. Therefore, we focused on pairs of children working on a learning assignment in their own school during regular school hours, without robot. The 22 Dutch Montessori elementary school students (50% female), aged 6–9 years old (the target age range of the EASEL project) participated in the study. Children in this age range show a variance of cognitive development, verbal, and non-verbal behaviour. As such, including children across this age range will give us a broad spectrum of insights on possible behaviours. The school is located in a residential area near the University of Twente, in the city of Enschede, The Netherlands.
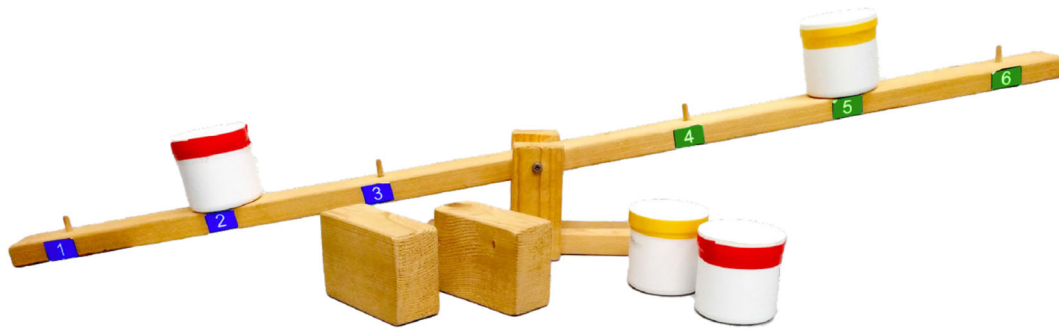
Fig. 2 The wooden balance with several differently weighted pots that can be placed over the pins and wooden blocks that can be placed underneath the balance

3.3 Procedure

Prior to the study, ethical approval was granted by the ethical boards of the Behavioural, Management and Social sciences (BMS) and the Electrical Engineering, Mathematics and Computer Science (EEMCS) faculties of the University of Twente. Permission forms were distributed to the children's parents through the school's regular communication channels.

The experiment started with an explanation to a pair of children about the assignment. The researcher showed them the learning materials and demonstrated how they worked. This introduction took approximately 5 min. After the introduction, the children would each receive the assignment sheets with five tasks from the experimenter. The children were given approximately 30 min to complete the assignment. When a pair finished the assignment, the experimenter would thank them for their participation and say goodbye. During the experiment, the experimenter was present so that children could ask for help if they had any questions.

The interactions were recorded by three cameras: two for recording actions of each of the two children and an overview camera to record the learning materials.

3.4 Balance Scale Materials

The physical balance scale instrument was originally designed by Inhelder and Piaget [30]. In our study, children used such a balance scale (see Fig. 2) to explore the moment of force. The moment of force (or torque) is a physical concept that describes the rotational force acting upon an object in relation to a reference point. To explore this concept, our children did small experiments in which they placed pots of various weights at various distances from the scale's central pivot point.

The two-armed wooden balance had a central pivot point and three pins at equal increasing distances on each arm. These pins were numbered 1 to 6 from left to right. Numbers 1 to 3 were given a blue label and numbers 4 to 6 were given

a green label. The children received four pots that could be placed over the pins: two heavy red pots and two lighter yellow pots. Two wooden blocks could be placed under the balance to secure it in a horizontal position.

3.5 Learning Assignment

Children received a total of five numbered tasks as part of their learning assignment. Task instructions were printed on paper and consisted of brief textual descriptions and illustrations.

Task 1, 2, 3, and 4 were similar in nature. In these tasks the children's instructions and actions were symmetrical and synchronous: each child had the same task instructions and could complete the same actions. Each task was structured according to the steps of the inquiry learning. In the preparation step, children were first asked to place a wooden block underneath each arm to prevent the balance from tilting prematurely. They were then asked to place two pots on the balance (pot color and position were given explicitly in each task). In the prediction step, children were asked to predict what they thought would happen when the wooden blocks were to be removed. They could select one of three options: (1) the blue side will go down; (2) the balance will stay in equilibrium; or (3) the green side will go down. In the experiment step, children were asked to remove the wooden blocks, allowing the balance to tilt freely. During the observation step they explicitly stated their observation about the position of the balance and explained whether this was in agreement or disagreement with their initial prediction. Finally, in the conclusion step they would give an explanation for their observation using their own words (explain why the balance tipped left, right, or stayed in equilibrium). They gave this explanation verbally to each other and the experimenter.

Task 5 was structured differently in order to possibly trigger more diverse collaborative behaviours. In this final task children worked asynchronously and were given asymmetrical task instructions: here it was necessary for the children to take consecutive steps, where each child took different

actions. First, child A would choose one of the pots and place it on the blue side of the balance. Then, child B would choose one of the pots and give it to child A. Subsequently, they had to answer the following question together: "where should the pot be placed on the green side to keep the balance in equilibrium?" Child A would then place the pot on the green side where they thought it should be placed to keep the balance in equilibrium. Finally, after removing the blocks, the children would explain to each other if their prediction was correct and why the balance was in this position.

4 Annotation Scheme

To further understand the children's educational, collaborational, and relational interactions and to understand how the children progress through the phases of the inquiry cycle, we manually annotated the video recordings using the ELAN tool.[2]

The annotation scheme was divided into five layers. To determine the label sets for these layers, we first inspected sample fragments from the recordings to get a sense of the kind of behaviours that could be found in them. We used this to choose the possible labels for annotating each of the three levels from the ECR model. In the fourth annotation layer we labeled video fragments according to the occurrence of the inquiry cycle phases. The fifth layer only had the label *Supplementary interactions*. This label was used to indicate fragments containing behaviour that we identified as important for the learning process, but did not fit cleanly in the existing label descriptions in our code book. These *Supplementary interactions* were subsequently also ordered along the three levels of the ECR model.

Using the developed annotation schemes, coder A labeled 100% of the data. To calculate the inter-annotator agreement score, coder B labeled an overlapping 20% of the data. On the combined annotations of the educational and collaborational levels, we computed a Cohen's Kappa of 0.11. Inspection of the disagreement data showed that this low score was mostly caused by one coder assigning a label to a certain video fragment while the other coder did not. Upon discussion of such fragments, it became clear that the majority of the behaviours that were missed by one of the coders, did indeed rightfully belong under the annotated label that the other coder assigned to it. When we excluded these annotation instances that were missed by one of the coders, Cohen's Kappa increased to 0.77. The Cohen's Kappa for all annotations in the inquiry process category was 0.61. After excluding the fragments that were missed by one of the coders, Cohen's Kappa increased

to 0.91. These inter-annotator agreement scores left us confident that the labels that we gathered from the annotation procedure were meaningful enough for further analysis, even though a number of instances of relevant interactions might not have been assigned a label.

The results of our annotations will allow us to identify *how* children express themselves when navigating various roles in relation to their learning partner. Such insights will form the basis for our behaviour design guidelines for a social robot that will function in similar roles. Additionally, these annotations will allow us to identify *when* children express certain behaviours, which gives us clues to when the robot might contribute. Finally, we will be able to assess whether our specific learning task is appropriate for the target user group, and whether it lends itself to investigate the full inquiry learning cycle.

The annotation schemes are described in more detail in the next five sections.

4.1 Label Set for Educational Interactions

During the course of an assignment, children often engaged in interactions that were educational by nature. That is, they engaged in discussion and knowledge construction around the learning content. On the one hand, we observed very explicit educational exchanges (for instance, child A asking child B for an explanation or opinion). On the other hand, children engaged in implicit educational interactions (for example, child B integrating an argument in their explanation that was previously made by child A). Additionally, we observed instances where children would engage in discussions where they agreed or disagreed with each other's answers or actions.

The following labels were used to identify the interactions on the educational level. These labels were partly based on work by van Dijk et al. [69], who investigated collaborative drawing behaviours.

Elicitation/asking questions When children want to extract knowledge or support from their partner, mostly in the form of a question. For example, "What do you think will happen?".

Explaining When a child gives any form of verbal explanation directed at the partner about the assignment or the principles of balance. This explanation may be spontaneous or elicited. For example, "This pot is at the same spot as the other, so I think it will stay in equilibrium".

Thinking aloud This label includes cases where a child thinks aloud about the construction of their argumentation, but this externalisation is not explicitly directed towards to the partner. Such utterances are often less elaborate than explanations. For example, "hmm... I think... hmm... the left side will go down".

Agreement or disagreement Instances where a child shows basic (dis)agreement with their partner's reasoning, without explicitly showing comprehension. This includes, for example, verbal utterances ("yes, I think so too" or "no, that's not right") and nonverbal gestures (nodding or shaking their head).

Integrating This label goes beyond basic agreement, when a child uses the content of a partners' explanation in their own line of argumentation. For example, child A says: "I think the balance tipped left because the pot on the blue side is placed more towards the end" and child B responds: "Because it is more towards the end of the balance the left side is heavier". In this case, the explanation of child A is labeled as *Explaining* and the response of child B as *Integrating*.

Elaborate critique When a child criticises or corrects the content of the explanation of the partner, going beyond basic disagreement. For example, child A explains "The yellow pot is heavier", and child B responds "No, it's lighter!" the response of child B is annotated as *Elaborate critique*.

4.2 Label Set for Collaborational Interactions

While working on their assignment, children had to collaborate in order to share the physical learning materials. Doing so effectively required them to align their interpretations of task objectives, communicate their approach, and agree upon actions. The following labels were used to identify the interactions on the collaborational level.

Reading aloud When a child reads a task or instruction out loud. Since it is often difficult to distinguish between reading out loud for oneself and reading for the partner, this label includes both cases.

Coordination When a child makes explicit suggestions or explanations about how to proceed with the task, or explicitly mentions their progress within the task. For example, "it's your turn to place a pot" or "I'm finished with this one".

4.3 Label Set for Relational Interactions

In order to collaborate effectively, partners needed to maintain mutually benevolent attitudes with respect to each other. Such attitudes were shaped and maintained through relational interactions, which took the form of shared experiences, emotions, and frustrations. We did not have a grounded annotation scheme for interactions that took place on the relational interaction level. Instead, we labeled salient episodes in this domain as part of the 'Supplementary interactions' category.

4.4 Label Set for Inquiry Processes

A typical inquiry learning cycle is composed of five processes: (1) preparation; (2) hypothesis generation; (3) experimentation; (4) observation; and (5) conclusion. To analyse how children progressed through these phases, we identified behaviours that represent these processes using the following labels:

Preparation When children prepare the physical learning materials in accordance to the task, e.g. adding pots on the balance and placing the wooden blocks.

Hypothesis generation In this phase, children state their hypothesis, e.g. by verbalising out loud to each other, or by selecting a hypothesis on their assignment sheet.

Experimentation Performing the actual experiment, i.e. removing the wooden blocks so that the balance tips over.

Observation In this phase children state their observations about the experiment, by verbalising or pointing (e.g. "That side went down!"), or by documenting the observation on their assignment sheet.

Conclusion When children draw conclusions about the current task, by explaining to each other. This phase only occurred verbally, e.g. by explaining whether their hypothesis was correct and why.

4.5 Supplementary Interactions

We observed many salient behaviours and interactions that appeared to be relevant to the above label sets, but did not fit cleanly into any of the available labels. These spanned a wide range of actions. For example, at one point a child grabbed and took materials from another child; this action happened only in one session, but was considered very salient for collaborational interaction. Therefore, to fully capture the richness of the social interactions we labeled such additional observed behaviours as *Supplementary interactions*. Behaviours with this label were then ordered along the three levels of the ECR model and are reported and discussed as such in the respective sections below.

5 Results of the Contextual Analysis

This section reports the content of the annotations. All the 11 pairs of children completed tasks 1, 2, 3, and 4, each of which followed a similar structure. However, in 36% of the tasks they needed some assistance from the researcher. In most cases, this simply meant the researcher reading the task out loud or rephrasing the task instructions in different words. In cases where that did not help, the researcher would ask thought-provoking open questions without revealing the answer.

Table 1 Frequency of educational interaction labels

Educational label	Frequency	Percentage
Elicitation/asking questions	116	25.3
Explaining	164	35.8
Thinking aloud	53	11.6
(Dis)agreement	97	21.2
Integrating	6	1.3
Elaborate critique	22	4.8
Total	458	100

Table 2 Frequency of collaborational interaction labels

Collaborational label	Frequency	Percentage
Reading aloud	23	30.3
Coordination	53	69.7
Total	76	100

Task 5 took on a different structure, which likely increased its complexity. During this final task, children required more support from the researcher; only 3 out of 11 pairs of children were able to complete the last task without any help or guidance. The children expressed confusion about the structure of this task, which was composed of different steps than the previous tasks. In particular, they often did not understand the different responsibilities that they needed to fulfill in order to complete the task. Still, during this task we did observe interactions related to learning.

5.1 Educational Interactions

In identifying the educational interactions between children, a total of 458 annotations were performed. Frequencies for individual labels on the educational level are shown in Table 1.

The most prevalent labels were *Elicitation*, *Explaining*, and *(Dis)agreement*, which indicates that children were intentionally helping each other by explicitly sharing knowledge. Their helping behaviours focused primarily on topics related to the process (e.g. "I don't understand what we have to do") and the learning content (e.g. "I don't understand what happened"). Additionally, learners unintentionally helped their partner during behaviours labeled as *Thinking aloud*, by implicitly sharing knowledge about the task or content. Finally, we sporadically found instances where children explicitly integrated or critiqued their partner's explanation or reasoning process, illustrating a deeper level of discussion and debate.

In addition to the annotations labeled according to our annotation scheme, we performed an explorative analysis of the annotations that were labeled as *Supplementary interactions*. Firstly, in these annotations we observed many instances of pointing and gesturing behaviour. This was often used as a nonverbal clarification or explanation (e.g. pointing at a pot or location on the balance, or gesturing the position of the balance, or pointing at other's assignment sheet to indicate the correct answer). Secondly, we found several implicit and explicit indications that children were compar-

ing and aligning their own answers to those of their partner. For instance, they would look at and sometimes copy the answer of their partner, they would discuss explicitly (e.g. "What are you going to fill in?"), or they would announce when they had a breakthrough of their own (e.g. "I know the answer!"). Thirdly, although most children accepted lessons learnt through observations of the embodied learning materials (e.g. "I have now of course seen it... so I have to believe it"), we occasionally saw instances of confirmation bias. For instance, some children would alter reality to fit their prediction (e.g. perhaps jokingly, they would say "See, I was correct, I said it would stay balanced", while artificially holding the balance in equilibrium with one hand). Others would adapt their original (incorrect) hypothesis to match their observations on the fly (e.g. "But I did think this would happen beforehand" or "You know, actually I knew this was going to happen"). Nevertheless, in some instances this prompted children to discuss an alternative hypothesis or propose an alternative experiment (e.g. "What would happen if we would place this pot here and the other here?"). Finally, some children showed curiosity about the balance unrelated to the assignment, which provoked them to investigate beyond the scope of the subject matter (e.g. they would try to figure out what was inside the weighted pots or they would place unrelated items on the balance).

From these additional observations we found that children discussed their observations and updated their knowledge while progressing through tasks. We also found that gestures and gaze behaviours played an important supplementary role in the collaborative inquiry learning interactions: children seemed to naturally take advantage of the physical nature of the embodied learning materials when explaining to their partner.

5.2 Collaborational Interactions

In identifying collaborational interactions between children, a total of 76 annotations were performed. Frequencies for individual collaborational labels are shown in Table 2. In general, we found behaviours corresponding to all of the criteria, requirements, and choreography of joint activities discussed in Sect. 2.3.2.

Behaviours that fell under the label *Reading aloud* included children reading (parts of) the assignment out loud.

Often, this role was taken on by the child who was most proficient in reading. However, when both children had difficulties reading the task instructions, the experimenter would read the assignment out loud. By reading instructions out loud, the children seemed to ensure that both team members were in agreement on the task objectives. Additionally, this naturally provided whoever was listening, the opportunity to ask clarifying questions, or suggest a plan of action.

Among behaviours labeled as *Coordination*, we mostly found interactions that were related to the planning and execution of individual actions. For instance, children would take turns manipulating the learning materials (e.g. "Now it's your turn to place the pots"), or they would share responsibility (e.g. "I'll place this pot, you do that one"). Furthermore, they would occasionally plan ahead based on the structure of the task (e.g. "Hang on, before we can remove the blocks we first have to give our answer"). Finally, we identified instances where children explicitly mentioned their progress in the assignment (e.g. "There, that was task 2 done, let's move on to the next!").

From the explorative analysis of the behaviours in the supplementary interactions category, we identified several additional situations that were relevant on the collaborational level. Firstly, in addition to the explicit verbal interactions coded under the *Coordination* label, we found several instances of successful implicit coordination. While children worked together on the physical learning materials, they would often maintain a shared focus of attention through mutual gaze and by observing each other's actions. Additionally, we noticed that children would often use behaviours, body language, pointing gestures, or gaze to signal their intentions to their partner. For example, children would seek for approval or permission while doing actions (e.g. very slowly/hesitantly starting an action, while searching for eye contact, giving their partner opportunity to intervene), or they would assist their partner's actions (e.g. holding the balance steady while the partner removes or places pots). Secondly, we observed instances where misalignment of actions or goals occurred. Children could often explicitly signal such a situation to their partner verbally (e.g. by stating "This pot is in the way" or "I think we did it wrong") or nonverbally (e.g. obstructing or overruling their partner's actions, for instance by forcefully taking or removing a pot). In such cases, partners would seek re-alignment through more explicit coordination. They would, for instance, seek agreement with their partner before performing actions (e.g. "Can I do this now?" or "Are you choosing that pot..? Ok good!") or they would explicitly time and synchronise their actions (e.g. "Now? Yes? ... Yes now!" or "3..2..1..Now!"). Finally, there were instances where misalignment of task progress occurred. For example, children would often check their partner's progress in the task by looking at and comparing each other's assignment sheets. If they noticed that

their partner was falling behind, children would often use either of the following two strategies: (1) they would either dawdle patiently while waiting for their partner to catch up, not directly contributing to their partner's progress (although some children asked the experimenter "Do we need to wait for each other?"). Or they would (2) directly interfere with their partner's progress by prompting the answer (e.g. "Fill this one in [pointing at answer sheet]"), or by taking on a more dominant role where they would give their partner explicit tasks or would explicitly direct the actions of their partner (e.g. "Now you do this" or "You have to do that").

5.3 Relational Interactions

Since we did not have a grounded annotation scheme for behaviours that constitute relational interactions, we looked at the interactions that fell under the supplementary interactions category. In doing this, we identified several situations that appeared relevant on the relational level. Firstly, in addition to mutual eye contact and shared focus of attention, children would often share mutual emotions while progressing through the tasks. For example, they would smile and laugh together when something (unexpected) happened, they would celebrate or cheer together when they had made a correct prediction (e.g. "Yaaay!"), or they would explicitly share their confusion or frustration when their predictions were proven wrong (e.g. "...HUH...??!!"). In addition, some partners would occasionally mimic actions of the other for comic effect, often resulting in shared laughter (e.g. task related actions such as placing pots in mirroring locations, or non-task related actions such as stacking papers). Secondly, on the one hand we observed examples of friendly courtesy towards each other (e.g. "[handing over a pot] if you want, you can place this one" or "you may remove the blocks if you like"). On the other hand, we saw signs of frustration about their partner's actions, often expressed as deep sighs or rolling their eyes. Finally, we observed cases where partners engaged in mutual "mischief" outside of the scope of the task, where they would reinforce each other's mischievous behaviours.

The behaviours that we identified on the relational level suggest that children worked towards establishing and maintaining a positive relationship. Mutual shared experiences (both positive and negative) contribute to a stronger social bond, which in turn can strengthen the feeling of being "allies in learning". Although we did observe instances of frustration, overall the children acted mutually benevolent towards each other. Additionally, we anecdotally identified behaviours that could be associated with building rapport. For example, children mirrored or mimicked each other's behaviours, gestures, and facial expressions, suggesting a level of interpersonal synchrony.

Table 3 Frequency of inquiry process labels

Inquiry process	Frequency	Percentage
Preparation	316	37.4
Hypothesis generation	94	11.1
Experimentation	196	23.2
Observation	150	17.7
Conclusion	90	10.6
Total	846	100

5.4 Inquiry Processes

In identifying the inquiry processes children engaged in, a total of 846 annotations were performed. Frequencies for individual labels are shown in Table 3.

These results verify that children go through all phases of the inquiry learning process when working with the balance task. However, we found an imbalance between the frequency of observations in the various phases: children more frequently engaged in preparation, experimentation, and observation behaviours. In fact, within a single task, children would often complete multiple quick, successive cycles of preparing and experimenting, skipping the hypothesis generation and conclusion phases. Preparation, experimentation, and observation phases were composed primarily of actions and verbalisations related to the practical or physical nature of the task (e.g. placing/removing pots and blocks, or remarking on the tilt of the balance). On the contrary, behaviours that were associated with the hypothesis generation and conclusion phases were more cognitive and reflective in nature and focus on constructing and organising knowledge (e.g. discussing or explaining about how to interpret the experiment outcome, or combining prior knowledge and new knowledge to construct an updated hypothesis).

5.5 Conclusion of the Contextual Analysis

In our analysis we have identified a rich collection of behaviours aligned with our ECR model that children exhibit in practice, in the context of our specific learning assignment. Furthermore, we found that children went through all phases of inquiry and displayed learning behaviours appropriate to each phase. These results show firstly, that our inquiry learning assignment is suitable for conducting research on collaborative interactions. Secondly, they show that our assignment is usable for the children in our target user group.

Based on these findings from the contextual analysis we would like to identify specific opportunities for when and how a social robot may make a meaningful contribution to the learning process. In the following section we present concrete design guidelines that focus on specifying potential roles and behaviours for the robot, as well as opportunities in the learning cycle where such a robot may potentially intervene.

6 Design Guidelines for Educational Robots

In this section we present four general Recommendations and fourteen concrete Design Guidelines (*DG 1-14*) for developing social robot behaviours, derived from our findings in the literature and our contextual analysis, and we offer concrete examples of possible robot behaviours and phrases that illustrate how each design guideline may materialise in practice.

Recommendation 1—Designing Educational Behaviours: The Robot Should Offer Adequate Content Support and Process Support

Learners who construct knowledge [55] through the process of inquiry learning [40] require a certain degree of scaffolding to achieve their next level of cognitive development [74]. In the observations from our contextual analysis we found many instances where children scaffolded each other's learning process. In some instances, we see behaviours that don't contribute much to learning, such as directly prompting answers or even blatantly copying answers. However, in most cases scaffolding takes place through explanations from one child to the other.

DG 1: The Robot Could Use the Physical Learning Materials to Supplement Its Scaffolding

From our observations of children interacting we found that, when explaining, they naturally used the physical learning materials as tools or props to supplement their explanations. A robot could take advantage of similar techniques when scaffolding the child's learning.

Both the robot and the learning materials are physically embodied in the environment of the child, which can be quite naturally leveraged to create more coherent scaffolding behaviour. By referring to objects or aspects of the learning materials, the robot can further supplement its explanations or increase the legibility of its instructions. Such supplementary didactic behaviours may take several forms depending on the robot's embodiment and physical capabilities.

For example, the robot may use nonverbal behaviours like gazing, gesturing, or pointing to reference locations, objects, or components of the learning materials. Such behaviours can be used while the robot is explaining or reading out loud, or when objects are referenced by the child or assignment. Furthermore, robots with fine motor skills may physically manipulate learning materials, to offer an illustration (e.g.

the robot may pick up a red pot and a yellow pot and say "this one seems heavier, what do you think?").

DG 2: A Robot Should Explicitly Take on a Certain Role

Based on our observations we found that an individual child might be more knowledgeable in one aspect of the task while they are less knowledgeable in another. Furthermore, we saw that children naturally and effortlessly switched between such roles as they progressed through the learning tasks together. In the current analysis we have not investigated in detail how children establish and regulate their respective roles of more knowledgeable and less knowledgeable team member—a process which seems to occur implicitly. However, we do see opportunities for a robot to support the learning process by explicitly navigating a certain role.

For a robot to effectively help children learn, the child must know how they can benefit from the robot's presence. In other words, the types and forms of support (e.g. actions, information, knowledge) that the robot will be able to offer them will depend on the role that the robot takes. Since we are unsure how the robot may navigate between roles implicitly, we suggest that the robot should have explicit means of presenting itself in a certain role, in order for children to more easily appraise the support that they may potentially receive from the robot.

For example, a robot could explicitly introduce itself in a certain way or the designer could provide the robot with a certain background story. On the one hand, a robot navigating the role of a less- or differently knowledgeable other, could for instance, introduce itself: e.g. "Hi, my name is [..]. I'm interested in learning about topic X, would you like to investigate it together?" or could be introduced by the designer with a background story: e.g. "This is [..], an exchange student from the robot school. Like you, it is studying topic X, maybe you can study together." On the other hand, a robot navigating the role of a more knowledgeable other, could for instance, introduce itself: e.g. "I have already learned something similar to this topic, maybe I can help you with it." or the designer could describe the background story of the robot: e.g. "The robot has practiced this task a lot of times before, so you can ask it for instructions." The choice for which role fits best for the robot will likely depend upon how the designer wishes to present and introduce the robot to the child, as well as the looks of the robot, the type of task, the learning context, and the research interests.

DG 3: A Robot Operating in the Role of Differently Knowledgeable or Less Knowledgeable Other Could Leverage the Interactive Explanation Paradigm

In our observations we often saw educational interactions by the more knowledgeable child being triggered by something that the less knowledgeable child said or did (e.g. making mistakes, looking confused, or asking). The more knowledgeable child would then use their expertise to provide support to the other child (e.g. correcting, explaining, or thinking/reading aloud) .

In such situations a peer-like robot can offer social facilitation (i.e. audience effect) through which it can leverage the interactive-explanation or the learning-by-teaching paradigm [56] by prompting task-related questions to provoke a (verbal) response. This may trigger the child to (re)organise specific elements of their knowledge to offer a coherent explanation.

For example, in some cases, the robot may ask questions that require the child to reflect on the process of the learning assignment (e.g. "What did we do in the last step?" or "What does the task say about the next step?" or "What was our initial hypothesis, was it correct or not?" or more generically, "What happened to the balance?"). In other cases, the robot may prompt for specific knowledge that children should have previously acquired, or highlight information that they need to answer the current task (e.g. "Can you remember which of the pots was heaviest?" or "Why is that pot further away from the centre than the other pot?" or "In the last exercise the balance tipped to the left, so why does it stay balanced now?" or "What do you think would happen if we moved one of the pots closer to the pivot?" or more generically, "Can you explain why the balance behaves like this?").

DG 4: A Robot Operating in the Role of a More Knowledgeable Other Could Offer Tutoring to the Learner

Our observations revealed many instances where children needed process and content support to understand the learning assignment. In many cases such support was offered by the more knowledgeable child. The experimenter would offer additional tutoring when both children were unable to continue with the task.

In such cases, a tutor robot may offer the child directions, advice, and explanations regarding the learning content and learning process. It may help them in specific areas where they are struggling or where they require additional information. Depending on the robot's reasoning capabilities it could either attempt to automatically recognise when a child requires tutoring or could respond to a child's explicit request for help. In some cases it would be sufficient for the robot to read the instructions out loud and in other cases the robot may need to employ more advanced forms of tutoring.

For example, regarding the inquiry learning process, the robot may offer additional background information about the steps of inquiry (e.g. "In the hypothesis step we must think about what we expect will happen" or "When giving a conclusion, try to think about whether your hypothesis was correct or not" or simply reading the assignment out loud).

Regarding learning content the robot may offer, for example, analogies that the child might be familiar with (e.g. "This balance is like a seesaw on the playground. What happens if a mouse and an elephant sit on either side of a seesaw?"), or it might refer to prior knowledge (e.g. "Remember, we saw last time that the red pot is heavier than yellow"), or might explain underlying conceptual principles (e.g. "When a pot is further away from the pivot it puts more force on that side of the balance").

Recommendation 2—Designing Collaborational Behaviours: The Robot and Child Should Engage in Joint Activities

Effective collaborational interactions require collaborators to work towards aligning their goals, objectives, strategies, abilities, and actions [18,41]. According to Klein and Feltovich [41] successful joint activities are achieved (among other things) through an interplay between the following five aspects: (a) an intention to collaborate; (b) a level of interdependence between collaborators; (c) common ground; (d) directability of one's actions by their partner; and (e) effective signalling behaviour. In our observations we saw each of these aspects materialised in one form or another in the children's behaviours and interactions. A social robot aimed at collaborating with a child should be designed to also navigate these aspects of joint activities. In the next five design guidelines we discuss each of Klein and Feltovich's requirements and translate them into more specific robot behaviours.

DG 5: To Express an Intention for Collaboration, the Robot Should Disclose Its Own Goals and Should Discuss Task Objectives, Strategies, and Joint Actions

Klein and Feltovich [41] describe the intention to collaborate as a mostly implicit and unspoken agreement between all team members. From our observations, however, we found examples where children agreed upon goals and strategies, and examples where children discussed intentions to perform individual actions. By doing so they explicitly expressed their intention to collaborate, in addition to any implicit unspoken agreement that might have preexisted.

Similarly, we suggest that a robot should express such intentions for collaborating. Since we did not investigate unspoken intentions, we can only offer suggestions for explicit robot behaviours. It is important to note that educational robots are usually created and programmed to accomplish a certain objective (e.g. to improve the child's learning in some form or another). Such an objective is not necessarily known to the child. However, to signal to the child that the robot has a benevolent intention to collaborate, it could disclose a different "personal goal" (e.g. with an introduction or a background story). While working on the learning assignment with the child, the robot should continue to express its intention for collaboration by discussing task objectives, strategies, and actions. The robot may take various approaches, depending on the type and content of the assignment, the learning context, and the role of the robot in relation to the child.

For example, a more knowledgeable or tutoring robot may introduce its personal goal to the child as: "teaching subject X". During the rest of the interaction, that robot could then act as a leader, tutor, or expert: reading the instructions out loud to the child (e.g. "The task says we should do: X..Y..and..Z"), referring explicitly to the task instructions when they are violated (e.g. "I think we're doing X wrong"), and proclaiming progress when the instructions are followed (e.g. "We finished X, now we do Y to move on to the next step"). However, a less knowledgeable or differently knowledgeable robot may tell the child that its personal goal is: "wanting to learn about subject X". That robot could then act like a peer-learner during the remainder of the interaction: asking the child to read the assignment out loud, asking explicitly whether the current step meets all task instructions (e.g. "Did we do everything that the task says?"), ask explicitly about the next actions (e.g. "What should we do next?"), and inquire about progression through tasks (e.g. "How far are we? When can we move on to the next step?").

DG 6: To Support a Level of Interdependence, the Robot and Child Should Facilitate Each Other's Actions

Interdependent activities, according to Klein and Feltovich [41], are interwoven actions between collaborators; their actions mutually influence and affect each other. Our learning assignment inherently imposed a certain level of interdependence. Since both children had to share the same physical learning materials, actions from the one, naturally influenced the other. In our observations we mostly saw examples of children working together (e.g. assisting each other with placing pots). However, we also saw instances of children working in separation (e.g. each doing "their own task" in sequence) and even obstructing each other (e.g. taking away pots that their partner had just placed).

A robot that aims to collaborate with children should aim to engage in interdependent interactions. Such interdependence in collaborative learning assignments may be introduced in different ways. Firstly, the designer of the learning assignment may choose to impose explicit interdependencies between the child's and the robot's actions. Secondly, the robot and child may assist each other's actions even if this is not imposed explicitly by the structure of the learning assignment.

For example, the assignment might require two different actions to be performed at the same time, or might require a sequence of actions (e.g. first the robot must complete action

A, then the child must complete action B, only then can the robot complete action C, and so on). In such situations the robot may either take on a dominant role (e.g. "It's your turn to do action X") or a submissive role (e.g. "What should I do next?"). In other cases, the robot may ask the child for assistance when it is incapable of performing a certain action (e.g. asking the child for an object that is outside the robot's reach). Furthermore, depending on the robot's physical capabilities it may also attempt to assist the child's actions (e.g. handing over/pushing objects to the child or manipulating other aspects of the learning materials).

DG 7: The Robot Should Build and Maintain Common Ground with the Child in the Context of the Learning Assignment

Klein and Feltovich [41] describe common ground between collaborators as the shared mutual knowledge, beliefs, and assumptions, which continuously evolve during interactions. In general, a strong common ground can result in more efficient communication and collaboration during joint activities. We speculate that, among other things, our children's implicit mutual gaze, attentive gaze, and shared focus of attention contributed to their common ground. Additionally, we saw instances where children sought explicit alignment by discussing task instructions and actions.

A strong common ground in child-robot interactions may contribute to an increased legibility and predictability of the robot, enabling the child to more accurately anticipate and infer the robot's future actions. In turn, this increased predictability may enable the child to align their own actions and react to the robot accordingly. Children from the same class might have a level of established common ground, allowing more implicit forms of alignment. However, we expect that not to be the case during initial child-robot interactions: the child may not be aware of the robot's knowledge, beliefs, and assumptions. Therefore, the robot should use explicit verbal statements and questions so that the child can build common ground, in addition to implicitly contributing to common ground through gaze behaviours.

For example, the robot should firstly seek explicit alignment on various aspects of the learning assignment: a) alignment on the learning objectives (e.g. "I think the idea of the assignment is to investigate topic A" or "What do you think topic A is about?'); b) the steps to take towards achieving these learning objectives (e.g. "To investigate topic A we need to do complete tasks X, Y, and Z" or "What information from task X do we need to use in task Y?"); c) the actions that are required within each step (e.g. "In this step we must first place 'object 1' here and then 'object 2' there"); and d) the references to physical objects and locations (e.g. "Is *this* the 'object' that we have to use in this task?" or "*There* is the 'location' where we have to put the object"). Secondly, depending on the available sensors and actuators, the robot

may attempt to emulate the gaze behaviours that we observed with our children: a) emulate paying attention by reactively gazing towards the child's actions; b) emulate mutual gaze by reactively gazing towards the child when they gaze towards the robot; and c) emulate shared focus of attention by gazing towards where the child gazes or points.

DG 8: The Robot's Actions Should be Directable by the Child and Should be Communicated as Such

According to Klein and Feltovich [41] directability of each partner's actions plays a key role in successful interdependent activities. Through mutual directability of their actions, collaborators signal to each other that their behaviours affect one another in a meaningful way. In our observations of children, we saw instances where children would mutually direct their actions, or where one child would overrule the actions of the other.

We imagine it can be challenging to design robot behaviours in such a way that they are *truely* directable by the child throughout an entire interaction. To exhibit *some* level of directability we suggest that designers focus on implementing directable actions that are isolated to specific aspects or moments of an interaction.

For example, asking the child for help (e.g. "Can you tell/show me how to do this?" or "What do I have to do next?"), explicitly asking for directions (e.g. "Which object would you like me to use?" or "Where would you like me to place this object?"), inviting the child to correct the robot's actions (e.g. "Am I doing this right?" or "Which answer should I pick?"), or synchronising timing of actions (e.g. "Shall I do this now?" or "Let's count down together!"). In addition to such explicit actions the robot could use implicit social cues to signal that its behaviours are directable. For instance by displaying hesitation or waiting for mutual eye contact before executing an action, or by allowing the child to interrupt the robot's speech and behaviours.

DG 9: The Robot Should Use Appropriate Signalling Behaviours for the Coordination of Joint Activities

Finally, Klein and Feltovich [41] state that to effectively plan and execute joint actions collaborators must clearly signal their own intentions, goals, and actions. In addition they must signal their recognition and interpretation of their *partner's* intentions, goals, and actions. Clear and legible signalling is particularly important in situations where common ground is weak (e.g. during first familiarisation), where something in the joint task has gone wrong (e.g. misunderstood task instructions), or when partners' goals or actions have become misaligned (e.g. working on different sub-tasks).

When using digital learning environments or sensorised physical learning materials, the system is capable of reliably

recognising the user's actions, enabling the robot to detect and respond to deviations from task instructions. However, reliable automatic detection, recognition, and interpretation of an interaction partner's underlying intentions and goals is quite challenging, which makes it all the more important for the robot to use appropriate signalling behaviours towards the child. The robot's signalling behaviour can be composed of explicit verbal statements and implicit non-verbal behaviours.

For example, in some cases, explicit verbal statements can be aimed at communicating the robot's goals and actions (e.g. "I am working on 'goal A' now" or "I think we should first focus on achieving 'goal A', then move onto 'goal B' afterwards" or "I'm going to do 'X' now" or "I'd like to do 'X', can you help me by doing 'Y'?"). In other cases, the robot may explicitly ask the child to clarify their goals and actions (e.g. "Are you working towards 'goal B'?" or "Are you going to do 'Y' next?"). Such explicit verbal statements of intention can be supported through implicit non-verbal behaviours. For example, through attentive gaze, mutual gaze, and shared focus of attention the robot signals to the child that it is "following", "paying attention to", and attempting to "recognise" their actions. Furthermore, through proactive behaviours and actions (e.g. by proactively handing over objects to the child, or by reciting information relevant to the task at hand) the robot signals that it is has interpreted the child's intentions and is anticipating their next action accordingly. Additionally, such proactive behaviours may be signals of the robot's benevolent intentions.

Recommendation 3—Designing Relational Behaviours: The Robot Should Work Towards Building and Maintaining a Positive Social Relationship with the Child

Effective collaborative learning requires a positive social relationship between learning partners. Partners need to judge each other's interests, abilities, skills, and knowledge, and must trust each other to act benevolently. Learning partners can take advantage of their social relationship to offer each other positive emotional support when faced with new experiences or setbacks.

The children who participated in our contextual analysis were from the same class and had already established interpersonal social relationships. Although we did not observe the formation of new social relationships, we did observe instances where children engaged in meaningful relational interactions. For example, children had mutual emotional experiences in response to successes and failures, they displayed prosocial behaviours towards each other, and they engaged in entertaining activities outside of the learning assignment. Designers of learning interventions using robots and other social agents, can take inspiration from such observations. Previous research has shown that children are prone to build relationships with robots (e.g. [6,33,35,65]). Fostering such positive child-robot relationships can indeed have beneficial effects for the child (e.g. [4,36,66]). The following design guidelines are suggestions on how to implement relational behaviours on an educational social robot.

DG 10: The Robot Could Display Emotional Responses to Learning Events

Children often display an emotional response when encountering learning events such as successes and failures. By designing the robot to display similar emotional responses it can attempt to emulate a mutual experience with the child.

For example, a basic version of such a robot may use prescripted responses to successes and achievements (e.g. "Yeah! We got it right!" or "Aha! I knew it!" or simply cheering and/or smiling). More advanced technological implementations may attempt to analyse the child's facial expressions, body poses, gestures, and affective speech to offer a tailored response, more accurately emulating a mutual emotional experience. If the intervention and the context of the learning assignment are appropriate, the robot may also emulate mutual experiences during failures (e.g. "Hrmpf.. Now I have to start all over again!" or "Grrr! I was so sure my answer would be correct!"). However, such a negative response may have a detrimental effect on the child's learning experience and motivation.

DG 11: The Robot Could Offer Emotional Support

In addition to mutual experiences, there are several opportunities for the robot to offer more explicit emotional support when children encounter frustrating setbacks that may leave them demotivated.

For example, firstly, while children work on a task, the robot may offer friendly motivational support, emphasising their task-related effort (e.g. "I see you're really doing your best!" or "You must have thought really hard on that answer!"). Secondly, the robot could focus on offering positive feedback when children experience successes (e.g. "Well done! Onto the next challenge!") as well as when they encounter failures (e.g. "That was a good effort, just a bit more practice and I'm sure we will make it!" or "I don't mind making mistakes, as I always learn something new!"). Finally, as children advance through various levels, assignments, or difficulties, the robot could offer praise regarding their progress (e.g. "Great! Thanks to your hard work we got a new highscore!" or "We really nailed that task, well done!" or "I see you learned a lot, let's try an even more difficult challenge next!").

DG 12: The Robot Could Show Prosocial Behaviour

In our contextual analysis we found instances of reciprocal, prosocial, benevolent behaviours between co-learners. In cases where it contributes to the learning experience, a designer may choose to translate specific instances of such prosocial behaviours to the robot.

For example, firstly, the robot could exhibit proactive behaviours that contribute to task progress (e.g. the robot offers help at appropriate moments, the robot hands over objects, fills out answers, or reads relevant task instructions). Secondly, the robot may offer the child the opportunity to perform the most amusing, comical, or entertaining actions themselves. What these actions are depends to a large extent on the task context; for our inquiry learning assignment children seemed to mostly enjoy manipulating the physical objects when preparing (e.g. placing pots on the balance) and initiating the actual experiment (e.g. removing wooden blocks from under the balance). Finally, on occasion the robot may directly offer the correct answer when the child is stuck (e.g. "Pssst..I think the answer should be 'A'!"), thus enabling them to continue with the remainder of the assignment. However, it is likely that directly prompting such answers doesn't contribute much to the child's knowledge. Therefore, we advise that the robot only uses this particular relational prosocial behaviour if the answer at hand is not directly related to the main learning goals of the intervention. Otherwise, if the answer *is* directly related to the learning goals, the robot should offer educational scaffolding instead.

DG 13: The Robot Could Engage with the Child in Activities Beyond Learning

We observed many positive relational interactions (e.g. mutual smiling and laughing) when children engaged in non-learning activities (e.g. neatly stacking papers or participating in mischief). More generally, past research has shown that engaging in off-topic activities with a social agent, such as social dialogues and smalltalk, can contribute to a positive relationship [10,11].

While we do not necessarily recommend that the robot promotes mischief, it still might be beneficial to the social relationship if the child and robot engage in non-learning activities. In practice, such activities would depend on the child's personal interests and the robot's (conversational) capabilities.

Examples of activities that children might like to do with a robot could include drawing and colouring, dancing, storytelling, or playing games.

Recommendation 4—The Robot Should Stimulate the Cognitive Processes of Inquiry Learning

In effective inquiry learning, students alternate between processes that are practical in nature (e.g. the process of preparing, experimenting, and observing) and cognitive in nature (e.g. the process of hypothesising and concluding) [39]. Learning interventions should focus on promoting an optimal balance between both types of processes. However, in our observations we found children's practical actions to be disproportionally more prevalent than behaviours associated with cognitive processes. Designers of learning interventions may address this by offering scaffolding to trigger such cognitive behaviours.

DG 14: The Robot Could Attempt to Elicit Explanations from the Child

Cognitive thought processes take place internally in the mind of the learner and can be difficult to observe externally. However, we observed that children would verbally express aspects of their reasoning when explaining their hypotheses and conclusions to each other, offering us a glimpse into their thought processes.

To promote more of such cognitive processes the robot could attempt to elicit explanations from the child. Such an approach could be similar to the interactive explanation paradigm [56] outlined above in DG 3.[3]

For example, the robot could ask the child to further clarify their hypotheses (e.g. "Why do you think that X will happen?"), explain the state of the learning materials (e.g. "Can you explain why the balance tilted this way?"), or formulate underlying physics principles in their own words (e.g. "How do you know which pot is heavier?" or "What happens when pots move closer/further from the pivot?").

7 Implementing the Guidelines to Facilitate Social Learning in CRI

To briefly summarise, the main goal of the work presented so far was to generate guidelines for designing socially appropriate educational robot behaviours. In the remainder of this paper we present our Computer Aided Learning (CAL) system that uses a social robot to facilitate the learning process of individual primary school children.

In the following sections we first briefly introduce the technical setup of our CAL system. Then, we discuss four variants of our CAL system that were implemented using

---

[3] DG 3: A robot operating in the role of differently knowledgeable or less knowledgeable other could leverage the interactive explanation paradigm.

Table 4 An overview of the specific design guidelines (DGs) that have been implemented in each of our four CAL system variants. For each variant we list the primary design guidelines (denoted with a black check mark) and supplementary design guidelines (denoted with a grey check mark) that were followed during the design and implementation of robot's behaviour and the learning assignment. Some design guidelines currently remain without implementation

	Implementation variants			
	I	II	III	IV
DG 1: The robot could use the physical learning materials to supplement its scaffolding	(✓)	(✓)		(✓)
DG 2: A robot should explicitly take on a certain role	✓	(✓)	✓	
DG 3: A robot operating in the role of differently knowledgeable or less knowledgeable other could leverage the interactive explanation paradigm	✓	(✓)		
DG 4: A robot operating in the role of a more knowledgeable other could offer tutoring to the learner	✓	(✓)	✓	(✓)
DG 5: To express an intention for collaboration, the robot should disclose its own goals and should discuss task objectives, strategies, and joint actions				
DG 6: To support a level of interdependence, the robot and child should facilitate each other's actions				
DG 7: The robot should build and maintain common ground with the child in the context of the learning assignment	(✓)	(✓)		(✓)
DG 8: The robot's actions should be directable by the child and should be communicated as such				
DG 9: The robot should use appropriate signalling behaviours for the coordination of joint activities				(✓)
DG 10: The robot could display emotional responses to learning events	(✓)	(✓)		(✓)
DG 11: The robot could offer emotional support	(✓)	(✓)	✓	✓
DG 12: The robot could show prosocial behaviour				
DG 13: The robot could engage with the child in activities beyond learning		✓		
DG 14: The robot could attempt to elicit explanations from the child	✓	(✓)		

our recommendations and design guidelines and we reflect on how children interacted (socially) with these variants during empirical studies. This will allow us to take a first step towards showing the potential value of our design guidelines for producing robots that are capable of supporting children's inquiry learning. Table 4 shows the specific design guidelines implemented in each variant of our CAL system.

7.1 Technical Setup

A platform for implementing the recommendations and design guidelines outlined in the previous section requires several (technical and pedagogical) components to be in place. Firstly, there must be an inquiry learning assignment and embodied learning materials that are suitable for the target user group. Secondly, there must be pedagogically appropriate exercises for the assignment. Thirdly, there must be a dialogue that can be made socially and educationally aware (i.e. using sensors for the child and task materials). Finally, there must be appropriate feedback and scaffolding to be offered to the child.

The modular architecture of the EASEL project was used as a foundation for our base CAL system [57]. This base CAL system setup consists of the following components: (1) dialogue manager; (2) sensors; (3) embodied learning materials; (4) social robot; and (5) tablet user interface. Figure 3
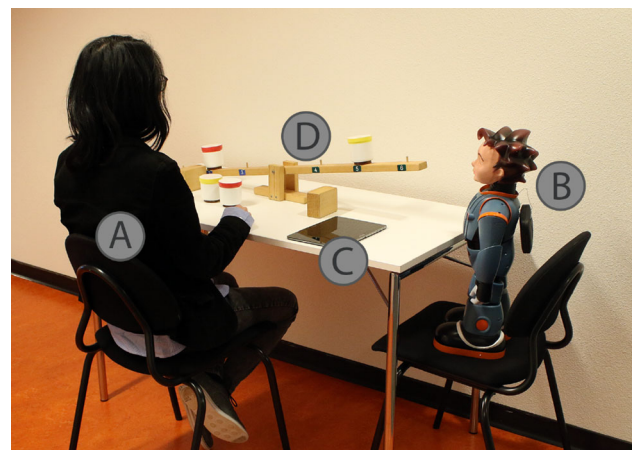

Fig. 3 A typical setup of our CAL system: (A) a child, (B) a robot, (C) a tablet device, and (D) embodied learning materials

depicts a typical setup of our base CAL system (note that computer hardware is hidden from the user where possible).

The dialogue manager forms the heart of our CAL system that ties together the various components. It is responsible for controlling the flow of the interaction between child, robot, and learning materials, and enables our CAL system to be responsive to the child. The dialogue manager uses sensors (or Wizard of Oz input) to follow the child's interactions with the robot and their progression through the learning assignment. This information is then used to determine an appropriate response from the CAL system according to the rules defined in our dialogue models. We use an Information State Update-based dialogue manager: dialogue models are specified in an early version of Flipper 2.0 [73] as a collection of rules and effects that are matched against an information state.

Data from external sensors is collected via network connections and incorporated into the information state. Such sensors are used to collect data about how the user interacts with the system. On the one hand, they measure social aspects of the child-robot interaction (e.g. proxemics). On the other hand, sensors can detect the child's actions with the embodied learning materials.

The embodied learning materials are based on existing inquiry learning principles and involve the exploration of physics phenomena. For example, we have built balance tasks to enable children to explore the moment of force and ramp tasks with which children explore principles of rolling resistance and potential energy. In early versions we used Wizard of Oz techniques to emulate detection of user interactions with the learning materials. Later versions of our CAL system used embedded sensors to automatically detect the state of the learning materials and the child's actions [22].

Our CAL system uses RoboKind's R25 Zeno.[4] robot to offer social support to the learner. Zeno's face can show simple facial expressions (such as smiling, gasping, crying, and frowning) and it can perform rudimentary gaze shifts. The robot can use either prerecorded speech phrases or a Text-To-Speech (TTS) system, both of which include a form of lip-synchronisation. Unfortunately, although Zeno's arms can be used for rough, imprecise movements, they do not have sufficient degrees of freedom for accurate pointing or gesturing. The robot's behaviours are specified in Behaviour Markup Language (BML) [42] which is translated by ASAPRealizer [58] into robot-specific control primitives.

Finally, the CAL system uses a tablet interface to display task instructions. The tablet can display text, images, and buttons. Children may provide their answers to tasks through the tablet. Their answers are then passed to the dialogue manager, enabling the robot to respond accordingly.

The CAL system base setup was used in four variants as described in the following sections. Each variant focused on implementing specific aspects of the educational interaction, for which we designed specific tasks and robot behaviours according to our design guidelines. For each variant we first briefly describe the underlying rationale and highlight the relevant design guidelines. Then, we present how the design guidelines influenced the implementation of the specific tasks and behaviours. Finally, where applicable, we reflect on insights from semi-structured interviews and on our anecdotal observations of educational, collaborative, and relational behaviours exhibited by children interacting with the robot and learning assignment.

7.2 Variant I: A Robot that Stimulates Explanatory Behaviours

From our observations of how children progressed through our inquiry learning assignment we found that children more frequently displayed practical actions than behaviours associated with cognitive learning processes. Therefore, we set out to implement a CAL system to stimulate the cognitive inquiry learning processes of hypothesising and concluding.[5] Our main rationale was to design a social robot that offered social facilitation to leverage the interactive explanation paradigm [56].

The following primary design choices were implemented to support this paradigm. We explicitly presented the robot to take on a certain role in relation to the child.[6] On the one hand, the robot acted as a less knowledgeable other[7] regarding the principles of balance: (1) the robot was given a background story and was introduced as being curious to investigate the concept of balance together with child (i.e. the robot offered no explicit content scaffolding); (2) it asked open questions at appropriate times during the inquiry phases in which the child formulates a hypothesis or draws a conclusion, inviting the child to give an explanation[8]; and (3) it gazed towards the child when asking questions to strengthen the social facilitation (audience effect), inviting the child to give more elaborate and deeper explanations. On the other hand, the robot acted as a more knowledgeable other[9] regarding the inquiry process: (1) as part of the robot's background story it was introduced as having experience in researching

---

[5] Recommendation 4—The robot should stimulate the cognitive processes of inquiry learning.

[6] DG 2: A robot should explicitly take on a certain role.

[7] DG 3: A robot operating in the role of differently knowledgeable or less knowledgeable other could leverage the interactive explanation paradigm.

[8] DG 14: The robot could attempt to elicit explanations from the child.

[9] DG 4: A robot operating in the role of a more knowledgeable other could offer tutoring to the learner.

physics phenomena; and (2) it could offer more elaborate explanations regarding the steps of the inquiry cycle (i.e. the robot offered explicit process scaffolding).

In addition, the following supplementary design choices were implemented throughout the interaction. Firstly, the robot emulated attentive gaze by following the child's interactions with the tablet and learning materials.[10] Secondly, it emulated mutual gaze and displayed emotions at particular moments in the interaction.[11] For example, just when the balance tipped to either side during the experiment phase, the robot would gaze towards the child and smile. Thirdly, the robot used didactic gaze to emphasise aspects of the learning materials.[12] For example, it would gaze towards the tablet as new instructions appeared and would briefly gaze towards the balance when explaining or reading task instructions aloud. Finally, the robot offered supportive encouragement and feedback as children progressed through the assignment.[13] For example, it praised the child's progress (e.g. "Well done!") and it always offered positive feedback without judging the child's performance (e.g. "That's interesting!" when hearing their explanations, regardless of whether they were correct).

The main study was conducted with children of 5–10 years old and its results are reported elsewhere as Wijnen et al. [75]. Here we reflect on the types of behaviours that we anecdotally observed, and discuss them in light of our ECR model.

On the educational level, we observed children giving rich explanations to the robot. This is interesting as it lends potential credence to a less knowledgeable robot as social facilitator to support the interactive explanation paradigm. Additionally, we found indications that children might have appreciated the process support of the robot; they would often request additional help and explanations from the robot when they didn't understand particular steps of the task, and younger children (who might not yet be proficient readers) often asked the robot to read tasks out loud. These observations were further supported by our exit interviews, during which some children described how the robot helped them during the tasks while others mentioned that they explained the tasks to the robot.

On the collaborational level, we anecdotally observed instances where children asked the robot to give its own input, thoughts, and interpretations to the task (e.g. "What is your opinion about this, robot?"). Additionally, we found that some children seemed to dislike the robot's repetitive

instructions. Due to its prescripted dialogue models, the robot would repeat the same instructions across different tasks when scaffolding the inquiry learning process. However, some children may have been annoyed by this repetition, as they stated that they already knew the order of the steps after the first task. Others seemingly took offence that the robot would "warn them" about following task instructions even though they had never once deviated from the task and always followed the instructions meticulously. More adaptive dialogue models could take the child's past actions into account when offering process support. For example, by only issuing elaborate instructions and reminders when first encountering new steps, or when children repeatedly skip certain steps.

7.3 Variant II: A Robot that Engages in Extracurricular Activities

In this variant we were interested in supporting interactions on the relational level.[14] From our observations of child-child interactions we found that children displayed mutually positive social emotions when they engaged in extracurricular activities. To stimulate such behaviours when interacting with the robot we implemented an additional activity besides the learning assignment.

This variant built upon the same robot, interaction, and learning assignment as in Variant I; as a result, it addressed the same design guidelines. The primary design effort in this variant, however, was focused on developing an additional intermediate extracurricular activity[15]. During this (hopefully) fun and engaging activity, child and robot attempted to make the same dance-like movements with their arms and legs. This shared experience was expected to influence how children perceive the robot on a social level. The robot would encourage the child to do their best, but would also mention that it is unable to perform all the movements itself.

This exploratory study was conducted with children of 5–11 years old. The main results are reported elsewhere in [21]. Here, we briefly reflect on our anecdotal observations of behaviours on the relational level. Firstly, we found indications that children seemed to enjoy the extracurricular activity with the robot, as we annotated more smiling during and after the activity. Secondly, through exploratory semi-structured interviews and picture association tasks we attempted to gather insights regarding how the children perceived the robot on a social level. These highly speculative

---

[10] DG 7: The robot should build and maintain common ground with the child in the context of the learning assignment.

[11] DG 10: The robot could display emotional responses to learning events.

[12] DG 1: The robot could use the physical learning materials to supplement its scaffolding.

[13] DG 11: The robot could offer emotional support.

[14] Recommendation 3—Designing relational behaviours: The robot should work towards building and maintaining a positive social relationship with the child.

[15] DG 13: The robot could engage with the child in activities beyond learning.

results seem to suggest that children considered the robot to be a social agent of sorts, of a similar age to themselves. In their descriptions of the robot most children mentioned that they thought it could be their friend, and they would describe hobbies, activities, and games they would like to do with the robot. Most children indicated they would like to sit next to the robot in class and that they would feel comfortable to entrust the robot with a secret. Finally, some children indicated that they were amused that they were able to do things that the robot could not. Others mentioned they were impressed by the robot's overall skill, in both the learning assignment and the dancing activity.

7.4 Variant III: A Robot that Automatically Offers Appropriate Scaffolding

Variant I and Variant II used prescribed fixed-pattern scaffolds that were triggered based on Wizard of Oz input. However, to support a fully autonomous interaction, the robot must be able to recognise at which point in time it should offer which type of support to the child. Therefore, we focused on automatically generating appropriate content and process scaffolding based on sensors in the learning materials.[16]

The learning materials were replaced with sensorised versions that used simple embedded sensors and an overhead camera to automatically measure the tilt of the balance (left/right/centered) and the location of pots (which weight on which pin) and supporting blocks (present/removed). The simple sensors captured the state of the materials regarding these aspects with an accuracy of 96%. The system used a screen to display tasks and children entered their answers by pressing physical buttons. The Zeno robot was replaced with a simplified static robot embodiment (a plastic teddy bear) to encourage children to focus primarily on its scaffolding behaviours as opposed to its social behaviours and appearance. Based on sensor data and the child's button presses, the robot would recognise which step of the inquiry cycle children were in. It would then automatically offer appropriate scaffolding and feedback. The robot was explicitly presented to the child as a more knowledgeable other that would help them complete the assignment.[17] Within that role it offered two forms of scaffolding.[18] Firstly, it offered process scaffolding by strictly monitoring and enforcing the task structure, and by issuing reminders if children skipped steps or forgot task instructions. Secondly, it offered content scaffolding by providing simplified analogies to the children when they were unable to formulate hypotheses or conclusions.

Additionally, it offered motivating, supportive, and friendly feedback while children progressed through the task.[19]

The system was tested with children of 5–10 years old. The design of the sensorised learning materials and corresponding scaffolding behaviour was presented as part of Davison et al. [22]. Here we present the relevant results from semi-structured interviews.

Regarding the educational level, children mentioned that the robot's support helped them judge the outcome of the tasks. They also mentioned that the robot helped them concentrate on the task and offered them useful instructions, explanations, and feedback.

Regarding the relational level, children praised the robot's niceness and his motivational and social support. Furthermore, we conducted an exploratory content analysis of their descriptions of the robot. We speculate that the robot's positive feedback and encouragement gave children a sense of achievement and success while progressing through challenges.

7.5 Variant IV: A Robot for Longitudinal in-the-Wild Deployment

Using the sensorised learning materials from Variant III we designed a learning intervention that could be deployed in an unsupervised manner in classrooms over extended periods of time. For this purpose, we focused on designing dialogues and behaviours that were robust enough for the robot to autonomously interact with the child without intervention or assistance from a researcher.

This variant was achieved by combining the Zeno robot's social repertoire from Variant I with the autonomous sensors and scaffolding from Variant III. The primary aim was to offer tailored emotional support and encouragement to the child as they progressed through several assignments of increasing difficulty.[20]

Most supplementary design choices were adopted from Variant I and were modified to work autonomously where necessary. Additionally, to signal that it could see and interpret what the child was doing, the robot would gaze reactively towards the child's actions as they manipulated the learning materials.[21] Finally, we made the system more robust so that it could be used unsupervised, in a real classroom setting. On the one hand, we simplified somewhat the interaction repertoire of the robot. For instance, it no longer asked for verbal explanations, but instead asked children to select their hypotheses, observations, and conclusions from a multiple-choice list on a tablet interface. On the other hand, we added

---

[16] Recommendation 1—Designing educational behaviours: The robot should offer adequate content support and process support.

[17] DG 2: A robot should explicitly take on a certain role.

[18] DG 4: A robot operating in the role of a more knowledgeable other could offer tutoring to the learner.

[19] DG 11: The robot could offer emotional support.

[20] DG 11: The robot could offer emotional support.

[21] DG 9: The robot should use appropriate signalling behaviours for the coordination of joint activities.

more variety to the exercises and gave children more freedom in how they could go through them. For instance, children could skip or repeat certain steps of the inquiry cycle as they preferred. In such cases the system reminded them about the proper order but did not necessarily prevent them from continuing.

This variant was deployed during a 4-month study in a classroom with 6–10 year old children. The results of this are still being analysed; we hope that they will show that the CAL system that we built up in the first three variants contributed to a varied, socially rich, and engaging interaction between the robot and the participating children.

8 Limitations and Future Work

In our contextual analysis we observed 22 children learning together in pairs in their usual school environment. Due to the relatively low sample size we may not have gathered an exhaustive collection of all possible educational, collaborational, and relational interactions. However, we did observe sufficiently rich and diverse behaviours for the purpose of specifying our set of relevant design guidelines.

Since the children participating in our contextual analysis all came from a single Dutch school we cannot easily generalise our results to different cultures. Moreover, our school followed the Montessori educational system, in which children are quite accustomed to collaborating together on embodied learning tasks. In particular, older, more knowledgeable children are regularly paired up with their younger, less knowledgeable classmates. This likely means that we may have observed different behaviours if we would have conducted our contextual analysis in schools following different educational systems. However, the implemented variants of our CAL system were tested with children from a wider selection of schools and daycare centres. We did not observe any behaviours with these children that would point to our design guidelines being biased towards the Montessori system.

The design guidelines presented in this paper followed from observations of 6–9 year old children. Children at either end of that age group are likely to have differences in their skills and cognitive development. Our guidelines do not specifically consider such differences between the younger and older children in an absolute sense. Instead, they offer suggestions from a relative point of view; in terms of a more knowledgeable child with a less knowledgeable partner or vice versa. Furthermore, we cannot generalise our results to children, students, or adult learners that fall outside our age group. In fact, we speculate that such more advanced learners working on more complex and challenging assignments would likely benefit from different forms of (social) scaffolding. Therefore, designers of learning interventions may

wish to adjust the difficulty of the task or the form of scaffolding to better cater to the needs of children in their target age group.

It is likely that the context and nature of our inquiry learning assignment has inherently shaped the interactions that we observed between children. This type of assignment affords a set of social learning behaviours, that is specific to inquiry learning contexts. We can therefore not generalise our recommendations and design guidelines to other types of learning paradigms.

While children interacted with our robots we anecdotally observed several types of behaviours that seem to align with the levels of our ECR model. Such anecdotal observations do not necessarily validate our model or our design guidelines. However, we speculate that these anecdotal observations do seem to suggest that children display signs of social learning while interacting with robots that were implemented according to our design guidelines.

In future work we suggest that designers continue to explore the applicability of our suggested guidelines in other educational interaction contexts. In particular, we see opportunities to address aspects of the ECR interactions that we have not yet investigated. Designing behaviours for more diverse types of robots may open up opportunities for other types of interactions. Our humanoid robot Zeno, for example, was unable to physically manipulate learning materials. Most of the collaborational interactions that we designed were therefore limited to gaze-shifts and speech. A robot with a different embodiment could be designed to make use of more intricate forms of collaborational behaviours. For example, a robot capable of manipulating and handing over physical objects (e.g. [32]) may be designed to facilitate joint actions in an interdependent task[22] and show prosocial actions towards collaborators.[23] Additionally, in several of our design guidelines we made suggestions related to the background story of the robot. Although we did not explore this aspect in detail, we found that children would occasionally make references to the background story of the robot or would use their imagination to describe such a story. For example, after playing with Variant II of our system (where children did an extracurricular dancing activity with the robot) they would imagine all sorts of other activities that the robot could do. We imagine that a robot could self-disclose its own background story as the interaction progresses (e.g. see the ALIZ-E project [70]) and may leverage this to disclose its own experiences and goals to the child.[24]

---

[22] DG 6: To support a level of interdependence, the robot and child should facilitate each other's actions.

[23] DG 12: The robot could show prosocial behaviour.

[24] DG 5: To express an intention for collaboration, the robot should disclose its own goals and should discuss task objectives, strategies, and joint actions.

In future work it could be interesting to investigate how such a background story may influence children's perceptions of the robot and how it could shape their collaboration efforts. Finally, in our various CAL system variants we created dialogue models that would generate appropriate feedback and scaffolding. However, our models were not yet fully adaptive to individual children's needs and preferences. During interviews we found indications that some children would have liked the robot's help to be less repetitive and preachy. To address this, future dialogue models could offer more personalised scaffolding. For example, by toning down the amount of instructions when a child performs well on a certain repetitive task.

9 Conclusion

This paper presented a contextual analysis from which we derived recommendations and design guidelines for designing educational social robot interventions.

First, we discussed underlying pedagogical literature and constructed the ECR model as a lens to look at three types of social learning interactions: (1) Educational interactions, (2) Collaborational interactions, and (3) Relational interactions, and we discussed related work on educational HRI.

Following this, we conducted a contextual analysis. We observed 11 pairs of real learners in a school while they were working on our specific collaborative inquiry learning assignment and learning materials. The video recordings of these sessions were annotated regarding the three ECR types of social interaction. These annotations were analysed to answer our first research question, "*Which typical behaviours and social interactions do we identify in learners working with our assignment and how do these behaviours align with theory?*". We found that our children often behaved as more-or less knowledgeable partners relative to each other while they worked together. In those roles they would scaffold each other's learning process by thinking aloud, by eliciting or offering explanations, and by agreeing or disagreeing with each other. Furthermore, we found that children coordinated their joint actions by taking turns, discussing goals and plans, and aligning their progress. When doing so, children would often signal their intentions verbally or through gaze and gestures. Misalignment of actions, goals, or task progress were often explicitly corrected verbally or nonverbally by either child. As they progressed through the assignment, the children showed signs of social bonding by displaying both positive and negative mutual emotions. These observed behaviours were aligned along the three levels of our ECR model.

We also analysed the same data regarding the physical and cognitive behaviours that children expressed in the various phases of the inquiry learning cycle, to answer the second research question, "*How do children progress through the various phases of our learning assignment and do the types of learning behaviour they exhibit vary with these phases?*". These analyses showed that our learners displayed a rich repertoire of behaviours and social interactions across all three types, but they were predominantly engaged in activities of a practical nature (i.e. preparation, experimentation, and observation), whereas the more cognitive activities (i.e. hypothesis generation and conclusion) were less pronounced.

Then, we looked at the third research question, "*What are opportunities for social robot interventions in regard to interaction patterns, specific behaviours and roles, to contribute to the child's learning in our assignment?*". To answer this, we used the above results to derive four generic high-level recommendations and fourteen concrete design guidelines for *when* and *how* a social robot may have a meaningful contribution to the learning process. These design guidelines can be used as a foundation for developing social robots that can navigate the various ECR types of interaction with the purpose of supporting a child's learning process.

Finally, we discussed four implemented variants of a Computer Aided Learning (CAL) system in which we translated our design guidelines into concrete robot behaviours. We briefly reflected on anecdotal behaviours that we observed from children interacting with the various variants of our CAL system. In doing so, we have taken the first steps towards showing the potential value of our design guidelines for producing robots capable of supporting children's inquiry learning.

Compliance with Ethical Standards

References

1. Banchi H, Bell R (2008) The many levels of inquiry. Sci Child 46(2):26–29
2. Baraglia J, Cakmak M, Nagai Y, Rao R, Asada M (2016) Initiative in robot assistance during collaborative task execution. In: ACM/IEEE international conference on human-robot interaction HRI'16, pp 67–74
3. Bell RL, Smetana L, Binns I (2005) Simplifying inquiry instruction. Sci Teach 72(7):30–33
4. Belpaeme T, Baxter PE, Read R, Wood R, Cuayáhuitl H, Kiefer B, Racioppa S, Kruijff-Korbayová I, Athanasopoulos G, Enescu V, Looije R, Neerincx M, Demiris Y, Ros-Espinoza R, Beck A,

Cañamero L, Hiolle A, Lewis M, Baroni I, Nalin M, Cosi P, Paci G, Tesser F, Sommavilla G, Humbert R (2012) Multimodal child–robot interaction: building social bonds. J Hum Robot Interact 1(2):33–53

5. Benitti FBV (2012) Exploring the educational potential of robotics in schools: a systematic review. Comput Educ 58(3):978–988

6. Beran TN, Ramirez-serrano A (2011) Can children have a relationship with a robot? In: Lamers MH, Verbeek FJ (eds) Human–Robot personal relationships. Lecture notes of the institute for computer sciences, social informatics and telecommunications engineering, vol 59. Springer, Berlin, pp 49–56

7. Beran TN, Ramirez-Serrano A, Kuzyk R, Nugent S, Fior M (2010) Would children help a robot in need? Int J Soc Robot 3(1):83–93

8. Bevan N, Macleod M (1994) Usability measurement in context. Behav Inf Technol 13(1–2):132–145

9. Beyer H, Holtzblatt K (1998) Contextual design: defining customer-centered systems, 1st edn. Morgan Kaufmann, Burlington

10. Bickmore T, Cassell J (2005) Social dialogue with embodied conversational agents. In: Kuppevelt JCJ, Dybkjær L, Bernsen NO (eds) Advances in natural multimodal dialogue systems. Springer, Dordrecht, pp 23–54

11. Bickmore TW, Picard RW (2005) Establishing and maintaining long-term human–computer relationships. ACM Trans Comput Hum Interact 12(2):293–327

12. Cameron D, Fernando S, Collins E, Millings A, Moore R, Sharkey A, Prescott T (2016) Impact of robot responsiveness and adult involvement on children's social behaviours in human-robot interaction. In: 5th international symposium on new frontiers in human–robot interaction

13. Chandra S, Alves-Oliveira P, Lemaignan S, Sequeira P, Paiva A, Dillenbourg P (2015) Can a child feel responsible for another in the presence of a robot in a collaborative learning activity? In: 2015 24th IEEE international symposium on robot and human interactive communication (RO-MAN), pp 167–172

14. Chang CW, Lee JH, Chao PY, Wang CY, Chen Gd (2010) Exploring the possibility of using humanoid robots as instructional tools for teaching a second language in primary school. Educ Technol Soc 13(2):13–24

15. Charisi V, Davison D, Wijnen F, van der Meij J, Reidsma D, Prescott T, van Joolingen W, Evers V (2015) Towards a child–robot symbiotic co-development: a theoretical approach. In: New frontiers in human–robot interaction. AISB

16. Chase CC, Chin DB, Oppezzo MA, Schwartz DL (2009) Teachable agents and the protégé effect: increasing the effort towards learning. J Sci Educ Technol 18(4):334–352

17. Chi MT, Bassok MM, Lewis MW, Reimann P, Glaser R (1989) Self-explanations: how students study and use examples in learning to solve problems. Cogn Sci 13(2):145–182

18. Clark H (1996) Using language. Cambridge University Press, Cambridge

19. Coninx A, Baxter P, Oleari E, Bellini S, Bierman B, Blanson Henkemans O, Cañamero L, Cosi P, Enescu V, Ros Espinoza R, Hiolle A, Humbert R, Kiefer B, Kruijff-Korbayovà I, Looije R, Mosconi M, Neerincx M, Paci G, Patsis G, Pozzi C, Sacchitelli F, Sahli H, Sanna A, Sommavilla G, Tesser F, Demiris Y, Belpaeme T (2015) Towards long-term social child–robot interaction: using multi-activity switching to engage young users. J Hum Robot Interact 5(1):32

20. Corsaro WA (1997) The sociology of childhood. Pine Forge Press, London

21. Davison D, Schindler L, Reidsma D (2016) Physical extracurricular activities in educational child–robot interaction. In: 5th international symposium on new frontiers in human–robot interaction (AISB-NFHRI)

22. Davison DP, Charisi V, Wijnen FM, Papenmeier A, van der Meij J, Reidsma D, Evers V (2016) Design challenges for long-term interaction with a robot in a science classroom. In: Proceedings of the RO-MAN2016 workshop on long-term child–robot Interaction. In: IEEE robotics and automation society

23. de Greeff J, Belpaeme T (2015) Why robots should be social: enhancing machine learning through social human–robot interaction. PLoS ONE 10(9):e0138061

24. Dillenbourg P (1999) What do you mean by collaborative learning? Collab Learn Cogn Comput Approach 1:1–19

25. Dolcos F, Iordan AD, Dolcos S (2011) Neural correlates of emotion–cognition interactions: a review of evidence from brain imaging investigations. J Cogn Psychol (Hove, England) 23(6):669–694

26. Fredrickson BL (2001) The role of positive emotions in positive psychology. The broaden-and-build theory of positive emotions. Am Psychol 56(3):218–226

27. Gordon G, Spaulding S, Westlund JK, Lee JJ, Plummer L, Martinez M, Das M, Breazeal C (2016) Affective Personalization of a Social Robot Tutor for Children's Second Language Skills. Thirtieth AAAI Conference on Artificial Intelligence

28. Hoffman G, Breazeal C (2009) Effects of anticipatory perceptual simulation on practiced human–robot tasks. Auton Robots 28(4):403–423

29. Hood D, Lemaignan S, Dillenbourg P (2015) When children teach a robot to write: an autonomous teachable humanoid which uses simulated handwriting. In: International conference for human–robot interaction HRI'15, pp 83–90

30. Inhelder B, Piaget J (1958) The growth of logical thinking from childhood to adolescence: an essay on the construction of formal operational structures, vol 84. Basic Books, New York

31. Janssen JB, van der Wal CC, Neerincx MA, Looije R (2011) Motivating children to learn arithmetic with an adaptive robot game. In: Mutlu B, Bartneck C, Ham J, Evers V, Kanda T (eds) ICSR 2011. Lecture notes in computer science, Lecture notes in computer science. Springer, Berlin, vol 7072, pp 153–162

32. Jung MF, DiFranzo D, Stoll B, Shen S, Lawrence A, Claure H (2018) Robot assisted tower construction—a resource distribution task to study human–robot collaboration and interaction with groups of people. arXiv:1812.09548 [cs.RO]

33. Kahn PH, Kanda T, Ishiguro H, Freier NG, Severson RL, Gill BT, Ruckert JH, Shen S, Kahn PH Jr (2012) Robovie, you'll have to go into the closet now: children's social and moral relationships with a humanoid robot. Dev Psychol 48(2):303–314

34. Kanda T, Hirano T, Eaton D, Ishiguro H (2004) Interactive robots as social partners and peer tutors for children: a field trial. Hum Comput Interact 19(1):61–84

35. Kanda T, Sato R, Saiwaki N, Ishiguro H (2007) A two-month field trial in an elementary school for long-term human–robot interaction. IEEE Trans Robot 23(5):962–971

36. Kanda T, Shimada M, Koizumi S (2012) Children learning with a social robot. In: Proceedings of the seventh annual ACM/IEEE international conference on human–robot interaction—HRI'12, ACM Press, New York, New York, USA

37. Kennedy J, Baxter P, Belpaeme T (2015) The Robot Who Tried Too Hard. In: Proceedings of the tenth annual ACM/IEEE international conference on human–robot interaction—HRI'15. ACM Press, New York, New York, USA, pp 67–74

38. Kjeldskov J, Graham C (2003) A review of mobile HCI research methods. In: Chittaro L (ed) Human–computer interaction with mobile devices and services. Springer, Berlin, pp 317–335

39. Klahr D (2000) Exploring science: the cognition and development of discovery processes. The MIT Press, Cambridge

40. Klahr D, Dunbar K (1988) Dual space search during scientific reasoning. Cogn Sci 12(1):1–48

41. Klein G, Feltovich PJ, Bradshaw JM, Woods DD (2005) Common ground and coordination in joint activity. In: Rouse WB, Boff KR (eds) Organizational simulation. Wiley, pp 139–184

42. Kopp S, Krenn B, Marsella S, Marshall AN, Pelachaud C, Pirker H, Thórisson KR, Vilhjálmsson H (2006) Towards a common framework for multimodal generation: the behavior markup language. In: International conference on intelligent virtual agents, IVA 2006, vol 4133. Springer, Berlin, pp 205–217

43. Kort B, Reilly R, Picard RW (2001) An affective model of interplay between emotions and learning: Reengineering educational pedagogy-building a learning companion. In: Proceedings—IEEE international conference on advanced learning technologies, ICALT 2001. IEEE Computer Society, pp 43–46

44. Kory J, Breazeal C (2014) Storytelling with robots: Learning companions for preschool children's language development. In: The 23rd IEEE international symposium on robot and human interactive communication, pp 643–648

45. Kyza EA (2009) Middle-school students' reasoning about alternative hypotheses in a scaffolded, software-based inquiry investigation. Cogn Instr 27(4):277–311

46. Lazonder AW (2014) Inquiry learning. In: Spector JM, Merrill MD, Elen J, Bishop MJ (eds) Handbook of research on educational communications and technology. Springer, New York, pp 453–464

47. Lemaignan S, Jacq A, Hood D, Garcia F, Paiva A, Dillenbourg P (2016) Learning by teaching a robot: the case of handwriting. IEEE Robot Autom Mag 23(2):56–66

48. Leyzberg D, Spaulding S, Toneva M, Scassellati B (2012) The physical presence of a robot tutor increases cognitive learning gains. In: Proceedings of the annual meeting of the cognitive science society, vol 34

49. Maguire M (2001) Methods to support human-centred design. Int J Hum Comput Stud 55(4):587–634

50. Millen DR (2000) Rapid ethnography: time deepening strategies for HCI field research. In: Proceedings of the conference on designing interactive systems processes, practices, methods, and techniques—DIS'00. ACM Press, New York, New York, USA, pp 280–286

51. Morse JM (1994) Emerging from the data: the cognitive processes of analysis in qualitative inquiry. In: Critical issues in qualitative research methods, pp 23–45

52. Mubin O, Stevens CJ, Shahid S, Mahmud AA, Dong JJ (2013) A review of the applicability of robots in education. Technol Educ Learn 1(1):1–7

53. Okita SY, Ng-Thow-Hing V, Sarvadevabhatla R (2009) Learning together: ASIMO developing an interactive learning partnership with children. In: RO-MAN 2009—the 18th IEEE international symposium on robot and human interactive communication, pp 1125–1130

54. Pellegrini A (2009) The role of play in human development. Oxford University Press, Oxford

55. Piaget J, Inhelder B (1969) The psychology of the child. Basic Books, New York

56. Ploetzner R, Dillenbourg P, Preier M, Traum D (1999) Learning by explaining to oneself and to others. In: Collaborative-learning: cognitive and computational approaches, pp 103–121

57. Reidsma D, Charisi V, Davison D, Wijnen F, van der Meij J, Evers V, Cameron D, Fernando S, Moore R, Prescott T, Mazzei D, Pieroni M, Cominelli L, Garofalo R, de Rossi D, Vouloutsi V, Zucca R, Grechuta K, Blancas M, Verschure P (2016) THE EASEL project: towards educational human–robot symbiotic interaction. In: Lecture notes in computer science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol 9793. Springer, Cham, pp 297–306

58. Reidsma D, van Welbergen H (2013) AsapRealizer in practice—a modular and extensible architecture for a BML realizer. Entertain Comput 4(3):157–169

59. Resnick M, Rosenbaum E (2013) Designing for tinkerability. In: Honey M, Kanter D (eds) Design, make, play: growing the next generation of STEM innovators. Routledge, New York, pp 163–181

60. Rogoff B (1998) Cognition as a collaborative process. In: Handbook of child psychology, chap. Cognition. Wiley, Hoboken, NJ, USA, pp 679–744

61. Rohrbeck CA, Ginsburg-Block MD, Fantuzzo JW, Miller TR (2003) Peer-assisted learning interventions with elementary school students: a meta-analytic review. J Educ Psychol 95(2):240–257

62. Saerbeck M, Schut T, Bartneck C, Janse MD (2010) Expressive robots in education. In: Proceedings of the 28th international conference on Human factors in computing systems—CHI'10, ACM Press, New York, New York, USA

63. Saldaña J (2014) The coding manual for qualitative researchers. 1. Sage, Thousand Oaks

64. Shin N, Kim S (2007) Learning about, from, and with robots: students' Perspectives. In: RO-MAN 2007—the 16th IEEE international symposium on robot and human interactive communication, pp 1040–1045

65. Tanaka F, Cicourel A, Movellan JR (2007) Socialization between toddlers and robots at an early childhood education center. Proc Natl Acad Sci USA 104(46):17954–17958

66. Tanaka F, Matsuzoe S (2012) Children teach a care-receiving robot to promote their learning: field experiments in a classroom for vocabulary learning. J Hum Robot Interact 1(1):78–95

67. Teasley SD (1995) The role of talk in children's peer collaborations. Dev Psychol 31(2):207–220

68. Topping KJ (2005) Trends in peer learning. Educ Psychol 25(6):631–645

69. van Dijk AM, Gijlers H, Weinberger A (2014) Scripted collaborative drawing in elementary science education. Instr Sci 42(3):353–372

70. van der Drift EJG, Beun R-J, Looije R, Blanson Henkemans OA, Neerincx MA (2014) A remote social robot to motivate and support diabetic children in keeping a diary. In: 2014 9th ACM/IEEE international conference on human–robot interaction (HRI), pp 463–470

71. van Joolingen WR, de Jong T (1997) An extended dual search space model of scientific discovery learning. Instr Sci 25(5):307–346

72. van Joolingen WR, de Jong T, Dimitrakopoulou A (2007) Issues in computer supported inquiry learning in science. J Comput Assisted Learn 23(2):111–119

73. van Waterschoot J, Bruijnes M, Flokstra J, Reidsma D, Davison D, Theune M, Heylen D (2018) Flipper 2.0: a pragmatic dialogue engine for embodied conversational agents. In: Proceedings of intelligent virtual agents (IVA 2018), ACM Press, Sydney, Australia

74. Vygotsky L (1978) Mind in society: the development of higher psychological processes. Harvard University Press, Cambridge

75. Wijnen FM, Davison DP, Reidsma D, van der Meij J, Charisi V, Evers V (2018) Now we're talking: Learning by explaining your reasoning to a social robot (submitted)

76. Wood D, Bruner JS, Ross G (1976) The role of tutoring in problem solving. J Child Psychol Psychiatry 17(2):89–100

77. You Z-J, Shen C-Y, Chang C-W, Liu B-J, Chen G-D (2006) A robot as a teaching assistant in an english class. In: Sixth IEEE international conference on advanced learning technologies (ICALT'06), pp 87–91

78. Zaga C, Lohse M, Truong KP, Evers V (2015) The effect of a robot's social character on children's task engagement: Peer versus tutor. In: Social robotics. Lecture Notes in Computer Science, LNCS vol 9388. Springer, Cham, pp 704–713

Daniel P. Davison has a Bachelor degree in Computer Science and a Master degree in Human Media Interaction from the University of Twente. As a PhD Candidate at the Human Media interaction group at the University of Twente, his research focused on educational social robots for primary school children. In the context of the EASEL project he was particularly interested in designing instrumented learning materials, expressive robot behaviours and rich interactive dialogues to optimally support a child's inquiry learning process. Daniel is currently involved in the DE-ENIGMA project, where he develops interactive game-like activities to help autistic children practice facial expressions and emotions with a social robot.

Frances M. Wijnen graduated in 2014 for her master Instruction, Learning and Development (Psychology) at the University of Twente. After her graduation she worked as a Junior Researcher at the department for Teacher Education at the University of Twente. She was involved in the project: Expressive Agents for Symbiotic Education and Learning (EASEL) in which she explored the interaction between children and a social robot in the context of inquiry learning. Frances currently works as a PhD-student on the project: Technology as a tool for 21st century learning. In this project she explores which factors stimulate and/or obstruct primary school teachers to use technology for stimulating higher-order thinking skills (such as: creative thinking, critical thinking, and problem solving) with children.

Jan van der Meij since March 2019 he is director of vocational education at Het Erasmus in Almelo. Until March 2019, he was programme director of the University of Twente teacher education, ELAN, department of teacher professional development. After finishing a teacher training in electronics (1994), graduating with a master's degree in educational technology from the University of Twente (1997), and working at the educational software design department of the national centre for the innovation of education and training (1997–2001), he received his PhD at the University of Twente in 2007 on supporting learning with multiple representations in simulation-based learning environments. From 2006–2011 he was staff member of the department of instructional technology and from 2011–2019 staff member of the UT teacher education.

Dennis Reidsma is Assistant Professor at the Human Media Interaction group and Lecturer at the Creative Technology curriculum at the University of Twente. After obtaining his PhD in 2008 on reliability and validity aspects of large annotated multimodal corpora for machine learning, he worked with Dr. Herwin van Welbergen on behaviour generation for coordinated interaction with social agents. Besides a number of high impact publications, this led to the release of AsapRealizer (formerly Elckerlyc), a state-of-the-art Open Source software platform for generating continuous interaction with Virtual Humans, social agents and robots. Currently he uses this platform to investigate the potential impact of social robots in the classroom.

Vanessa Evers is a full Professor of Human Media Interaction at the University of Twente. Her research focuses on Human Interaction with Autonomous Agents such as robots or machine learning systems and cultural aspects of Human Computer Interaction. She specifically likes to take theories on human behaviour from social psychology and see if similar processes occur when we interact with technology. She is best known for her work on social robotics such as the FROG robot (fun robotic outdoor guide), that can interpret human behaviour automatically and respond to people in a socially acceptable way. She received a M.SC. in Information Systems from the University of Amsterdam, and a Ph.D. from the Open University, UK. During her Master studies she spent two years at the Institute of Management Information Studies of the University of New South Wales, Sydney. After her Ph.D. she has worked for the Boston Consulting Group, London and later became an assistant professor at the University of Amsterdam's Institute of Informatics. She was a visiting researcher at Stanford University (2005–2007). She has published over 80 peer reviewed publications, many of which in high quality journals and conferences in human computer interaction and human robot interaction. She serves on Program Committees of HRI, CHI, CSCW and ACM Multimedia.