Fostering Relatedness Between Children and Virtual Agents Through Reciprocal Self-disclosure

Franziska Burger[1(✉)], Joost Broekens[1], and Mark A. Neerincx[1,2]

[1] Delft University of Technology, Delft, Netherlands
{f.v.burger,d.j.broekens,m.a.neerincx}@tudelft.nl
[2] TNO, Soesterberg, Netherlands

Abstract. A key challenge in developing companion agents for children is keeping them interested after novelty effects wear off. Self Determination Theory posits that motivation is sustained if the human feels related to another human. According to Social Penetration Theory, relatedness can be established through the reciprocal disclosure of information about the self. Inspired by these social psychology theories, we developed a disclosure dialog module to study the self-disclosing behavior of children in response to that of a virtual agent. The module was integrated into a mobile application with avatar presence for diabetic children and subsequently used by 11 children in an exploratory field study over the course of approximately two weeks at home. The number of disclosures that children made to the avatar during the study indicated the relatedness they felt towards the agent at the end of the study. While all children showed a decline in their usage over time, more related children used the application more, and more consistently than less related children. Avatar disclosures of lower intimacy were reciprocated more than avatar disclosures of higher intimacy. Girls reciprocated disclosures more frequently. No relationship was found between the intimacy level of agent disclosures and child disclosures. Particularly the last finding contradicts prior child-peer interaction research and should therefore be further examined in confirmatory research.

1 Introduction

Type 1 diabetes mellitus (T1DM) is an autoimmune disease of the pancreas that requires manual control of blood glucose levels. Strict adherence to a medical regimen is crucial to prevent many of the health risks associated with this chronic disease. T1DM accompanies diagnosed children and adolescents through various physical and mental stages of development. We develop a **P**ersonal **A**ssistant for a healthy **L**ifestyle (PAL[1]) with the aim of increasing the self-management skills of diabetic children (ages 7–14) by supporting them, their caregivers, and healthcare professionals in sharing responsibility. The PAL robot and its mobile avatar

[1] http://www.pal4u.eu/.

T. Bosse and B. Bredeweg (Eds.): BNAIC 2016, CCIS 765, pp. 137–154, 2017.
DOI: 10.1007/978-3-319-67468-1_10

are companion agents intended to function as a pal for the children, helping them to accomplish their diabetes-related goals through person- and time-adaptive, engaging interactions.

Since children are susceptible to novelty effects, and the PAL solution can only be effective when children continue to engage with the system, ways of sustaining their motivation are highly desirable. We are interested in exploring the possibilities and limitations of creating a bond between diabetic child (8–12 years) and the PAL virtual companion agent through self-disclosure with the goal of increasing the motivational capacity of the agent. To this end, we conducted a two-week exploratory field study in which 11 children were given the opportunity to interact with the self-disclosing mobile PAL avatar by either disclosing as well or by simply listening.

2 Theoretical Foundation

Companion agents are developed for long-term use. A key challenge in the field is thus the maintenance of motivation when novelty effects wear off.

Social relationships often play a large motivational role in our behaviors. According to Self Determination Theory (SDT), successful establishment of a social bond between human and agent leads to sustained motivation both to interact with the agent and to engage in activities that the agent proposes. SDT [7] argues that the basic psychological needs for *autonomy*, *competence*, and *relatedness* must be satisfied by the social environment for humans to feel motivated to attempt a task. Relatedness here refers to the feeling that one is accepted and cherished by another individual or community. It comes into play when the intrinsic motivation to engage in an activity is low. More simply put: if we like or want to be liked by someone, we feel more inclined to do what they suggest, even if we are not too fond of the activity itself.

An important mechanism by which such a bond could be established is described by Social Penetration Theory (SPT) [1]. It proposes a directional development of interpersonal relationships whereby the involved individuals first share and explore each others personalities at a superficial level before disclosing more intimate information. Disclosing proceeds along two dimensions: breadth and depth, with *breadth* describing the number of different topics that are disclosed about and *depth* describing the personal value these topics have. Finally, an important determinant of self-disclosure is reciprocity. This describes the tendency to self-disclose as a result of being disclosed to. Reciprocal disclosures in successfully progressing relationships are usually on a similar level of intimacy.

3 Related Literature

One of the key interests in human-human self-disclosure research has been the close link between disclosure and liking. Specifically, three persistent disclosure-liking relations have been identified [6]: (a) the more someone intimately discloses to us, the more we like that person, (b) the more we like someone at the outset

of the interaction, the more we will disclose, and (c) the more intimately we disclose to someone, the more we like that person.

When children were asked what a friend is and what differentiates a friend from a non-friend, children older than nine indicated that friends take an interest in each others problems and care for their friends' emotional well-being. Additionally, cooperation and the insight that each child should contribute equally to the interaction can be expected in this age group [15]. In line with this, in a study conducted in the United States, it was found that 6th grade children's liking of another child was influenced by that child's ability to match the intimacy level of a disclosure while that of 4th graders was not [13].

Support for the disclosure-liking effect has also been found in the domains of human-computer (HCI) and child-robot (cHRI) interaction. In [11], a computer first disclosed some information about itself before asking the user (all university undergraduates) an interview question. As hypothesized, interviewees shared more intimate information with the computer that told personal information about itself but only if this personal information would gradually increase in intimacy throughout the interview. However, the liking for the computer only depended on the sharing of personal information and was not influenced by the intimacy strategy. When a robot was used to elicit self-disclosures from children (aged 10–14), those who were prompted to disclose to the robot described the robot significantly more often as a *friend* than children in the control condition [10]. In [9], a two-month study was conducted in an elementary school with a relational robot capable of identifying children (aged 10–11) and calling them by name, showing more varied behavior with time, and disclosing personal information as a function of a child's interaction time. It was found that children's desire to be friends with the robot at the end of the study was positively correlated with the interaction time.

To the best of our knowledge, there has been no empirical investigation of whether and how the sharing of disclosures between user and system contributes to sustaining user motivation over longer periods of time. The here described research presents a first step in closing this knowledge gap. We developed the initial prototype of a dyadic disclosure dialog module (3DM, Sect. 4) to gain insights into how and how readily diabetic children respond to self-disclosures of an embodied conversation agent (ECA) and to learn about the possibilities of sustaining children's motivation in this way. We were particularly interested in the following research questions, summarized by the relationship *disclosure* → *relatedness* → *motivation*:

RQ1 Can the relatedness that the child feels towards the avatar be predicted from (1) the amount of disclosures that the child heard from the avatar, (2) the amount of disclosures the child made to the avatar, and (3) the relatedness the child felt at the outset of the study?

RQ2 Is relatedness a good indicator for children's motivation to use the application?

Furthermore, we were interested in learning about how children respond to a self-disclosing avatar. Studies on self-disclosure reciprocity in child-child

interaction have been conducted mainly in North America several decades ago (compare [5,12,13]). It was therefore uncertain whether insights transfer to today's children in Europe or to child-robot interaction. We thus also pursued the following research questions:

RQ3 How do children respond to the disclosures of the avatar?
- (a) Is there a relationship between the avatar intimacy and children's responsiveness?
- (b) Do children match the intimacy level of the avatar disclosure when they respond?
- (c) What role do age and gender of the children play in how children respond to the avatar?

A situated approach was taken by integrating the module (described in the following section) into a mobile application for diabetic children to be used in an uncontrolled environment for a period of two weeks. In so doing, we found that while children did not match the intimacy of disclosures from the ECA, those children who replied more actively to the disclosures also felt more related to the avatar. Furthermore, children were more likely to reciprocate a disclosure when it was of lower intimacy or when the child was a girl.

4 Dyadic Disclosure Dialog Module

The first prototype of the dyadic disclosure dialog module (3DM) was developed. While it is the ultimate goal of the module to manage the sharing of personal information between agent and child in an adaptive and engaging manner, the first prototype, developed for this research project, only served the purpose of exploring the disclosure behavior of the children when interacting with a self-disclosing ECA. For this, content that the ECA could disclose was needed. Below we therefore briefly touch on the steps taken to develop the disclosure database and its structure. This is followed by a description of how the module functions and how it is integrated into application of the PAL project.

4.1 Content of 3DM

To design suitable disclosures for the ECA, three preliminary steps had to be taken. First, a personality for the avatar was crafted by determining sensible traits for the given domain (e.g. the ECA should be conscientious because this is important in diabetes self-management and we would like the ECA to provide positive examples of self-discipline for the children). The Murphy-Meisgeier Type Indicator for Children[2] was employed for finding a suitable type to integrate these initial traits into one coherent personality. Second, a background story was written for the robot by the lead researcher from which consistent disclosures at various intimacy levels could be derived. Here, the goal was to obtain a story that is both in line with the fact that robots are not human and in line with a character that children can embrace[3]. Third, a scaling method (rating scale) to

[2] https://www.capt.org/.
[3] http://latd.tv/Latitude-Robots-at-School-Findings.pdf.

design agent disclosure statements at various intimacy levels and to assess the depth of children's disclosures was developed [4].

The current disclosure database consists of approximately 150 English disclosures for the avatar at all four intimacy levels of the rating scale. These were written by the lead author taking into consideration the personality of the robot and its background story. In designing the rating scale, a selection of the statements were evaluated as a set by ten participants with regard to believability and consistency ($mean = 4.3$ on 5-point Likert-scale) with the designed personality and background story. The disclosure statements are organized into the four categories *food*, *school*, *social*, and *sports*. They also have valence labels so as to be matched to the child's affective state if available. Since two of the partner hospitals of the PAL project are in the Netherlands and the study was carried out with Dutch children, all disclosures have Dutch translations. Translations were first made by the lead author, a non-native speaker. They were then double checked and edited by two native Dutch speakers, including the second author. Finally, the Dutch translations (as these were to be used in the experiment) were investigated for their age appropriateness by the developmental psychologist, also a native Dutch speaker, involved in the PAL project. The disclosures are stored as instances of the Disclosure class—a class in the associated ontology described below.

4.2 Functionality of 3DM

Within the PAL-project, knowledge is represented in ontologies. Specifically for the module, a small ontology was therefore made. There are three main classes in the ontology for 3DM: Disclosure, Prompt, and Closer. These correspond to the three types of statements that 3DM relies on. All disclosures have the parameters intimacy level, valence, and topic. Agent disclosures additionally have an associated prompt. Prompts are said by the agent to elicit a disclosure from the child. Closers are used to end the off-activity chat and return to the main activity: a *positive* closer is said when the child chooses to disclose something, a *negative* closer is said otherwise[4]. Since the module is not yet capable of comprehending a child's disclosure, closers are very general statements that make no reference to the disclosure content. The ontology is specified in RDF[5].

The flow of the disclosure module follows a loop. From the perspective of the user this proceeds as illustrated in Fig. 1. While inactive, 3DM waits for a trigger event from the interface. When it receives this, it selects a disclosure and sends it with a gesture to the avatar for rendering. Upon execution, it follows up with the prompt. The interface then provides a pop-up asking the child whether it would like to respond. If the child chooses not to (*passive interaction*), a

[4] It is important to note that *positive* and *negative* here are not synonymous with rewarding or punishing the child. An example for a positive discloser can be found in the example dialog. An example for a negative closer is "That's alright. Maybe next time! In any case, thanks for listening.".

[5] https://www.w3.org/.

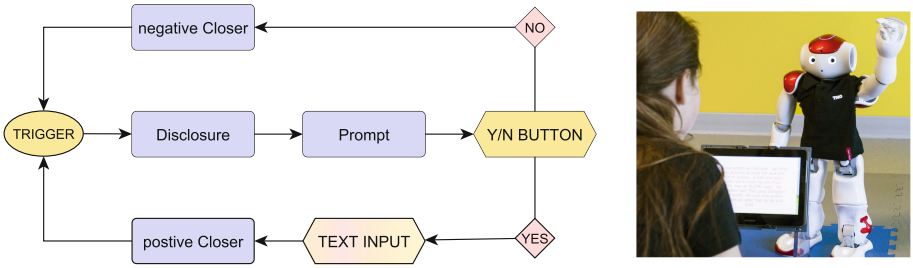

Fig. 1. *Left.* Illustration of the 3DM functionality. Interface actions are hexagonal, agent actions are rectangular, and child actions are diamonds. The trigger event (opening of the diabetes diary) has a circular shape. *Right.* A diabetic child interacts with the PAL robot. Photo courtesy of Rifca Peters.

negative closer command is sent to the avatar. If the child wants to respond (*active interaction*), it can do this in a second pop-up that allows it to type some text. Once the module has received the text, it sends a positive closer command to the avatar. It then simply waits for the next trigger event. In the first prototype, the trigger event was chosen to be the opening of the diabetes diary area of the app. Both closer sentences and prompt sentences contain a placeholder for using the name of the child. It is randomly decided whether to use the name in the prompt, in the closer, or not at all.
An example dialog of the agent (**A**) with a fictional child (**C**) called Maria may look like this:

C selects diary feature of application to access the diabetes diary. Before diary opens:
A(disclosure): "I also go to school! Together with all the other robots at the hospital. Our teachers are doctors and nurses."
A(prompt) : "Enough about me! Tell me something interesting about yourself!"
Interface : *Would you like to tell NAO something?* yes/no
C(selecting) : yes
Interface : *Please provide your response below.* text input field
C(typing) : "I had a lot of fun at school today. We played hide and seek during the break. No one found me!"
A(p. closer) : "Thanks for sharing that with me, Maria!"
Diabetes diary opens

5 Method

To investigate the relationship *disclosure* → *relatedness* → *motivation*, a two-week, exploratory field study was conducted.

5.1 Participants

Participants in the study were 11 diabetic children between the ages of 8 and 12 ($Mean_{age} = 9.91\,years$, $SD_{age} = 1.08\,years$, 6 girls). All participants had previously participated in an evaluation of the MyPal application at home for 2–4 weeks in May of 2016 and were recruited for this through the two partner

hospitals in the PAL project. Only children that had been diagnosed with diabetes at least six months prior to the evaluation in May were included to avoid any influence of effects (e.g. psychological, lifestyle, family relations) of a recent diagnosis. Children were not reimbursed for the study, which is why efforts were made to ensure that neither children nor parents perceived participation as a burden (e.g. the researcher visited families at home, so they would not need to travel; children should use the application only as frequently as long as they found it enjoyable).

5.2 Materials and Measures

Children were provided a Lenovo tablet computer running Android and with the MyPal Application installed. The app had three main functionalities— a quiz game, the diabetes diary, and an overview of current and achieved diabetes-related objectives of the child. When children chose to open the diabetes diary, the avatar started the disclosure loop only when the child was not using the application offline. The avatar disclosure came from one of two sets, depending on the number of interactions the child already had. The first set contained six *background* disclosures of low intimacy that provided general information about the avatar. These were required to understand some of the disclosures from the second set. When the child had heard all disclosures from this small set of *background* disclosures, avatar disclosures would be randomly chosen from the second larger set, containing a balanced amount of disclosures of low, medium, high, and very high intimacy. One week into the experiment, we found that children were barely using the application, so that few children ever heard disclosures of higher intimacy. To obtain more dislcosure data to study the intimacy, we therefore decided to use the physical NAO robot in the second appointment to also disclose to the children in an introduction round before playing a hangman game with the children. Here one disclosure of each intimacy level was covered. Just like in the app, children were prompted to disclose and could choose not to. Also, as explained in the procedure section, the robot used the words in the hangman game as disclosure triggers. All disclosures for this interaction were drawn from the second set of disclosures, and thus contained no background disclosures. From here on after we will therefore systematically refer to the *ECA* when drawing on data obtained from both robot and avatar and refer to the *avatar* when considering data obtained only from the application.

A total of three questionnaires were used. The first questionnaire contained a subset of the questionnaire from the prior evaluation and was aimed at assessing children's relatedness and motivation at the outset of the study. The second questionnaire asked for children's opinion of the app, the new avatar within the app, technical difficulties, and how much they were using the application approximately. The third questionnaire was the first questionnaire augmented with additional questions for determining relatedness (see *Relatedness* section below).

Age and gender of the children were already recorded in the prior evaluation. The four concepts of interest *disclosures*, *relatedness*, *motivation*, and *intimacy* were operationalized as follows. *Disclosures* were counted. Disclosures

could be passive (child only heard a disclosure from the avatar) or active (the child disclosed to the avatar in return). Disclosures made by the physical robot were only used to determine how children respond to disclosures of various intimacy levels, but were not used in the relatedness and motivation analysis. *Relatedness* we had originally intended to measure exclusively with a subset of the questionnaire from the prior evaluation. However, ceiling effects were obtained on all questions concerning relatedness. As a result, the pre-intervention relatedness measure could not be determined. For the post-intervention measure, the subscales *Companionship* (how much the child enjoys spending time with the avatar), *Reliable Alliance* (how trustworthy the avatar is in terms of disclosure), and *Closeness* (how attached the child feels to the avatar and how much the child believes that the avatar reciprocates this connection) from the Friendship Qualities Scale [3] were added as additional questions to the post-questionnaire. It must be emphasized that making such alterations was only accepted because of the exploratory nature of the study. *Motivation* to use the system was assessed through (1) *usage*: the amount of content a child added to the app while interacting (sum of played quiz questions, diary entries made, and active disclosure interactions) and (2) *consistency*: the percentage of days on which the child engaged with the app. *Intimacy* of child disclosures was determined in a post-analysis using our own disclosure intimacy rating scale [4].

5.3 Procedure

Children and their parents were contacted by phone in the second week of June 2016 to inform them of the purpose of the study, to explain the details of the procedure, and to invite them to participate again. If interested, parents were asked for their email address to receive an information letter and to then schedule an initial appointment.

The first appointment took place in the homes of the children. After parents and children had signed the consent form, children were interviewed using the initial questionnaire. At the end of the interview, children were given the tablet computer and it was explained to the child that the app now contained a new robot with a different name (Robin). Other than that, the functionalities were the same as in the prior evaluation and they could use it without further instructions. Children were not given any guidelines as to how much they should use the application per day, because we were interested in the intrinsic motivation to use the application. The children were then left to their own devices for one week, after which parents received an intermediate questionnaire by email. After two weeks, the first author again visited the children at home to administer the final questionnaire in the form of a semi-structured interview and pick up the tablet computers again. After the interview, the child was given a chance to play a hangman game with the physical robot. This game consisted of an introduction round, in which the robot told a bit about itself and then disclosed to the children four times at all four intimacy levels. Children were prompted to respond in return. Children could play up to four rounds of hangman with the robot. At the end of every round, the robot would again use the hangman solution to

disclose to the children. In total, children could thus hear up to eight further disclosures from the robot. The lead researcher was present during these interactions, noting the children's disclosures. The final interaction with the robot thus served three purposes: to add to the dataset of disclosure interaction between ECA and child, to provide a form of closure for the children and reward them for their participation. The first appointment took approximately 30 min, the second approximately 60 min. All interactions with the children were conducted by the lead researcher.

6 Results

All analyses and plots were made using R-Cran version 3.2.4. We adopted $\alpha = .05$ as the significance threshold. Given the small number of of participants in the study, we strongly advise to take all analyses conducted on variables measured per child with caution (for these we provide post-hoc power analyses results as given by G*Power [8]). We have conducted these analyses to detect trends rather than to confirm hypotheses.

6.1 Disclosure and Relatedness

As described in Sect. 1, Social Penetration Theory posits a strong link between liking and disclosure. It was hence of interest whether the disclosure activity of children was indicative of the relatedness they felt with the avatar at the end of the evaluation period.

To determine the reliability of the relatedness measure in this study, Cronbach's α was computed separately for each of the employed subscales of the Friendship Qualities Questionnaire ($\alpha_{\text{COMP}} = .73$, $\alpha_{\text{RA}} = -.41$, $\alpha_{\text{AB}} = .84$, $\alpha_{\text{RApp}} = .91$). The two items of the *Reliable Alliance* subscale were found to negatively correlate ($r = -.18$). As this should not be the case and the reliability of said subscale is low, we chose to drop the one of the two items ("If there is something bothering me, I can tell my friend about it even if it is something I cannot tell to other people") that increased the overall reliability of the scale (from $\alpha = .89$ to $\alpha = .90$). Active and passive disclosure counts were standardized for each child with the total number of days that it used the application.

To obtain insight into how the two different disclosure behaviors (active vs. passive) relate to the bond between child and avatar, the correlations between the variables could be determined separately. These are illustrated in Fig. 2. However, these correlations do not control for the overall activity of children. We thus modeled the relationship between disclosure behavior and relatedness using linear regression with the predictors *total number of disclosures* (active and passive) and *percentage of active disclosures*. Active disclosures are those where the child actively responds to the avatar with a disclosure of its own, in passive disclosures it does not. The former predictor thus also corresponds to the number of disclosures the child heard from the avatar. The model is given

by the equation:

$$Relatedness = \theta + \beta_1(Disclosures) + \beta_2 \left(\frac{Active.Disclosures}{Disclosures} \right)$$

The two predictors were not correlated ($\rho_S(9) = .10, p = .75$). The total amount of disclosures was not found to be a significant predictor in the model ($b_1 = .05, t(8) = 1.854, p = .10$). The ratio of active disclosures to total disclosures did however significantly predict relatedness ($b_2 = 1.49, t(8) = 2.480, p < .05$). This means that of two children that are interacting with the disclosure module equally often, the child that responds more actively feels more related to the avatar ($R^2 = .529$, $f^2 = 1.122$, $1 - \beta = .866$).

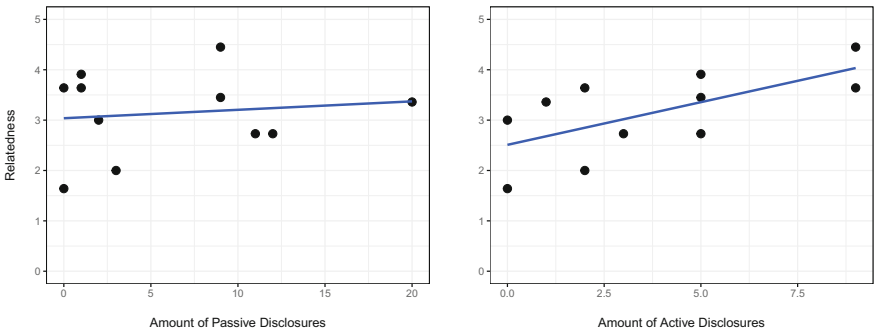

Fig. 2. The relationship between the absolute amount of passive (*left*) and active (*right*) disclosures of children within the application and their relatedness as indicated on the final questionnaire.

6.2 Relatedness and Motivation

Self-Determination Theory argues that relatedness plays a role in motivation. To determine whether the data of this evaluation constitute supportive evidence, the relatedness was correlated with children's overall consistency (how often they used the application) as well as their overall activity (how much they used the application). Using a one-tailed Spearman's rank order correlation, a significant relationship was found between the relatedness and the consistency with which children used the application ($\rho_S(9) = .59$, $p = .03$, $1-\beta = .723$) and the average daily activity ($\rho_S(9) = .64$, $p = .019$, $1-\beta = .816$). To further get an impression of whether there were differences in how much more related and less related children used the application over time, we artificially divided the children into two equally sized ($n_{related} = 6, n_{unrelated} = 5$) groups based on the overall post-evaluation relatedness mean. The evaluation period was divided into two halves for each child and their average daily activity (number of active contributions—diary entries, quiz questions, active disclosures—to the application per day) was calculated for each half. The results are shown in Fig. 3. Since group sizes were

small (5 to 6 children), we believe it to be more informative to inspect the data than to subject them to statistical analyses. The interaction plot shows that children in the more related group were more active in both evaluation halves, but their activity levels decreased substantially between the first and the second week nonetheless and much more so than those of children in the less related group.

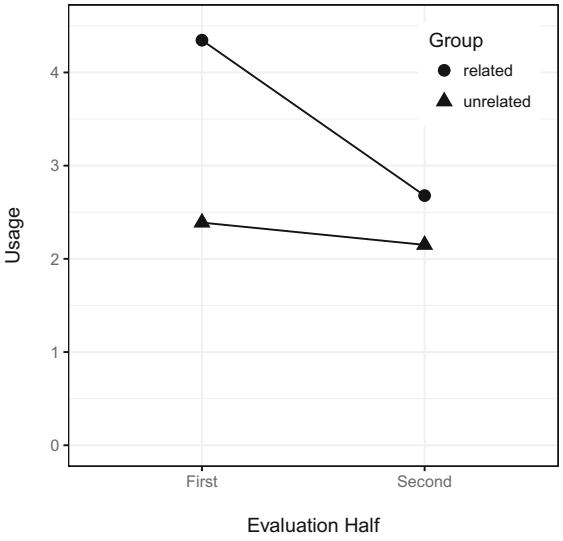

Fig. 3. Average number of activities per evaluation half across children that were artificially split into the two groups related ($n = 6$) and unrelated ($n = 5$) based on their indication of Relatedness on the final questionnaire.

6.3 Intimacy

Three main questions were of interest: (1) does the intimacy level of the avatar disclosure influence whether a child chooses to respond or not (2) if the child responds, does the intimacy level of the prior ECA disclosure predict the intimacy level of the response (3) what role do age and gender of the children play in the former two questions. Since two different types of ECA were used in collecting the active disclosures of children (robot and avatar), we included the ECA type as an additional predictor in the second model for response intimacy described below.

Response Choice. Children were given the choice whether to disclose to the avatar in response to a disclosure from the avatar. It was therefore also of interest to investigate whether their choice to reciprocate depended on the intimacy

level of the disclosure, their age, and their gender. The interaction term between intimacy and time (how much percent of the total experiment time had elapsed when the disclosure occured) was included because the background disclosures caused disclosures of lower intimacy from the avatar to coincide with the beginning of the evaluation period. Due to the binary nature of the response, we use a logistic regression model and since choices are again nested within children, a mixed logistic regression was first fit, allowing intercepts to vary across children. This was again nearly unidentifiable and did not fit the data significantly better than the non-multilevel equivalent ($\chi^2(1) = 3.12$, $p = .08$). We thus chose the simple effect model.

The model for measurements $i = 1, \ldots, n$ is given by the equation:

$$logit(E[Reciprocation_i]) = \theta + \beta_1(Avatar.Intimacy_i) + \beta_2(Child.Age) + \beta_3(Child.Gender) + \beta_4(Avatar.Intimacy_i * Time_i)$$

Figure 4 illustrates the effect of each predictor separately on the binary variable *Reciprocation*. The results from fitting the model match with the visual impression. Both the intimacy level of the avatar disclosure and the gender of children significantly predict whether children choose to respond. As can be seen in Table 1, the odds of disclosing decrease for higher levels of intimacy (OR = .45). Furthermore, the odds of boys disclosing are 7.86 times lower than those of girls. It must be noted here that the confidence interval for this latter effect is large.

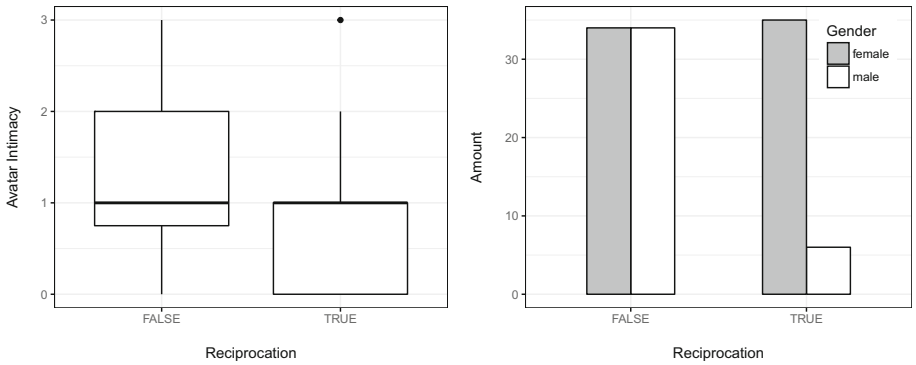

Fig. 4. The relationship between the significant predictors, avatar intimacy (*left*) and gender of child (*right*), and the outcome variable *Reciprocation* in the logistic regression model of whether a child chooses to respond.

Intimacy Prediction. The intimacy of ECA and child disclosures was rated on a four point scale with higher values indicating higher intimacy. A weighted Cohen's kappa which squares the deviance between ratings (extent of disagreement) was used to check for rater agreement. For the disclosures made by the

Table 1. Results of fitting the logistic regression model to the response choice of children within the application.

Predictor	Coefficients					Odds ratio		
	b	z	p	CI		OR	CI	
				2.5%	97.5%		2.5%	97.5%
Avatar Intimacy	−.81	−2.07	.039	−1.61	−.07	.45	.20	.93
Age	.18	.73	.465	−.29	.65	1.19	.74	1.93
Gender	2.06	3.06	.002	.82	3.49	7.86	2.28	32.89
Avatar Intimacy x Time	.00	.02	.99	−.02	.01	1.00	.98	1.01

ECA and the children, agreement was substantial with $\kappa = .707$, $n = 63$ and $\kappa = .697$, $n = 88$ respectively. We averaged the ratings of both raters and used the ceiling function to not inflate the number of to-be-predicted classes.

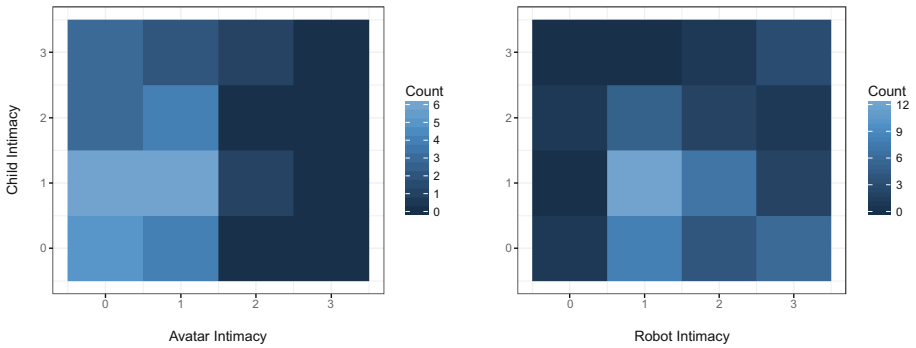

Fig. 5. This shows the contingency matrix of avatar (*left*) and robot (*right* disclosure intimacy and respective child disclosure intimacy as heatmaps. The bottom left corner represents the number of child disclosures of intimacy level 0 that were made in response to agent disclosures of level 0. Intimacy values were based on the combined ratings of both raters (ceiling of average).

From this, it follows that ECA and child intimacy are ordinal variables and should be analyzed with a cumulative link model[6]. Furthermore, the data is hierarchical with disclosures nested within children. Consequently, a cumulative link mixed model was first fit to account for random effects. However, since this model was nearly unidentifiable (condition number of the Hessian = 52790.17)

[6] Cumulative link models are an extension of logistic regression to more than two categories. Thus, where logistic regression determines $logit(P(Y_i = j))$ with $J = 2$, the cumulative link model determines $logit(P(Y_i \leq j))$ with j falling in one of J categories. The model thus accumulates the probabilities of a response being smaller than or equal to a certain category.

and since the multilevel model did not fit the data significantly better than a non-multilevel one ($\chi^2(1) = .07$, $p = .79$), we opted for the latter.

For this analysis, several predictor variables are of interest, the most important being the intimacy level of the ECA disclosure that preceded the child disclosure (Fig. 5). This is followed by the type of ECA (avatar or robot) that made the disclosure. The related literature indicates children's disclosure intimacy may depend on their age and gender, these variables were also included in the model. The predictors of interest were therefore: ECA.Intimacy, ECA.Type, Child.Age, and Child.Gender.

The model is given by the following equation:

$$logit(Child.Intimacy_i \leq j) = \theta_j - \beta_1(ECA.Intimacy_i) - \beta_2(ECA.Type_i) \\ -\beta_3(Child.Age_i) - \beta_4(Child.Gender_i)$$

with $i = 1, \ldots, 88$ (disclosures) and $j = 0, \ldots, 3$ (intimacy categories). None of the independent variables showed any significant relationship with the intimacy of child disclosure. The results are displayed in Table 2.

Table 2. Results of fitting the cumulative link model to predict children's disclosure intimacy from the preceding ECA disclosure intimacy, the type of ECA, the age, and the gender of the child. The first five columns show the log-odds and significance tests. The next set of three columns show the likelihood ratio if the respective predictor is dropped from the model as compared to the full model. The final three columns show the cumulative odds ratios and respective confidence intervals.

Predictor	Coefficients					Likelihood ratio			Odds ratio		
	b	*z*	*p*	*CI*		*AIC*	$\chi^2(1)$	*p*	*OR*	*CI*	
				2.5%	97.5%					2.5%	97.5%
ECA Intimacy	−.13	−.63	.528	−.53	.27	267.53	.40	.528	.88	.59	1.31
ECA Type	−.11	−.42	.673	−.62	.40	267.31	.17	.672	.90	.54	1.49
Age	−.19	−1.05	.294	−.56	.16	268.24	1.10	.294	.82	.57	1.18
Gender	.49	1.08	.282	−.40	1.37	268.30	1.17	.280	1.63	.67	3.95

7 Discussion

The data analysis resulted in several interesting and partially unexpected findings. In this section, we therefore regard the results in light of the larger context of the study and its theoretical background.

Of interest was the chain of *disclosures* → *relatedness* → *motivation*. For the link between disclosure and relatedness, we found that the percentage of active disclosures that children make can be regarded as an indicator for how related they feel towards the agent. While this is not in line with the finding that we typically like those more who disclose to us more, it may match with the finding that the more we like someone at the outset, the more we disclose [6]. Since the

initial questionnaire that we administered to children was not sensitive enough to capture their relatedness at the outset of the study, causal inferences cannot be made, i.e. it is unclear whether disclosing more led the children to feel more related or whether they disclosed more because they already felt more related. This should be investigated again in a controlled experiment.

When regarding the link between relatedness and usage, we find that while more relatedness is associated with more, and more consistent usage, the usage of the related group decreased from the first to the second evaluation half. This is in-line with Self-Determination Theory. Relatedness is a factor in motivation, but not sufficient for it. The application as a whole may not have been attractive enough for the children. Especially the magnitude of the decrease in usage in the related group in comparison to the unrelated group is disconcerting. It is possible that children who felt more related to the avatar had high expectations concerning the avatar's capabilities or the app in general that were then disappointed.

We found avatar intimacy to be a significant predictor in whether children choose to respond with children being more responsive to disclosures of lower intimacy than disclosures of higher intimacy. Although low intimacy disclosures coincided with the novelty of the application due to the background disclosures, *time* did not prove to be a moderator in the relation. With the small amount of data, however, it cannot be excluded with confidence. Other possible explanations for the link are that children may felt overwhelmed by disclosures of higher intimacy ("too much information") or they wanted to match the intimacy but did not know anything of higher intimacy to share. However, in the prior evaluation as well as in the focus group of the ALIZ-e project [2], parents and children stated that they would appreciate a "buddy" robot with whom children can talk about their troubles. It is therefore unlikely that children are entirely untroubled, especially when taking into consideration that they are chronically ill. Instead their troubles may not be salient enough when interacting with the app, they may not trust the avatar sufficiently despite saying so in questionnaires, or the avatar may be too limited in responsiveness. A future study could be conducted to systematically discern these possibilities. Another significant predictor in children's decision to disclose was the gender of the child with boys making substantially fewer disclosures to the avatar than girls. Three of the five participating boys barely used the application. Of the two boys that engaged with MyPal, both disliked the module, one because he could not get directly to the diary, the other because he did not want to talk to the avatar. For the six girls, two also showed very little usage. However, all girls expressed their liking of the module in questionnaires. Since the sample was very small, it is not clear how this generalizes to larger populations. Before drawing conclusions, the gender effect should be re-examined in a confirmatory study.

Finally, when children responded to the ECA, no pattern could be found regarding prior intimacy of the ECA's disclosure, the type of ECA, the gender, or the age of children. This contradicts prior results from child-peer disclosure behavior, in which children in the same age range as in the current study either

relatively or absolutely matched the intimacy of the discloser when reciprocating [14]. From the heat maps, it appears that children are conservative in their replies, tending more towards the lower two intimacy levels regardless of the ECA's intimacy level. This result must be considered with caution, since it is based on sparse, unbalanced data. Furthermore, a problematic influence in the interactions may have been the lack of privacy given to the child when disclosing. In interactions with the physical robot, the experimenter was present and due to the spatial arrangement of some of the children's homes, it was not always possible to isolate the children from parents or siblings or ensure that no disturbances (such as family members coming home) would occur. It is also possible that children experienced similar lacks of privacy when interacting with the application or that some of the disclosures occurred in the context of children demonstrating the application to others.

The data does not paint a coherent picture with children disclosing more actively to disclosures of lower intimacy but not following any particular pattern when they do disclose. The external validity of results is not given because of the small sample size of both children and disclosures as well as the unequal distribution over different intimacy levels. Furthermore, the nature of the study led to potential influences of confounding variables. Particularly since the latter result does not match with prior findings from child-peer interaction, it is important to investigate again whether it is attributable to the replacement of the human peer with an artificial one or if other variables influenced children's true intimacy tendency.

8 Directions for Further Research

The nature of the study required flexibility and some adaptations had to be made to the protocol. Nonetheless, several interesting results were found. It appears that the amount of disclosures that children make towards the avatar is an indicator of how related they feel towards it. No support could be found that children feeling more related to the avatar maintain their initially high usage over time but they use the application more than less related children.

An important goal of this research was the generation of new research questions. These questions can be derived from both the significant and the insignificant results of this study:

nRQ1 What is the causal link between active disclosing and relatedness in the context of long-term child-avatar interaction?
nRQ2 In an interface that clearly supports autonomy and competence, what role does relatedness play in motivating children?
nRQ3 Do children feel more related to a more responsive avatar in the context of long-term interaction?
nRQ4 Are children less likely to respond to more intimate avatar disclosures? If so, why?
nRQ5 Is there a general or child-dependent strategy that the ECA should follow in terms of intimacy development over time to obtain more active disclosures from children?

nRQ6 Do boys disclose less to an avatar than girls? If so, why?
nRQ7 Do children also not match the intimacy level of an ECA when they are given complete privacy?
nRQ8 Is there a difference in how children match disclosure intimacy depending on whether a physical ECA, virtual ECA, or another child is disclosing first?
nRQ9 Is there a difference between diabetic and healthy children in their disclosure behavior towards an ECA?

These research questions should be addressed in confirmatory studies with larger populations of children. The module in itself is flexible and could easily be integrated into another software as well to gather more data. In its current state, however, it is still too limited to provide engaging dialog interactions for children. Hence, a second prototype should be developed.

Several points of improvement for the module became evident during the study. For one, as already identified in Sect. 7, not all children appreciated the placement of the module within the app. This is something that seems to clearly be a personal preference and thus should be personalized. The application was also very limited in its dialog capabilities and from the responses of children it is clear that they figured this out soon (e.g. children attempted to ask the avatar questions several times). In a similar vein, 8 of 11 children had the impression that the avatar knew them better as a consequence of their disclosure. It would be nice for future iterations of the module if the avatar could also show this. To this end, the PAL user model should be augmented with information filtered from the dialog and means should be found to incorporate knowledge from the user model again into the dialog. All in all, this can be summarized as a need for more intelligent behavior of the module.

9 Conclusion

Due to the lack of recent research in the areas of child-peer and child-robot bonding, we conducted an exploratory field study using the first prototype of the dyadic disclosure dialog module. The purpose of the study was two-fold: on the one hand, we wanted to learn about diabetic children's behavior towards a self-disclosing virtual agent. On the other hand, we were interested in possibilities and limitations of creating a bond between child and agent to increase children's motivation in using the application. More related children both disclosed more actively and used the application more than less related children. Children were less likely to respond to disclosures of higher intimacy and boys disclosed less than girls. Future research will need to investigate whether there is truly a difference between ECA and human as conversational partner for children. We thus conclude that the current project presents a promising starting point for further research.

References

1. Altman, I., Taylor, D.: Social Penetration Theory. Rinehart & Mnston, New York (1973)
2. Baroni, I., Nalin, M., Baxter, P., Pozzi, C., Oleari, E., Sanna, A., Belpaeme, T.: What a robotic companion could do for a diabetic child. In: The 23rd IEEE International Symposium on Robot and Human Interactive Communication, pp. 936–941. IEEE (2014)
3. Bukowski, W.M., Hoza, B., Boivin, M.: Measuring friendship quality during pre- and early adolescence: the development and psychometric properties of the friendship qualities scale. J. Soc. Pers. Relat. **11**(3), 471–484 (1994)
4. Burger, F., Broekens, J., Neerincx, M.A.: A disclosure intimacy rating scale for child-agent interaction. In: Traum, D., Swartout, W., Khooshabeh, P., Kopp, S., Scherer, S., Leuski, A. (eds.) IVA 2016. LNCS, vol. 10011, pp. 392–396. Springer, Cham (2016). doi:10.1007/978-3-319-47665-0_40
5. Cohn, N.B., Strassberg, D.S.: Self-disclosure reciprocity among preadolescents. Pers. Soc. Psychol. Bull. **9**(1), 97–102 (1983)
6. Collins, N.L., Miller, L.C.: Self-disclosure and liking: a meta-analytic review. Psychol. Bull. **116**(3), 457 (1994)
7. Deci, E.L., Ryan, R.: Overview of self-determination theory: an organismic dialectical perspective. In: Handbook of Self-Determination Research, pp. 3–33 (2002)
8. Faul, F., Erdfelder, E., Lang, A.G., Buchner, A.: G* power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behav. Res. Methods **39**(2), 175–191 (2007)
9. Kanda, T., Sato, R., Saiwaki, N., Ishiguro, H.: A two-month field trial in an elementary school for long-term human-robot interaction. IEEE Trans. Robot. **23**(5), 962–971 (2007)
10. Kruijff-Korbayová, I., Oleari, E., Bagherzadhalimi, A., Sacchitelli, F., Kiefer, B., Racioppa, S., Pozzi, C., Sanna, A.: Young users' perception of a social robot displaying familiarity and eliciting disclosure. ICSR 2015. LNCS, vol. 9388, pp. 380–389. Springer, Cham (2015). doi:10.1007/978-3-319-25554-5_38
11. Moon, Y.: Intimate exchanges: using computers to elicit self-disclosure from consumers. J. Consum. Res. **26**(4), 323–339 (2000)
12. Rotenberg, K.J., Chase, N.: Development of the reciprocity of self-disclosure. J. Genet. Psychol. **153**(1), 75–86 (1992)
13. Rotenberg, K.J., Mann, L.: The development of the norm of the reciprocity of self-disclosure and its function in children's attraction to peers. Child Dev. **57**, 1349–1357 (1986)
14. Rotenberg, K.J., Sliz, D.: Children's restrictive disclosure to friends. Merrill-Palmer Q. **34**, 203–215 (1988). (1982)
15. Youniss, J.: Parents and Peers in Social Development: A Sullivan-Piaget Perspective. University of Chicago Press, Chicago (1980)