When Children Teach a Robot to Write: An Autonomous Teachable Humanoid Which Uses Simulated Handwriting

Deanna Hood, Séverin Lemaignan, Pierre Dillenbourg
Computer-Human Interaction in Learning and Instruction Laboratory (CHILI)
École Polytechnique Fédérale de Lausanne (EPFL)
firstname.lastname@epfl.ch

ABSTRACT

This article presents a novel robotic partner which children can teach handwriting. The system relies on the *learning by teaching* paradigm to build an interaction, so as to stimulate meta-cognition, empathy and increased self-esteem in the child user. We hypothesise that use of a humanoid robot in such a system could not just engage an unmotivated student, but could also present the opportunity for children to experience physically-induced benefits encountered during human-led handwriting interventions, such as motor mimicry.

By leveraging simulated handwriting on a synchronised tablet display, a NAO humanoid robot with limited fine motor capabilities has been configured as a suitably embodied handwriting partner. Statistical shape models derived from principal component analysis of a dataset of adult-written letter trajectories allow the robot to draw purposefully deformed letters. By incorporating feedback from user demonstrations, the system is then able to learn the optimal parameters for the appropriate shape models.

Preliminary in situ studies have been conducted with primary school classes to obtain insight into children's use of the novel system. Children aged 6-8 successfully engaged with the robot and improved its writing to a level which they were satisfied with. The validation of the interaction represents a significant step towards an innovative use for robotics which addresses a widespread and socially meaningful challenge in education.

1. INTRODUCTION

Handwriting difficulties in children at an early age often negatively affect the academic performance of the students [5], in addition to their self-esteem being adversely affected [14], causing them to shy away from expressing what they know [16]. Successful interventions for children with handwriting difficulties involve the student in many sessions where they are engaged in physically practising the skill [10]. However, the link between handwriting difficulties and low self-efficacy [6] results in children who are unmotivated to participate in such sessions, potentially leading to a developmental arrest in the acquisition of the skill.

The *learning by teaching* paradigm, which engages the target student in the act of teaching another, has been shown to produce motivational, meta-cognitive, and educational benefits in a range of disciplines [18]. To our best knowledge, the application of the paradigm to handwriting intervention remains, however, unexplored. One reason for this may due to the requirement of an appropriately unskilled peer for the target child to tutor, as this may present a logistical constraint if the target child is the lowest performer in their class. In some cases, it may be appropriate for a peer or teacher to simulate a naïve learner for the target child to teach. For handwriting, where one's skill level is visually evident, however, this acting is likely to be eventually detected. As such, there is motivation for the use of a teachable agent which can be configured for a variety of skill levels, and for which children do not have preconceptions about its handwriting ability.

We present the development of a novel teachable agent that intentionally makes mistakes typical of children learning handwriting. Through this capability, the robot can be taught by children, who themselves may learn through their teaching.

Within this article, Section 3 presents the novel work in the area of artificial intelligence to develop a learning algorithm suitable for a teachable agent in the context of handwriting. Section 4 details the extension of this algorithm to an embodied robotic learning agent, including the new approach for achieving simulated fine motor skills on commercially affordable humanoid robots such as the NAO. Section 5 explores the contributions made to the study of human-robot interaction, in discussing the use of the system with primary school children and its potential as a tool for addressing wider pedagogical research questions in education. Finally, Section 6 addresses the challenges which are faced in extending this system to a level suitable for long-term studies, and Section 7 concludes by reiterating the impact of the article's contributions.

2. RELATED WORK

Teachable computer-based agents have previously been used to encourage the "protégé effect", wherein students invest more effort into learning when it is for a teachable agent than for themselves [4]. As we are concerned with learning of a physical skill, the learning agent developed is embodied in a humanoid robot which is capable of physically demon-

HRI '15, March 02 - 05 2015, Portland, OR, USA

ACM 978-1-4503-2883-8/15/03 ...$15.00.
http://dx.doi.org/10.1145/2696454.2696479

strating handwriting trajectories to its child learning partner. This is supported by the potential for motor mimicry to yield significant improvements in handwriting interventions in which letter formations are demonstrated to participants [2]. Furthermore, when compared to screen-based agents, robotic partners have been shown in some contexts to increase users' compliance with tasks [1], maintain more effective long-term relationships [11], and produce greater learning gains when acting as tutors [12].

Robots have been used as teachers or social partners to promote children's learning in a range of contexts, most commonly related to language skills [9], and less often to physical skills (such as calligraphy [15]). Looking at the converse (humans *teaching* robots), Werfel notes in [22] that most of the work focuses on the robot's benefits (in terms of language [19] or physical [17] skills, for example) rather than the learning experienced by the human tutor themselves. Our work concentrates on this latter aspect: by demonstrating handwriting to a robot, we aim at improving the *child's* performance. Note that our work must be distinguished from "learning from demonstration" approaches to robots learning physical skills, as the agent we present is only simulating fine motor skills for interaction purposes.

A robotic learning agent which employs the learning by teaching paradigm has previously been developed by Tanaka and Matsuzoe [21]. In their system, children learn vocabulary by teaching the NAO robot to act out verbs. The robot is tele-operated and mimics the actions that the children teach it, but with no long-term memory or learning algorithm in place. Our project significantly extends this line of work in two ways. First, by investigating the context of children's acquisition of a challenging physical skill (handwriting), and second by proposing a robotic partner which is fully autonomous in its learning.

3. A LEARNING AGENT IN THE CONTEXT OF HANDWRITING

A parameterisation of letters and their deformities is used such that different quality shapes can be generated, depending on the parameters input to the letter models. This allows us to configure the system to improve its writing by modifying the parameters based on feedback from the reinforcement learning partner (Section 3.3).

3.1 Shape Modelling of Letters

We use statistical shape modelling for generating a shape model which can appropriately represent realistic variations in shapes. Statistical shape modelling is an application of principle component analysis (PCA), where a linear transform which de-correlates data vectors is found [20] and allows for dimensionality reduction.

PCA is performed on a set of letter paths captured from a digital pen, using the UJI Pen Characters 2 dataset [13] with 120 instances of each letter (2 repetitions from 60 adult users). While it may be appropriate in future work to identify the location of salient features of the shapes which are robust to unanticipated user input (such as shapes drawn backwards), the features are currently taken as $n = 70$ uniformly spaced points along the shape path. The points are arranged into an observation vector presented in (1), where x_i and y_i represent the coordinates of each of the points along the path. The observation shapes are normalised to have a unit maximum dimension and $\mathbf{0}$ mean.

$$\mathbf{x} = [x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n]^T \tag{1}$$

Equation (2) represents the projection from the original $2n$-dimensional feature space to a reduced N-dimensional space, where $\mathbf{p}$ contains the coordinates in the N-dimensional space with $\mathbf{0}$-origin, calculated as in (3). $\mathbf{\Phi}$ is an orthogonal $2n \times N$ matrix composed of the eigenvectors $\mathbf{v}_i$ corresponding to the largest N eigenvalues (λ_i) of the covariance matrix of the observations [20], as shown in (4). If there is correlation between the points in the observations, there will be eigenvalues of the covariance matrix which are close to zero. As such, removing the associated eigenvectors from $\mathbf{\Phi}$ allows for dimensionality reduction with minimal impact.

$$\tilde{\mathbf{x}} = \bar{\mathbf{x}} + \mathbf{\Phi}\mathbf{p} \tag{2}$$

$$\mathbf{p} = \mathbf{\Phi}^T(\mathbf{x} - \bar{\mathbf{x}}) \tag{3}$$

$$\mathbf{\Phi} = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_N]^T \tag{4}$$

PCA is performed on all of the paths of a particular allograph in the dataset individually, to reduce the $2n$-dimensional space for that shape to one with $N = 10$ dimensions. Each shape is then approximated by the mean shape of the allograph plus a sum of the top 10 eigenvectors, weighted by the parameter vector $\mathbf{p}$.

Equation (2) may also be used to generate new shapes based on the parameters $\mathbf{p}$ which are used. $\mathbf{p} = \mathbf{0}$ will yield the mean shape, and variations to each of the N values in $\mathbf{p}$ will cause a change in the shape represented by the corresponding eigenvector (Figure 1). For the dataset presented in Figure 1, the eigenvectors associated with the 3 largest eigenvalues explain 78.5% of the variance in the dataset, illustrating the capability of the statistical shape modelling approach to produce compact parameterisation of shapes.

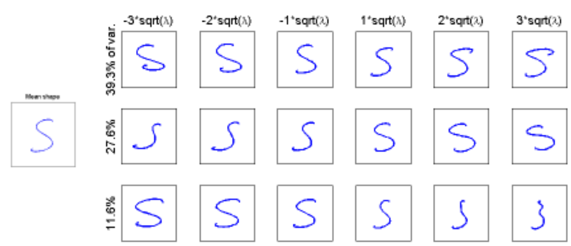

Figure 1: The mean shape (left) and the effect (right) of varying the first three parameters (each row) of a shape model. Parameter variation is dependent on the eigenvalue λ corresponding to the parameter's eigenvector. The percentage of the total variance in the dataset explained by each parameter is shown to the left of the corresponding row.

Interestingly, although the parameters are the result of an *unsupervised* shape analysis, they still represent variations which could have been intuitively identified by a manual parameterisation. For example, for the model shown in Figure 1, the parameters may represent the height of the top half of the letter compared to the bottom half, the width of the overall shape, etc. The ability to generate varied levels of deformations which may be ascribed descriptive interpretations (not just numerical) is an advantage of this method, given its intended use with humans. It is, for instance, possible for a teacher to create letters – which will be used as a starting point for the system – with a particular feature (a wide ‘s’ or a ‘d’ with a large loop, for example).

3.2 Generating Poor Letters

As explained, new letters can be generated by varying the parameter values for a shape model in accordance with (2). By choosing parameter values which lie within the observed range in the dataset, it is possible to produce letters which are more likely to be reasonable looking. When the parameter values are outside of the range observed in the dataset, they are less likely to represent shapes from the dataset of adult-written letters, and as a result are more likely to represent poor shapes. Figure 2 illustrates sample letters generated from the models of ‘e’ and ‘g’ by selecting random values for the first 5 parameters from a distribution with standard deviation of $3\sqrt{\lambda_i}$, rather than the $\sqrt{\lambda_i}$ standard deviation observed in the dataset.

Figure 2: Sample letters generated from the PCA shape model on ‘e’ (left) and ‘g’ paths (right), generated randomly from parameters with $3\times$ the standard deviation observed in the dataset.

In [3], Chandra found that children aged 4-6 years participating in a handwriting peer tutoring pilot study most often made mistakes qualitatively classified as *internal proportions* (inappropriate proportion of the different strokes within a letter), or *global deformations* (overall deformation in the appearance of the letter). As exemplified in Figure 2, the shapes generated by the system exhibit the same kind of deformities. Chandra identifies other, less common, mistakes which involve topological changes, such as letters being broken into subparts or mirrored. Using a database of children’s letters when available may yield potential for better parameterising these other mistakes. However, as an initial approximation, the shape models generated from PCA on a dataset of only adults’ writing have shown to be well-suited to generate ‘poor’ letters that children were able to identify as such and successfully improve.

3.3 Responding to Feedback

In addition to generating letters by varying input parameters, the statistical shape model of Section 3.1 may also be used to determine a particular letter’s parameters, given the model. The parameters of user-drawn letters may therefore be used in order implement a learning algorithm which adapts to the user’s feedback via demonstration letters.

The statistical shape model is used to determine the parameters of a demonstration shape by projecting the features of the observed shape into the lower-dimensional space determined by the model. Mathematically, the parameters associated with a demonstration $\mathbf{x}_{demo}$ are determined as in (3) with $\mathbf{x} = \mathbf{x}_{demo}$, and will reconstruct the closest approximate shape.

The method we employ for responding to user demonstrations is to move the learning algorithm’s parameters towards those of the demonstration. In the results presented in this work, the linear update equation shown in (5) is used, where $\mathbf{p}$ is the learned parameter vector at time step k, and α is the learning rate, between 0 and 1.

$$\mathbf{p}^{(k+1)} = \mathbf{p}^{(k)} + (\mathbf{p}_{demo} - \mathbf{p}^{(k)}) \times \alpha \tag{5}$$

Figure 3 illustrates the response of the system to demonstrations from a child for the letter ‘s’ using a learning rate of $\alpha = 1/2$. Observe that even poorly-written demonstrations allow the system to improve.

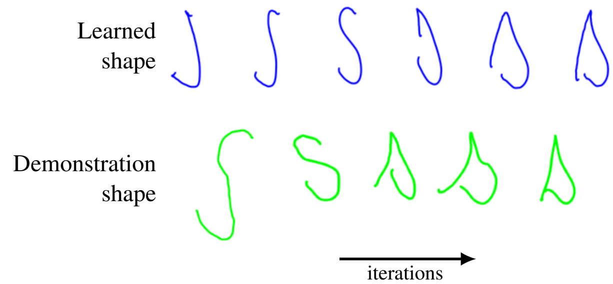

Figure 3: Example of the learning algorithm responding (top) to user demonstration of shapes (bottom) for the letter ‘s’ (demonstrations received from two 7-8 year-old children taking turns).

It is possible that parameters $\mathbf{p}^{(k)}$ and $\mathbf{p}_{demo}$, which individually yield acceptable shapes, produce parameters $\mathbf{p}^{(k+1)}$ which yield an unacceptable shape. This is especially true if the demonstration shape is of a different style to that learnt at time k (see Figure 4), as there are no restrictions imposed on parameter values. However, the proposed method for adapting to the demonstration shapes has the advantage of being able to recover from such a situation: with further demonstrations of the same letter, the system would eventually approach the demonstration shape. As a result, the event of poor-looking intermediate letters would not limit the interaction later proposed in Section 5 in a technical sense, but it may influence the user’s *perception* of the learning agent. It remains to be seen if it is necessary to avoid such an event mathematically.

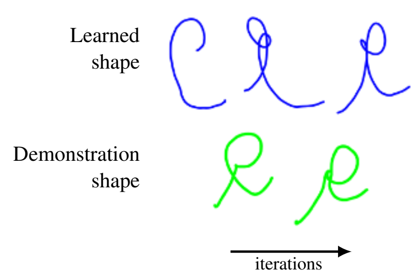

Figure 4: Example of the learning algorithm responding (top) to user demonstration of shapes (bottom) for the letter ‘e’, passing through a parameter state which yields a poor letter (demonstrations received from a 7-8 year-old child).

4. EMBODIMENT OF THE LEARNING AGENT WITH THE NAO HUMANOID ROBOT

In order to develop a teachable agent that is appropriate for engaging a child in the learning by teaching paradigm, we have established capabilities for the robot to engage in handwriting and interactive turn-taking.

The NAO V4 humanoid robot, which has been purposely designed by Aldebaran Robotics to look approachable [8], is used for this work. It is a commercially affordable biped robot, 58cm tall, with 25 degrees of freedom, two cameras, speech capabilities and the ability to autonomously execute a range of tasks.

Precise control over what the robot is writing is necessary in the proposed application of a handwriting. Because of the limited fine motor skills possible with such an affordable robot, in addition to the absence of force feedback and other technical necessities, we have configured the NAO to use simulated handwriting with a synchronised tablet to achieve this level of control.

The development of the necessary components for embedding the handwriting learning algorithm presented in Section 3 in the humanoid agent are presented in the sections which follow.

4.1 Robot Trajectory Following Movements

Using simulated handwriting provides an opportunity for the robot's writing to appear smoother than would be achievable with a writing instrument. However, the robot's motions must still sufficiently match the displayed trajectory in order capture the engagement of the child participant in the action. Aldebaran's NaoQi API is used for the inverse kinematics of the trajectory following. The Robot Operating System (ROS)[1] is used for integration of the NAO with external reference frames, such as the tablet's location, using the tf transformation library [7].

When using simulated handwriting, it is no longer necessary that the robot engages in the typical style of handwriting of using a writing instrument at a desk. Having the robot point at a vertical writing surface to cause the trajectory to appear (as in Figure 5) has several advantages:

- The working space of the robot increases, both in the technical sense and the interaction sense: someone can, in theory, show the tablet to the robot from across the room and have it still respond, without needing the tablet to be within arm's reach.

- Concerns about whether or not the child would start mimicking the robot's incorrect writing form (*e.g.* pen grip) are mitigated.

- Perhaps most significantly, the accuracy of the matching of the robot's motion with the trajectory displayed on the tablet is not as critical. This is because a pen tip would be expected to touch the tablet exactly at the trajectory point, while a fingertip may not.

We have therefore designed the system in such a way that the robot is simulating handwriting by pointing at the tablet[4]. As interacting with a tablet with one's finger is not uncommon, this may aid the acceptance of the writing style by users.

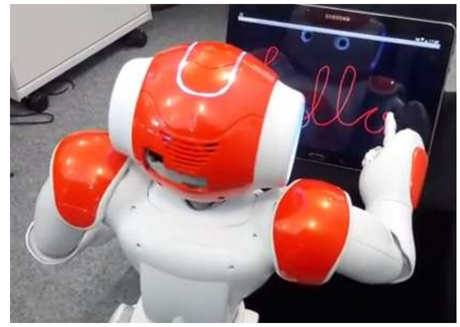

Figure 5: A demonstration of the robot simulating the writing of a word with its finger. The motion of the robot is synchronised with the display of the tablet, communicating over ROS.[3]

Because motion planning is performed with respect to the hand of the robot, rather than its fingertip, one or two of the orientation degrees of freedom of the hand are fixed to keep the finger approximately perpendicular to the writing surface, depending on the desired accuracy. The remaining free orientation(s), coupled with the whole-body motion control available, allow for a sufficient working space for writing on the entire tablet.

4.2 Synchronisation with the Tablet Trajectory Display

To enable the robot's 'writing' to display while the robot is tracing trajectories, ROS is used for the communication between the devices, including the Android tablet[5]. As a result, aspects of the networking between the tablet and the robot, such as the overheads associated with connections, ports, etc. have been simplified.

An Android application has been developed to receive the trajectory message over a ROS topic and display the trajectory as an animation. Synchronisation between the tablet and the robot is achieved by using NTP servers to synchronise device clocks; passing only the necessary number of points (7) to the robot's motion planner to improve timing accuracy; and not running computationally expensive tasks on the robot (such as camera publishing) while it is writing.

To instruct the robot where to write, the robot has been configured to detect a particular fiducial marker, a *chilitag*[6], with the camera located in its head, and to use that to determine the relative position of the writing surface (Figure 6). When used in an interaction involving a participant, this allows a user to move the tablet as required for the interaction. The tablet is assumed to be stationary during the

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

[3] See `https://www.youtube.com/watch?v=2qWFSJRxCU0` for a video of the synchronised writing demonstration.

[4] Teachers interviewed for their feedback on the system advised that children are asked to draw letters in the air in a similar manner as part of their handwriting education. The behaviour is hence not unfamiliar to children.

[5] For more information about ROS on Android devices see `http://wiki.ros.org/android`

[6] See `https://github.com/chili-epfl/ros_markers` for more information on the fiducial markers used.

writing process as detecting the tablet interferes with the robot's adherance to writing synchronisation.

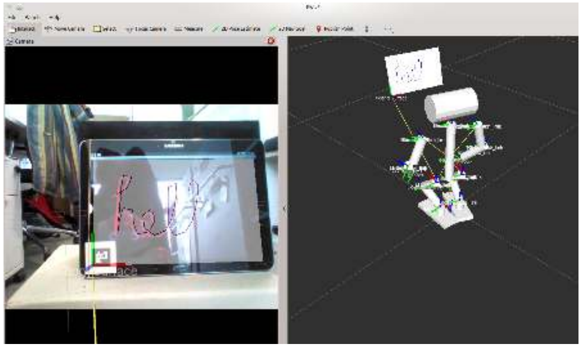

Figure 6: Detection of the tablet using a fiducial marker to represent the origin of the writing surface frame, visualised in RViz. The robot's camera image is on the left, with the text trajectory overlay visible.

4.3 Integration into a Teachable Robotic Agent

The fusion of the embodied handwriting agent developed with the handwriting learning algorithm presented in Section 3 involves the integration of three components: the robot, the tablet, and a central controller (Figure 7). The robot and Android tablet application present the writing process/result to the user, as explained in the previous section. The tablet application has been extended to act as the primary medium for capturing participant input, and submits the user's demonstrations when they are satisfied with their writing.

The user demonstrations are received by the interaction controller running on a desktop computer. It is responsible for getting the NAO to prompt and respond appropriately to feedback received using a finite state machine to manage the interaction stage and various system inputs. In the context of learning handwriting, an additional input to the system is a word from which the letters are to be learned, which is detected by a fiducial marker on the card displaying the word. The controller provides inputs to the learning algorithm including the word to learn and the user demonstrations, by inferring the letter which the demonstrations are intended for based on their position on the tablet. The output shapes from the learning algorithm are then sent again to the devices which write them.

The source code for the teachable robotic handwriting partner has been made available at `https://github.com/chili-epfl/cowriter_letter_learning`.

5. A TOOL FOR SOCIAL AND PEDAGOGICAL INVESTIGATIONS

In addition to constituting a technically novel system, the presented teachable robotic agent represents a tool which may be used for investigating social and pedagogical research questions. For example, one such question is what impact the addition of such a teachable robotic agent would have on the outcomes of a typical handwriting intervention. Preliminary studies at two schools in the Geneva area, involving over 50 children, have been conducted to evaluate the feasibility and technical soundness of the interaction system

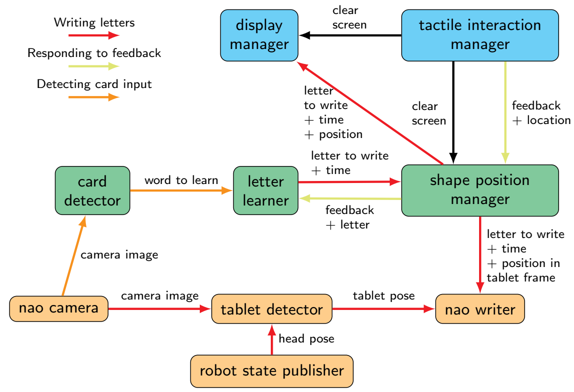

Figure 7: Overview of the system. Components in the top row run on the tablet, those in the middle row on the central controller, and those in the bottom row on the robot.

proposed as a tool for such investigations.

5.1 Interaction Context

Figure 8 illustrates an example interaction sequence between the participant and the robot which consists of the following stages:

1. The participant shows the robot one of seven different 3-letter words to write, made up of 7 possible letters ('c', 'e', 'n', 'o', 's', 'u', 'w'). Fiducial markers which are printed on the word cards allow them to be detected with the robot's camera.

2. The robot responds to the word request verbally and writes the letters according to the method described in Section 4.2.

3. The robot asks for feedback from the participant and they demonstrate how to write the letter which they feel needs to be corrected. The tablet may be moved into the most appropriate position for the child to write on it with the stylus. Only one letter may be demonstrated at a time and the position of the participant's demonstration on the tablet encodes the letter that it is a demonstration for. The participant can remove and repeat their letter if they are unhappy with it. When the participant is satisfied with the demonstration, they press a button on the tablet which signals that it is the robot's turn.

4. The robot writes an adapted letter in response to the participant's feedback, and the interaction iterates, with participants taking turns to interact with the robot if necessary. When the participant(s) is/are satisfied with the robot's performance on all letters, they may use the "test" card and an additional word for which they will verbally evaluate the robot's performance.

5.2 Outcomes of Preliminary Study 1

A pilot study at the first school consisted of four groups of approximately 8 english-speaking children each, aged 6-7 years. The children interacted with the system for a total of 65 minutes, with the robot writing 96 letters. 49 of

(a) The user shows a card to the robot with a word to write.

(b) The robot writes the word seen on the card and asks for feedback.

(c) The user provides feedback on the letters written via demonstration.

(d) The robot responds to the feedback, until the user is satisfied.

Figure 8: A user engaging with the robot in the *learning by teaching* interaction, using demonstrations as feedback.

these letters were in response to demonstrations from the children and the remaining were from when new words were requested.

As a result of the pilot study, we acted on two key observations. The first was that children appeared to have a difficult time providing demonstrations to the robot in the same place as previously-written letters. At the time, the system required the children to write on top of a letter of the same type as the one which they were demonstrating, and the children seemed to find this counter-intuitive and would occasionally just trace the robot's letter as it appeared. As such, the technical components of the system were extended to allow the children the opportunity to write around previously-written letters instead of on top.

The second key point which came from the pilot study was that children were observed giving advice to the child designated as the letter demonstrator, potentially giving rise to a higher level paradigm of learning by *teaching to teach*. As a consequence of this observation, the study which followed at the second school was designed to further observe the effect of the number of children interacting with the robot.

5.3 Outcomes of Preliminary Study 2

The study at the second school involved 21 french-speaking students aged 7-8 years. 7 children interacted with the robot individually and 7 sessions included the remaining 14 students interacting with the robot in pairs. Initial parameters were drawn from a range purposely chosen to generate shapes for letters 'e' and 's' which would elicit correction. For the other letters, parameters were fixed, generating shapes which some groups still chose to correct (e.g. 13/14 for 'n' and 2/14 for 'c'). The duration of the sessions was between 8 and 15 minutes, with an average of 11.4 minutes (SD = 2.3).

We have concluded following the second study that the system has been validated as a technically sound autonomous interaction. The interaction setup including the teachable robotic agent withstood the interaction which lasted for a total of 160 minutes. During this time the robot wrote 335 letters, 152 of which in response to demonstrations received from the 21 children. Technical intervention was only required for the three instances that the robot fell later in the day. Otherwise, the technical components of the system operated autonomously and as expected over the sessions.

Furthermore, no child indicated that they did not believe that the robot was writing by itself. There were, at times, questions about the robot's writing method at the beginning of the interaction, but when advised that the robot "tells the tablet what it wants to write," this was accepted by the children. On the rare occasion that the robot's writing was not correctly synchronised with the tablet, this did not appear to influence the children's impression. If older children participate in the interaction study – which may be likely as children with lasting handwriting difficulties are included as participants – it may become more important to invest time into the believability of the robot's writing scheme. However, for 6-8 year olds the proposed setup appears sufficient.

Regarding the engagement of the children in teaching the robot, an average of 10.9 demonstration letters (SD = 4.4) were provided to the robot for each session during the interaction. In 9 out of the 14 sessions (64%), the robot received demonstration letters even *after* reaching the test stage of the interaction. The participants' teaching after the test word had been written and evaluated – the only purposefully imposed external motivation – may suggest that by that time the participants had become intrinsically motivated to engage in the interaction, as we anticipated.

6. TOWARDS LONG-TERM STUDIES

A conclusion drawn in a systematic review of handwriting intervention studies [10] is that any of the studies considered which involved fewer than two practice sessions per week and fewer than a total of 20 practice sessions, including homework, were not found to demonstrate effective results. This highlights the necessity to engage students in an interaction which will be sustainable over the long-term if we want to address research questions which involve the measurement of learning gains. Several challenges are raised in developing such long-term capabilities for the system.

In terms of the interaction experience, the current experimental setting, while technically autonomous, can not robustly recover from situations outside of the nominal protocol presented in Section 5, and consequently still requires the supervision of an experimenter. The interaction finite state machine would require extension in non-trivial ways to allow for long, fully autonomous interactions with children.

In the current system, the robot can ask questions and prompt participants, but it cannot engage in discussions with the participants. It is clear that work is necessary to develop the conversational agent in the interaction so that the presence of an experimenter is not required for a captivating and continuing engagement. While there is the possiblity to focus the interaction design on group-based interaction with the robot in order to alleviate the necessity of a conversational agent, we find reason to believe that constructing such a social interaction is not a trivial task. Anecdotes from the preliminary studies have shown that some children may criticise another's demonstrations to the robot, which may or may not be as damaging to a child's self-efficacy as when they are criticised in a typical educational context.

How the children's perception of the robot as a learning agent may change over the long-term remains to be seen. On one occasion during the preliminary studies, a child's response to whether or not the NAO could write its own name (not previously demonstrated) included that it may have problems with the 'n', as the child had been correcting the robot on this letter. We believe that the user was projecting human-like learning features, such as forgetfulness, onto the robot, although they were not technically present in the system. This may need to be capitalised on when considering how to extend the interaction for long-term use, as the present system – with a learning rate such that progress is evident to the user – will cause convergence for a letter within a few iterations.

We expect that incorporating a database of letters drawn by children into the shape modeling process will facilitate generating shape models which capture a wider range of mistakes typical of children learning handwriting. However, the current system has conceptual – a PCA-only approach can not generate or learn a different shape topology – and technical – no support is currently implemented for shapes which require pen lifting between strokes – limitations which would need to be overcome.

If the system is extended to allow for a wider range of mistakes, a further topic for exploration then is how the handwriting error generation of the system may be abstracted to a higher level of control so that a teacher may configure it to work with a child on a particular type of mistakes based on the child's performance. Where would the balance lie between developing autonomous capabilities for the system to determine the child's difficulties and empowering the teaching staff to decide for themselves instead?

Addressing these challenges will take us further towards answering if the addition of a teachable robotic agent to handwriting interventions would benefit the participants' self-esteem, motivation, and learning gains.

7. IMPACT AND CONCLUSION

We believe that this article introduces three noteworthy contributions: an innovative application of data processing and artificial intelligence for the learning of hand-written letters suitable for educative purposes; a robotic system which was able to provide scaffolding for complex human-robot interactions (teacher-learner social interactions, learning by demonstration, simulated robotic fine motor skills) during two preliminaries studies; and an initial experimental investigation of what appears to be a new role for robots in education.

Specifically, the technical challenges involved in developing a teachable robotic agent in the context of handwriting which have been addressed in this work include:

- developing capabilities for a robot with limited fine motor capabilities, in particular the NAO robot, to engage in the act of handwriting in a way which is believable for interacting with children. This is accomplished by leveraging simulated handwriting with a synchronised tablet communicating via ROS;
- developing an algorithm capable of incorporating user feedback and demonstrations in order to adapt artificially generated handwriting quality so as to simulate a teachable agent, which has been implemented by maintaining a learning algorithm in the parameter space of the PCA-based shape models and converging towards the parameters of user demonstrations; and
- integrating the system into a working interaction suitable for engaging children in the learning by teaching paradigm, accomplished by fusing the robotic drawing capabilities and the learning algorithm for handwritten letters established with a central controller which manages the flow of the interaction, turn taking and integration of the connected devices.

However, we believe that the strongest impact of this work is for the human-robot interaction community and relates to the very *nature* of the interaction fostered by this research. The work presented here investigates a particular role for a robot in the education of handwriting: not only is the robot actively performing the activity by drawing letters, but it does so in a way that engages the child in a very specific social role. The child is the teacher in this relationship and the robot is the learner: the child must engage in a (meta-) cognitive relationship with the robot to try to understand why the robot fails and how to help it best. Here, the robot is more than just an activity facilitator or orchestrator – its physical presence and embodiment induce agency and anthropomorphising, and cognitively engage the child into the learning activity, which we predict will lead to higher learning efficacy.

Also notable, the robot is not used in the usual context of robotics or computer education, but instead in an activity – handwriting – which requires fine physical skills. In such activities, the embodied nature of the robot is appropriate as in interventions where motor mimicry is elicited [2] the arm motion for instance is, *by itself*, part of the teaching. Furthermore, when facing a child with school difficulties, robots can play the role of a naïve learner which neither adults nor peers – because of the social effects it would induce – can convincingly play. Along these lines, we hope to see more research on non-STEM educational applications of robotics.

The strong social impact of early educational problems makes continued research in this field an undoubtedly meaningful challenge for robotics and human-robot interaction.

Acknowledgments

This research was supported by the Swiss National Science Foundation through the National Centre of Competence in Research Robotics.

8. 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. *International Journal of Social Robotics*, 3(1):41–52, 2011.

[2] V. W. Berninger, K. B. Vaughan, R. D. Abbott, S. P. Abbott, L. W. Rogan, A. Brooks, E. Reed, and

S. Graham. Treatment of handwriting problems in beginning writers: Transfer from handwriting to composition. *Journal of Educational Psychology*, 89(4):652, 1997.

[3] S. Chandra. Co-writer: Learning to write with a robot, Dec. 2013. [Internal report, EPFL].

[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] B. Engel-Yeger, L. Nagauker-Yanuv, and S. Rosenblum. Handwriting performance, self-reports, and perceived self-efficacy among children with dysgraphia. *The American Journal of Occupational Therapy*, 63(2):182–192, 2009.

[7] T. Foote. tf: The transform library. In *Technologies for Practical Robot Applications (TePRA), 2013 IEEE International Conference on*, pages 1–6. IEEE, 2013.

[8] 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.

[9] J. Han. Robot-Aided Learning and r-Learning Services. In D. Chugo, editor, *Human-Robot Interaction*. InTech, 2010.

[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] C. D. Kidd and C. Breazeal. Robots at home: Understanding long-term human-robot interaction. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 3230–3235. IEEE, 2008.

[12] D. Leyzberg, S. Spaulding, M. Toneva, and B. Scassellati. The physical presence of a robot tutor increases cognitive learning gains. In *Proceedings of the 34th Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society*, 2012.

[13] D. Llorens, F. Prat, A. Marzal, J. M. Vilar, M. J. Castro, J.-C. Amengual, S. Barrachina, A. Castellanos, J. Gómez, et al. The UJIpenchars database: a pen-based database of isolated handwritten characters.

[14] T. Malloy-Miller, H. Polatajko, and B. Anstett. Handwriting error patterns of children with mild motor difficulties. *Canadian Journal of Occupational Therapy*, 62(5):258–267, 1995.

[15] A. Matsui and S. Katsura. A method of motion reproduction for calligraphy education. In *Mechatronics (ICM), 2013 IEEE International Conference on*, pages 452–457. IEEE, 2013.

[16] J. Medwell and D. Wray. Handwriting – a forgotten language skill? *Language and Education*, 22(1):34–47, 2008.

[17] 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.

[18] 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*, 95(2):240–257, 2003.

[19] J. Saunders, C. L. Nehaniv, and C. Lyon. Robot learning of lexical semantics from sensorimotor interaction and the unrestricted speech of human tutors. In *Proc. 2nd Int. Symp. New Frontiers Human–Robot Interact. ASIB Convent*, 2010.

[20] M. B. Stegmann and D. D. Gomez. A brief introduction to statistical shape analysis. *Informatics and Mathematical Modelling, Technical University of Denmark, DTU*, 15, 2002.

[21] 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.

[22] J. Werfel. Embodied teachable agents: Learning by teaching robots. In *Intelligent Autonomous Systems, The 13th International Conference on*, 2013.