Issues in Human/Robot Task Structuring and Teaching

Joe Saunders, Nuno Otero and Chrystopher L. Nehaniv

Adaptive Systems Research Group, University of Hertfordshire, College Lane, Hatfield, Herts AL10 9AB, UK
e-mail : J.2.Saunders|N.Otero|C.L.Nehaniv@herts.ac.uk

Abstract—Teaching a robot new skills may require that the teacher scaffolds the teaching experience appropriately. However, due to inherent assumptions made by a human teacher the scaffolding process may in some circumstances fail to effectively teach the robot. Here we illustrate this issue in two simple robot teaching exploratory studies and examine the assumptions made by the teacher when teaching the robot. In the first study the human teacher had to reason about robot perceived states in order to provide suitable teaching. In the second study the human teachers had to understand the perceptual constraints of the robot based on the instructions given beforehand by the experimenter and subsequently adapt the guidance given.

The results suggest that although the two tasks are quite distinct in their level of complexity a common thread can be observed: people tend to underspecify their teaching. It seems that steps of the explanation are assumed to be known and skipped or not even considered at all. We reflect on the possibility that one of the major challenges in designing robots that are capable interaction partners in these teaching situations is to be able to make them communicate their internal state and current capabilities effectively. Furthermore, we also reflect on designing appropriate behavioral primitives for the robot, corresponding implications on the level of task description and for benefiting from human teaching.

I. INTRODUCTION

In our previous work we have developed, implemented and validated a robot social learning architecture on various physical robotic platforms [1]. The architecture allows a possibly naïve human teacher to teach a robot skills and competencies in a flexible, extendable and open-ended way. The human trainer teaches the robot via the process of 'putting through' and constructs hierarchies of skills based on the developmental psychological idea of task scaffolding combined with an ecological idea of environmental scaffolding. Part of this research also considered how a teacher might approach the problem of teaching a robot [2] and we have observed that, in some circumstances, the teacher may fail to teach the robot effectively (see also [3] for a complimentary approach to this issue). We have suggested that this may be due to inherent assumptions made by a human teacher during the teaching process regarding the capabilities of the robot and the fact that people are not very explicit in the instructions that they give to the robot. In order to explore this phenomenon two simple robot teaching exploratory studies were carried out. These serve to examine the assumptions made by the teacher when teaching the robot.

Prior to describing these studies and in order to place these studies in context we provide, in section II, some background information on the relationships between the learning sciences and human robot interactions. We then describe, in section III, two different user studies involving human robot teaching. The first explores how, when teaching a robot a very simple state based task, training episodes or states can be missed by the human teacher. In this task the teacher trains the robot using the learning architecture described in [1]. The second study considers a more complex teaching task which further reinforces this idea of task 'underspecification'. In this study the teacher attempts to show the robot how to "lay the table". In this latter study the robot is assumed to be learning from observation and responses from the robot are enabled using wizard-of-oz techniques [4]. This section also includes detailed results from both of these studies. In section IV we discuss the results of both studies in the broader field of the learning sciences and make suggestions on how the issues arising from the studies might be addressed in the future.

II. BACKGROUND

Many of the challenges facing HRI and learning/teaching tasks have strong resemblances to issues raised in 'AI in education' or the learning sciences in general. Greeno [5] describes a research framework, which he terms *situative learning*, that investigates how complex social organizations (involving teachers, learners, materials and the physical environment) facilitate (or hinder) the emergence of learning phenomena. According to Greeno: "Situative analyses include hypotheses about the principles of coordination that support communication and reasoning in activity systems, including construction of meaning and understanding." (p. 79). Proponents of the distributed cognition approach also advocate the same type of framework (for example, see [6]). By considering the importance of the interactional processes between social agents and the inter-connections with the physical environment, these perspectives are in line with overall conceptualization under discussion in this paper. To design human-robot interactions we will have to analyze the discourse and interactions happening during learning/teaching episodes and also how the actual actions exerted in the physical environment are indicative of the informative exchanges between the teaching/learning partners. Active and well-chosen scaffolding of the environment by the teacher [1,2] or by the designer of human-robot interactions can lead to more efficient learning by the robot.

III. The User Studies

A Experiment 1 - Teaching a robot a simple stimulus-action task

In this study human teachers were asked to train a robot using the process of 'putting through'. This is where the teacher moves the robot through the actions that need to be taught whilst the robot learns the appropriate motor actions in its given environmental state. The actions and state are recorded. When executing the taught behaviour the robot finds the state closest to that previously taught and executes the corresponding action. This experiment required the execution in certain circumstances of simple actions from a set of primitives that had to be taught to the robot in a very restricted environmental state. The actions were limited to "move the gripper up" or "move the gripper down". The states were restricted to "a light is on" or "a light is off". There was no feedback during training from the robot to the trainer, however the robot could be tested to see if it had learned the task immediately after training had finished.
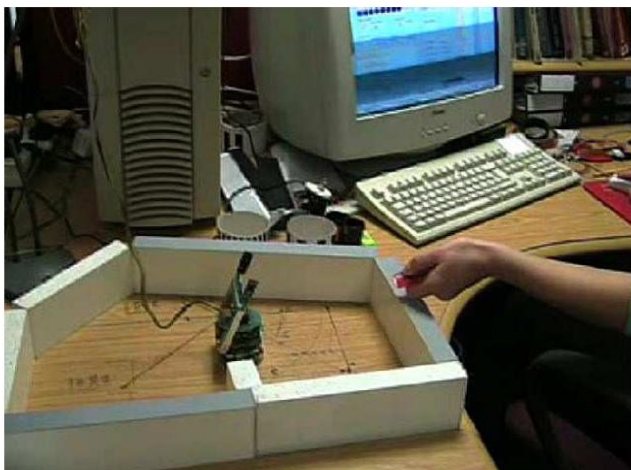

Figure 1. Teaching a robot a simple stimulus-action task. The robot is 'put through' its movements via a computer interface allowing direct manipulation by the teacher whilst simultaneously recording its state. When tested, the robot executes the action associated with the nearest learned state.

Our expectation in carrying out this experiment was that the teacher would have no difficulties in training the robot to carry out this task. However, this proved not be the case.

B. Instructions and Experimental Setup

For this study the robotic architecture described in [1] was used on miniature 5cm diameter non-holonomic Khepera robots equipped with a gripper arm. Four test subjects (3 male, 1 female) acting as teachers were asked to train the Khepera to behave as follows:

"the robot should normally have its gripper arms in a raised position unless a light is shone on it. The robot should then lower its arms".
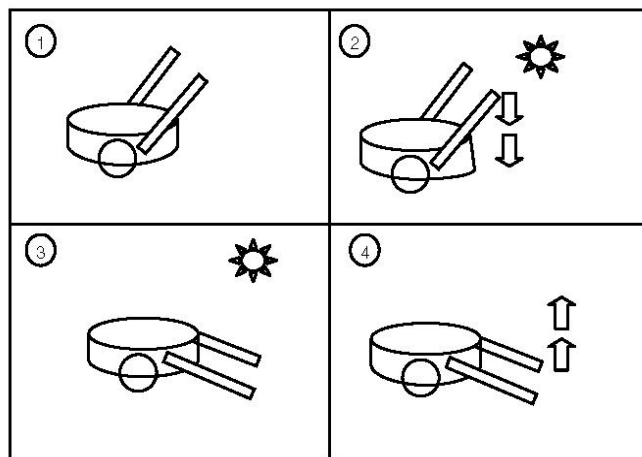

Figure 2. Teaching a robot a simple stimulus-action task. Each trainer carried out the same procedure and taught the robot situations 2 and 4 above. However, situations 1 and 3 were ignored by all of the teachers.

C. Teaching a robot a simple stimulus-action task - Results

All four of the test subjects then consistently carried out the same training procedure (see fig. 2) as follows:

1. learning started, robot placed in an 'arm up' position
2. trainer shines light on Khepera (see box 2, figure 2), robot 'put through' to an 'arm down' position
3. light turned off, robot 'put through' to an arm up position (see box 4, figure 2)
4. learning terminated

The learning architecture allows immediate execution by the robot of the competence taught to it by the teacher. The trained behaviour was started from the 'arm up' position and the trainer(s) asked to test the behaviour by instructing the robot to execute the taught competence.

Testing proceeded with the trainer shining the light on the Khepera. The robot in *all* cases correctly responded by moving to the 'arm down' position. However the robot then immediately moved its arm back to the 'arm up' position. If the light remained on the robot would repeat this sequence. The behaviour of the robot was a surprise to all of the test subjects who had expected the gripper arm to remain in the down position if the light remained on it. In fact, the robot's unexpected behaviour occurred because the 'move arm up' action was the result after finding the robot learned state most similar to the current state. The robot had only been trained to move to the 'arm down' position when the light was on and move to the 'arm up' position when the light was off, *it had not been trained to keep its arm down when the light was on and inversely had not been explicitly trained to keep its arm up when the light was off (see table 1)*. These extra training states are what we call *missed states*, which can be thought of as 'not obvious'. Typically these training episodes only become obvious when discovered from testing the robot.

Table 1. Results from stimulus-action Study

| Robot State | | Actions | |
Arm	Light	Action that should be Taught	Action actually Taught
Up	Off	Arm up	Not taught
Up	On	Arm down	Arm down
Down	Off	Arm up	Arm up
Down	On	Arm down	Not taught

The results indicated that assumptions are made by the teacher on how the robot learns. The training given to the robot by the teacher(s) could be due to the teachers(s) only considering the positive results of the training experience i.e. only those instructions which achieve the desired state. Alternatively the teacher may be assuming that there is an implicit 'while' or 'maintain' condition (which would be the case when training another human – perhaps an unspoken rule that one should 'do nothing' in cases not covered by the training), e.g. teacher says "when light is on put your arm down" however what the teacher means is "when light is on ALWAYS keep your arm down". Training which operates to *keep* the robot in a desired state does not appear to be considered in advance. Alternatively, the nature of the task (which here was highly discrete) may be problematic since a state in which no teaching has occurred but no action is required might be close to one in which an action is taught, however this experiment was designed to be simple in order to illustrate situations that can arise in teaching a robot behaviours. In other more complex teaching studies that we have carried out similar problems occurred [2]. In these cases the states were more numerous and far less discrete in nature. Proximity of states in which no action should be taken to those in which an action was explicitly taught could not alone allow them to be distinguished from nearby states in which generalization of the action should occur.

D. Experiment 2 - Demonstrating how to lay a table

In this experiment a more open scenario was used which illustrated further aspects that arise during a robot training session. In this case the robot control/learning system is dispensed with completely and only the actions of the teacher are analysed. Here the human trainer attempts to teach a robot to "lay a table". Instructions on how the human should go about this are given in advance and stress aspects such as clarity of the demonstration, use of speech and gestures and speed of movement. The study concentrated on how the human trainer went about teaching the robot. In this study there was no attempt by the robot to demonstrate that it had learnt the behaviour. The only feedback to the trainer was driven by an external operator sending positive and reassuring speech commands (indicating that the robot had understood the teacher) from the robot to the trainer. The experiment is described in detail below.

The "How to Lay a Table" study followed a within-subjects design where the participants were told to demonstrate how to lay a table for one person utilizing either of two strategies:

1. Using gestures only
2. Combining gestures and speech

For this paper only four participants' demonstrations in the second condition are discussed in a qualitative way since the aim is to provide a set of exemplars that can highlight a range of styles of demonstrating the task. The overall sample of the study run, however, consisted of 4 female and 8 male participants.

Figure 3 - The "Robot House" and the robot's view of a participant showing a plate before transporting it to the table.

E. The physical setting and material utilized

The study took place at the University of Hertfordshire "Robot House". This is a normal apartment where studies in HRI have been run to approach more ecologically valid conditions regarding the way people might interact with robots in their own home. The following additional materials were used to set the scene:

- one cup, plate, and cutlery in order to demonstrate laying a table for one person,
- a camera for video-recording
- a robot (an ActivMedia Peoplebot) that gave positive verbal feedback to the participants whenever an object was laid on the table

The production of the feedback was controlled by one of the experimenters using the Wizard of Oz technique [4]. The robot said things such as: "I understand", "I am following", "OK".

F. Instructions given to the trainer prior to the experiment.

At the beginning of the session, the participants were briefed about the study and specific instructions were given. The participants were told that their performance was not being assessed and that there was no right or wrong way of doing things. They were asked to demonstrate how to lay a table for one person:

a. pick up a plate, a cup and cutlery from the kitchen and lay the table in the living room.
b. put the objects back into the kitchen.
c. they should only move one object at a time and they should first show the object and then demonstrate where to put it.
d. the experimenter also noted that they should also try to ensure that the robot could see the actions, not to cover the objects, and to perform actions slowly. Pointing was also encouraged.

Finally, depending on the condition we explained that only gestures or speech and gestures could be used in the demonstration.

G. The coding scheme

This experiment was video-taped and the actions of the human trainer subsequently analysed. The coding scheme utilized here to classify the observed gestures is in line with the multi-functional classification system proposed by Nehaniv et al. [7]. The coding scheme has been developed, validated and successfully used in other studies [see, 8, 9]. Nehaniv et al. [7] propose the following six classes of gestures:

- *Manipulative gestures* - gestures that involve the displacement of objects.
- *Symbolic gestures* - these are gestures that follow a conventionalized signal. Its recognition is highly dependent on the context (both current task and cultural milieu).
- *Interactional Gestures* - this category identifies gestures used to regulate interaction with a partner.
- *Deictics* - the gestures that fall into this category are gestures used to indicate objects or loci of interest.
- *Side Effect of Expressive Behaviour* - these are gestures that occur as side-effects of people communicative behaviour.
- *Irrelevant* - these are gestures do not have a primary communicative or interactive function.

H Demonstrating how to lay a table - Results

The results below are separated into *explicit style* where the participant gave a detailed teaching demonstration to the robot, *inexplicit referencing* where no indication is given by the participant of the relative positions of the objects on the table and finally, *inexplicit explanation* where the participant gave even fewer indications to the robot in the demonstration regarding how to perform the task.

1) Explicit style

Table 2 shows some descriptive statistics regarding the demonstration of participant 2. Participant 2 gave a detailed and explicit demonstration. The utterances produced are

clearly separable in two types:

a. identifying the object being picked (e.g., "Pick up the plate") and
b. indicating the location of where to place it (e.g., "Put the knife on the right hand side of the plate on the table").

Furthermore, the location of the object on the table is given by relative coordinates to explicitly mentioned objects.

Table 2 - Frequencies and duration for the different gestures coded the video recording of participant's 2 demonstrations.

Type of gesture	Freq.	Duration (seconds)		
		Mean	Min.	Max
Utterances	9	NA	NA	NA
Interactional	6	2.4	1.44	2.76
Deictics	5	1.79	0.32	3.08
Manipulative	8	4.25	1.52	6.96

In terms of accompanying gestures to the utterances, the naming of the objects is linked, in terms of time of occurrence and content, to manipulative gestures clearly intended to show the object (following one of the rules presented in the beginning of the session). Verbal reference to the location on the table happens at the same time as the manipulative gesture of transporting the object to its place on the table. So, all the utterances occur while the participant is showing or transporting the object. Furthermore, he directs his eye gaze towards the robot (interactional gesture) when he speaks.

Within this multi-function classification system, the showing of the different objects (manipulative gestures) was also classified as cases of deictics. There was one occasion when the participant pointed to the object and its final location on the table.

2) Inexplicit referencing

Table 3 - Frequencies and duration for the different gestures coded from the observation of the video recording for the participant's 4 demonstration.

Type of gesture	Freq.	Duration (seconds)		
		Mean	Min.	Max.
Utterance	8	NA	NA	NA
Interactional	5	0.6	0.4	0.8
Deictics	8	1.2	0.92	1.28
Manipulative	4	6.8	5.9	6.92

To some extent the demonstration of participant 4 is similar to participant 2. The utterances also distinguish between naming the objects and naming the place where they should be laid (e.g., "this is the plate" and "you put that in the middle"). However, the relative coordinates of the location of the object on the table stated are not as explicit as for participant 2 since there is no mentioning of the other objects that could serve as points of reference. Furthermore, the manipulative gestures showing the object and showing the location are also

conflated. Although it is possible to observe that the participant did show the object and then placed it on the table the overall movement is continuous – in other words we do not see the same clear separation of the two functions as with participant 2. The video coding reflects this difference, instead of having two separate manipulative gestures "grasp" (corresponding to a deictic gesture) and "transportation" we have "grasp and transportation". The deictic classification appears four times at some point of the "grasp and transportation" manipulative gesture.

3) Inexplicit explanations

The demonstrations of the following two participants (1 and 3) go a step further to less detailed and less explicit demonstrations. Tables 4 and 5 indicate that both participants did not show examples of deictic gestures. In the utterances produced by participant 1 the reference to the object and its place on the table are mingled. For example, the participant says in a continuous form "The glass on the right hand side". Moreover, all the utterances happen at the end of "grasping and transporting" (manipulative gestures) the different objects to the table. No clear indication regarding showing the object is given. However, this participant briefly looked at the robot while speaking (interactional gesture). Looking at his utterances one can also observe that the relative coordinates are far from informative since no reference points were identified.

Table 4 - Frequencies and duration for the different gestures coded from the observation of the video recording for the participant's 1 demonstration.

Type of gesture	Freq.	Duration (seconds)		
		Mean	Min.	Max.
Utterance	4	NA	NA	NA
Interactional	7	0.4	0.2	1
Manipulative	4	5.47	4.6	6.6

Table 5 - Frequencies and duration for the different gestures coded from the observation of the video recording for the participant's 3 demonstration.

Type of gesture	Freq.	Duration (seconds)		
		Mean	Min.	Max.
Utterance	4	NA	NA	NA
Interactional	4	1.0	0.8	1.2
Manipulative	4	5.1	4.7	5.5

Participant 3, in her turn, gave the least detailed and explicit demonstration of the four participants considered here. In terms of the verbal productions, she started with one utterance naming the object and using the word "here" as a marker to the corresponding manipulative gesture that showed where to place the object - "Put the plate here". In the following utterances, however, only the objects were referred to giving some implicit indication of the order (which may tacitly have served to specify relative placement) e.g. "And then the fork...".

IV. DISCUSSION AND CONCLUSIONS

We have described above two different, but we believe complimentary experiments, which demonstrate that human teachers tend to be inexplicit in their teaching demonstrations and make assumptions about the capabilities of the robot.

Investigating human learning and building supportive artifacts has a long tradition in many disciplines. Our present focus, humans teaching robots, somehow seems to shift the perspective but we would argue that it maintains the general research challenges. For example, du Boulay [10] reviewing the progress of artificial intelligence in education points out challenges facing the teacher, some of which seem quite appropriate to be considered here. He identifies difficulties on understanding the level of consonance between the teacher's and learner's goals. In terms of human and robot teaching these could be construed as follows:

- how much time is a teacher really willing to spend teaching a robot?
- how easily can the teacher clearly formulate his requirements to the learner?
- how can the teacher give a good explanation?
- how can the teacher know when to re-state an explanation?
- when can the teacher provoke a response in order to assess the state of the learner?

More research is needed to understand which strategies to use in order to design robots that are truly able to make the teaching interactions between humans and robots efficient, efficacious and satisfactory. Let us take a closer look at the results of the studies reported here.

In considering the first experiment, the human teacher firstly had to clearly formulate his requirements to the robot and secondly, in order to succeed, would have had to understand what the limitations of the robot were in respect to the learned states. The difficulty here is two-fold. The teacher makes the, not unnatural, assumption that the robot interprets instructions as would another human and is not helped by the rather strange behaviour of the robot that results during testing. Here it would be useful for the robot to indicate that it has "holes" in its taught competencies. Finding these problem areas might be similar to a compiler finding unbalanced "if-then-else" constructs in a programming language. Alternatively, the design of the robot learning system might be modified to cope with the human assumptions more effectively if they could be formulated and operationalized.

The results obtained in the "How to Lay a Table" study suggest support to previous findings from other studies we conducted regarding the paucity of explicit pointing and symbolic gestures in similar demonstrations [8, 9]. Furthermore, two of the rules regarding the way the demonstration should be segmented were not unambiguously followed by three of the participants being described here ("You should try to follow two stages: (a) show the object and (b) demonstrate where to put it" and "Try to define the positions of the objects with clear indications. For example, pointing").

In fact, it seems clear from these four descriptions of the participants' teaching/explanations that there is quite a wide range of styles in this particular type of task: from a detailed approach following the rules given (participant 2), to a quite loose demonstration where no explicit indications were supplied (participants 1 and 3).

A general question is: to what extent does a robot need to cope with these different teaching styles? Considering a home use scenario one can easily imagine that different styles might coexist: from different members of a family, to people outside the family that work at the home. These issues suggest the need to find ways to cope with variation of styles. For the moment two distinct (but not exclusive) strategies come to mind. One possibility is to invest in the modeling of the robot's feedback in order to make it communicate its current internal state and understanding of the task taught. For example, the robot could communicate that its sensors were only being able to capture such-and-such gestures and that the user could try to slow down his movements and/or emphasize them. Furthermore, by explicitly stating how the task was apprehended the robot will be providing valuable information on which bits of the overall explanation may need further elaboration. Moreover, it could be also important that the robot could recognize the parts of a demonstration or explanation that are not clear enough and ask for clarification. The issues put forward here highlight the need to consider that teaching is a complex interaction between teacher and learner. However, a relevant HRI issue is: to what extent will people accept the demand to be more explicit in their explanations of tasks that they might consider simple? And, the point of view of respecting a humane role for technology, to what extent is it really necessary to make such demands on humans?

Another strategy, in addition (or complementary) to improving the robot's feedback, is to reconceptualize the way the overall activities need to be performed in order to minimize under-specifications and cope with technological constraints - segmenting tasks based on people's way of "chunking" the actions and understanding the level of familiarity and automaticity for the human. We are proposing to try to balance, on the one hand, the need to make the interaction between humans and robots as natural as possible and, on the other hand, to make the activity technologically feasible, given the current state of the art. In other words, we suggest that the design of typical tasks could be done at a level where people usually plan (or consciously reflect on) the task. As an example: in the "How to Lay a Table" scenario, instead of asking people to show the object and place it on the table we can ask people to first show all the objects and then show where they should be placed.

We expect that such research carefully examining the processes of embodied interactive behaviour will continue to play a key role in gaining an understanding of the prospects of designing different interaction strategies. Recognizing and coping with the diversity of user strategies will be crucial for the development of HRI design. Humans in interaction tend to design their utterances and behaviour in a way so as to take into account how their interaction partners are likely to in-

terpret it (*recipient design [11]*), and the learner robot can do this as well by making its internal state and current capabilities apparent to the human teacher. HRI designers too may have to take into account general communicative strategies of humans analogous to Grice's conversational maxims describing pragmatics of human language use in interaction [12], including "be relevant", "make your contribution as informative as is required for the current purposes of the exchange", "do not make your contribution more informative than is required", in allowing robots and humans to benefit from human instruction given to robots.

V. ACKNOWLEDGEMENTS

VI. REFERENCES

[1] J. Saunders, C. L. Nehaniv, and K. Dautenhahn, "Self-imitation and Environmental Scaffolding for Robot Teaching," *International Journal of Advanced Robotics Systems*, vol. 4, pp. 109-124, 2007.

[2] J. Saunders, "Observational Imitation, Self-Imitation and Environmental Scaffolding in Robotic Systems," University of Hertfordshire, Ph.D. Thesis 2006.

[3] A. L. Thomaz, G. Hoffman, and C. Breazeal, "Reinforcement Learning with Human Teachers: Understanding how people want to teach robots," in Proc, IEEE RO-MAN, 2006, pp.. 352-357.

[4] N. Dahlbäck, A. Jönsson, and L. Ahrenberg, "Wizard of Oz studies -- why and how," *Knowledge Based Systems*, vol. 6, pp. 258-266, 1993.

[5] J. G. Greeno, "Learning in Activity," in The Cambridge Handbook of the Learning Sciences, R. K. Sawyer, Ed. Cambridge: Cambridge University Press, 2006, pp. 79-96.

[6] G. Salomon, "Distributed Cognition: Psychological and educational considerations " in Learning in Doing: Social, Cognitive, and Computational Perspectives. Cambridge: Cambridge University Press, 1993.

[7] C. Nehaniv, K. Dautenhahn, J. Kubacki, M. Haegele, C. Parlitz, and R. Alami, "A methodological approach relating the classification of gesture to identification of human intent in the context of human-robot interaction," in Proc. IEEE RO-MAN, 2005, pp. 371-377.

[8] N. Otero, S. Knoop, C. Nehaniv, D. S. Syrdal, K. Dautenhahn, and R. Dillman, "Distribution and Recognition of Gestures in Human-Robot Interaction," in Proc. IEEE RO-MAN, 2006, pp. 103-110.

[9] N. Otero, C. Nehaniv, D. S. Syrdal, and K. Dautenhahn, "Naturally Occurring Gestures in a Human-Robot Teaching Scenario," in Proc. IEEE RO-MAN, 2006, pp. 533-540.

[10] B. du Boulay, "Fallible, Distractible, Forgetful, Willful and Irrational Learners.," in Computers as Cognitive Tools: no more walls., vol. Volume Two, S. P. Lajoie, Ed. NJ: Lawrence Erlbaum, 2000, pp. 339-376.