

Building Successful Long Child-Robot Interactions in a Learning Context

Alexis Jacq^{1,2}, Séverin Lemaignan¹, Fernando Garcia¹, Pierre Dillenbourg¹, Ana Paiva²

¹CHILI Lab, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

²INESC-ID & Instituto Superior Técnico, University of Lisbon, Portugal

Abstract—The CoWriter activity involves a child in a rich and complex interaction where he has to teach handwriting to a robot. The robot must convince the child it needs his help and it actually learns from his lessons. To keep the child engaged, the robot must learn at the right rate, not too fast otherwise the kid will have no opportunity for improving his skills and not too slow otherwise he may lose trust in his ability to improve the robot's skills. We tested this approach in real pedagogic/therapeutic contexts with children in difficulty over repeated long sessions (40-60 min). Through 3 different case studies, we explored and refined experimental designs and algorithms in order for the robot to adapt to the troubles of each child and to promote their motivation and self-confidence. We report positive observations, suggesting commitment of children to help the robot, and their comprehension that they were good enough to be teachers, overcoming their initial low confidence with handwriting.

Index Terms—Child-robot Interaction; Learning by teaching; Protégé Effect; Children Self-confidence; Extrinsic Motivation; Robotic Handwriting Learning

I. INTRODUCTION

Children facing difficulties in handwriting integration are more exposed to troubles during the acquisition of other disciplines as they grow up [5]. The CoWriter activity introduces a new approach to help those children [9]. While traditional successful interventions involve children in long intervention (at least 10 weeks) focused on *motor* skills [10], CoWriter is based on *learning by teaching* paradigm and aims to repair self-confidence and motivation of the child rather than his handwriting performance alone.

Learning by teaching is a technique that engages the students to conduct an activity as the teachers in order to support their own learning process. This paradigm is known to produce motivational, meta-cognitive and educational benefits in a range of disciplines [20]. The CoWriter project is the first application of the learning by teaching paradigm applied to handwriting with a robot.

The effectiveness of our learning by teaching activity builds on the “protégé effect”: the teacher feels responsible for his student, commits to the student's success and possibly experiences student's failure as his own failure to teach. Teachable computer-based agents have previously been used to encourage this “protégé effect”, where students invest more effort into learning when it is for the benefit of a teachable agent than for themselves [4]. We rely on this cognitive mechanism to reinforce the child's commitment into the robot-mediated handwriting activity.

We assume here that the key of such a relationship between the child and the robot relies on the credibility of the robot: the more the robot convinces the child that it is a beginner in handwriting who needs help – therefore initiating a “protégé effect” – the deeper the child will engage in the interaction. We focus hereafter on two aspects that are instrumental in building a credible teaching situation: how to generate the initial state of the learner-robot, and how to design its learning behavior.

We previously used a limited approach [?] in which letters had to be written as a single stroke (no pen lifting) and that covered typical mistakes of adults extracted from an handwritten letters database. Our initial experiments were conducted in school, involving either group of children doing the activity together or children in short individual sessions. These studies were conducted to evaluate the feasibility and technical soundness of the interaction system. Because of the group effect and the briefness of the interactions, no conclusions could be reached about the positive effect of the interaction. Participating children were randomly chosen in school classes and had no specific difficulties in handwriting. This made it difficult to observe any remediation of low self-esteem or motivation.

As a follow up, this article reports on further experimental investigations. We explore different algorithmic and staging approaches built on top of the original system in order to figure out intricate aspects of long child-robot interactions in a pedagogical context. We solved previous technical limitations of robot's letter learning and generation, and we introduce new algorithmic approaches that makes the behavior of the robot more convincing. Through three experiments, we involved children with actual handwriting troubles or low self-esteem in repeated long sessions (four times about one hour). We used different measures, both qualitative and quantitative, to express the impact of those interaction with the CoWriter robot on the child.

After a presentation of the work related to this topic, this paper is divided in four sections. In the first section we give technical details of our setup, such as how the different modules are connected together and which algorithms are used for the robot to learn and generate letters. The following three parts report our the three experiments and results: two case studies specifically designed to be adapted to one child IV V; one user study conducted with 8 children separately VI.

II. RELATED WORK

Concerning learning non-physical skills, the protégé effect has been used in the past by computer-based agents [4]. Robots maintain better long-term relationship [12] and contribute to obtain more learning gains [15] than with screen-based agents in pedagogical interactions. Specifically, when learning physical skills, robotic partners have been showed to increase users' compliance with the tasks [1].

Many studies have been conducted with language skills acquisition [8], less often involving physical skills (such as calligraphy [17]). Regarding *learning by teaching* paradigm with robots, Werfel notes in [23] that studies tend to focus on the ability of the robots to learn (in terms of language [21] or physical [18] skills, for example) rather than the beneficial impact on the teaching for the human. To contrary, our work minimized the robot's skills while we concentrate on the possible improvement of children self-confidence and motivation promoted by the behavior of the robot.

The usage of tutor robots in educative activities with children is a sensible point. A bad choice of the robot's behavior can have negative impacts on the learning [11], that are consequent in long-term interactions [13]. Peer robot partners seems more efficient than tutor [24], but no such study relates advantages of learning by teaching a robot. However, Tanaka and Matsuzoe [22] explored this paradigm with a Nao robot learning vocabulary from children, and Chandra [3] used a peer robot leading a learning by teaching activity performed by two children.

A remaining difficulty in Child-Robot-Interaction concerns the evaluation of the interaction. Using questionnaires with children can lead to contradictions between the actual behavior of the child during interaction and his answers during the interview [14]. One reason is that children have the tendency to try to please the experimenter, rather than answer truthfully to survey questions [2]. Various metrics can be used to describe the behavioral aspects of the interaction (duration of interaction, proximity...) and learning gains (pre/post tests), but it is much harder to obtain measure of psychological impacts without a very large sample size providing significant results [2].

Measurement of the children engagement must be based on a rigorous model. O'Brien and Toms [19] provided such a framework and listed different attributes that can provide information about the engagement. In our context of Human-Robot-Interaction, we can make distinction between three kind of engagement : social engagement, task engagement and social-task engagement [6]. Along this paper, we focus on the persistence of the "protégé effect": we aim to play with the children's perception of the robot in order to create motivation. In that way, we base our observations and results on metrics of social-task engagement.

III. EXPERIMENTS DESIGN

A. Interaction overview

Figure 1 illustrates our general experimental setup: a face-to-face child-robot interaction with an autonomous Alde-

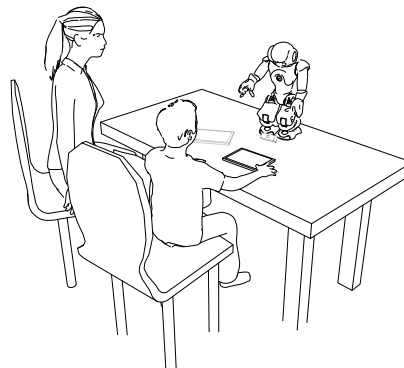


Fig. 1. Our experimental setup: face-to-face interaction with a NAO robot. The robot writes on the tactile tablet, the child then corrects the robot by directly overwriting its letters on the tablet with a stylus. An adult (either a therapist or an experimenter, depending on the studies), remains next to the child to guide the work (prompting, turn taking, etc.). For some studies, a second tablet and an additional camera (lightened) are employed.

baran's NAO robot.

A tactile tablet (with a custom application) is used by both the robot and the child to write: in each turn, the child requests the robot to write something (a single letter, a number or a full word), and pushes the tablet towards the robot, the robot writes on the tablet by gesturing the writing (but without actually touching the tablet). The child then pulls back the tablet, corrects the robot's attempt by writing himself above or next to the robot's writing, and "sends" his demonstration to the robot by pressing a small button on the tablet. The robot learns from this demonstration and tries again.

Since the child is assumed to take on the role of the teacher, we had to ensure he would be able to manage by himself the turn-taking and the overall progression of the activity (moving to the next letter or word). In our design, the turn-taking relies on the robot prompting for feedback once it is done with its writing (simple sentences like "What do you think?"), and pressing on a small robot icon on the tablet once the child has finished correcting. We found that both approaches were easy to be understood by children.

B. Generating and learning letters

Since our approach is based on teaching a robot to write, generating (initially bad) letters and learning from demonstrations is a core aspect of the project. The initial state of the robot and his ability to learn in an obvious way from demonstrations of the child is the key to lend credibility to the activity and to induce the "protégé" effect.

The technical idea is simple: allographs of letters are encoded as a sequence of 70 points in 2D-space and can be seen as vectors with 140 elements $(x_1, \dots, x_{70}, y_1, \dots, y_{70})$. We arbitrary chose a set of allograph that define the initial state of generated letters. After the child provided a demonstration of a letter, the algorithm generates a new letter corresponding to the middle point between the last state and the demonstration.

In the following sections, we present various techniques to create the initial state, and different metrics used to compute progression of the robot, tested as hypothesis within our three experiments.

1) *Generation of initial allographs*: The first question relates to the construction of the initial set of allographs. In previous experiments presented in [?], we built a subspace based on principal component analysis (PCA) of a standard dataset of 214 adult letters (the UJI Pen Characters 2 dataset [16]). We used the first n eigenvectors (in these experiments, $3 < n < 6$) of the covariance matrix generated from PCA to create a subspace. To create new letter shapes, we chose random coordinates close to the origin of this subspace. Each eigenvector provided the direction of a principal deformation of the allograph in human handwriting [9]. But generated “imperfections” of letters were not representative of children deformations: they were reflecting typical defects when adults are writing to fast. Over the following studies, we explored three different ways to generate samples closer to beginners. In our first case study (section IV), we used homework of the child previously provided by his mother, to exaggerate by hand his main defects. This way, the child was going to correct his own kind of mistakes. In the second study (section V), the child was suffering from visuo-constructive deficits. Since it was difficult for him to improve already recognisable allographs, we decided under the guidance of his occupational therapist to make the robot start from simple vertical stroke for all letters. In the third study VI we chose to use the middle point between a vertical stroke and correct letters as a starting point for the robot.

2) *Metrics used for the learning curve of the robot*: The second question focuses on the learning algorithm. In [9], we were projecting children’s demonstrations in PCA’s subspace in order to compute the middle between that point and the previous state of the robot. Then, we generated the allograph in middle way as the new state of the robot. For the experiments introduced in this paper, we explored two other ideas: In the first study (section IV) we generated a PCA subspace from a small set of allographs we drew arbitrary. Each time the child was providing a demonstration, we added that demonstration to the small set and re-built the PCA subspace. That way, the principal eigenvectors obtained progressively tended to encode the main deformations of letter done by the child. The algorithm 1 explains the successive steps of this approach.

From our perspective, this dynamic subspace was more adapted to the progression of the child, and the sequence of tries performed by the robot looked smoother. However using metrics in subspace can make the learning algorithm too slow in some cases, because consecutive projected demonstrations can sometimes be too far from each other in subspace while they appears similar in Cartesian space. In other studies, we decided to put aside the PCA approach and to always use the middle point in Cartesian space, in order to have a better control over the convergence of the robot tries to the demonstrations.

Algorithm 1: Learning from demonstration in an adaptive PCA subspace

```

Generate initial dataset  $D$ 
Generate initial subspace  $S$  by PCA of  $D$ 
Generate initial robot state  $r$  (random point in  $S$ )
if robot receives a demonstration  $d$  then
    Add  $d$  to dataset:  $D' \leftarrow D \cup d$ 
    Recompute subspace  $S'$  by PCA of  $D'$ 
    Compute coordinates  $r'$  of  $r$  in  $S'$ 
    Compute coordinates  $d'$  of  $d$  in  $S'$ 
    Learn the demonstration:  $r = \frac{(r'+d')}{2}$ 
end

```

C. *Robotic Implementation*

The actual implementation on the robot requires the coordination of several modules (from performing gestures and acquiring the user’s input to the high-level state machine), spread over several devices (the robot itself, one laptop and up to four tactile tablets for certain studies we conducted). We relied on ROS to ensure the synchronization and communication between different devices.

Our system is embodied in an Aldebaran’s NAO (V4 or V5, depending on the studies) humanoid robot. This choice is motivated by its approachable design [7], its size (58cm) and inherently safe structure (lightweight plastic) making it suitable for close interaction with children, its low price (making it closer to what school may afford in the coming years) and finally its ease of deployment on the field.

Robotic handwriting requires precise closed-loop control of the arm and hand motion. Because of the limited fine motor skills possible with such an affordable robot, in addition to the absence of force feedback, we have opted for *simulated handwriting*: the robot draws letters in the air, and the actual writing is displayed on a synchronised tablet.

The overall architecture of the system (Figure 2) is therefore spread over several devices: the NAO robot itself, that we address via both a ROS API¹ and the Aldebaran-provided Naoqi API, one to four Android tablets (the main tablet is used to print the robot’s letter and to acquire the children’s demonstrations; more tablets have been used in some studies, either to let the child input words to be written, or for the experimenter to qualitatively annotate the interaction in a synchronized fashion), and a central laptop running the machine learning algorithms, the robot’s handwriting gesture generation and high level control of the activity.

Since the system does not actually require any CPU-intensive process, the laptop can be removed and the whole logic run on the robot. Due to the relative difficulty to deploy and debug ROS nodes directly on the robot, the laptop remains however convenient during the development phase and we kept using it in our experiments.

¹The ROS stack for NAO is available at http://wiki.ros.org/nao_robot.

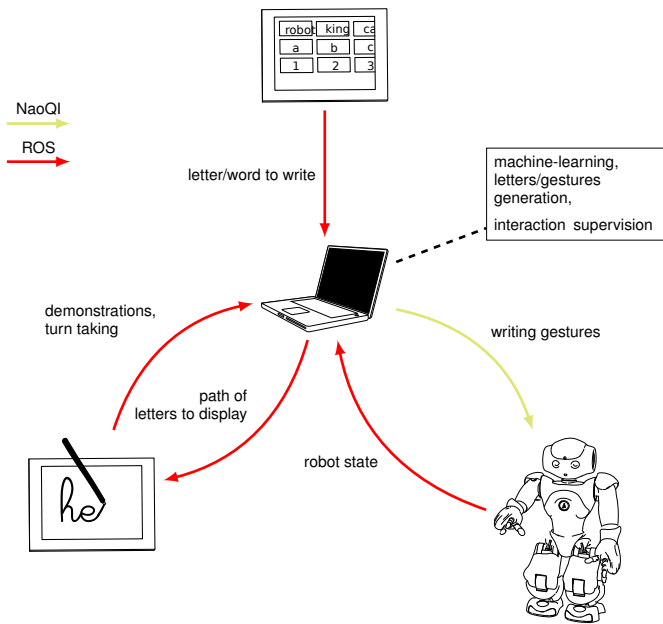


Fig. 2. **Overview of the system.** In total, the system runs about 10 ROS nodes, distributed over the robot itself, a central laptop and Android tablets.

Most of the nodes are written in Python, and the whole source code of the project will be made available online².

IV. CASE STUDY 1: VINCENT

A. Context

Vincent³ is a five year-old child. At school, he has difficulties to learn writing, particularly with cursive letters. From our perspective, Vincent is shy and quiet. He suffers from poor self-confidence much more than any actual writing problem. The experiment was conducted without any therapist, in our laboratory. A parent was here to accompany the child, but she did not intervene during interactions. Children’s personalities, conditions and state evaluation were reported by the parent.

B. Hypothesis

The CoWriter activity needs a child engaged as interaction leader. With this study we consider the problem of long-term interactions. We hypothesize that with an appealing scenario children can maintain motivation in doing a handwriting activity for an hour over 4 sessions.

C. Experimental design and methodology

Our goal was to provide Vincent with an environment that would enable him to sustain engagement over four one-hour sessions, one session per week. We decided to introduce a scenario to elicit a strong “protégé effect” and such induce a stronger commitment. While the child came with low motivation in writing exercise for himself, our idea was to use this

²The primary repository is https://github.com/chili-epfl/cowriter_letter_learning.

³The names of children has been changed.

effect to promote a new extrinsic motivation: improving letters in order to help the robot.

In our scenario we used two Nao robots: a blue one (called Mimi) and an orange one (called Clem). Mimi was away for a scientific mission, and the two robots had to communicate by mails. But they decided to do it “like humans”, with handwritten messages. While Mimi was good in handwriting, Clem had strong difficulties and needed Vincent’s help.

Mimi’s mission was to explore a mysterious hidden base. Each week, a postal mail contenting a picture of a curious object it found and a few handwritten words about its discoveries. The picture showed itself exploring a dark room of the hidden base (that was actually our laboratory’s workshop).

During the three first sessions, Clem (the robot interacting with the child) was waiting for Vincent with the received mail. It let Vincent take a look at the picture and the object, and then it asked him to read the message. Finally, Vincent formulated a response and helped the robot to write it.

The fourth and last session was set as a test: Mimi, the “explorer” robot, came back from its mission and challenged Clem in front of Vincent: “*I don’t believe you wrote yourself these nice letters that I received! Prove it to me by writing something in front of me!*” In this situation +we pressed the protégé effect: Clem is going to be judged on its writing skills by Mimi, but Vincent is here to give a last help and to encourage his student.

To complement the motivation of helping a robot to communicate with another one, we gradually increased the complexity of Vincent’s task to keep it challenging and interesting (first week: demonstration of single letters; second week: short words; third week: a full message – Figure 3).

Vincent had to tell the robot what to write with small plastic letters. A third person was here to send the formed word to the robot via the computer.

During the experiment, we recorded writings of the child and the robot on the tablet into log files. We also recorded the time date when the child started and finished a demonstration.

D. Measures

We measured the commitment of the child with the number of demonstration he provided. We also measured the duration of sessions. During the two last sessions, we recorded the time taken by the child to write each demonstration.

After the experiment we interviewed the parent of the child. She was asked if she observed any impact of our activity on the child.

E. Analysis

We compared the number of demonstrations provided by Vincent along the 4 sessions (reported on Table I) and we summed the time spend by the child to write demonstration during the 2 last sessions.

F. Results

Overall, Vincent provided 155 demonstrations to the robot. We can see in Table I that the number of demonstrations

TABLE I
NUMBER OF DEMONSTRATIONS PROVIDED BY VINCENT OVER THE FOUR SESSIONS.

Session	S1	S2	S3	S4	Total
Number of demonstrations	23	34	52	46	155

provided by Vincent was globally increasing along sessions while the difficulty of the activity was also increasing. Interestingly, as the number of demonstration decreased from session 3 to session 4, the total time spend to write demonstrations is similar: 41.6s in session 3 ($\sim 0.8s$ per letter) and 41.1s in session 4 ($\sim 0.89s$ per letter). A explanation of this result could be that since the difficulty was increasing the child spent more time to write his demonstrations.

After the first week, he showed confidence when playing with his “protégé” and he built affective bonds with the robot over the course of the study, as evidenced by some cries on the last session, and several letters sent to the robot *after* the end of the study (one of them 4 months later) to get news. This represents a promising initial result: we can effectively keep a child committed into the activity with the robot for a relatively long periods of time (about 4 hours).

From the parent’s perspective, Vincent was actually showing a new motivation in improving his handwriting. He took pleasure to work with the robot and to accomplish his teacher’s mission. She confirmed that an affection of the child for the robot took root within the experiment. Finally she saw an improvement of his handwriting and explained that the child “passed from a mix of script and cursive writing up to a full-cursive writing”.

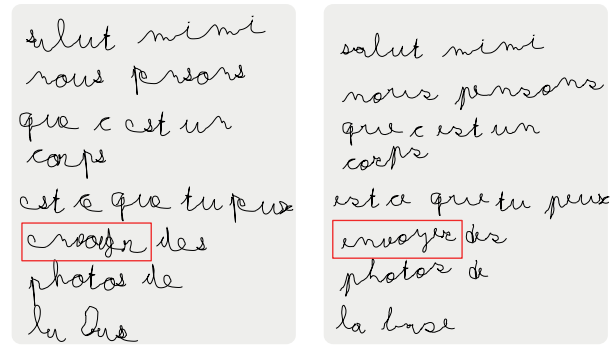
But no conclusion can be drawn in terms of actual handwriting remediation: we did not design this study to formally assess possible improvements. However, as pictured on Figure 3, Vincent was able to significantly improve the robot’s skill, and he acknowledged that he had been able to help the robot: in that regard, Vincent convinced himself that he was “good enough” at writing to help someone else, which is likely to have a positive impact on his self-esteem.

V. CASE STUDY 2 : THOMAS

A. Context

Thomas, 5.5 years old child, is under the care of an occupational therapist. He has been diagnosed with visuo-constructive deficits. He was frequently performing random attempts and then was comparing with the provided template. According to the therapist, Thomas is restless and careless: he rarely pays attention to advice and does not take care of his drawing movement when he is writing. He is quickly shifting his attention from one activity to another.

Thomas was working on number allographs with his therapist. During a prior meeting, the therapist provided us with a sequence of numbers written by Thomas. one of the observed problems was drawing horizontally-inverted allographs, mainly for “5”. The experiment was conducted with Thomas’ therapist.



(a) Initial letter, generated by the robot (b) Final letter, after training with Vincent

Fig. 3. (French) text generated by the robot, before and after a one hour long interaction session with the child. As an example, the red box highlights the changes on the word “envoyer”.

B. Hypothesis

We want to see if the CoWriter activity can be adapted to a pedagogical context in order help a child with diagnosed deficits to learn handwriting.

We believe that small modifications of the activity adapted to Thomas problems (visuo-constructive deficits and inattention) could help to keep him focused on the activity during forty-minutes sessions, and to evidence to the child that the robot is progressing by dint of his demonstrations.

C. Experimental design and methodology

The experiment was conducted in the therapist’s office (four sessions spanning over 5 weeks). We assumed that a scenario like the one we used for Vincent would not be usable with Thomas. We just introduced the robot and quickly said that it was seeking help to train for a robot handwriting contest.

In order to integrate our work with that of the therapist, we decided to adapt the CoWriter activity to work with numbers.

Since Thomas was frequently drawing horizontally-inverted numbers, or even unrecognisable allographs, the learning algorithm of the robot was converging to meaningless scrawls. To fix this problem, we programmed the robot to refuse allographs that were too distant to a reference with a threshold we arbitrary fixed. In that way, the child was forced to take care of his demonstrations for the robot.

According to the therapist, it was easier for Thomas to memorize the way to draw a number if it was always done is the same trajectory, *e.g.* if the “5” was always drawn from the top-right tip down to bottom. Therefore we programmed the robot to refuse as well a good allograph drawn in a wrong trajectory. But in order to reassure Thomas about the right final allograph’s shape, we made the robot able to recognize such a drawing, and, when it occurred, to use the phrase: “Oh, this is exactly the shape of the number I want to learn, but can you show me how to draw it in the opposite trajectory?”

Also, to make the robot’s progresses evident, we modified the initialization step of the learning algorithm to start with a




Fig. 4. Demonstrations provided by Thomas for the number “6” (top row) and corresponding shapes generated by the robot. After eight demonstrations, Thomas decided that the robot’s “6” was good enough, and went to another character: in that respect, he was the one leading the learning process of the robot.

roughly vertical stroke instead of a deformed number (round 0 on Figure 4).

In this setup, we added a second tablet with one button per number. It was used by the child to chose a new number to teach to the robot. It also provided the possibility to enter letters or words, and to switch to another activity (robot telling a story if the child needs a short break).

D. Measures

We recorded all the demonstration performed by the child and by the robot. The duration of sessions and the time spend by demonstration were also recorded by the logs of the tablet.

E. Analysis

It was difficult to make comparison between different sessions since the child did not work on the same numbers. But we could study the evolution of the quality of Thomas’ demonstration when he was working on a given number (Figure 6). To show how Thomas led the robot to reach his level we plotted on the same graph the evolution of the quality of Thomas’ demonstrations and the robot’s trials (Figure 5). We also reconstructed and displayed the drawn allographs of the number 6 to visualize the impact of the lessons of Thomas on the robot (Figure 4).

F. Results

Despite his attention deficit, Thomas was able to remain engaged in the activity during more than forty minutes in each session. In total, 55 allographs out of 82 demonstrated by the child were acceptable considering our threshold (with a progressive improvement from 13 out of 28 in the first session up to 26 out of 29 in the last session).

As soon as Thomas understood that the robot was only accepting well-formed allographs, he started to focus on it and he would typically draw 5 or 6 times the number before actually sending to the robot (the tablet lets children clear their drawing and try again before sending it). According to the therapist, it was the first time that Thomas corrected himself in such a way: he mad the effort to take into account how *another agent* (the robot) would interpret and understand his writing. Figure 6 shows how he gradually improved his demonstrations for some numbers, according to the metric we used to make the robot accept/refuse samples.

Since the robot’s handwriting started from a simple primitive (a stroke), each time Thomas succeeded to have his

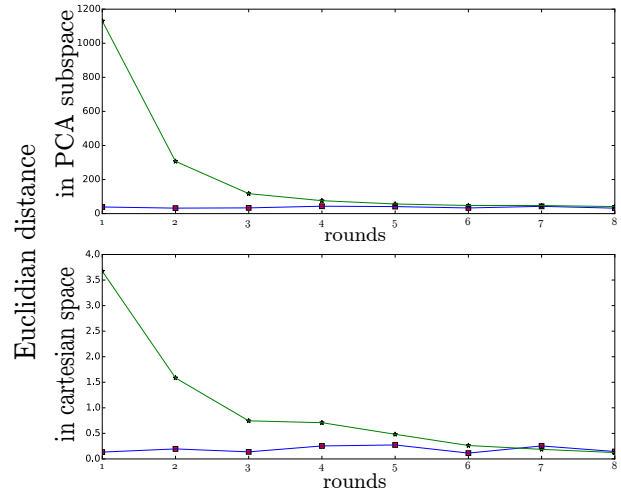


Fig. 5. Two metrics to assess the handwriting progresses: Euclidean distance between demonstrations and templates in the subspace of the number dataset (top figure) or in Cartesian space (bottom figure). Green lines represent the robot performance, blue lines performance of the child. The round IDs correspond to the demonstrations pictured on Figure 4.

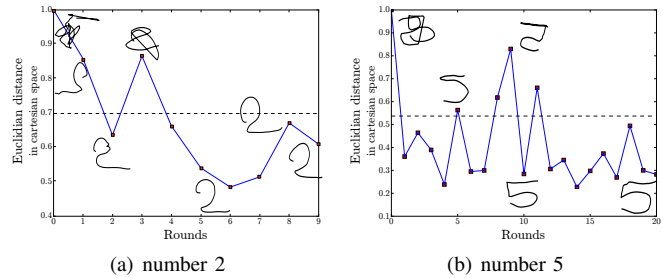


Fig. 6. Improvement of Thomas demonstrations for some numbers: a) the number 2 and b) the number 5. Thomas progressively took care of the demonstrations he was providing to the robot for those numbers. We used for this figure the same metric than the one used for the acceptance algorithm to measure distance between demonstration and templates. Distances are normalized with respect to the biggest value. The dashed line correspond to the threshold of robot’s acceptance.

demonstrations accepted by it, the robot’s improvement was clearly visible (as measured in Figure 5). This led to a self-rewarding situation that effectively supported Thomas’ commitment.

VI. CASE STUDY 3: WHEN CHILDREN EVALUATE THE ROBOT

A. Context

Each of previous studies was specifically adapted to a particular child: we relied on two different designs in order to sustain each child’s commitment. In this new experiment, we conducted a study with eight children using a single experimental design. The children all have in common difficulties to learn cursive writing but the nature and magnitude of these troubles are significantly different from one child to another. Valerie (7 years old), Antoine (6.5) and Johan (7) are under the

care of an occupational therapist. Emilien (8) and Mathieu (7) are repeating their school year because of writing. Marie (6) and Adele (8) are bottom of their respective classes in writing activities. Nicolas (7) is under the care of a neurologist, and has been diagnosed with specific language impairment. Given their school year, all of these children would be expected to know the shape of cursive letters. The experiment was conducted in collaboration with an occupational therapist.

Our goal was to study the perception of the robot's progress in children. We wanted to know how easily children were able to take the role of teachers and to detect improvements or eventual degradations of the robot's letters.

B. Hypothesis

Children understand their role and find motivation to teach the robot. They are able to perceive the progress of the robot, and their evaluations correlates with its handwriting performance.

C. Experimental design and methodology

This experiment took place in an occupational therapist clinic in Normandy, France. Over a period of two weeks, each child came three times for a one hour long session (except Adele and Marie who only attended one session). An experimenter was present to explain the rules of the game and tablet usage. As in the previous experiments, children were provided with two tablets: one to choose a word (or a single letter) to teach, one used by both the child and the robot to write. We also provided printed templates for the letters if the child asked for them.

The initial shapes used by the robot when writing were the same for all children: we used the average of a simple vertical stroke and the reference letter. In this study, we wanted the robot to be only influenced by the demonstrations provided by the child, so we did not project allographs in a subspace. The new samples generated by the robot were simply computed as the average (in Cartesian space) between the last demonstrations and the previously generated samples.

The robot was programmed to accept all demonstrations, endowing the child with the full responsibility of a teacher.

Besides, we added two buttons to the tablet interface: a green one with a "thumbs up", and a red one with a "thumbs down". Those buttons could be used by the children to evaluate the robot (the green one was for positive feedback while the red one was for negative feedback). We used it as a measure of the perception of the robot by the child: the more the child used evaluation buttons, the more he was adopting the role of the teacher, judging the robot instead of himself. Children were free to use the buttons whenever they wanted during the experiment.

D. Measures

As in previous studies, we recorded the timestamps of all demonstrations, the duration of demonstrations and we measured the overall commitment of the children as the number of demonstrations provided per session. We also logged all

the evaluations provided by the children. The awareness of children for the robot progress is measured as the correlation between children evaluations and distances between the robot's letters and reference templates.

E. Analysis

Since sessions took place over only two weeks, we did not attempt to study possible handwriting remediation in children, and we focused instead on the correlation between the children's evaluations and the robot's progression. We estimated the robot's progression as the difference between an initial score (score of the first robot's attempt when the children have chosen a new word/letter to work on) and the current robot's score (after being taught by the child). The score is calculated as the average of the euclidean distance between the robot's generated letter and the reference allograph for each of the letters of the word. The reference letters were manually created beforehand, based on typical cursive letters template⁴. At every turn, we associate two values: the current score of the robot, and the child's immediate feedback (+1 if the child pressed the green button, -1 if he pressed the red one, 0 if he did not press any button during the round). We only keep rounds with feedback (*i.e.* a non-zero grade) and computed a Pearson's correlation between the robot score and the child feedback.

F. Results

All children maintained their engagement during all the sessions. They provided on average 42 demonstrations per session. All children made use of the evaluation buttons and had preference to reward the robot (in total, 99 positive feedbacks and 33 negative ones were recorded). Interestingly, the time spent by the children to draw the demonstrations systematically increased from one session to the other. We interpret this result as the children being more careful and demonstrating the correct gestures to the robot in a slower fashion.

We found that five children out of the eight provided evaluations that significantly correlated with progress of the robot. The coefficients of correlation r_{robot} are reported in Table II.

We also computed the correlation between the children's evaluations and their own progress. The analysis was conducted in the same way, using distances between the children's demonstrations and reference allographs as a progress score. The evaluations of three out of the five children whose evaluations correlated with the robot's progress, were also significantly correlated with their own progress (r_{child} in Table II). For those children, it seems that the robot was reflecting their own performances, and while they were judging the robot positively (three times more positive feedback than negative feedback), they were actually evaluating themselves.

⁴<http://www.education.com/slideshow/cursive-handwriting-z/>

TABLE II

FEEDBACK FROM THE CHILDREN TO THE ROBOT. #Demo DENOTES THE AVERAGE NUMBER OF DEMONSTRATIONS PER SESSION PROVIDED BY THE CHILD; #Pos AND #Neg THE TOTAL NUMBER OF POSITIVE (RESP. NEGATIVE) FEEDBACKS THEY PROVIDED. r (ROBOT) IS THE CORRELATION COEFFICIENT BETWEEN THE FEEDBACK PROVIDED BY THE CHILDREN AND THE PERFORMANCE OF THE ROBOT. r (CHILD) IS THE CORRELATION COEFFICIENT BETWEEN THE FEEDBACK PROVIDED BY THE CHILDREN AND THEIR OWN PROGRESS.

Child	# Demo	# Pos	# Neg	r_{robot}	r_{child}
Valérie	42	24	6	0.25 **	0.14 ns
Émilien	74	20	9	0.06 ns	0.02 ns
Mathieu	43	10	3	0.23 **	0.21 **
Nicolas	38	16	4	0.31 ***	0.20 **
Johan	32	10	5	0.10 ns	0.03 ns
Antoine	27	10	3	0.20 *	-0.02 ns
Adèle	35	4	2	0.28 *	0.30 **
Marie	40	5	1	-0.02 ns	0.13 ns

VII. CONCLUSION

This paper provides the first results of long-term experiments conducted with CoWriter apparatus.

Through adaptations of the interaction design and the learning algorithm, our goal was to induce a protégé effect in order to promote an extrinsic motivation to help the robot.

The first case study evidences that it is possible to conduct a long-term interaction with the child and the robot, while maintaining a high-level of child engagement. According to our observations and interviews, we successfully induced a protégé effect by building an affective bond between the child and the robot. We evidenced the engagement of the child from his increasing number of demonstrations and time spent to write demonstrations.

The second study shows that the activity can be adapted to real therapeutic scenarios, and that a child with severe attentional troubles can be motivated to train the robot over extended period of time.

The third study aimed at evaluating how children adopt their “robot teacher” role and how they perceive the robot’s progress. By presenting the children with an interface to evaluate the robot performance over the course of the experiment, we were able to assess their commitment as teachers. If all children appeared intrinsically motivated by the activity, we can assume that those who understood and well played their role of teacher (given our measure) found their motivation by playing this role. In their case, the fact that evaluations were mostly positive reveals the protégé effect induced by the activity.

The on-line evaluation of the robot by the child also provides insights on the child’s satisfaction with the learning process (both the robot’s ability to learn and the child’s own ability to teach). We believe that this information could be taken into account by the robot in order to improve the quality of the interaction. As an example, it could be used at two levels: • it is possible to get relevant information about the perception of the robot’s progress by the child. • We can reinforce the learnt letters when the robot receives a good

evaluation, or make it forget the letter when it receives a bad evaluation.

REFERENCES

- [1] W. A. Bainbridge, J. W. Hart, E. S. Kim, and B. Scassellati. The benefits of interactions with physically present robots over video-displayed agents. *ICSR*, 2011.
- [2] T. Belpaeme, P. Baxter, J. De Greeff, J. Kennedy, R. Read, R. Looije, M. Neerinx, I. Baroni, and M. C. Zelati. cHRI: Perspectives and challenges. In *Social Robotics*, pages 452–459. Springer, 2013.
- [3] S. Chandra, P. Alves-Oliveira, S. Lemaignan, P. Sequeira, and P. Paiva, A. Dillenbourg. Can a child feel responsible for another in the presence of a robot in a collaborative learning activity? In *Ro-Man*, 2015.
- [4] C. C. Chase, D. B. Chin, M. A. Oppezzo, and D. L. Schwartz. Teachable agents and the protégé effect: Increasing the effort towards learning. *Journal of Science Education and Technology*, 18(4):334–352, 2009.
- [5] C. A. Christensen. The Role of Orthographic/Motor Integration in the Production of Creative and Well-Structured Written Text for Students in Secondary School. *Educational Psychology*, 25(5):441–453, Oct. 2005.
- [6] L. J. Corrigan, C. Peters, G. Castellano, F. Papadopoulos, A. Jones, S. Bhargava, S. Janarthanam, H. Hastie, A. Deshmukh, and R. Aylett. Social-task engagement: Striking a balance between the robot and the task. In *Embodied Commun. Goals Intentions Workshop ICSR*, volume 13, pages 1–7, 2013.
- [7] D. Gouaillier, V. Hugel, P. Blazevic, C. Kilner, J. Monceaux, P. Lafourcade, B. Marnier, J. Serre, and B. Maisonnier. The NAO humanoid: a combination of performance and affordability. *CoRR*, 2008.
- [8] J. Han. Robot-Aided Learning and r -Learning Services. In D. Chugo, editor, *Human-Robot Interaction*. InTech, 2010.
- [9] D. Hood, S. Lemaignan, and P. Dillenbourg. When children teach a robot to write: An autonomous teachable humanoid which uses simulated handwriting. In *HRI*, 2015.
- [10] M. M. P. Hoy, M. Y. Egan, and K. P. Feder. A systematic review of interventions to improve handwriting. *Canadian Journal of Occupational Therapy*, 78(1):13–25, Feb. 2011.
- [11] J. Kennedy, P. Baxter, and T. Belpaeme. The robot who tried too hard: Social behaviour of a robot tutor can negatively affect child learning. In *HRI*, 2015.
- [12] C. D. Kidd and C. Breazeal. Robots at home: Understanding long-term human-robot interaction. In *Intelligent Robots and Systems*, IEEE, 2008.
- [13] I. Leite, G. Castellano, A. Pereira, C. Martinho, and A. Paiva. Empathic robots for long-term interaction. *ICSR*, 2014.
- [14] S. Lemaignan, J. Fink, F. Mondada, and P. Dillenbourg. You’re doing it wrong! studying unexpected behaviors in child-robot interaction. In *ICSR*, 2015.
- [15] D. Leyzberg, S. Spaulding, and B. Scassellati. Personalizing robot tutors to individuals’ learning differences. In *HRI*, 2014.
- [16] D. Llorens, F. Prat, A. Marzal, J. M. Vilar, M. J. Castro, J.-C. Amengual, S. Barrachina, A. Castellanos, J. Gómez, et al. The UJPenchars database: a pen-based database of isolated handwritten characters.
- [17] A. Matsui and S. Katsura. A method of motion reproduction for calligraphy education. In *Mechatronics (ICM)*, 2013.
- [18] K. Mülling, J. Kober, O. Kroemer, and J. Peters. Learning to select and generalize striking movements in robot table tennis. *The International Journal of Robotics Research*, 32(3):263–279, 2013.
- [19] H. L. O’Brien and E. G. Toms. What is user engagement? a conceptual framework for defining user engagement with technology. *Journal of the American Society for Information Science and Technology*, 59(6):938–955, 2008.
- [20] C. A. Rohrbeck, M. D. Ginsburg-Block, J. W. Fantuzzo, and T. R. Miller. Peer-assisted learning interventions with elementary school students: A meta-analytic review. *Journal of Educational Psychology*, 2003.
- [21] J. Saunders, C. L. Nehaniv, and C. Lyon. Robot learning of lexical semantics from sensorimotor interaction and the unrestricted speech of human tutors. In *New Frontiers Human-Robot Interact.*, 2010.
- [22] F. Tanaka and S. Matsuzoe. Children teach a care-receiving robot to promote their learning: Field experiments in a classroom for vocabulary learning. *Journal of Human-Robot Interaction*, 1(1), 2012.
- [23] J. Werfel. Embodied teachable agents: Learning by teaching robots. In *Intelligent Autonomous Systems*, 2013.
- [24] C. Zaga, M. Lohse, K. P. Truong, and V. Evers. The effect of a robots social character on childrens task engagement: Peer versus tutor. In *Social Robotics*, pages 704–713. Springer, 2015.