A Humanoid Robot's Effortful Adaptation Boosts Partners' Commitment to an Interactive Teaching Task

ALESSIA VIGNOLO, Istituto Italiano di Tecnologia, Italy
HENRY POWELL, Institute of Neuroscience and Psychology, University of Glasgow, Scotland
FRANCESCO REA and ALESSANDRA SCIUTTI, Istituto Italiano di Tecnologia, Italy
LUKE MCELLIN, Department of Psychology, University of Warwick, United Kingdom
JOHN MICHAEL, Department of Cognitive Science, Central European University, Austria

We tested the hypothesis that, if a robot apparently invests effort in teaching a new skill to a human participant, the human participant will reciprocate by investing more effort in teaching the robot a new skill, too. To this end, we devised a scenario in which the iCub and a human participant alternated in teaching each other new skills. In the *Adaptive condition* of the *robot teaching phase*, the iCub slowed down its movements when repeating a demonstration for the human learner, whereas in the *Unadaptive condition* it sped the movements up when repeating the demonstration. In a subsequent *participant teaching phase*, human participants were asked to give the iCub a demonstration, and then to repeat it if the iCub had not understood. We predicted that in the *Adaptive condition*, participants would reciprocate the iCub's adaptivity by investing more effort to slow down their movements and to increase segmentation when repeating their demonstration. The results showed that this was true when participants experienced the *Adaptive condition* after the *Unadaptive condition* and not when the order was inverted, indicating that participants were particularly sensitive to the changes in the iCub's level of commitment over the course of the experiment.

CCS Concepts: • **Computer systems organization** → **Robotics**; • **Human-centered computing** → **User studies**;

Additional Key Words and Phrases: Commitment, human-robot interaction, kinematics, non-verbal communication, movement understanding

1 INTRODUCTION

As robots become increasingly prevalent in many domains of everyday life, such as healthcare, education, and manufacturing [3, 5, 7, 9, 22], researchers are devoting ever more attention to developing new ways of optimizing **human-robot interaction (HRI)**. Over the years, this has inspired efforts to design robots that look or act in ways that elicit some of the same responses from humans as other humans do (i.e., to design robots that humans perceive and respond to as though they were human) [19]. For example, some studies have documented evidence of "automatic imitation" [20] or 'co-representation' [25] on the part of humans interacting with a robot. Similarly, some studies have focused on establishing that robot behaviours can elicit trust from humans [8, 28]. Building on these advances, in many (e.g., industrial or healthcare) contexts, there has been a shift from the use of robots deployed as fully pre-programmed devices performing single predefined tasks towards the adoption of co-bots, able to learn new tasks from expert human demonstrators [16]. This shift brings about a new challenge: to boost human interactants' willingness to invest time and effort when interacting with a robot partner and in persisting in an interaction as the robot adapts to a new context and a new human partner. While there is a risk of human interactants becoming frustrated or impatient when a robot is slow to adapt, the potential benefits of adaptation are high insofar as they can maximise a robot's ability to contribute to new tasks with new partners.

To address this challenge, Powell and Michael [12] (cf. also[13]) have recently proposed that a potentially effective and low-cost strategy could be to develop design features that serve to maintain a human's sense of commitment to an interaction with a robot. In this context, they define the sense of commitment as a regulative mechanism which sustains the motivation to perform an activity despite temptations and distractions, and which is boosted by the perception or belief that a partner is relying on one to perform the activity. This way of conceptualising the sense of commitment highlights the possibility that the motivation to contribute to a partner's goals or to shared goals can be modulated by cues that one's partner values the joint action and may be relying on one to make one's contribution [15]. On this basis, recent research has begun to investigate the cues and situational factors that may trigger/modulate such a sense of commitment. Two recent studies [14, 33], for example, revealed evidence that a high degree of spatiotemporal coordination within joint action (i.e., the two agents form a chain in cleaning up a pile of sand, with one agent scooping and the other pouring) can engender a greater sense of commitment than would be present if there were only a low degree of spatiotemporal coordination (i.e., the same two agents do not form a chain but, rather, work in parallel). This leads observers of the joint action to expect both agents to remain engaged in the joint action for a longer time and to be more likely to persist until the goal is achieved. In a related strand of research, one recent study [26] found evidence in support of the hypothesis that the perception of a partner's effort elicits a sense of commitment to joint action, leading to increased persistence in the face of a temptation to disengage (see also [4]). Extending this research into the context of HRI, Székely and colleagues [27] have recently found evidence that the perception of a robot partner's apparent investment of cognitive effort boosted people's persistence on a boring task which they performed together with a robot.

Building upon these previous findings, we hypothesised that, if a robot invests effort in adapting to a human partner in a context in which the robot is teaching the human a new skill, she or he

will experience a sense of commitment, leading her or him to reciprocate by investing more effort and patience in a subsequent task in which she or he teaches the robot a new skill. In testing this hypothesis, the current study also draws upon research in developmental psychology which has been identified as having considerable potential for social robotics. Specifically, a wealth of research has shown that human infants benefit from the spontaneous use of *motionese* on the part of caregivers—i.e., caregivers slow down their movements and introduce more segmentation [2]. In the context of human-robot interaction, Vollmer et al. [34] found that human participants produce *motionese* in demonstrations directed towards a robot learner, and Nagai and Rohlfing [17] showed that a robot observer could be designed to pick up on, and extract information from, *motionese* produced by a human. The current study aims to extend this previous research by showing that a robot can implement *motionese* in teaching motion sequences to a human, and that doing so may facilitate interaction with a human partner in at least two ways. First, it may help to scaffold learning insofar as a system that can evaluate the performance of an end-user in real time and tailor its *motionese* to the specific learning needs of that end user could be particularly useful for real-life teaching scenarios. Second, it may facilitate interaction by eliciting a sense of commitment on the part of human learners, which may increase their willingness to reciprocate by investing more effort and patience in a subsequent task in which she or he teaches the robot a new skill.

To test our hypothesis, we devised a scenario in which the humanoid robot iCub [11, 21] and a human participant alternated in teaching each other new skills. The method in which a human and a robot iteratively switch roles has been successfully employed in [18], and it has been shown that cross-training provides improvements in team performance, in the perceived robot performance, and in the human's level of trust when compared to standard reinforcement learning techniques.

2 METHODS

The experiment was comprised of a *robot teaching phase* and a *participant teaching phase*, which alternated over the course of the experiment. During the robot teaching phase, the iCub demonstrated a sequence of movements which the participant was required to observe, memorise, and repeat. After the robot's demonstration, the participant was required to reproduce the sequence in the correct order. For each sequence, the participant was given two chances to perform the movement sequence correctly (meaning that if the participant did not understand the sequence the first time, the robot demonstrated it a second time). Across separate blocks, the iCub demonstrated the movement sequences with either a high or a low level of effort, thereby exhibiting either a high or a low degree of commitment to the teaching task. In both conditions, the robot first demonstrated the movement sequence at a baseline speed, and would repeat the sequence if the participant asked for a second demonstration. In the *Unadaptive condition*, the iCub repeated the sequence at a higher speed. This manipulation was based on our own informal observation that when people are uncommitted to a task, they often try to complete it as quickly as possible regardless of how negatively the task may be affected by their rushing through it. In the *Adaptive condition*, the robot also initially demonstrated the movement sequence at a baseline speed. Then, however, if the participant asked the robot to repeat the sequence, the iCub repeated it at a slower speed. Potentially, such behaviour may be perceived by the human as reflecting the iCub's genuine intention to facilitate the partner's task. Ideally, this should have made it easier for the participant to understand and reproduce the sequence but at the cost of requiring a greater investment of time. In the participant teaching phase, the participant was instructed to teach the iCub a word (which was displayed on the screen behind the robot) by drawing it in the air with her or his right hand. The participants were told that if the iCub did not understand the word the first time they would have been asked to repeat the demonstration. In reality, the 'failure' event was implemented at the same point in the experiment for every participant. This was done in order to ensure that all the

participants had to repeat the same words a second time—i.e., if it had been random, participants would have had to repeat different subsets of words, implying an uneven exposure to the stimuli for different participants.

We reasoned that, if the robot's investment of effort in adapting its teaching demonstration to the human learner elicited a sense of commitment on the part of the human, then this may lead the human to reciprocate the iCub's investment of effort in adapting her or his teaching demonstration to the iCub—i.e. it may lead the human to invest more effort in tailoring the demonstration to the robot in the *Adaptive condition* compared to the *Unadaptive condition*. Theoretically, we expect commitment to manifest as a willingness to invest time or effort despite temptations or distractions. In other words, commitment was operationalised in terms of effort invested by the human in order to ensure a successful interaction. Crucially, we implemented two strategies in the design of the task to be able to rule out any general mirroring effect. First, we were not interested in whether participants would demonstrate the sequence more slowly in one condition than the other, but whether they would slow down more when repeating the sequence a second time within each trial (i.e., adapting more from the first to the second repetition within each trial). Second, the movements that the human taught to the robot were of a different kind to the movements which the robot demonstrated to the human; this enabled us to ensure that participants were not simply mirroring specific movements. Instead, it enabled us to probe whether human participants would use the same general strategies as the robot (slowing down and introducing segmentation) flexibly and as appropriate in a different domain. We therefore predicted that when repeating a demonstration of a word which the robot did not understand (i.e., in the *participant teaching phase*), participants would reduce their velocity more in the *Adaptive condition* than in the *Unadaptive condition* (P1), and introduce longer segmentation pauses in between letters (P2).

We also predicted that the participants' responses to questionnaires would reveal more positive attitudes towards the iCub and towards the interaction in the *Adaptive condition* than in the *Unadaptive condition*. More specifically, we predicted that in the *Adaptive condition*, participants would report a greater sense of having taught the iCub successfully (P3), of having learned successfully from iCub (P4), of having helped the iCub to learn (P5), that they had a good relationship with the iCub (P6), and that they had been coordinating with the iCub (P7). We also predicted that they would report a greater willingness to wait for the iCub in a hypothetical social situation in which iCub was running late for a meeting with them (P8) and that they would perceive the iCub as more anthropomorphic (P9), animate (P10), likeable (P11), and intelligent (P12).

2.1 Experimental Setup

The pre-registration for this experiment can be accessed at: http://aspredicted.org/blind.php?x=7xz2rh. The experiment aimed to replicate an ecologically valid situation in which a robot and a human were involved in a teaching activity. Such situations could feasibly be realised in contexts such as schools or personal care. Participants carried out the task by standing in front of the iCub. Behind the robot was a television monitor on which the words used to teach the iCub during the *participant teaching phase* were displayed. A button was placed in front of the participant which was pressed by the participant in order to progress the experiment from one phase to the next. In order to monitor the experiment, we placed a hidden RGB-D (Primesense Carmine) camera in the lab. The feed from these was then fed through to a laptop hidden behind a curtain from where the experimenters could not only observe the participants carrying out the task but also control the task itself. The task was controlled by leveraging the robot middleware **YARP (Yet Another Robot Platform)** [10]. We created a YARP module comprised of a finite state machine that allowed us to remain in control of the iCub's behavior—specifically the onset and termination of its speech
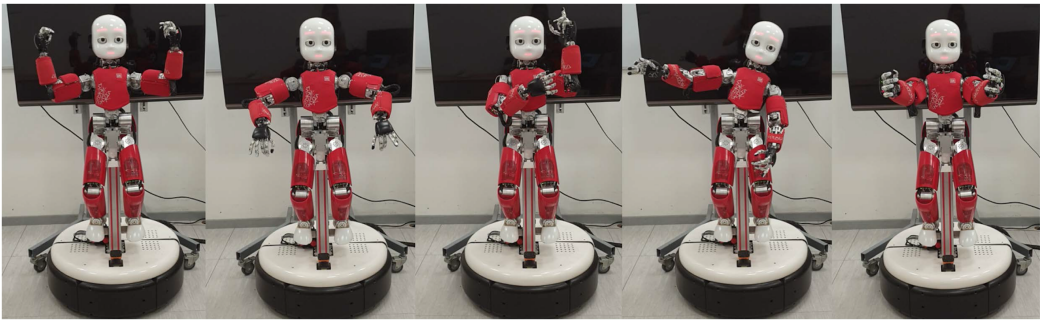
Fig. 1. *Robot teaching phase*: Example of sequence of movements.

and motion. During the *participant teaching phase*, participants would press the button in front of them in order to make a word appear on the screen. On a second key press, the robot would provide feedback as to whether or not it had understood the word that the participants had attempted to write. During the *robot teaching phase*, state transitions were triggered by the experimenter using the connected laptop. The main control of the robot involved the experimenter pressing the "Y" key on the laptop keyboard if the participant repeated the taught sequence of movements correctly and "N" if they did not. Each button pressed triggered the corresponding positive or negative feedback from the iCub. If the participant again failed to accurately reproduce the movement, the same "Y"/"N" procedure was implemented. A "Wizard of Oz" [24] paradigm was adopted to give the impression of an autonomous robot, whereby the hidden experimenter was needed to start the interaction at the beginning of the experiment and to assess the correctness of the sequence of movements repeated by the participant, triggering the behaviour of the robot accordingly. Although, some of the robot's behaviours, like the face tracking system, were autonomous.

Kinematic (x, y, and z position) data were captured using an Optotrak system with four active markers placed in the following way: two on participants' index fingertip, one on the pinky finger knuckle and one on the wrist. As the markers were active and thus involved wired connections to a receiver there was a risk that participant's movements would be constrained. To ensure this was not the case, we had each participant move their arms around freely and asked whether they still had full range of movement and whether they were comfortable while being connected to the Optotrak system. Any participants who reported discomfort or restricted movement had their markers reattached to the point that they were satisfied that this was not the case.

2.2 Robot Stimuli

We preprogrammed each of the movement sequences (an example on Figure 1) that the iCub taught the participants. This was done by first programming 22 distinct individual movements, each comprising some combination of the robot's head, torso and arms. From these 22 movements, we then created 12 sequences of 5 individual movements with each sequence being composed of randomly selected individual movements. There was a risk that the movement sequences would be improperly balanced with respect to their difficulty. This was worrying insofar as we wanted participants to struggle enough with each sequence so that they would have the chance to see the iCub make the variation in its effort. In order to avoid this as much as possible, we ensured that none of the individual movements had any obvious and commonly used communicative content (such as waving or giving the thumbs up), because communicative content may make certain movements easier to memorise. Second, we made sure that none of the sequences contained any repeated movements, as repetition could also simplify the task.

In the *Unadaptive condition* of the robot's teaching task, the robot performed six of the pre-recorded sequences wherein each movement i had duration T_i. When asked to repeat the movement, the iCub would demonstrate it a second time but with a duration of $0.75T_i$ for each movement. In the *Adaptive condition*, the robot did the same but repeated the sequence with a duration of $1.39T_i$ for each movement. These coefficient values were selected (after the piloting phase) because they allowed the robot's movements to be perceived with different velocities, to be not too long in the *Adaptive condition*, and not too fast and dangerous in the *Unadaptive condition*. For example, for a sequence with a duration of 15 seconds, the second demonstration would be of 11 ($\approx 15 \cdot 0.75$) seconds in the *Unadaptive condition* and 21 ($\approx 15 \cdot 1.39$) seconds in the *Adaptive condition*. We validated that slowing down is indeed more effective in teaching than rushing, and that it is perceived as more helpful [32].

A further variation was included in the robot's behaviour between the two conditions. In order to ensure that each individual movement in each sequence was clearly delineated, we programmed the iCub to return to a "home"/neutral position after each of the individual movements. In the *Adaptive condition,* the iCub took longer to reach the home position than in the *Unadaptive condition* when asked to repeat the sequence. This not only made the sequence in the *Adaptive condition* easier to learn in virtue of clearer segmentation of the individual movements but also gave the impression that the iCub was investing more effort than in the *Adaptive condition* because it was investing more time to ensure that the participant had an easier time learning the sequence.

In order to make the interaction more natural, the iCub was programmed to portray a set of social skills throughout the duration of the experiment. Specifically, a face-tracking module made the robot look at the participant's face during the course of the interaction. Second, we programmed the iCub's eyelids to close intermittently to simulate blinking. Finally, we included some random slight movements of the arms that gave the impression that the robot was breathing and making subtle natural variations in its posture in the same way that humans naturally do during social interactions. The iCub's speech was synthesised using a speech module and came out of an external speaker. We also had the iCub's speech come up on the left side of the television monitor to ensure that participants understood exactly what the iCub was saying.

2.3 Experimental Design

Participants and the iCub had two asymmetrical teaching tasks: in the *robot teaching phase*, the iCub taught participants sequences of five movements (see an example of sequence in Figure 1), while in the *participant teaching phase* (Figure 2, left), participants taught the iCub a word of five letters by drawing them in the air with their right index finger. During the *robot teaching phase*, participants had to memorise the movements of the sequence in the correct order and try to repeat them after the demonstration. The iCub then gave feedback by informing the participants whether they accurately recreated the sequence or not. Upon a failed attempt, the iCub would repeat the sequence once again in two different ways depending on the condition. In the *Adaptive condition,* the iCub would slow its movements down in order to make them more salient to the participants. In the *Unadaptive condition,* the iCub would speed up its movements as if it was trying to get through the task as quickly as possible. During the *participant teaching phase*, the participant was instructed to draw the word that appeared on the TV screen behind the iCub with their right hand, and to repeat the drawing once again if the robot informed the participant that it did not understand.

The experiment consisted of two initial sessions (baseline and familiarisation sessions), and two core blocks (with the two different conditions: Unadaptive and Adaptive blocks which were counterbalanced between participants). Before the experiment, participants were asked to fill out a *Godspeed questionnaire* [1] and to answer to the *Commitment Question* explained below in the
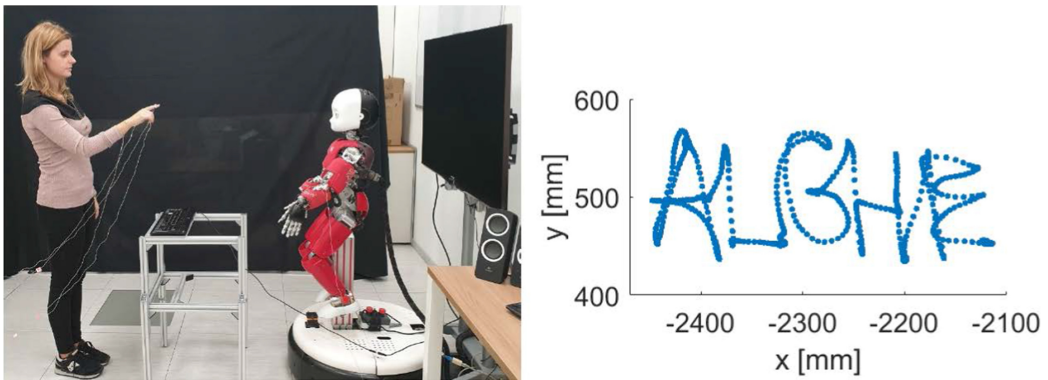
Fig. 2. Representative image of *Participant teaching phase*[1] (left) and an example of drawing trajectory (right).

list of questions, with the main goal being to assess their *a priori* feelings toward the iCub. In the familiarisation session, the participant was familiarised with the robot and the task by observing and trying to repeat three single movements individually and one sequence of three movements (as they would be presented in the experiment), and by teaching the iCub three single letters and a word of three letters in length. In the baseline block, participants had to draw a word which appeared on the TV screen in the same way as they would be doing in the experiment. The two core blocks consisted of six trials each (Figure 3, above). In each trial (Figure 3, below), iCub taught a movements and the participant taught a word. In both blocks, for one trial, the iCub informed the participants that it understood the word after the first demonstration (to give the impression that it could understand it the first time and in order to avoid participants thinking that the first attempt was always useless). For the remaining trials, the iCub asked the participants to repeat the word again. After the second repetition, the robot would thank the participant and say that it understood the word better. Between the two blocks and at the end of the experiment, the participants were asked to answer again the *Godspeed questionnaire* and the following questions:

> "Did you feel you were teaching something to the iCub?" (*Teaching Success Question*; 1, Not at all - 5, A lot) "Did you feel like you learned something from the iCub?" (*Learning Success Question*; 1, Not at all - 5, A lot) "Did you get the impression that you helped iCub when it found it difficult to understand a word?" (*Teaching Effort Question*; 1, Not at all - 5, A lot) "Describe the impressions you had about iCub during the session just ended. Did you feel that you had a good relationship with it?" (*Relationship Question*; 1, Not at all - 7, A lot) "Describe how you think iCub felt during the last session. Did it seem to it that your actions and words were coordinated with its words and actions?" (*Perceived Coordination Question*; 1, Not at all - 7, A lot) "Imagine that you and iCub were going to see a concert, which starts in 10 minutes and iCub is late. How long would you wait for iCub before entering to watch the concert?" (*Commitment Question*; 1, 5 minutes - 6, 30 minutes).

The whole duration of the experiment was approximately 1h and 15min.

2.4 Participants

We determined that 35 participants would be needed in order to observe a small- to medium-sized effect (80% power and an alpha level of .05), and therefore pre-registered a sample size of 35. Accordingly, 35 participants (mean age 32 years ± 11 SD) were recruited for the experiment; 18 were female, 17 male. The regional ethics committee ("Comitato Etico Regionale della Liguria")
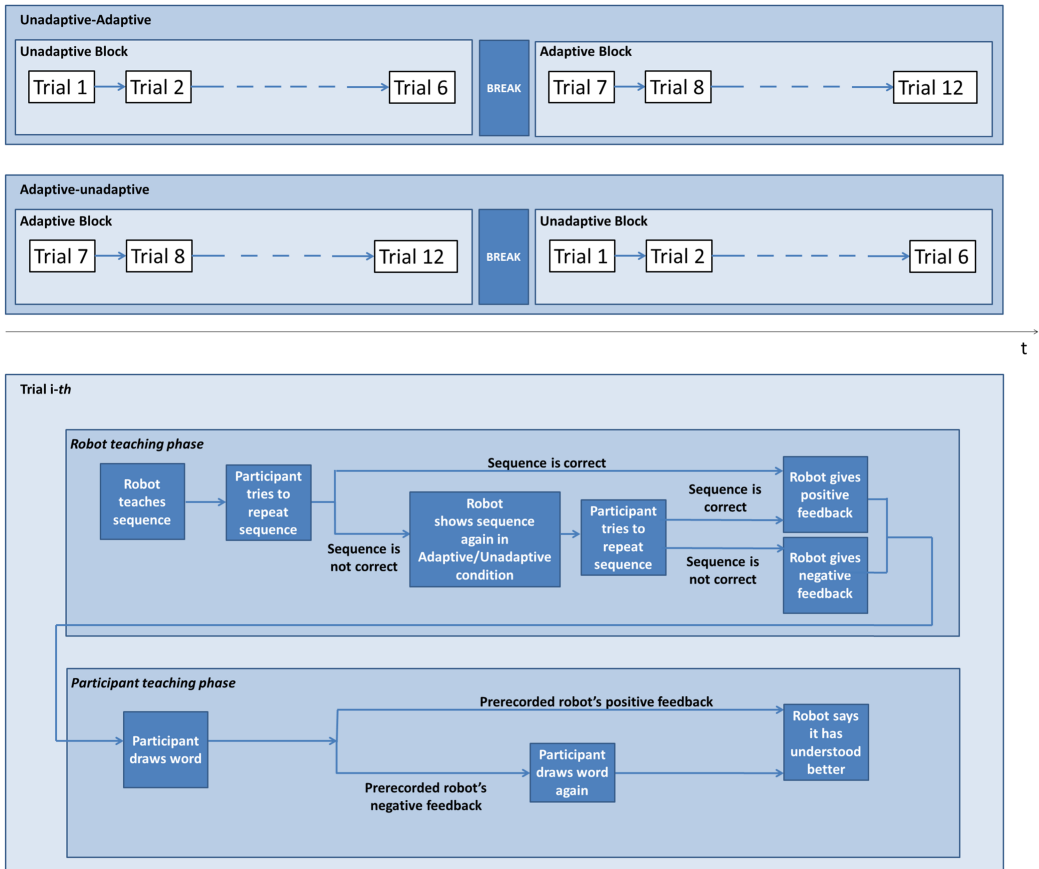
Fig. 3. Overview of the procedure.

2.5 Data Analysis

The kinematic data relating to each word was roughly segmented during the experiment by the participants' pressing of the button. Participants were asked to push the button immediately before and after each drawing trial. A short program written in MATLAB was then used to more precisely define when the beginning and ending of each movement occurred to ensure that the measure was consistent among participants. We defined a velocity threshold computed as 5% of the maximum velocity of each participant. Movement onset was defined as the point in time at which the movement exceeded this threshold, while the end of the movement was defined as the point at which the movement fell below this threshold after the last letter of the word had been finished. In some cases, participants forgot to press the button when their hand was in the starting position, making it difficult to determine the precise start and/or end point of those movements. To adjust for this, we determined the point at which the movement began as the minimum value of the position on the y axis after the previous movement had ended and when the marker returned to this minimum position after the word had been spelled. This was in order to capture when the participant's hand was resting on the button (which was the lowest point on the y-axis

for each movement) and when it returned to this position after spelling the word. An example of a segmented drawing is displayed in Figure 2, right.

For each kinematic feature, defined as a parameter of movement that is potentially relevant for the interaction, we computed the difference between the second and the first demonstration. Then, for each participant, we averaged the differences for each block and then averaged the results obtained for all the participants. Velocity was computed as the velocity on the vertical axis (as participants wrote the word on the vertical plane). Pause time was computed as the time in seconds that the velocity on the vertical axis was lower than a certain threshold. The threshold th was computed as: $th = [0.05(max(v_y) - min(v_y)] + min(v_y)$ where v_y is the velocity in the vertical direction.

Kinematic data as well as questionnaire answers were submitted to Mixed Model Anovas with adaptivity (Unadaptive or Adaptive) as repeated-measures factor and block order (Unadaptive-Adaptive or Adaptive-Unadaptive) as between-groups factor, followed by Tukey *post hoc* tests.

3 ANALYSES AND RESULTS

We analysed the kinematic data and the questionnaire answers of 35 naïve participants. The participants were divided into two groups: 18 were exposed to the *Unadaptive condition* as the first block and then to the *Adaptive condition* as the second block, while the other 17 were exposed to the opposite order.

We operationalised participants' commitment to teaching the robot by terms of the following parameters: the velocity of the movement and the segmentation between strokes (duration of pause time).

We analysed these variables for the two groups of participants, i.e., for the group that was exposed first to the *Adaptive condition* and then to the *Unadaptive condition* and for the group that was exposed to the conditions in the opposite order, considering separately the two blocks of trials.

3.1 Kinematic Data Analysis Results

Each participant saw the repetitions of the robot's movements for at least two trials in each block (this means that all participants performed the sequence incorrectly after the first demonstration at least twice), ensuring that all participants indeed experienced the manipulation. In particular, a Mixed Model Anova on the number of times participants failed to reproduce the sequence after the first demonstration, shows no effect of adaptivity condition (F(1,33) = 0.31, p = 0.582, η^2 = 0.006), of order (F(1,33) = 0.65, p = 0.426, η^2 = 0.011) or of the interaction between adaptivity condition and order (F(1,33) = 2.26, p = 0.144, η^2 = 0.042). Therefore, participants in the two groups were equivalent in this respect, as they failed to reproduce the sequence after the first demonstration in each condition and for each group a similar number of times: participants in the group who had the order Unadaptive-Adaptive were exposed to the second demonstration a similar number of times in the *Unadaptive condition* (M = 5.80, SD = 0.41) and in the *Adaptive condition* (M = 5.33, SD = 1.18); participants in the group who had the order Adaptive-Unadaptive were also exposed to the second demonstration a similar number of times in the *Unadaptive condition* (M = 5.29, SD = 0.91) and in the *Adaptive condition* (M = 5.50, SD = 0.65).

For each kinematic variable, we computed the difference between the first demonstration and the second, starting from the kinematic data of the finger registered with the Optotrak.

A Mixed Model Anova on the velocity differences (Figure 4, left) shows that there was no main effect of adaptivity (F(1,33) = 0.70, p = 0.410, η^2 = 0.005), and no order effect (F(1,33) = 1.44, p = 0.238, η^2 = 0.029), while the interaction between adaptivity and block order was significant (F(1,33) = 8.16, p = 0.007, η^2 = 0.061). A Tukey *post-hoc* test shows a marginal effect (t(33) = 3.66, p = 0.065, 95% CI
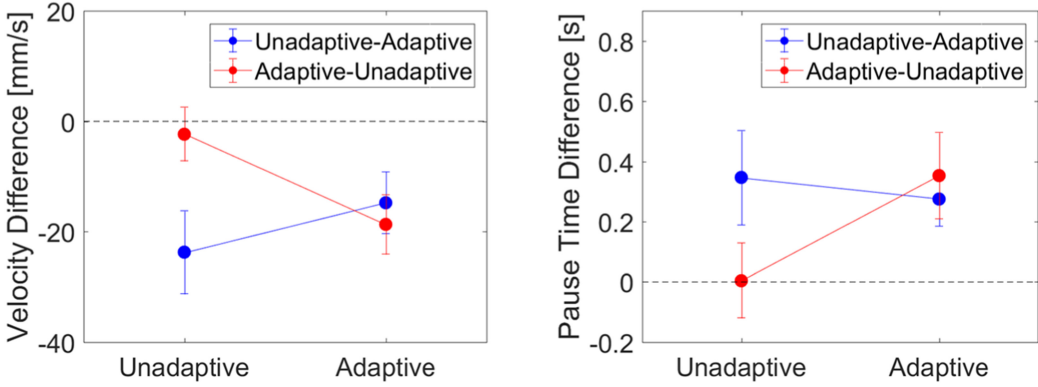
Fig. 4. Velocity and pause time differences. Values were derived by computing the difference between the first demonstration and the second, averaging the difference per block and then averaging it across participants that had the same block order. Error bars represent standard errors.

[−43.85 0.96]): in the *Unadaptive condition*, participants who had the order Unadaptive-Adaptive (*M* = −23.74, *SD* = 31.82) slowed down more in the *Unadaptive condition* than participants who had the order Adaptive-Unadaptive (*M* = −2.30, *SD* = 20.76).

A Mixed Model Anova on the pause time differences (Figure 4, right) shows a marginal effect of adaptivity (F(1,33) = 3.12, *p* = 0.087, η² = 0.015) and an effect of the interaction between adaptivity condition and order (F(1,33) = 7.14, *p* = 0.012, η² = 0.035) while there was no order effect (F(1,33) = 0.59, *p* = 0.448, η² = 0.014). A Tukey post-hoc test shows no significant effect within *Unadaptive condition* between people who had a different order of blocks (*t*(33) = 2.64, *p* = 0.262, 95% CI [−0.15 0.84]), within *Adaptive condition* between people who had a different order of blocks (*t*(33) = 0.61, *p* = 0.973, 95% CI [−0.57 0.42]), within order Unadaptive-Adaptive between the two adaptivity conditions (*t*(33) = 0.56, *p* = 0.979, 95% CI [−0.4 0.56]), order Adaptive-Unadaptive between the two adaptivity conditions (*t*(33) = 2.66, *p* = 0.256, 95% CI [−0.85 0.15]).

To further explore the interaction effects which we observed for pause time and velocity, we conducted a series of one-sample *t*-tests comparing the differences with 0 (Figure 5).

The one-sample *t*-tests show that in the first block, both participants who had the order Unadaptive-Adaptive (*t*(17) = −3.17, *p* = 0.006, *d* = −0.75, 95% CI [−39.57 −7.92], *M* = −23.74, *SD* = 31.82) and the order Adaptive-Unadaptive (*t*(16) = −3.39, *p* = 0.004, *d* = −0.82, 95% CI [−30.40 −7.01], *M* = −18.70, *SD* = 22.75) slowed down significantly in the second demonstration. In the second block, participants who had the order Unadaptive-Adaptive continued to slow down (*t*(17) = −2.64, *p* = 0.017, *d* = −0.62, 95% CI [−26.57 −2.94], *M* = −14.75, *SD* = 23.75), while this was not the case for participants who had the order Adaptive-Unadaptive (*t*(16) = −0.46, *p* = 0.654, *d* = −0.11, 95% CI [−12.97 8.37], *M* = −2.30, *SD* = 20.76).

The same pattern was observed for pause time: the one-sample *t*-tests reveal that in the first block, both participants who had the order Unadaptive-Adaptive (*t*(17) = 2.21, *p* = 0.041, *d* = 0.52, 95% CI [0.02 0.68], *M* = 0.35, *SD* = 0.66) and participants who had the order Adaptive-Unadaptive (*t*(16) = 2.40, *p* = 0.029, *d* = 0.58, 95% CI [0.04 0.67], *M* = 0.35, *SD* = 0.61) segmented the movement significantly more in the second demonstration. In the second block, participants who had the order Unadaptive-Adaptive continue to segment more in the second demonstration (*t*(17) = 3.05, *p* = 0.007, *d* = 0.72, 95% CI [0.08 0.47], *M* = 0.28, *SD* = 0.38), while this was not the case for those participants who had the order Adaptive-Unadaptive (*t*(16) = 0.04, *p* = 0.969, *d* = 0.01, 95% CI [−0.27 0.28], *M* = 0.01, *SD* = 0.53).
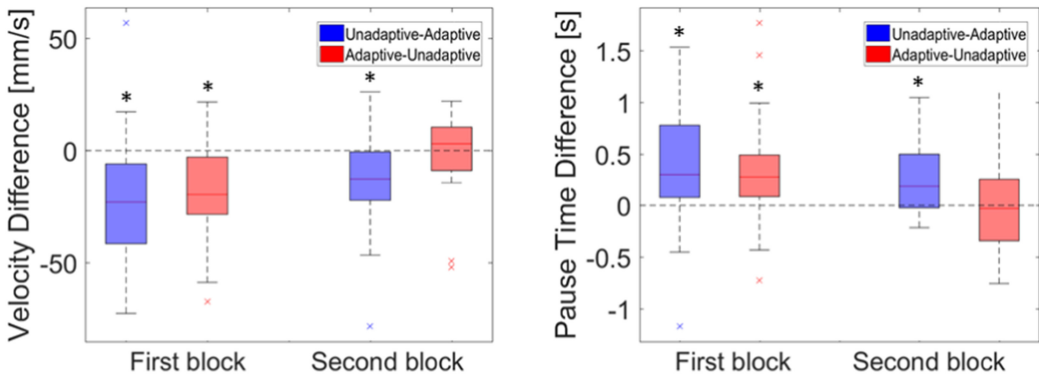
Fig. 5. Box plot of velocity and pause time differences in the two experimental blocks. Values were derived by computing the difference between the first demonstration and the second, averaging the difference per block and then averaging it across participants that had the same block order. The horizontal line is the median.

The results show that in the first block, independently of the condition, when demonstrating the word for a second time, participants slowed down the movement and introduced greater segmentation between the letters by making longer pauses. In the second block, participants who experienced an increase in the robot's effort between blocks (from Unadaptive to Adaptive) continued to slow down the movement and make longer pauses in the second repetition of the word, while participants who experienced a decrease of the robot's effort between blocks (from Adaptive to Unadaptive) stopped this behaviour—i.e., the velocity and pause time no longer differed significantly between the first and second repetitions (Figure 4).

We also performed an ex-post statistical power analysis: the analysis indicated that, given the sample size of 35 participants, this test had 20% power for detecting an effect size Cohen's f = 0.1, 81% for Cohen's f = 0.25, and 99% for Cohen's f = 0.5. This implies that the current study had sufficient power to detect medium to large effects, but was not sufficiently powered to detect smaller effects.

3.2 Questionnaire Results

We performed a series of Mixed Model ANOVAs with adaptivity as a within factor and block order as a between factor on the answers given by the participants after each experimental block (Figure 6, Figure 7).

No significant difference was found for the *Teaching Success Question* or for the *Teaching Effort Question*. This indicates that, while participants varied their level of effort in response to that of the iCub, they were not conscious of this.

The analysis of responses to the *Learning Success Question* revealed an interaction effect (F(1,33) = 7.80, p = 0.009, η^2 = 0.036) and a main effect of adaptivity (F(1,33) = 5.63, p = 0.024, η^2 = 0.026). The Tukey post-hoc test did not show any significant difference: participants who had the Unadaptive-Adaptive order judged that they had learnt more from iCub in the Adaptive block (M = 3.28, SD = 1.13) than in the Unadaptive block (M = 2.56, SD = 0.92), but the difference was not significant (t(33) = 3.12, p = 0.143, 95% CI [−1.61 0.16]).

The analysis of responses to the *Relationship Question* revealed an interaction effect between adaptivity condition and order (F(1,33) = 9.23, p = 0.005, η^2 = 0.032) and a main effect of adaptivity (F(1,33) = 5.23, p = 0.029, η^2 = 0.018). The Tukey post-hoc test did not show any significant difference: participants who had the Unadaptive-Adaptive order perceived a stronger relationship with
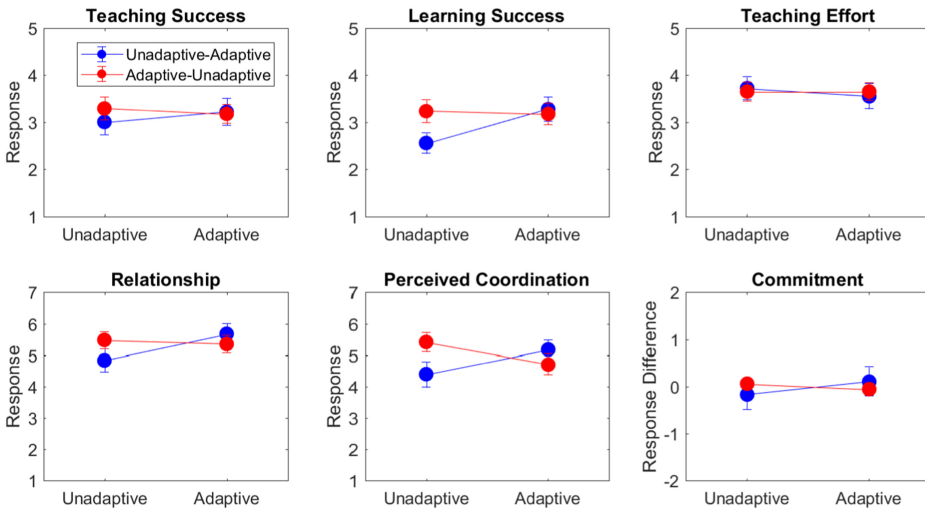
Fig. 6. Responses to the questionnaires. For the *Commitment Question*, the response difference was derived by computing the difference between the response given after each block with the response given in the baseline (i.e., response given prior to the experiment). Values were derived by averaging the response or the response differences per block and then averaging it across participants that had the same block order. Error bars represent standard errors.
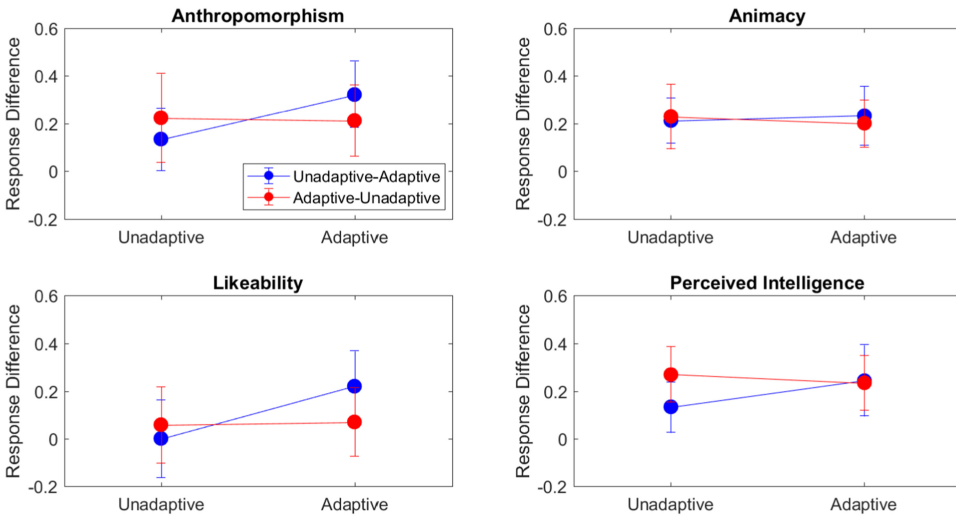
Fig. 7. Responses to the *Godspeed questionnaire*. The response difference was derived by computing the difference between the response given after each block with the response given in the baseline (i.e., response given prior to the experiment). Values were derived by averaging the response or the response differences per block and then averaging it across participants that had the same block order. Error bars represent standard errors.

iCub after the Adaptive block ($M = 5.67$, $SD = 1.46$) than in the Unadaptive block ($M = 4.83$, $SD = 1.54$), but the difference was not significant ($t(33) = -0.83$, $p = 0.434$, 95% CI [-2.01 0.34]).

The analysis of responses to the *Perceived Coordination Question* revealed an interaction effect between adaptivity condition and order ($F(1,33) = 9.82$, $p = 0.004$, $\eta^2 = 0.064$). The Tukey post-hoc

test did not show any significant difference: participants who had the Unadaptive-Adaptive order perceived iCub thinking to have more coordinated actions and words with them after the Adaptive block (M = 5.17, SD = 1.34) than in the Unadaptive block (M = 4.39, SD = 1.72), but the difference was not significant (t(33) = 1.66, p = 0.639, 95% CI [−2.09 0.54]).

For the *Commitment Question* (which was posed before the experiment as well as after each block), responses given prior to the experiment were treated as a baseline; we computed the difference between the response given after each block with the baseline. An interaction effect between adaptivity condition and order was found (F(1,33) = 6.12, p = 0.019, η^2 = 0.010). The Tukey post-hoc test did not show any significant difference: participants who had the order Unadaptive-Adaptive increased their commitment more after the *Adaptive condition* (M = 0.11, SD = 1.32) than after the *Unadaptive condition* (M = −0.17, SD = 1.34), but the difference was not significant (t(33) = 1.19, p = 0.833, 95% CI [−1.17 0.61]).

For the *Godspeed questionnaire* (which was posed before the experiment as well as after each block), responses given prior to the experiment were treated as a baseline; we computed the difference between the response given after each block with the baseline. For *Anthropomorphism*, a marginal interaction effect between adaptivity condition and order was found (F(1,33) = 2.51, p = 0.123, η^2 = 0.007). The Tukey post-hoc test did not show any significant difference: participants who had the order Unadaptive-Adaptive increased *Anthropomorphism* responses more after the *Adaptive condition* (M = 0.32, SD = 0.595) than after the *Unadaptive condition* (M = 0.13, SD = 0.55), but the difference was not significant (t(33) = 1.30, p = 0.796, 95% CI [−0.75 0.37]). For *Animacy*, there was no main effect of the adaptivity, no order effect and no interaction effect. For *Likeability*, a marginal interaction effect between adaptivity condition and order (F(1,33) = 3.24, p = 0.081, η^2 = 0.007) and a marginal adaptivity condition effect (F(1,33) = 4.00, p = 0.054, η^2 = 0.009) were found. The Tukey post-hoc test did not show any significant difference: participants who had the order Unadaptive-Adaptive increased *Likeability* responses more after the *Adaptive condition* (M = 0.22, SD = 0.62) than after the *Unadaptive condition* (M = 0.00, SD = 0.69), but the difference was not significant (t(33) = 1.51, p = 0.71, 95% CI [−0.78 0.34]). For *Perceived Intelligence*, there was no main effect of the adaptivity, no order effect, and no interaction effect.

4 DISCUSSION

The current study probed whether the iCub's effortful adaptation to the human would lead the participant to experience a sense of commitment, and consequently to reciprocate by investing additional effort to adapt her or his teaching demonstration to the robot during a subsequent *participant teaching phase*. In relation to the first prediction (P1), the results of the velocity analysis did not reveal a main effect of the adaptivity condition but did reveal a significant interaction between the adaptivity condition and the order in which participants were exposed to each condition. As Figure 4 shows, both groups of participants (i.e., those who were first exposed to the *Adaptive condition* and then to the *Unadaptive condition*, and those who were first exposed to the *Adaptive condition* and then the *Unadaptive condition*) tended to slow down their movements when repeating word demonstrations in the *participant teaching phase*. However, the group that first experienced the *Adaptive condition* and then the *Unadaptive condition* stopped doing this during the second (i.e., the Unadaptive) block. In contrast, the group that first experienced the *Unadaptive condition* and then the *Adaptive condition* continued to slow down their movements when repeating word demonstrations in the second (i.e., the Adaptive) block. The results from the *t*-tests (see Figure 5) confirm that this decrease in velocity when repeating demonstrations was significantly different from zero. Insofar as slowing down demonstrations requires a teacher to adapt her or his teaching style to the learner, it constitutes an investment of cognitive effort in the teaching interaction. Insofar as it entails an increase in the overall movement duration, it also constitutes an

investment of biomechanical effort. These findings therefore support the hypothesis that participants' sense of commitment was more stable when the robot increased its effort over the course of the experiment than when the robot's effort level decreased, and thereby had a positive impact on their willingness to invest effort in teaching the robot. It is worth emphasizing that participants' slowing down is unlikely to have been due to a general mirroring effect. There are two reasons for this. First, we did not probe whether participants would demonstrate the sequence more slowly in one condition than the other, but whether they would slow down more when repeating the sequence a second time within each trial (i.e., adapting more from the first to the second repetition within each trial). Second, the movements that the human taught to the robot were of a different kind to the movements which the robot demonstrated to the human; this enabled us to ensure that participants were not simply mirroring specific movements. Instead, it enabled us to probe whether human participants would use the same general strategies as the robot (slowing down and introducing segmentation) flexibly and as appropriate in a different domain.

In relation to the second predictions (P2), the results of the pause-time analysis reveal a marginal effect of adaptivity: the extent to which participants increased the segmentation between letters from the first to the second demonstration was greater in the *Adaptive condition* than in the *Unadaptive condition*. This supports the hypothesis that the robot's apparent investment of effort in adapting its demonstration to facilitate learning elicited a sense of commitment on the part of the participants, leading them to reciprocate by investing more effort to adapt their demonstration to the robot learner. As with the velocity analysis, we also observed an interaction between effort and the order in which participants were exposed to each condition. As Figure 4 shows, both groups of participants (i.e., those who were first exposed to the *Adaptive condition* and then to the *Unadaptive condition*, and those who were first exposed to the *Unadaptive condition* and then the *Adaptive condition*) tended to increase segmentation when repeating word demonstrations in the *participant teaching phase*. However, the group that first experienced the *Adaptive condition* and then the *Unadaptive condition* stopped doing this during the second (i.e., the Unadaptive) block. In contrast, the group that first experienced the *Unadaptive condition* and then the *Adaptive condition* continued to do so in the second (i.e., the Adaptive) block. The results from the *t*-tests (see Figure 5) confirm that this increase in pause time when repeating demonstrations was significantly different from zero. Insofar as increasing pause time requires a teacher to adapt her or his teaching style to the learner, it constitutes an investment of cognitive effort in the teaching interaction. Insofar as it requires the teacher to stop and re-start movements, this also constitutes an investment of biomechanical effort. Thus, these findings suggest that participants' sense of commitment was more stable when iCub increased its effort than when it decreased its effort from the first block to the second block. This impacted upon their willingness to invest effort in teaching the robot.

The results from the questionnaires reveal a pattern that is broadly consistent with the results from the motion data. Responses to the *Learning Success Question* indicate a trend that participants judged that they were better able to learn in the *Adaptive condition* than in the *Unadaptive condition* (P4). Similarly, responses to the *Relationship Question* revealed a trend that participants perceived a stronger relationship with the iCub after the Adaptive block (P6), and responses to the *Perceived Coordination Question* indicate a trend that participants experienced a higher degree of coordination with the iCub in the *Adaptive condition* (P7). For responses to the —i.e., with respect to the differences between responses given after each block and prior to the experiment—we did not observe a main effect of condition, but we do see a marginal effect of the interaction, with responses given after the second block of the experiment indicating higher levels of commitment for those participants who experienced the *Adaptive condition* after having experienced the *Unadaptive condition* (P8). One reason why we may not have observed a clear effect for the *Commitment Question* is that the scenario (in which the participants imagines having to wait for the iCub to show up to

a social engagement) may not have been perceived as realistic; future research should attempt to improve on this. These results are consistent with the pattern observed for the motion data: when the iCub's adaptivity decreases over the course of the experiment, participants' commitment is less stable than when its adaptivity increases over the experiment. Finally, it is interesting to note that in the responses to the *Teaching Success Question* and *Teaching Effort Question*, we do not observe any significant differences between conditions, suggesting that participants were not aware of their own differential responsiveness to the *Adaptive condition* and the *Unadaptive condition* (P3, P5), and also regarding the *Godspeed questionnaire* we do not observe any significant differences between conditions (P9, P10, P11, P12).

5 CONCLUSION

Our findings build upon a recent body of research investigating how movement kinematics can be adapted to increase mutual understanding and intuitive collaboration in the context of HRI [29–31]. Indeed, the potential to make robots' movements more easily legible to human interactants is a crucial goal of current and future research in social robotics [6, 23]. Our findings go beyond this by highlighting the possibility that the adaptation of movement kinematics may be used not only to increase legibility but also to enhance human interactants' persistence, effort and patience within human-robot interactions. In this respect, the current study links research on legibility with research on the sense of commitment [14], which may offer considerable potential in the context of social robotics. Indeed, Székely et al. [27] have already shown that by eliciting human interactants' sense of commitment to an interaction with a robot, their persistence and patience can be enhanced. In the current study, the manipulations of adaptivity were more implicit and based only on kinematic features. Combining our results with those of Székely and colleagues, it would be valuable to explore whether human participants would also persist longer in interactions with an adaptive robot than with an unadaptive robot.

Our findings also complement previous efforts to implement *motionese* in the context of human-robot interaction. First, they do so by showing that *motionese* may help to scaffold learning insofar as a system that can evaluate the performance of an end-user in real time and tailor its *motionese* to the specific learning needs of that end user could be particularly useful for real-life teaching scenarios. Second, our results show that *motionese* in may facilitate HRI by eliciting a sense of commitment on the part of human learners, which may increase their willingness to reciprocate by investing more effort and patience in a subsequent task in which she or he teaches the robot a new skill. In teaching more complex action sequences, the potential to use movement kinematics not only to optimise teaching in HRI but also to generate and maintain a sense of commitment has important implications in such contexts as physiotherapy, exercise classes, or other skill training programs. In particular, if movement kinematics can be used as an effective and inexpensive strategy for boosting human learning from robots and for building up a sense of commitment to the interaction, then humans may not only find it easier to learn, but may also be more motivated to do so. Extending this line of research further, researchers should also explore whether manipulation of movement kinematics may be especially useful if it is implemented such that specifically the most difficult parts of a sequence of movements are highlighted. In addition, the manipulation of kinematics may be complemented by other techniques, such as verbal encouragement or patterns of eye gaze, in order to further bolster a human's commitment to an interaction with a robot (in the current study, other such modalities were kept constant across conditions in order to maintain experimental control).

It would be valuable for future research to replicate our results—particularly the interaction between Adaptivity and Block Order, as this was not predicted *a priori*. It would also be important to investigate other contexts in which robot *motionese* may facilitate human learning, such as in

producing or using novel tools or machines. It would be useful to actually try to train the robot to read the words and to give real feedback depending on the behaviour of the participant. In the current study, the iCub's behaviour merely appeared as adaptive to the observers—i.e. it did not actually sense human responses and accordingly decide how to change its movements but, rather, executed a pre-programmed series of actions designed to exemplify the experience of adaptation. Future work will be dedicated to enabling robots to be sensitive to the degree of commitment exhibited by a human partner in order to adjust its movements accordingly. Finally, it would also be important to investigate to what extent the skills or information learned with the help of robot *motionese* are recalled after several weeks or months—in other words, to probe whether robot *motionese* also facilitates the automatisation of new skills or the encoding of new information in long-term memory.

REFERENCES

[1] C. Bartneck, D. Kulic, E. Croft, and S. Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics* 1(1), 71–81.

[2] R. J. Brand, D. A. Baldwin, and L. A. Ashburn. 2002. Evidence for 'motionese': Modifications in mothers' infant-directed action. *Developmental Science* 5, 1 (2002), 72–83.

[3] C. Breazeal, A. Brooks, J. Gray, G. Hoffman, C. Kidd, and H. Lee. 2004. Humanoid robots as cooperative partners for people. *Journal of Humanoid Robots* (2004), 34.

[4] M. Chennells and J. Michael. 2018. Effort and performance in a cooperative activity are boosted by perception of a partner's effort. *Scientific Reports*.

[5] A. Clodic, H. Cao, S. Alili, V. Montreuil, Alami R., and Chatila R. 2009. SHARY: A supervision system adapted to human-robot interaction. *Experimental Robotics. Springer Berlin Heidelberg* (2009), 229–38.

[6] A. D. Dragan, K. C. T. Lee, and S. S. Srinivasa. 2013. Legibility and predictability of robot motion. *8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*.

[7] E. C. Grigore, K. Eder, A. G. Pipe, C. Melhuish, and U. Leonards. 2013. Joint action understanding improves robot-to-human object handover. *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2013), 4622–4629.

[8] J. J. Lee, B. Knox, and C. Breazeal. 2013. Modeling the dynamics of nonverbal behavior on interpersonal trust for human-robot interactions. *Trust and Autonomous Systems: Papers from the 2013 AAAI Spring Symposium.*

[9] C. Lenz, S. Nair, M. Rickert, A. Knoll, W. Rosel, and J. Gast. 2008. Joint-action for humans and industrial robots for assembly tasks. *IEEE* (2008), 130–5.

[10] G. Metta, P. Fitzpatrick, and L. Natale. 2006. YARP: Yet another robot platform. *International Journal of Advanced Robotic Systems* 3, 1 (2006), 43–48.

[11] G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, and L. Montesano. 2010. The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks* 23 (2010), 1125–1134.

[12] J. Michael and H. Powell. 2019. Feeling committed to a robot: Why, what, when, and how? *Philosophical Transactions of the Royal Society B: Biological Sciences* (2019), 374 (1771).

[13] J. Michael and A. Salice. 2017. The sense of commitment in human-robot interaction. *International Journal of Social Robotics* 9, 5 (2017), 755–763.

[14] J. Michael, N. Sebanz, and G. Knoblich. 2016. Observing joint action: Coordination creates commitment. *Cognition* (2016), 106–113.

[15] J. Michael, N. Sebanz, and G. Knoblich. 2016. The sense of commitment: A minimal approach. *Frontiers in Psychology* 6, 1968 (2016).

[16] C. Mueller, J. Venicx, and B. Hayes. 2018. Robust robot learning from demonstration and skill repair using conceptual constraints. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 6029–6036.

[17] Y. Nagai and K. J. Rohlfing. 2009. Computational analysis of Motionese toward scaffolding robot action learning. *IEEE Transactions on Autonomous Mental Development* 1, 1 (2009), 44–54.

[18] S. Nikolaidis and J. Shah. 2013. Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 33–40.

[19] E. Oztop, D. W. Franklin, T. Chaminade, and G. Cheng. 2005. Human–humanoid interaction: Is a humanoid robot perceived as a human? *International Journal of Humanoid Robotics* 2, 04 (2005), 537–559.

[20] C. Press, G. Bird, R. Flach, and C. Heyes. 2005. Robotic movement elicits automatic imitation. *International Journal of Humanoid Robotics* 25, 3 (2005), 632–640.

[21] G. Metta, G. Sandini, and D. Vernon. 2007. The iCub cognitive humanoid robot: An open-system research platform for enactive cognition. *50 Years of Artificial Intelligence* (2007), 358–369.

[22] A. Sciutti, A. Bisio, F. Nori, G. Metta, L. Fadiga, and T. Pozzo. 2012. Measuring human-robot interaction through motor resonance. *International Journal of Social Robotics* 4, 3 (2012), 223–234.

[23] A. Sciutti, M. Mara, V. Tagliasco, and G. Sandini. 2018. Humanizing human-robot interaction: On the importance of mutual understanding. *IEEE Technology and Society Magazine* 37, 1 (2018), 22–29.

[24] Aaron Steinfeld, Odest Chadwicke Jenkins, and Brian Scassellati. 2009. The oz of wizard: Simulating the human for interaction research. In *4th ACM/IEEE International Conference on Human Robot Interaction (HRI'09)* (La Jolla, California). ACM, 101–108. https://doi.org/10.1145/1514095.1514115

[25] A. Stenzel, E. Chinellato, M. A. T. Bou, A. P. del Pobil, M. Lappe, and R. Liepelt. 2012. When humanoid robots become human-like interaction partners: Corepresentation of robotic actions. *Journal of Experimental Psychology Human Perception and Performance* 38, 5 (2012), 1073–1077.

[26] M. Székely and J. Michael. 2018. Investing in commitment: Persistence in a joint action is enhanced by the perception of a partner's effort. *Cognition* 174 (2018), 37–42.

[27] M. Székely, H. Powell, F. Vannucci, F. Rea, A. Sciutti, and J. Michael. 2019. The perception of a robot partner's effort elicits a sense of commitment to human-robot interaction. *Interaction Studies* 20, 2, 234–255.

[28] R. Van den Brule, R. Dotsch, G. Bijlstra, D. H. Wigboldus, and P. Haselager. 2014. Do robot performance and behavioral style affect human trust? A multi-method approach. *International Journal of Social Robotics* 6, 4 (2014), 519–531.

[29] F. Vannucci, G. Di Cesare, F. Rea, G. Sandini, and A. Sciutti. 2018. A robot with style: Can robotic attitudes influence human actions? In *IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)* (Beijing, China).

[30] F. Vannucci, A. Sciutti, H. Lehman, G. Sandini, Y. Nagai, and F. Rea. 2019. Cultural differences in speed adaptation in human-robot interaction tasks. *Paladyn, Journal of Behavioral Robotics* 10, 1 (2019), 256–266.

[31] A. Vasalya, G. Ganesh, and A. Kheddar. 2018. More than just co-workers: Presence of humanoid robot co-worker influences human performance. *PloS one* 13, 11 (2018).

[32] A. Vignolo, H. Powell, L. McEllin, F. Rea, A. Sciutti, and J. Michael. 2019. An adaptive robot teacher boosts a human partner's learning performance in joint action. In *28th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN 2019)* (New Delhi, India).

[33] A. Vignolo, A. Sciutti, F. Rea, and J. Michael. 2019. Spatiotemporal coordination supports a sense of commitment in human-robot interaction. In *11th International Conference on Social Robotics (ICSR'19)* (Madrid, Spain).

[34] A. L. Vollmer, K. S. Lohan, K. Fischer, Y. Nagai, K. Pitsch, J. Fritsch, K. J. Rohlfing, and B. Wrede. 2009. People modify their tutoring behavior in robot-directed interaction for action learning. In *IEEE 8th International Conference on Development and Learning* (2009), 1–6.