Helping robots imitate

Acknowledgment of, and adaptation to, the robot's feedback to a human task demonstration

Aris Alissandrakis,* Dag S. Syrdal† and Yoshihiro Miyake*
*Dept. of Computational Intelligence & Systems Science Interdisciplinary Graduate School of Science & Engineering Tokyo Institute of Technology, Japan and †Adaptive Systems Research Group School of Computer Science University of Hertfordshire, UK

Towards informing the design of more 'natural' (non-expert) human–robot interaction (HRI) systems, this chapter presents two studies, focusing on a scenario where a human has to teach a robot how to perform an everyday household task.

The first study considered the humans' acknowledgment of feedback from a robot, to which they had to demonstrate how to do a task. The research questions were if (and how) the participants would modify their teaching instructions when the robot stated a misunderstanding, and whether they would remain consistent for the remainder of the task.

In the second study, the robot feedback increased in complexity, in order to elaborate on what the human had demonstrated. The research question was if the participants, based on the robot feedback, would be able to resolve any misunderstandings by modifying their teaching instructions.

Besides illustrating that people express a wide variety of teaching behaviors when faced with the task of demonstrating to a robot, the results from the second study support our design suggestion that robots should actively and accurately advertise their cognitive and manipulative capabilities via form (appearance) and function (behavior) to improve the dialogic interaction.

1. Introduction

As robots integrate into human society, expanding from their traditional use in the industry into assisting and collaborating with humans in public and domestic spaces, new paradigms need to be developed towards more 'natural' forms of human–robot interaction, trying to avoid the need for technical or scientic expertise from the part of the human.

Social learning (in the sense of taking advantage of the presence of other agents), compared to individual learning (like trial-and-error), allows for task learning and skill transfer in more ecient and adaptive ways. Imitation (a form of social learning) has been studied by psychologists and ethologists for over a century; it is defined and interpreted in many ways [Zentall and Galef, Jr., 1988, Zentall, 1996, Zentall, 2001, Galef and Heyes, 1996]. Mitchell considers it a conceptually coherent phenomenon with various manifestations, resulting from processes at different but hierarchically related levels; a program producing imitation at any level has control over programs in preceding levels. As evidence of imitation, the following requirements are to be satised: something C (copy) is produced by an organism or machine, where C is similar to something else M (model), registration of M is necessary for the production of C, and C is designed to be similar to M [Mitchell, 1987].

Here we consider imitation in terms of the *agent-based perspective* (cf. [Dautenhahn and Nehaniv, 2002]), and within this focus we address mainly the question of *what* [information is transmitted by teaching], and to a lesser degree the *how* [to reproduce the task]. The answers to the remaining questions of *who* [to imitate], *when* [to imitate], and *how to evaluate the reproduction* are here assumed as given, or not directly addressed.

Our broader research goal is investigating social learning in the context of human–robot interaction (HRI) and as part of this, how people explain routine home tasks to robots, and how these human–robot interactions unfold by using speech and gestures. Using that as a starting point, we want to examine the possible role of feedback by a robot acknowledging the instructions/demonstrations of a human teacher. Are people willing to accept the robot's requests regarding the way a task can be performed? What form should these requests take if the human has to adapt their teaching in order to help the robot understand better? This role would need to be systematically investigated, and examine how it affects the cycle of human–robot interactions.

We would like to also highlight the dialogic nature of the problem and try to consider two perspectives (illustrated in Figure 1). By examining the human perspective, we can better inform the system perspective; this line of research will enable researchers to design robots that can pick 'close enough initial metrics of similarity' depending on context, kick-starting the robot's understanding of taught instructions and process of common ground negotiation with the human. For a more detailed discussion of three lines of research (observed human behavior, formal metrics and implemented technical solutions) related to the issue of helping robots imitate people, please see [Alissandrakis et al., 2010].

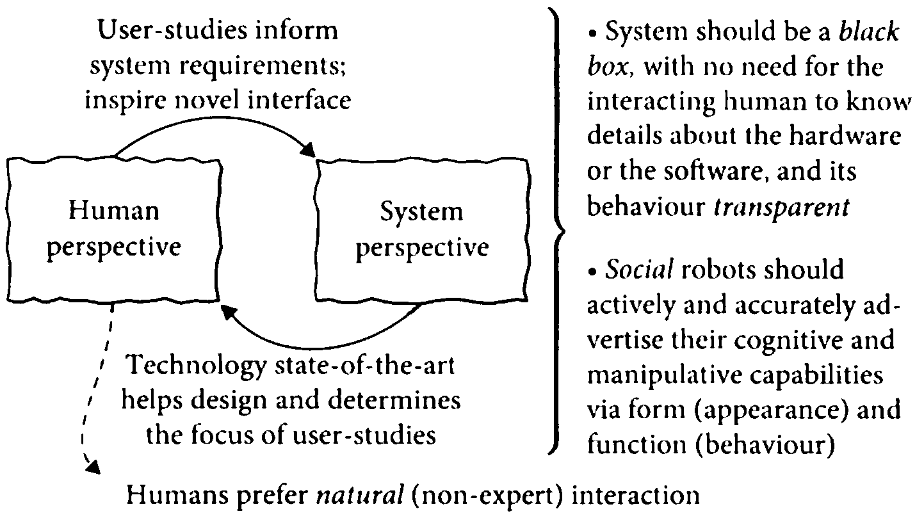

Figure 1. The human and system perspectives design cycle in human–robot interaction
a. a *human perspective* – how people do teaching (more-or-less) naturally, and
b. a *system perspective* – how can hardware and sensor requirements be met and extensive (or even impossible) assumptions and pre-knowledge about the world be addressed

The rest of the chapter is structured as follows: Section 2 gives an overview of two studies, one on the "acknowledgment of" (Section 3) and one on the "adaptation to" (Section 4) the robot's feedback by the human participant. Section 5 offers discussion and Section 6 some possible future work directions.

2. Studies overview

We have been conducting a series of user studies in order to understand how people naturally demonstrate some 'basic' home tasks to a robot. The sessions for both studies presented here were conducted at the Miyake Lab, Tokyo Institute of Technology, in Japan (see Figure 2). Previous studies were conducted in the UK and were presented in [Otero et al., 2008b, Otero et al., 2008c, Otero et al., 2008a]. Up to the present, our overall strategy in addressing this problem has taken a step-by-step approach: we chose to introduce small changes (reaction/actions of the robot towards the human demonstration) from one study to the next whilst maintaining the overall experimental instruction requests to the human participants. It is worth emphasizing that we are not actually evaluating or developing a specific robot system here; this line of studies is intended to be informative towards the iterative design of more natural HRI interfaces. Therefore, the results and discussion are presented in a qualitative and mostly descriptive manner.

Figure 2. Participants interacting with the Wakamaru robot. Photos from the "acknowledgment of feedback" (left) and the "adaptation to feedback" (right) studies

More specically, in the two user studies described in the next sections, we highlight the nature of human demonstrations of a specific routine home task: how to lay a table. This particular task was chosen because, while relatively simple and familiar[1] to most people, the number of tableware involved, their order and their possible congurations on the table can vary, depending on personal preferences.

Note that due to technological issues (current state-of-the-art robots are not yet able to detect and understand *unconstrained* behavior by humans), plus the fact that the robot did not actually need to detect or respond to the actual participant's behavior (as responses were predetermined to a great extent), in both studies presented here the robot was not autonomous, but controlled using a Wizard-of-Oz (WOZ) methodology (cf. [Dahlbäck et al., 1998]).

3. "Acknowledgment of feedback" study

The main motivation of this study was to conrm whether the initial observations of a previous study (presented in [Otero et al., 2008a]) on the participant's acknowledgment of the robot's feedback, were valid in a Japanese cultural setting, this time using a humanoid[2] robot, and to inform follow-up studies like the one in Section 4.

3.1 Research questions

Four research questions were considered in this study:

Q1: Do people change their demonstration when a robot declares its inability to understand the participant's gestures and/or speech?

Q2: What is the nature of the change of the demonstration (if any)?

Q3: To what extent are people consistent (maintain any changes for the remaining of the teaching task)?

Q4: Does the familiarity with the task to be taught (compared to one performed only occasionally) have any influence on the previous questions?

3.2 Design and methodology

In each session, a participant, after completing a background questionnaire, was asked to demonstrate to the humanoid communication robot Wakamaru (Mitsubishi Heavy Industries) the task of laying a table, twice, using a dierent set of objects each time – Japanese, and 'non-Japanese' eating utensils (the set order was reversed for each successive participant).[3] The utensils were placed on a side-table, from which they should be picked up by the participants and placed on another table, in front of the robot. A video camera was used to capture the participant's demonstrations and interaction with the robot.

An example setup can be seen in Figure 3. As Wakamaru is only 1m tall, a base was used to allow the robot to reach the table. The side table, where the utensils were initially placed, is off-camera (in front of the robot). The 'wizard' (human operator of the robot) was situated out-of-sight, behind the partition, monitoring the session via the camera (seen mounted on the top of the partition).

The sample consisted of 11 native Japanese participants (6 female and 5 male), all of them university students in their early twenties. None of them had any computer science or robotics background, or (more importantly) had any previous experience of interacting with a robot, and could therefore be considered *naive* in respect to their expectations of a robot behaving in a social manner.

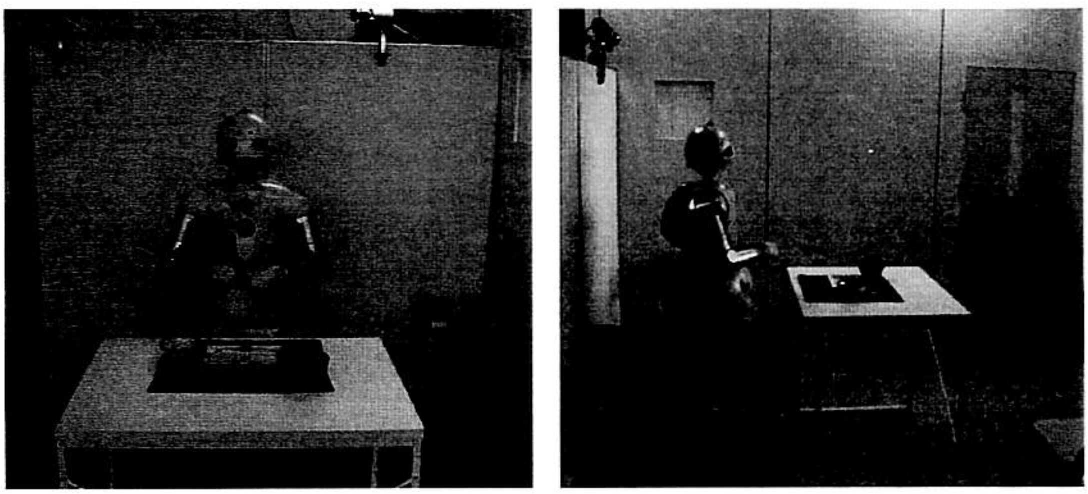

Figure 3. Study setup (front, side)

In order to capture the most natural responses and behavior by the participants, no specific instructions were given to the participants as to how to interact with the robot, besides that they could use gestures and/or speech, and the single restriction that they must demonstrate only one object at a time (but in any order), and hold it using one hand;[4] as there were four utensils, each task was decomposed to four sub-tasks.

The participants were told that their own performance was not to be evaluated, and that there was no right or wrong way to complete each task, but instead the purpose of the study was the evaluation of the robot's learning capabilities. If they wanted, they could stop the session at any time and still receive full compensation.

In every session, the robot would always provide 'positive' feedback to the participant's instructions after the first, third and fourth sub-task demonstration, by nodding its head forwards and saying "Hai, wakarimashita! ⟨Yes, I understand!⟩". Similarly, irrespectively to what the second sub-task demonstration was, the robot would tilt its head sideways and say "Suimasen, wakarimasen... ⟨Sorry, I don't understand... ⟩", thus giving 'negative' feedback. If the participant chose to repeat or modify their second demonstration, the feedback would then revert to positive for that sub-task (see illustration in Figure 4).

Finally, each participant was asked in a post-task semi-structured interview to subjectively recall the events as they occurred, and comment on their impression of their interaction with the robot.

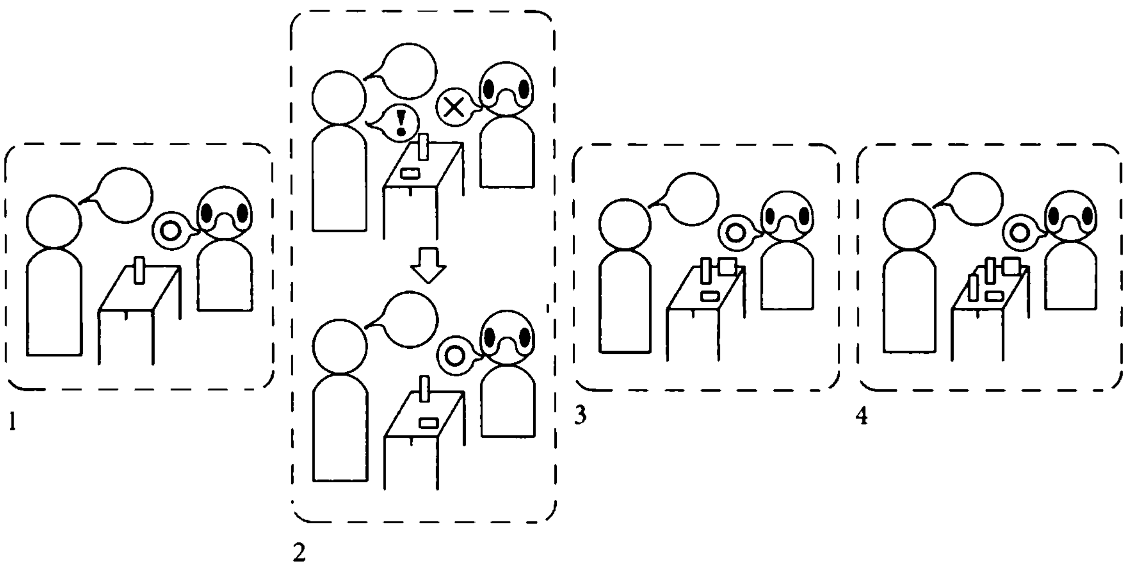

Figure 4. Robot's feedback in the "acknowledgment of feedback" study. The robot would state that it understands the first, third and fourth sub-task demonstration, but would initially state to not understand the second (reverting back to understanding, if the participant clarified or repeated their demonstration)

3.3 Summary of results

Using the study design described in 3.2, the change from positive to negative feedback between the first and second demonstrations allows to examine Q1 (and, in the cases that Q1 happens to be true, Q2); the change back to positive feedback for the third and fourth demonstrations allows to examine Q3, assuming Q1 occurred. Q4 is examined by comparing the interaction during the Japanese and non-Japanese tasks.

The conceptual framework presented in [Nehaniv et al., 2005] can be used to capture requirements for contextual interpretation of body postures and human activities for the purposes of HRI; it defines five functional classes of gestures, and a classication scheme based on it was used both in this study and [Otero et al., 2008a]. Other approaches to characterize the interaction, besides the classification of gestures, include the frequency of gestures and/or utterances, along with the classification of object localization (the latter being related, from the system perspective, to effect metrics for robot imitation, cf. [Alissandrakis et al., 2006]).

All[5] of the participants did acknowledge the robot's feedback on their demonstration of teaching the task of laying a table (Q1); they appreciated receiving positive feedback (indicating that the robot was paying attention to them), but most importantly they responded to the negative feedback, by attempting to clarify their demonstration.

Unfortunately, the participants used very few distinct deictic gestures,[6] and, as the negative feedback was not very informative about the cause of the misunderstanding, the majority assumed that e.g. they simply had to speak louder, so they did not greatly modify their demonstrations (Q2).

However, in the few cases they did [in particular, change the object localization], they remained consistent for the rest of the task (Q3) by keeping the same style of object localization ('absolute' or 'relative', see Section 4.2.3) – this observation was more promising than the findings of the [Otero et al., 2008a] study.

Concerning the influence of task familiarity on the current task teaching (Q4), there were no clear results from this study, however it can be noted that the participants added details about the object's intended use and function in their verbal instructions for the Japanese ("more familiar") utensils only.

For more details, please see [Alissandrakis and Miyake, 2009].

4. "Adaptation to feedback" study

The participants of the "acknowledgment of feedback" study (see Section 3) had a positive impression of the robot and the teaching interaction; however, they (quite understandably) noted that the feedback was too vague to explain what the misunderstanding was, in order for them to correctly modify their demonstration. In addition, they noted that (except when it was providing feedback), the robot was too 'passive'.

Therefore, informed by the results and the post-session participant interviews, the robot's feedback increased in complexity. The robot would also actively regulate the interaction with a variety of interactional[7] gestures, in addition to welcoming the participants at the beginning, and thanking them at the end of each task.

4.1 Research question

Extending the theme of the research questions from Section 3.1, this study looks at the hypothetical scenario of interacting with a robot that is able to understand only a specific way of being taught. Could that robot successfully advertise its cognitive capabilities to its human teacher by means of its feedback, so that she modifies any subsequent demonstrations accordingly?

In our approach (as outlined in Section 1) towards developing more natural HRI interfaces, we believe that a robot should be like a 'black box', with no need for the interacting human to know details about the underlining hardware or the software, and its behavior should be 'transparent'. Therefore, the main research question here was whether the participants could adapt to the robot's object localization, one that is implicitly expressed by its gestures and speech during the confirmation feedback of the just observed demonstration, rather than explicitly by identifying said (e.g. absolute or relative) metric in technical terms.

4.2 Design and methodology

The sample consisted of 18 native Japanese participants (13 female and 5 male); the majority were parents (early-to-late twenties) and some worked as kinder-garden teachers. None of them had any computer science or robotics background; a few of them had interacted previously with the same robot, but in an unrelated study that did not include teaching.

In each session, a participant, after completing a background questionnaire, was asked to demonstrate to the robot how to lay the table three times; however, in this study, we did not include a 'non-Japanese' task variation, therefore the utensils in all cases were a plate, a bowl, a cup and chopsticks. As reason for repeating the same task three times, the participants were told that we wanted to try dierent learning algorithms with the robot, and that they could change their own demonstration order and layout, if they wanted, each time. Like previously, we emphasized that there was no right or wrong way to complete each task, but instead we only wanted to evaluate the robot's learning capabilities. They were free to use any combination of speech and/or gestures they wanted.[8] After the completion of each task, they were asked to evaluate their interaction experience with the robot, and recall what happened in their own words, in a brief semi-structured interview.

The participants had to pay attention and decide if the robot's feedback accurately corresponded what they just demonstrated; the robot would not explicitly state to 'not understand', but instead ask the participant to conrm that its understanding of the demonstration was correct. Each sub-task was effectively composed of two stages: the participant's demonstration and the robot's feedback.

4.2.1 *Participant's demonstration*

The robot would initiate every demonstration stage: "[Saisho no/tsugi no] shokki wo dozo ⟨The [first/next] dish, please⟩". During the demonstration, every time there was a pause in the participant's instructions, the robot would nod and say "Hai ⟨Yes⟩!", indicating paying attention. If there was a long pause after the last utterance/gesture, the robot would tilt its head sideways and ask "Sore deii desu ka ⟨Are you finished⟩?". Unless the participant replied that they had more instructions, the demonstration stage would end, with the robot tilting its head towards the table and saying "Kangaeteimasu ⟨I am thinking about it⟩". This action on the part of the robot would serve a dual purpose: give to the participant the impression that the robot had to 'think' and 'understand' what was just demonstrated, but also most importantly, give to the wizard enough time to construct the appropriate response for the feedback.

4.2.2 *Robot's feedback*

The robot would point to a table location (see Figure 5), verbally identify the current object and suggest an object localization (see Figure 6), followed by a request for confirmation "Koreha tadashii desu ka ⟨Is this correct⟩?".

Compared to the simplistic statement of understanding (or not) of the demonstration given as feedback in the previous study, the robot now provided correct or incorrect feedback in terms of the object localization; object itself would always be correctly identified. The object localization would be either factually[9] *correct*, with the robot pointing at the appropriate part of the table where the object was placed (and in addition, in the case of relative reference, providing a valid reference to an existing referencing object – or to the robot itself, if there were no objects yet on the table), or *incorrect*, with the robot pointing at a wrong location on the table and/or giving a non-valid reference.

If the participant accepted and confirmed the feedback as correct, the robot would then ask her to proceed to the next object. There was also an option for the

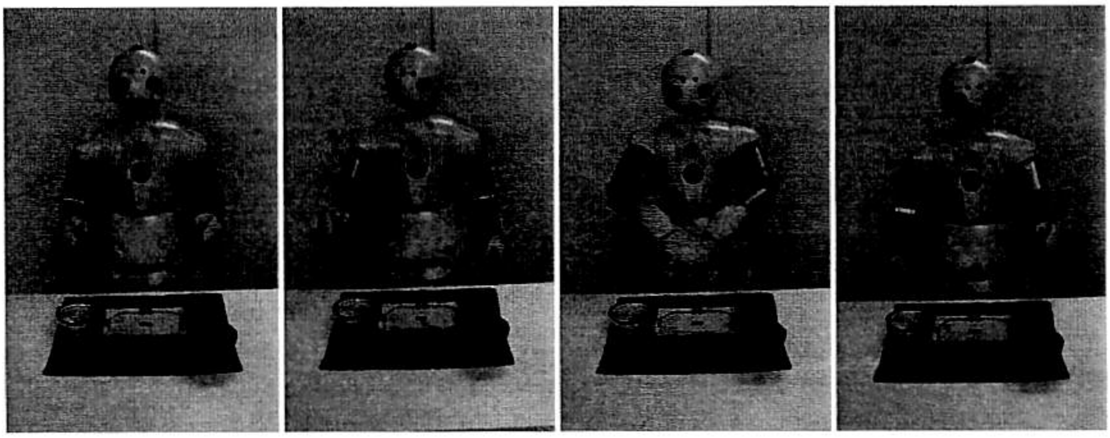

(a) Neutral (b) Right (c) Middle (d) Left

Figure 5. Robot feedback gestures. (a) neutral posture and gazing at the participant during the demonstration, (b) pointing and gazing to the right, (c) pointing to the middle and gazing downwards, (d) pointing and gazing to the left

(**Object** / *o sara* ⟨plate⟩ / *o wan* ⟨bowl⟩ / *o chawan* ⟨cup⟩ / *o hashi* ⟨chopsticks⟩) *wo* ⟨the⟩ *kokoni* ⟨here⟩ *okitaii* ⟨place⟩.

Abs[olute] object localization: "The (object) should be placed here."

(Object / *o sara* ⟨plate⟩ / *o wan* ⟨bowl⟩ / *o chawan* ⟨cup⟩ / *o hashi* ⟨chopsticks⟩) *wo* ⟨the⟩ (**relative object** / *o sara* ⟨plate⟩ / *o wan* ⟨bowl⟩ / *o chawan* ⟨cup⟩ / *o hashi* ⟨chopsticks⟩ / *Wakamaru*) (**reference** / *no migini* ⟨right of⟩ / *no hidarini* ⟨left of⟩ / *no maeni* ⟨before⟩ / *no okuni* ⟨behind⟩) *okitaii* ⟨place⟩.

Rel[ative] object localization: "The (object) should be placed (reference) the (relative object)."

Figure 6. Robot feedback phrases. Depending on the object localization type (top: Abs[olute], bottom: Rel[ative]), two types of phrases were used by the robot. In combination with the gestures shown in Figure 5, this allowed the robot to communicate what it 'understood' from the demonstration, advertising its cognitive capabilities

robot to repeat its feedback, in case the participant wished it to do so. If the participant responded that the feedback was not correct, the robot would instead ask her "Mou ichido shokki wo misete kudasai ⟨Please do that again⟩"; if this was repeated more than five times, the robot would instead say "Wakarimasen deshita ⟨Sorry, I could not understand⟩" and ask the participant to proceed to the next object.

4.2.3 Object localization classication

The object localization, both for the participant and the robot, could be classified as either[10] one of

a. Absolute (Abs) – no relation to reference points (other task objects, environmental landmarks) is given or implied e.g. by pointing gestures. A typical example would be "[This object] should be placed here" (which can be accompanied by a deictic gesture to the final destination). Alternatively, no verbal instructions at all are given, and the object is simply transported and placed at its intended location.
b. Relative (Rel) – the relation to other task objects (or environmental landmarks) is given, possibly combined with pointing gestures.

During the sessions, the participants sometimes used variations of these object localizations. For example, they would use only gestures to indicate where an object should be placed, or would not say (or point) which other object (or landmark) it should be "to the right of"; these examples could be considered as more implicit or 'incomplete' versions of Abs and Rel, respectively. As the robot could not modify its feedback beyond the gestures and phrases shown in Figures 5 and 6 respectively, in order to determine the appropriate response to the demonstration during the sessions, and for the discussion of the results here, any such variations were grouped together as either Abs or Rel.

4.3 Choice of object localization

Depending on the choice of localization on the part of the human (demonstration), the choice of localization on the part of the robot (feedback), and also whether this feedback would be correct (or not) was determined for each of the three tasks as illustrated in Figure 7.

During the first task, the robot would always give correct feedback, using the same object localization as the one used by the participant (for each demonstration), either absolute or relative. In this task, the participants were meant to familiarize themselves with the teaching interaction, and the robot's capabilities.

During the second task, the robot would use only the object localization that was used more frequently by the participant during the first task, and the feedback would be correct only if the participant also used it for the current demonstration. Assuming some consistency on the part of the participants, for the majority of cases, no misunderstandings should occur.

During the third task, this preference would be inversed, with the robot only using the object localization that was less (if at all) used by the the participant during the first task. In this case, it is almost certain that misunderstandings will occur, and the participants will have to modify their instructions to reach a common

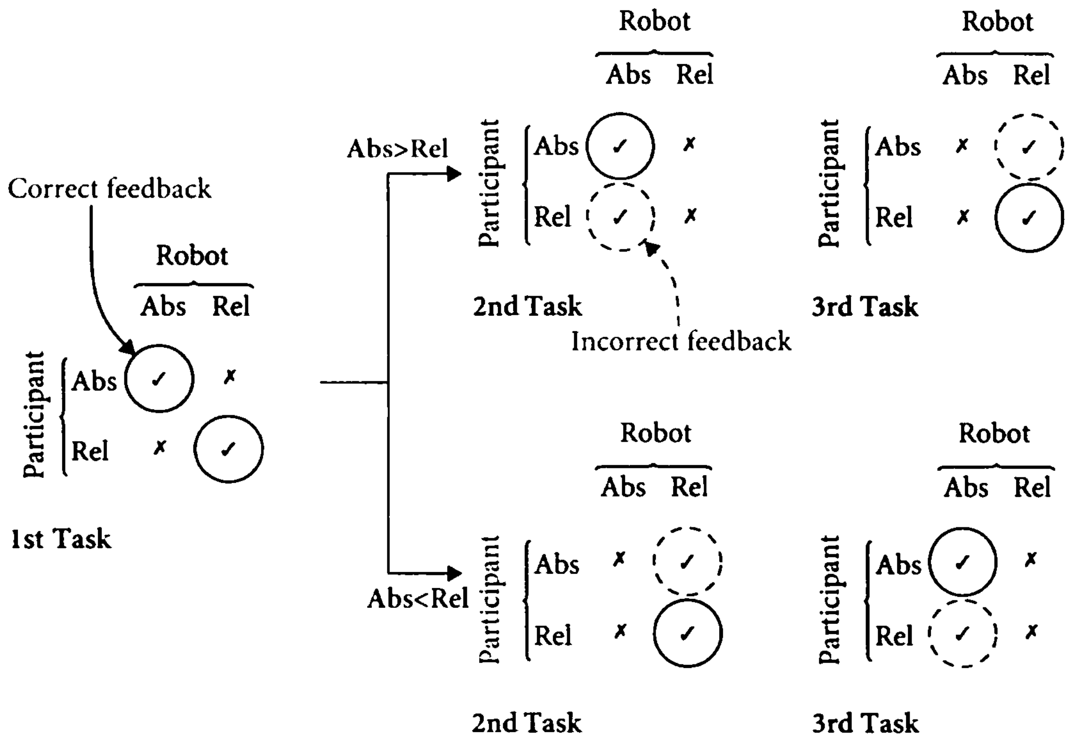

Figure 7. Choice of object localization

ground with the robot (i.e. adapt and use another object localization themselves) in order to successfully transmit the task knowledge.

A short video highlighting different interactions from this study can be found on the first author's homepage[11].

4.4 Results

4.4.1 *Overview*

As mentioned above, there was no right or wrong way to do the task and therefore evaluate performance in such a manner; each participant's teaching behaviour also varied significantly compared to the rest, as expected. Therefore, in the qualitative discussion of the results here (see the overview in Table 1), we classify the participants as those who adapted, and those that did not. Adapted were those participants that, in the third (or in one case, second) task, changed their object localization (in a repeated demonstration) in response to incorrect feedback given by the robot. We use this definition here loosely; the change of object localization (to the one that the robot used for its feedback) needs only to have occurred once. For the majority of participants, a number of attempts (if not a number of entirely unsuccessful sub-tasks) was required until they established that their choice of object localization was the cause of the misunderstanding. In addition, in subsequent

Table 1. Overview of "adaptation to feedback" sessions. For every participant (identied as p_x), the number of Abs and Rel object localizations used for the demonstrations during the 1st task, and (depending on which was used more and which less) the object localizations used exclusively for the robot feedback, for the 2nd and 3rd tasks respectively, plus notes

	1st Task		2nd Task		3rd Task	
	Abs	Rel	feedback	notes	feedback	notes
p_1	0	4	Rel	-	Abs	na, fn
p_2	0	4	Rel	-	Abs	na
p_3	0	4	Rel	-	Abs	-
p_4	1	3	Rel	-	Abs	a, fn
p_5	4	0	Abs	-	Rel	a, c
p_6	2	2	*Abs*	a, c	*Rel*	-
p_7	4	0	Abs	-	Rel	a, c
p_8	0	4	Rel	-	Abs	na, fn
p_9	4	0	Abs	-	Rel	na, fn
p_{10}	0	4	Rel	-	Abs	na, fn
p_{11}	4	0	Abs	-	Rel	a, c
p_{12}	1	3	Rel	-	Abs	a
p_{13}	4	0	Abs	-	Rel	a
p_{14}	4	0	Abs	-	Rel	na, fn
p_{15}	0	4	Rel	-	Abs	a, c, fn
p_{16}	0	4	Rel	fp	Abs	na
p_{17}	1	3	Rel	-	Abs	na, fn
p_{18}	2	2	*Abs*	fn	*Rel*	a, c

Notes legend: - = both participant and robot used the same object localization throughout the task; a = adapted; c = consistent; na = not adapted; fn = false negative (accepted incorrect feedback); fp = false positive (not accepted correct feedback)

sub-tasks, some (adapted) participants remained *consistent*, using the same object localization for at least one subsequent demonstration.

During the study sessions, there were both cases of *false negatives*, where participants accepted incorrect feedback by the robot as correct, but also one case of *false positive*, where a participant did not (initially) accept the correct feedback; however, after repeating the demonstration and receiving the same (correct) feedback, he realized that the robot was using its own perspective and not his.

In Figure 8 (top right) the participants are shown grouped according to the object localization they used for the first task demonstrations. One third of the participants were absolute for every demonstration, roughly another third were relative, while the remaining used a mixture (once absolute and the rest relative, or twice absolute and twice relative).

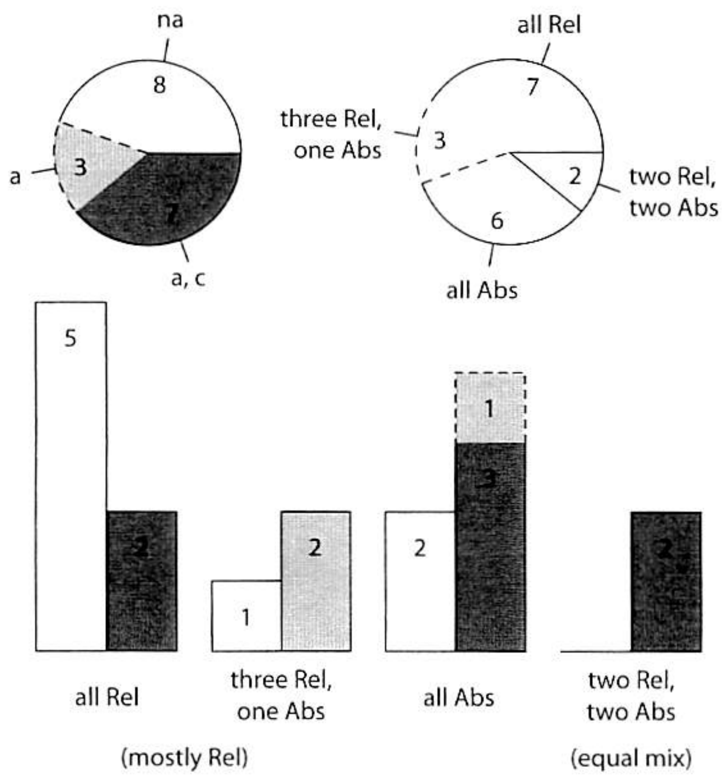

Figure 8. Overview of "adaptation to feedback" sessions. Top right: participants according to the object localization they used during the first task. Top left: participants that did not adapt (*na*), those who did (*a*), and those who adapted and were also consistent (*a,c*). Bottom: adaptation and consistency (colors as top left), grouped by object localization preference during first task. See also Table 1

Figure 8 (top left) shows how many participants adapted to the robot's feedback (during the third and/or second task, depending on the context) or not, and if they did, whether they remained consistent or not. More than half of the participants did adapt, and the majority of these were consistent.

Figure 8 (bottom) shows the adaptation of the participants, grouped by first task object localization. For the discussion the next subsections, 'all relative' and 'three relative, one absolute' are grouped as 'mostly relative', and 'two relative and two absolute' are referred to as 'equal mix'.

4.4.2 *Mostly relative*

Seven participants were relative for every demonstration, and three were absolute once (and in the other demonstrations relative).

Among the adapted, we include the special case of participant p_3. Although she was always relative throughout the first and second tasks, during the third task she spontaneously switched to absolute, from the first demonstration. In her post-task interview she mentioned that (presumably after growing condent with the interaction in the previous two tasks) she wanted to try whether the robot would still be able to understand her if, instead of calling each utensil by its name she simply said "this", and instead of giving a reference she said "here". She assumed that this would be more difficult for the robot, and was amazed that the robot could (by accident of the experimental procedure) understand her.

Participant p_1, according to her post-task interview, assumed that the robot was capable of responding only to a single phrase and tried to (unsuccessfully) find which one was suitable.

Participant p_2 noted that the robot was using "here" for the absolute feedback, but assumed that the robot was not using a reference point, like before, because her own demonstration was somehow wrong.

Participant p_{10} during the second task tried to place the chopsticks separately (ignoring the study instructions), on either side of the plate, but after a few attempts she realized that the robot can only give feedback for one location. During the third task, she accepted reluctantly every incorrect feedback after unsuccessful attempts to correct the first demonstration; she later noted that she "just could not make the instruction simpler" for the robot to understand.

4.4.3 *All absolute*

From the six participants that were absolute for every demonstration, only two did not adapt.

Participant p_9, although he did acknowledge and responded by "yes" and "no" to the robot's questions, did not provide any spoken instructions throughout, but simply placed the utensils in front of the robot, (with slower motions when asked to demonstrate again), and using no additional pointing gestures.

The remaining four participants all noted in their post-task interviews that they realized that the conrming feedback in the third task now included relative references, and therefore they changed their demonstrations accordingly.

4.4.4 *Equal mix*

Two participants were equally absolute and relative for their first task demonstrations. For both participants the choice of object localization for the robot's feedback in the second and third task was absolute and relative, respectively.

Participant p_6 adapted to absolute in the second task, and was relative, from the start, during the third task. Participant p_{18} adapted to relative in the third task, but for every sub-task during the second task, she accepted the incorrect (absolute) feedback to her (relative) demonstrations.[12]

4.5 Other observations

Overall, the teaching interaction was well perceived. Even those participants who were uncertain on the demonstration-feedback turn-taking, quickly got used to the regulation of the interaction by the robot. Five participants (p_1, p_6, p_7, p_{15}, p_{18}) stated (in various degrees) that they found the freeform nature of the interaction with the robot somehow dicult and would prefer a more structured format, possibly having to use specific, predefined, sentences; however, three of them were still able to adapt.

Compared to what was observed in our previous studies, the current participant demonstrations were, in general, much more salient in respect to the sub-task segmentation. Usually, every object was initially identified (either verbally or by a presentation-type gesture), and only after receiving acknowledgment from the robot, the participant proceeded to demonstrate its placement on the table. When designing the verbal component of the robot's feedback for this study, it was important that there would be no ambiguity about the cause of the misunderstanding; therefore, the robot would always correctly identify the utensil, even when the localization was incorrect. However, the participants, upon having to demonstrate a subtask again, would almost always repeat the object identication in their instructions, instead of correcting the object localization alone. This was not explored in the post-session interview, but a possible reason for repeating both identification and localization could be an assumption on their part regarding the robot's capabilities, based on the fact that the feedback was not segmented (single utterance and gesture, combining identification and localization).

In their interviews, the participants made a number of assumptions regarding the robot's capabilities. Participant p_7 initially expected the robot to move the utensils itself, according to her demonstration. Participants p_6 and p_{15} noted that they were quite surprised that a robot could actually make a mistake. On the other hand, participant p_{11} noted that she enjoyed when the robot was making mistakes, cause then she then had the opportunity to teach it how to do the task the right way. Participant p_{10} noted that "[...] when I teach Wakamaru something he doesn't know, Wakamaru just takes what I said instead of what I intended, I guess that is the best it can do". Participant p_2 made assumptions about the robot's visual capabilities, based on the fact that the color or shape of the objects was not mentioned in the feedback, but just the utensil name. Participant p_{15} made the assumption that the robot uses vision and not audio, since it sometimes gave different responses to what she said; therefore, during the third task she tried to do the demonstration gestures more slowly. Similarly, participants p_6 and p_{18} noted that they were uncertain if the robot better understands audio or visual cues; p_{18} noted that he made the assumption that the robot must rely on vision, after the robot gave relative feedback from its perspective, while he had given the reference for the (also relative) demonstration from his own perspective.

Comparing the adaptation ratios of the relative and absolute groups (in Figure 8, bottom), we see that the proportion of participants that were absolute (in their demonstrations during the first task) and did change their object localization to relative (in response to the robot's incorrect feedback during the third task), is greater than the proportion of participants that were relative and changed to absolute. This result is not entirely unexpected, as the two types of object localization (as classied here) are not 'interchangeable'; the participants' personal preferences can be considered as an influencing factor, in combination with the human cognitive processes involved. For example, as noted in Section 4.4.2, participant p_{10} stated that she could not think of how to make her (relative) instructions 'simpler'. On the other hand, in order to test the robot's capabilities, participant p_3 (also relative by preference) used "this" and "here" instead of labeling and referencing, as doing so (i.e. effectively changing to absolute) would be more 'complex'. Within the context of this study we also hypothesize that the (for some participants, low) saliency of the robot's feedback gesture might have been an additional factor. The gesture was more important in the cases of absolute feedback, as it disambiguated where "here" was, than in the cases of relative feedback that already included a reference point in the utterance. In order to reduce the complexity of the WOZ system (and compensating for the robot's technical limitations in respect to the movement range of the arms) one of the (only three possible) pointing gestures shown in Figure 5 was used to indicate the object location during each feedback, in an encompassing manner, along with the head gazing direction. However, five participants (p_4, p_6, p_9, p_{10}, p_{13}) mentioned in their post-session interviews that they would like more precision (they also noted that the robot does not have any fingers for pointing). From the interview of participant p_8, it became clear she did not notice that the robot was pointing at all, causing her to question the (as a result, rather vague) feedback. The low saliency of the feedback gestures possibly also contributed to the false negatives of participants p_{14}, p_{17} and p_{18} (as in all these cases, the incorrect feedback was accepted without any correction attempts).

4.6 Questionnaire results

The questionnaire used to asses the participant's subjective impressions of the robot (after completing each task) was based on the one used and evaluated by [Kanda et al., 2001].[13] Table 2 shows the individual items along with the factor names,

Table 2. Questionnaire item and factor names. The naming used is similar to the one used in [Kanda et al., 2001]

Factor	Items
Familiarity	accessible, altruistic, distinct, favorable, frank, kind, pretty, safe
Enjoyment	exciting, friendly, interesting, likable, pleasant, pretty
Activity	active, cheerful, showy
Intelligence	intelligent

Factor	Mean (SE)		
	1*st* Task	2*nd* Task	3*rd* Task
Familiarity	1.97(.01)	1.90(.13)	2.54(.15)
Enjoyment	2.06(.15)	1.89(.15)	2.43(.21)
Activity	2.92(.11)	2.72(.11)	2.72(.14)
Intelligence	2.00(.21)	2.05(.21)	3.15(.22)

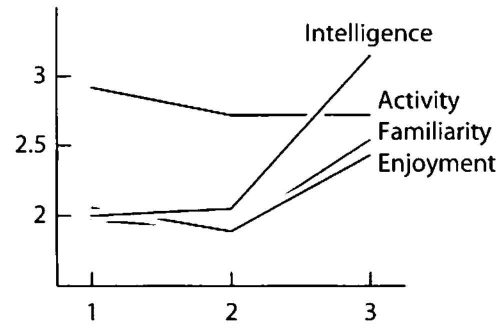

Figure 9. Questionnaire overall results. Mean scores with standard errors for each factor, across the three tasks

and Figure 9 shows the mean scores for each factor, after each task. Note that a high score indicates disagreement with the item/factor.

For Familiarity, there was a signicant main effect for task number (F(2,18) = 6.527, p = .007); post-hoc tests using the Bonferroni correction revealed that the score after the 3rd task was signicantly larger than those after the 1st (Mean Di = .58, SE = .16, p = .006) and 2nd task (Mean Di = .64, SE = .18, p = .005) suggesting that participants overall viewed the robot displaying less familiarity after the 3rd task.

Enjoyment also had a signicant main effect for task number (F(2,18) = 6.59, p = .007); post-hoc test using the Bonferroni correction revealed that there was a signicant dierence between the 2nd and 3rd task (Mean Diff = .54, SE = .17, p = .015), suggesting that participants viewed the interaction with the robot more enjoyable in the 2nd task than in the 3rd.

There was no main effect for task number and Activity (F(2,18) = 2.420, p = .12), suggesting that participant's impression of the robot along this dimension did not change over the trials.

There was a signicant result for Intelligence in terms of task number (F(2,18) = 9.724, p = .001). Participants viewed the robot as more intelligent in the 1st (Mean Di = 1.15, SE = .25, p = .001) and 2nd task (Mean Di = 1.10, SE = .28, p = .003) than in the 3rd task.

4.6.1 *Adaptation and impressions*

Figure 10 shows the mean scores for each factor, after each task, depending on whether the participants adapted or not.

There was no signicant interaction effect for adaptation (whether the participant adapted or not) and task number for Familiarity. There was a main effect for adaptation approaching signicance (F(1,13) = 3.661, p = .078) wherein participants who adapted in the 3rd task viewed the robot as more familiar across all trials. While this trend was present across all three tasks, it was most marked in the 3rd task.

There were no signicant interactions or main effects for adaptation and task number for Enjoyment.

There was no signicant interaction for task number and Activity. There was a main effect approaching signicance for Activity (F(1,13) = 3.408, p = .088), in which participants who adapted to the robot in the 3rd task viewed the robot as more active.

For Intelligence, there was no signicant interaction between task number and adaptation. There was however, a signicant main effect for Intelligence and adaptation (F(1,13) = 6.394, p = .025), where participants who adapted to the robot signicantly rated the robot as more intelligent across all three tasks.

	Mean(SE)					
Factor	1st Task		2nd Task		3rd Task	
	a	na	a	na	a	na
Familiarity	1.85(.14)	2.10(.20)	1.66(.15)	2.08(.21)	2.48(.21)	3.00(.30)
Enjoyment	1.92(.21)	2.10(.30)	1.65(.21)	2.03(.30)	2.15(.28)	3.00(.40)
Activity	2.73(.15)	3.00(.21)	2.46(.16)	3.00(.23)	2.40(.16)	2.80(.22)
Intelligence	1.60(.20)	2.20(.28)	1.70(.23)	2.20(.33)	2.70(.26)	3.80(.37)

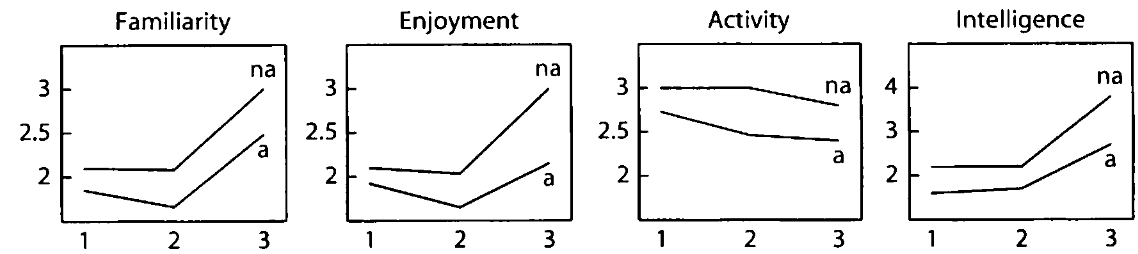

Figure 10. Adaptation and impression results. Mean scores with standard errors for each factor, across the three tasks, for participants that adapted (*a*, including consistent and non-consistent) and those that did not (*na*)

4.6.2 *Object localization preference and impressions*

Figure 11 shows the mean scores for each factor, after each task, depending on the participants preference of object localization during the initial task.

There was no main overall signicant effect for object localization preference, either relative ("mostly relative") or absolute ("all absolute"), on Familiarity. There was however, an interaction effect, for object localization preference and task number of Familiarity ($F(2,10) = 4.054$, $p = .05$). This effect suggested that, while relative participants rated the robot as more familiar during the 1st and 2nd task, the relationship was reversed after the 2nd task, where absolute participants rated the robot as more familiar.

There was no main overall signicance effect for object localization preference on Enjoyment. There was an insignicant trend mirroring that of the interaction effect between object localization preference and task number for Familiarity, but this trend was not signicant for Enjoyment in this sample size ($F(2,10) = 2.34$, $p = .146$, observed power = .37).

There was no main or interaction effect for object localization preference and task number on Activity.

There was a signicant interaction effect for object localization preference and task number on Intelligence ($F(2,10) = .4854$, $p = .03$). This effect was similar as to that described for Familiarity and Enjoyment.

Factor	Mean (SE) 1*st* Task		2*nd* Task		3*rd* Task	
	R	A	R	A	R	A
Familiarity	1.78(.16)	2.10(.20)	1.56(.15)	1.95(.19)	2.91(.24)	2.45(.30)
Enjoyment	1.67(.21)	2.33(.26)	1.58(.24)	1.93(.30)	2.71(.34)	2.27(.43)
Activity	2.79(.17)	3.00(.22)	2.75(.20)	2.67(.26)	2.50(.17)	2.80(.22)
Intelligence	1.63(.23)	1.80(.29)	1.63(.26)	2.00(.33)	3.50(.28)	2.80(.35)

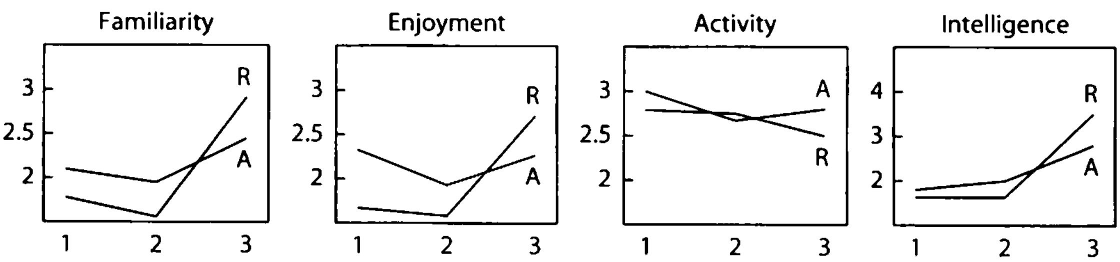

Figure 11. **Object localization preference and impression results.** Mean scores with standard errors for each factor, across the three tasks, for participants that were determined to be relative (*R*, "mostly relative") and those were absolute (*A*, "all absolute") from the initial task. The "equal mix" group was not considered due to the small sample size

4.6.3 *General comments on the participant's impression of the robot*

Overall positive impression (in terms of Familiarity, Activity and Intelligence) is related to willingness, or possibly, capability, to adapt to the robot's object localization preference. Moreover, these positive impressions for the group that adapted cannot (as shown by the lack of signicant interaction effects between adaptation and task number for these impressions) be considered to be the result of the successful adaptation in the third task, but rather as a causal factor that preceded it.

This suggests that adaptation was most likely caused by an existing good impression of the robot, most likely through one of two possible avenues. The first is that the participant's overall good impression of the robot might have led them to invest more into the interaction, and part of this investment was the willingness put in the extra eort to change their own object localization preference.

The second, and more likely, avenue is that the participants' overall higher scores along the Intelligence dimension meant that these participants viewed the robot as more intellectually capable, and as such, adopting new strategies to communicate with the robot would be considered more worthwhile than for those who rated the robot as being less intelligent. The comment by participant p_{10} regarding not being able to "make the instructions simpler", as well as the approach to teaching the robot by participant p_2, suggest that a view of the robot as less sophisticated intellectually did negatively impact the ability to adapt.

Irrespective of which avenue was most instrumental in causing this effect, these results do suggest that in this particular short-term teaching scenario, *a priori* mental models based on initial interactions [Syrdal et al., 2007], may have repercussions for the entirety of the teaching interaction, including its eventual outcome.

The interaction effects between object localization preference and task number for the different impression dimensions suggest that the task-based mental model (i.e. whether relative or absolute is understandable to the robot) impacts the more higher-level mental model of the robot (as measured by the impression questionnaire), and there is an asymmetry between the two object localization preferences in terms of how these two levels of mental models interact, specically that the change from a task-based mental model in which relative is used is more costly in terms of a favorable higher-level mental model than the change from a task-based mental model using absolute.

5. Discussion

We consider the two studies presented here as part of ongoing work towards systematically investigating the role of feedback (by a robot that is shown how to do a task by a human), and how this affects the cycle of human–robot interactions. The two 'thought experiments' behind the studies' design were:

i. If a robot states to not understand the demonstration of a task, would the human teacher acknowledge this and try to do something about it?
ii. Assuming that a robot, due to limitations of hardware and/or software, is only capable of detecting a particular type of object localization, would it be able to make this known to the human in the course of the teaching interaction?

The results from these and previous studies suggest that people show a wide variety of behaviors, actions and gestures when faced with the task of demonstrating to a robot. The dialogical perspective that we are advocating seems to be a way forward but the details of the whole process are still very open. Nevertheless, it appears that the initial publicizing by the robot of its basic skills and abilities, together with some degree of consistency regarding the course of action, might facilitate the human teacher to adapt her own demonstration. In turn, this human demonstration should be effectively used by the robot in order to build up its abilities.

Finally, considering the issue of feedback on a related note, it is worth mentioning that Thomaz and Breazeal make some observations about the way people tend to administer their own feedback when teaching a Reinforcement Learning agent [Thomaz and Breazeal, 2008]:

a. they use the reward channel not only for feedback, but also for future-directed guidance;
b. they have a positive bias to their feedback, possibly using the signal as a motivational channel; and
c. they change their behavior as they develop a mental model of the robotic learner.

These last points (about people's feedback) are outside the current scope of the work presented here (focusing more on the robot's feedback), but very much of interest in the broad context.

6. Future work

As the work and results presented here aim to inform the design and development of more natural HRI interfaces, possible next steps include the extension of another system (targeted as a companion[14] robot) by implementing our design suggestion of advertising that particular system's capabilities. Doing this in a more autonomous (and less WOZ) system would also allow to conduct studies where any improvement of performance (from the part of both the human and the robot) can be measured in more quantitative ways.

In both studies the participants were relatively young; it is possible that slightly different results can be obtained in a study that considers elderly persons instead, as attitudes towards technology and expectations will differ. Differences can be also be expected conducting a study in dierent cultural settings; see for example footnote 5 (dissimilarities), but also footnote 6 (similarities).

The results as shown in Figure 8 (top right) suggest that a variety of object localizations can be used 'spontaneously' by the demonstrating humans. The influence of the human cognitive abilities related to object localization (as well as any personal preferences) on the disproportionate adaptation rate, as well as any influence of the feedback on the demonstration segmentation, require additional studies.

Another possible future direction is looking at the response timing of the feedback (time between the end of the participant's utterance and beginning of the robot's utterance); this could be controlled as a function of the duration of the participant's utterance (similar to the approach in [Muto et al., 2009]). This timing synchronization would facilitate a *co-creation* process (cf. [Miyake, 2003]), meaning the co-emergence of real-time coordination by sharing subjective time and space between different persons (and robots).

Acknowledgments

This work was supported by a FY2007(P07369) Postdoctoral Fellowship for Foreign Researchers from the Japan Society for Promotion of Science (JSPS) and was partly supported by EU Integrated Project, LIREC (LIving with Robots and interactivE Companions), funded by the European Commission under FP7-ICT under contract FP7-215554. The first author would like to acknowledge the hosting by Prof. Yoshihiro MIYAKE in Tokyo Institute of Technology. The authors would like to thank Yukiko Thuresson and Yumiko MUTO for assistance organizing and conducting the user studies, Shohei YOSHIDA, Kengo OKITSU, Tatsunori NISHI and Shoji TAKASUGI for technical assistance with the robot.

References

Alissandrakis, A., & Miyake, Y. (2009). Human to robot demonstrations of routine home tasks: Acknowledgment and response to the robot's feedback. In *New Frontiers in Human-Robot Interaction, symposium at AISB 2009 convention* (pp. 9-15), 6-9 April 2009, Edinburgh, Scotland.

Alissandrakis, A., Nehaniv, C.L., & Dautenhahn, K. (2006). Action, state and effect metrics for robot imitation. In *The 15th IEEE International Symposium on Robot and Human Interactive Communication, 2006 (ROMAN 2006)*, (pp. 232–237).

Alissandrakis, A., Otero, N., Saunders, J., Dautenhahn, K., & Nehaniv, C.L. (2010). Helping robots imitate – Metrics and computational solutions inspired by human–robot interaction studies. In J. Gray & S. Nefti-Meziani (Eds.), *Advances in Cognitive Systems* (pp. 127–167). IET Press.

Calinon, S., Guenter, F., & Billard, A. (2007). On learning, representing and generalising a task in a humanoid robot. *IEEE Transactions on Systems, Man and Cybernetics, Part B. Special Issue on robot learning by observation, demonstration and imitation*, *37*(2), 286–298.

Dahlbäck, N., Jönsson, A., & Ahrenberg, L. (1998). Wizard of Oz studies – Why and how. In *Readings in intelligent user interfaces* (pp. 610–619). San Francisco CA: Morgan Kaufmann Publishers Inc.

Dautenhahn, K. (2007). Socially intelligent robots: Dimensions of human–robot interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1480), 679–704.

Dautenhahn, K., & Nehaniv, C.L. (2002). An agent-based perspective on imitation. In K. Dautenhahn, & C.L. Nehaniv (Eds.), *Imitation in animals and artifacts* (pp. 1–40). Cambridge MA: The MIT Press.

Galef, B.G., & Heyes, C.M., (Eds.). (1996). *Social learning in animals: The roots of culture*. San Diego CA: Academic Press.

Kanda, T., Ishiguro, H., & Ishida, T. (2001). Psychological analysis on human-robot interaction. In *Proceedings of the 2001 IEEE International Conference on Robotics and Automation, ICRA 2001*, May 21–26, 2001, Seoul, Korea, (pp. 4166–4173).

Mitchell, R.W. (1987). A comparative-developmental approach to understanding imitation. In P.P.G. Bateson & P.H. Klopfer (eds.), *Perspectives in Ethology 7: Alternatives* (pp. 183–215). New York NY: Plenum Press.

Miyake, Y. (2003). Co-creation in man-machine interaction. In *Proceedings of the 12th IEEE International Workshop on Robot and Human Interactive Communication (ROMAN 2003)* (pp. 321–324).

Muto, Y., Takasugi, S., Yamamoto, T., & Miyake, Y. (2009). Timing control of utterance and gesture in interaction between human and humanoid robot. In *The 18th IEEE International Symposium on Robot and Human Interactive Communication, 2009 (RO-MAN 2009)* (pp. 1024–1028).

Nehaniv, C.L., Dautenhahn, K., Kubacki, J., Haegele, M., Parlitz, C., & Alami, R. (2005). A methodological approach relating the classification of gesture to identification of human intent in the context of human-robot interaction. In *IEEE International Workshop on Robot and Human Interactive Communication, 2005 (ROMAN 2005)* (pp. 371–377).

Otero, N., Alissandrakis, A., Dautenhahn, K., Nehaniv, C.L., Syrdal, D.S., & Koay, K.L. (2008a). Human to robot demonstrations of routine home tasks: Exploring the role of the robot's feedback. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction* (pp. 177–184), Amsterdam: ACM.

Otero, N., Nehaniv, C.L., Syrdal, D.S., & Dautenhahn, K. (2008b). Naturally occurring gestures in a human–robot teaching scenario. *Interaction Studies*, *9*(3), 519–550.

Otero, N., Saunders, J., Dautenhahn, K., & Nehaniv, C.L. (2008c). Teaching robot companions: The role of scaffolding and event structuring. *Connection Science*, *20*(2&3), 111–134.

Saunders, J., Nehaniv, C.L., Dautenhahn, K., & Alissandrakis, A. (2007). Self-imitation and environmental scaffolding for robot teaching. *International Journal of Advanced Robotics Systems, Special Issue on Human-Robot Interaction*, *4*(1), 109–124.

Syrdal, D.S., Walters, M.L., Koay, K.L., Woods, S.N., & Dautenhahn, K. (2007). Looking good? Appearance preferences and robot personality inferences at zero acquaintance. In *AAAI Spring Symposium 2007, Multidisciplinary Collaboration for Socially Assistive Robotics* (pp. 86–92).

Thomaz, A.L., & Breazeal, C. (2008). Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence, 172*(6–7), 716–737.

Zentall, T.R. (1996). An analysis of imitative learning in animals. In B.G. Galef & C.M. Heyes (eds.), *Social learning in animals: The roots of culture* (pp. 221–243). New York NY: Academic Press.

Zentall, T.R. (2001). Imitation in animals: Evidence, function and mechanisms. *Cybernetics and Systems, 32*(1–2), 53–96.

Zentall, T.R., & Galef, Jr., B.G. (eds.). (1988). *Social learning: Psychological and biological perspectives*. Mahwah NJ: Lawrence Erlbaum Associates.

Zukow-Goldring, P., & Arbib, M.A. (2007). Affordances, effectivities and assisted imitation: Caregivers and the directing of attention. *Neurocomputing, 70*(13–15).