A model for inferring the intention in imitation tasks

Bart Jansen and Tony Belpaeme

Abstract— **Robot imitation comprises a number of hard problems, one of the most difficult problems is the extraction of intention from a demonstration. A demonstration can almost always be interpreted in different ways, making it impossible for the imitating robot to find out what exactly it is that was demonstrated. We first attempt to set out the problem of intention reading. Next, we offer a computational model which implements a solution to intention reading. Our model needs repeated interactions between the demonstrator and the imitator. Through keeping a score about which interactions where successful, the imitating robot gradually builds a model which "understands" what the intent is of the demonstrator.**

I. INTRODUCTION

The subject of imitation has received increasing attention in robotics [1]. The main reason for this is that imitation holds the promise of interacting with robots and teaching of robots in ways which come very natural to lay users. No programming would be needed to teach the robot how to vacuum the sofa; only a quick demonstration would be needed for the robot to understand what a sofa is and how it should be cleaned. However, between this rather utopian illustration and the current realities of robot imitation stand a number of very hard problems, problems for which only a partial solution or no solution at all exists. These issues are traditionally cast into five questions: who, when, what and how to imitate and how to judge if imitation was successful. The "how" question for instance investigates how to map observed actions to a —sometimes very dissimilar— body. When a human demonstrates a task to a robot, the robot needs to map the human's actions to its own motor skills. This mapping [2] and the degree to which sub-goals should be matched [3] have been the subject of study of a large part of research in robot imitation. For an overview of issues in imitation see [4].

This paper focuses on the problem of understanding the intention of a demonstration. Understanding *what* a demonstration is about comes deceptively easy to humans, but replicating this in an artificial system proves to be a challenge. If a human observes someone picking up a broom and sweeping the floor, one will immediately infer that the intent of that person is to clean the floor. However, for a naïve imitator it is hard to know what aspect should be imitated: should only the swishing sound of the broom be imitated? Would only using a stick do? Should one be standing in exact the same position? Wearing the same clothes?

Bart Jansen is with the Electrical Engineering and Informatics group of the Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium. bjansen@etro.vub.ac.be

T. Belpaeme is with the School of Computing, Communications and Electronics, University of Plymouth tony.belpaeme@plymouth.ac.uk

Imitation, and especially imitating the intent of a demonstration, is remarkably difficult, as it is unclear to a naïve observer —in our case: a robot— what the precise goal is of a demonstration.

We propose a model for recognising the *goal* of a demonstration, and imitating that goal. The model has been implemented and tested in simulation. We give details on the model, and demonstrate its behaviour. In the discussion we elaborate on the difficulties of porting the model to a physical robot.

II. RELATED WORK

Plan recognition, an area in human-computer interaction, is also concerned with extracting goals from a temporal series of events [5], [6]. However, plan recognition differs from goal-extraction in imitation in two aspects. First, plan recognition tries to classify observed action into one of several predefined classes. In imitation, these classes —or goals— are not predefined, instead they should be learnt from imitative interactions. Second, plan recognition's ultimate goal is to recognise or classify the goal of the demonstration *before* the demonstration ends. The applications in which plan recognition is used aim at automating tasks. For example, when the user is typing a letter, his word processor (and the plan recognition behind it) should recognise this as early as possible and offer advice or help in finishing the letter. In imitation, the user is expected to finish the demonstration before the robot will deduce the goal[1].

III. A COMPUTATIONAL MODEL

Our model contains a number of components. Some of these are rather trivial, such as mapping perception onto an internal representation or executing actions, and will not be dealt with in detail. However, there are two components which are essential to inferring the intent of the instructor. One is *perspective taking*, the other is maintaining a *model of belief* of the instructor.

- **Perspective taking** For the imitating agent to able to infer the intent of a demonstration, it will need to reverse its perspective to that of the instructor. The imitator imagines being the demonstrator and for each goal in its repertoire it investigates how it could lead to the observed blocks world. In language understanding, mutual knowledge and perspective-taking are crucial to correctly interpret linguistic utterances [8], [9]. This also holds for imitation. In order to understand what the goal

[1] Although there are some cases of immediate imitation [7].

is of an imitation the imitating agent will need some form of perspective-taking to solve ambiguities [10].
- **Model of belief** The imitating agent will need to maintain a model of belief of what the instructor is trying to achieve with the demonstration. When humans imitate each other, they can rely on a vast amount of world knowledge which helps them to disambiguate the demonstration. When building an imitating agent, two orthogonal approaches can be taken. Either world knowledge is pre-programmed, for example in the form of constraints, or the agent autonomously constructs a belief model of the goals of the instructor. In this paper, we will look at the latter —in our opinion, more interesting— option.

A. The approach

The model of intentional imitation which is outlined in this paper, is an agent-based computer simulation. Every agent acts on a blocks world in which it can manipulate blocks and visually observe changes in the configuration of the blocks. Rather than studying the exact copying of actions, we assume the agents to be able to perform and observe a set of actions. The exact vision and motor algorithms enabling these capabilities are not studied. In the case study which is used throughout the text, the blocks world is highly trivial: it is a two dimensional grid containing blocks. The agents can move the blocks one cell to any of the four directions, as long as the blocks remain on the board.

In its most simple set-up, two agents are placed around such a blocks world. One agent acts as a demonstrator, the other one as the imitator. The imitative interactions are organized by a strict scenario: *the imitation game*. The imitation game is a very simple interaction: the demonstrator performs some behaviour which is observed by the imitator. Upon completion of the demonstration, the imitator tries to imitate the observed behaviour. The demonstrator carefully observes the imitation and decides on the success of the imitative attempt. The demonstrator sends feedback about the outcome of the game to the imitator. Consequentualy, the imitator can use this feedback as a source of information for improving future imitation quality.

The kind of behaviour which is learnt by playing repeated imitation games is not shaped by the interaction protocol. Imitation games are just an abstract model for the learning of categories through repeated interactions. They were previously used in the imitation of speech sounds [11], gestures [12], sequences of actions [13] and intentional behaviour as in this paper [14].

In the case study which is used throughout the article, the agents imitate spatial relations over block configurations: the agents can express relations over pairs of blocks using for instance the relations *left-of?(A,B)* and *above?(A,B)*. Relations can be combined, such that for instance *left-of?(A,B)* $\wedge$ *above?(B,C)* denotes a configuration of the board in which a block A is at the left of a block B and block B is above a block C.

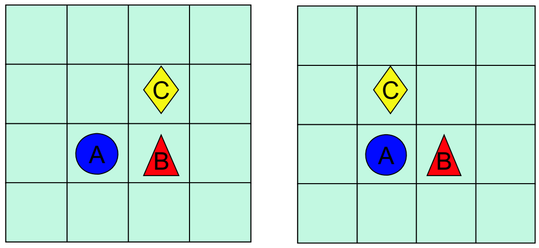

Fig. 1. Although the demonstrator performed actions transforming the board on the left into the one on the right with a clear goal in mind, the imitator can typically not uniquely determine that goal.

These are the spatial relations over configurations of blocks which are to be imitated by the agents. The imitation of such goals is non-trivial as the mapping of actions onto goals as well as the inverse mapping are non-unique: a given goal (f.i. *left-of?(A,B)*) can often be obtained in several ways (f.i. by moving block *A* to the left, or by moving *B* to the right). On the other hand, the same sequence of actions can be performed in order to obtain different goals, even when starting from the same configuration and thus leading to the same final configuration. Consider for instance the board's configuration as shown in figure 1 on the left. Performing the actions *Left(C)* and *Up(B)* will result in the configuration shown on the right in figure 1. As the imitator observes these actions being performed, he could not uniquely identify the demonstrator's goal. The demonstrator might have pursued the goal *left-of?(C,B)* or *above?(B,A)* or even the goal *left-of?(C,B)* $\wedge$ *above?(B,A)*, as all three goals were obtained by performing those two actions. Moreover, suppose that the imitator does not try to resolve this ambiguity and simply copies the final blocks' configuration. Although imitation succeeds, the imitator can not learn any goal. Thus, the sequence of actions performed by the demonstrator or the final blocks' configuration do not provide enough information to detect the goal of the demonstrator's behaviour.

Typically, various kinds of domain knowledge are used to facilitate the mapping of actions onto goals. In our case, four important assumptions implement this domain knowledge. We assume that all agents have the same set of action primitives (or at least know the primitives used by the others), that all agents have primitive predicates with the same expressive power (otherwise agents will not be able to learn a subset of all possible goals), that agents have a repertoire to store goals in, together with usage and success counters. As a fourth assumption, we stipulate that there is no deception: agents do not perform any action which could confuse the others, in other words, only actions which are directly relevant for the pursued goal are performed. By consequence, this entails that the planners used by the agents —although each agent might use a different planner— are optimal.

B. The game

Every agent A_k learns and maintains a repertoire R_k of goals. Each goal in the repertoire has a usage counter and a success score associated with it, so $R = \{(g, u, s)_i\}$. A goal g itself consists of a conjunction of predicates $g = q_1 \wedge q_2 \wedge ... \wedge q_n$ where all predicates $q_i \in P_k$.

The imitation game is defined as follows:

1) The demonstrator randomly selects a goal g from its repertoire, builds a plan p for it and executes the plan p.
2) The imitator observes this sequence of actions and finds the best matching goal g' from its own repertoire. If no suitable goal can be found in its repertoire, the game fails and a new goal is constructed.
3) The imitator builds a plan p' for reaching the goal g' in its own environment and executes the plan p'.
4) The demonstrator observes this sequence of actions and verifies whether its initial goal g holds in the blocks world arranged by the imitator. If that is the case, the game succeeds, in all other cases it fails.
5) The demonstrator informs the imitator about the success of the game. He sends a single bit of information to the imitator.
6) The imitator adapts its repertoire.

C. Goal categorization

The key issue in attributing intentions to a sequence of actions is implemented in the second step of the game. As the agents always categorize action sequences as a goal already in their repertoire, intention attribution is accomplished by finding the best matching goal in the agent's repertoire.

In order to find the best match, the imitator investigates every goal in its repertoire. For every goal which holds in the blocks world after the demonstration, he takes the perspective of the demonstrator and verifies what kind of actions the demonstrator would have performed if its goal were the goal currently considered by the imitator. Thus, for every goal in its repertoire, the imitator builds a plan as if it were the demonstrator and starting from its initial blocks configuration. By assuming optimality in the planners of the agents, all candidates which do not require exactly the same amount of actions as the original action sequence performed by the demonstrator, can be discarded.

This process of *perspective taking* heavily reduces the space of possible goals. From the remaining candidate goals (this set is called hypothesis set), the one with best score is taken as the intention of the demonstrated behaviour.

D. Learning

Goal categorization relies on the existence of a model of beliefs of the goals of the other agent, implemented as a repertoire of goals which have scores. The imitator starts without such a belief model and gradually constructs the model during the repeated demonstration cycles. The imitator has several learning mechanisms available: if it fails to categorize an observed sequence of actions as a goal in its repertoire, it can construct a new goal for it. Secondly, the scoring mechanism can be used to learn preference over goals for several action sequences. Moreover, using the scoring mechanism goals which are permanently unsuccessful can be identified and removed from the repertoire.

A new goal is created whenever the categorization process fails. This can happen if the repertoire of the imitator is empty or when no goal in its repertoire will lead to a configuration in which the same relations hold as in the state which was obtained by the demonstrator in the same number of actions. The goal creation process investigates the observed sequence of actions and constructs a random goal out of it such that the observed sequence of actions will be categorized as that goal by the agent's categorization mechanism. The goal creation process is explained in more detail in [14].

The scoring mechanism adapts the scores of the goals at the end of every game. First of all, the imitator increases the score of the goal it used. If the game succeeds, the score of that goal is increased, while the scores of all other goals in the hypothesis set are decreased by the same amount δ. Scores are limited between zero and one. All goals which reach a score which is below a given threshold ϵ are removed from the repertoire.

E. Example

Below three small examples are introduced which will provide better insight in the imitation game and the agents' goal recognition process. For the ease of the argument, a four by four blocks world was used. We suppose that both agents have only three blocks in their environment.

Suppose that the demonstrator wants to demonstrate the goal *left-of?(C,A)* to an imitator which has the following goals in its repertoire: *left-of?(C,B)*,*left-of?(C,A)* and *left-of?(A,B)*. The initial board of the demonstrator is drawn in figure 2 on the left. The planner of the demonstrator proposes the following sequence of actions: Left(C), Left(C), Left(C). The configuration after pursuing the goal is the one drawn in figure 2 on the right.

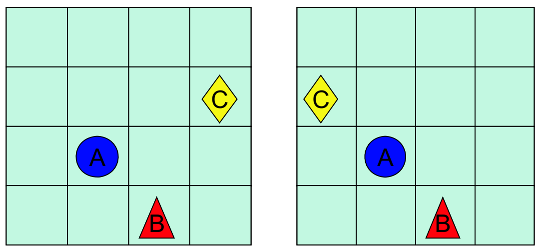

Fig. 2. The configuration of the demonstrator's environment before and after pursuing the goal *left-of?(C,A)*.

The imitator observes this sequence of actions being performed and investigates its repertoire. Obviously, *left-of?(A,B)* could not have been the goal intented by the demonstrator as it was already valid in the demonstrator's initial configuration. The imitator now builds plans for the other goals in its repertoire, as if it were the demonstrator himself. For the goal *left-of?(C,A)* the planner for instance proposes the following sequence of actions: Left(C), Right(A), Right(A). Thus, the goal *left-of?(C,A)* can thus be reached in exactly three steps, just as the demonstrator reached its goal. For the goal *left-of?(C,B)* the planner for instance proposes the following sequence of actions: Left(C), Right(B). As only two actions are required, this could not have been the demonstrator's intended goal. The only goal that remains in the hypothesis set is *left-of?(C,A)*, which will be the goal the imitator beliefs the demonstrator performed.

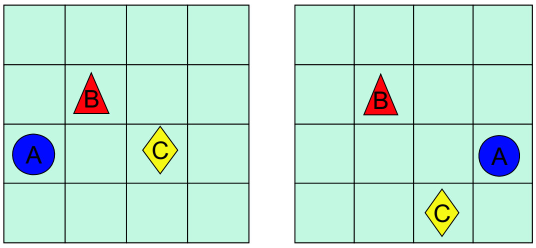

Fig. 3. The configuration of the imitator's environment before and after pursuing the goal *left-of?(C,A)*.

The imitator pursues this goal in its own blocks world (Figure 3 on the left). The planner proposes Down(C), Right(A), Right(A), Right(A) which results in the board depicted in figure 3 on the right.

In the final part of the game, the demonstrator verifies whether its original goal holds in the imitator's blocks world. As this is indeed the case, the game succeeds.

Now suppose that the imitator did not have the goal *left-of?(C,A)* in its repertoire. In that case, it would not have been able to categorize the observed sequence of actions, as it rejected the other goals in its repertoire as valid candidates. In that case, it would have constructed a new goal.

If the goal demonstrated by the demonstrator was *left-of?(C,A)* $\wedge$ *left-of?(C,B)*, it would also have been categorized as $left-of?(C,A)$ by the imitator. If the planner came up with the same plan as in the previous example, then the original goal of the demonstrator is not valid in the imitator's resulting configuration (Figure 3 on the right).

F. Measures

The imitation games are monitored with three complementary measures: *imitative success*, *repertoire size* and *average category distance*. The imitative success is simply a running average of the fraction of successful games. The repertoire size is a running average of the number of goals in the repertoire of the current imitator. The category distance (CD) provides a quantitative measure for the similarity between the repertoires of the demonstrator and the imitator and is defined in equation 1, in which capitals denote the repertoires of the agents, while small letters denote goals.

$$CD(A,B) = \frac{\sum_{a \in A} \min_{b \in B} d(a,b) + \sum_{b \in B} \min_{a \in A} d(a,b)}{|A| + |B|} \tag{1}$$

$d(x, y)$ is a binary difference measure on goals defined in equation 2.

$$d(a,b) = \begin{cases} 0 & \text{a = b} \\ 1 & \text{a} \neq \text{b} \end{cases} \tag{2}$$

IV. RESULTS

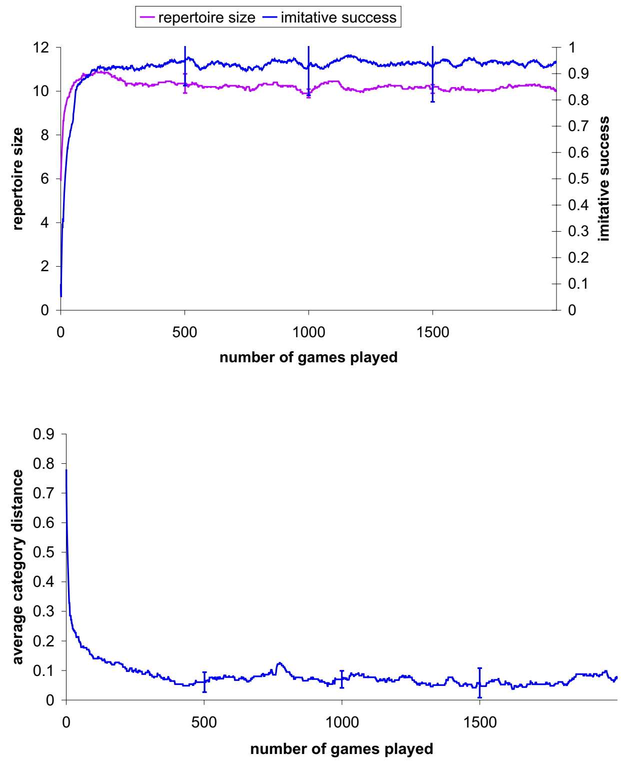

Fig. 4. A single demonstrator and a single imitator engage in 2000 imitation games. On top: imitative success and repertoire size. The category variance is plotted at the bottom.

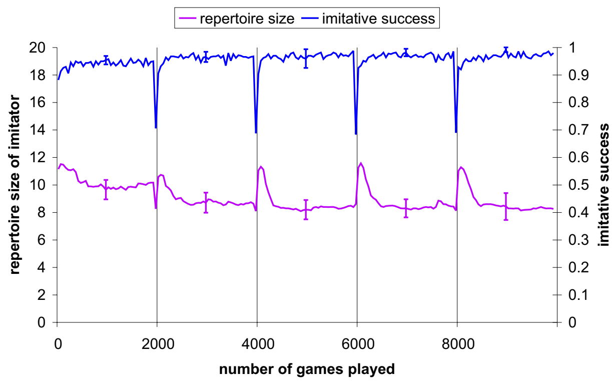

Fig. 5. Results of an iterative learning experiment in which every 2000 games the demonstrator is replaced by the imitator and a fresh imitator replaces the original imitator.

We performed a transmission experiment in which a single demonstrator is endowed with ten random goals. The imitator on the other hand starts the experiment without any goals in its repertoire. Both agents engage in 2000 imitation games. δ was set to 0.025 while epsilon was set to 0.1. The entire experiment was repeated ten times, 95% confidence intervals are plotted. Results are plotted in figure 4 and show how the repertoire of the demonstrator is transferred very quickly and reliably to the imitator: the repertoire size of the imitator stabilizes around ten goals while the category distance approaches zero. Imitative success immediately climbs towards 100%, so the imitator succeeds in synchronizing its repertoire with the demonstrator's one very fast. After several hundreds of games, the demonstrator's repertoire is stable, leading to successful imitation.

This simple experiment demonstrates how the imitator learns the intentions of the demonstrator by trying to imitate its actions. As this learning involves building up a repertoire of goals, this helps the imitator to classify each demonstration, and consequently the imitations become better and better. Detailed parameter studies on this simple vertical transmission experiment can be found in [15].

generation 0	
$Above?(C,A)$	1.0
$Above?(A,C) \wedge Left\text{-}of?(B,A)$	1.0
$Left\text{-}of?(C,B)$	1.0
$Above?(C,B) \wedge Left\text{-}of?(B,A)$	1.0
$Above?(A,B)$	1.0
$Above?(A,C)$	1.0
$Left\text{-}of?(A,C)$	1.0
$Above?(C,B)$	1.0
$Left\text{-}of?(B,C)$	1.0
$Above?(C,A) \wedge Left\text{-}of?(A,B)$	1.0
generation 3	
$Left\text{-}of?(A,B) \wedge Above?(C,A)$	1.0
$Left\text{-}of?(B,A) \wedge Above?(C,B)$	1.0
$Above?(A,B)$	1.0
$Left\text{-}of?(B,C)$	1.0
$Left\text{-}of?(A,C)$	1.0
$Above?(A,C) \wedge Left\text{-}of?(B,A)$	1.0
$Left\text{-}of?(C,B)$	1.0
$Above?(C,B)$	1.0
$Above?(C,A)$	1.0
generation 5	
$Above?(C,A) \wedge Left\text{-}of?(A,B)$	1.0
$Above?(C,B) \wedge Left\text{-}of?(B,A)$	1.0
$Left\text{-}of?(C,B)$	1.0
$Left\text{-}of?(B,C)$	1.0
$Left\text{-}of?(A,C)$	0.975
$Above?(A,B)$	1.0
$Left\text{-}of?(B,A) \wedge Above?(A,C)$	0.975
$Above?(C,A)$	0.35

TABLE I

THE ENTIRE REPERTOIRE OF THE AGENTS IS PLOTTED AT GENERATION ZERO (I.E. THE ORIGINAL REPERTOIRE OF THE DEMONSTRATOR), AT GENERATION THREE AND AT GENERATION FIVE. THIS PRINTOUT CONFIRMS THAT THE REPERTOIRE OF THE DEMONSTRATOR CAN INDEED BE TRANSFERRED RELIABLY.

In a second experiment, we investigate whether the initial repertoire of goals of the demonstrator can be transferred over multiple generations. An experiment was performed which starts just as the previous experiment. However, after 2000 games, a new imitator without any goals in its repertoire enters the experiment, while the current imitator replaces the demonstrator. Experiments run over five iterations of 2000 games to investigate the reliability of transmission over multiple generations. From figure 5 it is clear that the imitative success remains high throughout the entire experiment. After every iteration, the success drops instantly and then reaches its original high level very soon. After the fifth generation, the repertoire of the agents contains slightly less goals then ten. This is because sometimes the initial repertoire of the demonstrator contains two or more goals which are highly similar. For instance, if it has $Above(A,B)$ and $Above(A,B) \wedge Below(B,C)$ in its repertoire, the more specific one will often survive and replace the more general one.

In table I a printout of the original repertoire of the demonstrator is given (top) and of the imitator at the third and fifth generation. In the printout, the process of generalization can be observed clearly: At the third generation, the goal $Above?(A,C)$ and the goal $Above?(A,C) \wedge Left\text{-}of?(B,A)$ seem to have merged. Also between the third and fifth generation, the goals $Above?(C,B)$ and $Above?(C,B) \wedge Left\text{-}of?(B,A)$ are merged.

V. DISCUSSION

A. Insights from cognitive science

Traditionally robot imitation has often been influenced by work in psychology, biology and neuroscience on imitation. Most of this research has focused on the mapping of observations to actions and how both are linked through one system: the mirror system, which is the seat for action imitation in monkeys (and humans). Recently Fogassi et al. [16] showed neurophysiological evidence for intention reading: different neurons in the inferior parietal lobule fired when performing actions with a specific goal: either grasping for eating or grasping for placing. The same neurons showed specificity to not only the action, but also to the *goal* of the action. The same neurons were active when observing both goal-directed actions, showing that neurons were active in understanding the goal of an action. We only wish to mention this and do not speculate on how our model maps onto these findings. Our model clearly is a far cry from how the brain is organized, and only provides an implementation of learning the intention behind actions through repeated imitative interactions.

The relation of imitation to language has been hinted at before, as mirror neurons might play a role in action imitation as well as in vocal imitation. Rizzolatti and Arbib [17] suggest that the mirror neuron system is a homologue to Broca's area in humans, a cortical area crucial to speech production and comprehension. It has been speculated that language also draws heavily on perspective taking [8], [9]. The same has been proposed for imitation — e.g. [18]. This only strengthens the case that language and imitation might have roots in or might even be drawing on the same cognitive functions.

B. Future directions

The model implements goal-directed imitation through repeated imitative interactions. The sheer number of imitative interactions, which exponentially increases with the number of states, makes it difficult to implement the model onto a physical robot. In figure 4 for example, agents need a few hundred interactions to transfer ten goals from the demonstrator to the imitator. If a physical system would be able to demonstrate a goal every minute, then this approach would be rather slow. However, it is not unreasonable to assume that children go through a similar process. Children however manage to build on previous experiences and construct their skills hierarchically. The possibility of building new goals by concatenating existing goals seems promising, and could be a way to speed up goal-directed imitation learning.

Our computational model now assumes that the states and actions form a deterministic Markov decision process. A possible extension would be to work away this limitation. Allowing continuous states and actions, instead of discrete ones, might be done through discretisation. Allowing non-deterministic states and actions might be more difficult, as significant changes might be needed to the learning and planning algorithms. Especially the planner, which in the current deterministic implementation assumes that each plan is optimal to deduce the goal behind a plan, will need to be revised.

VI. CONCLUSION

We identified two essential elements in constructing goal-level imitation: *perspective-taking*, seeing the events and effects on the world through the perspective of the instructor, and *belief modeling*, constructing and maintaining a model what goals the instructor wishes to demonstrate. We have presented and demonstrated a computational model, in which perspective-taking was implemented by letting the imitator construct a plan from the perspective of the instructor. This allows the imitator to disambiguate possible conflicting goals in a demonstration. The belief modeling was implemented by letting the imitator try a possible goal and keep a memory of how the instructor judged the imitation. Experimental results on a simulated model based on perspective taking show that it is indeed possible to attribute intentions to a set of behaviours. However, the model is also limited in the sense that it does not implement one-shot imitation learning, but instead requires a substantial number of training interactions.

VII. ACKNOWLEDGMENTS

Bart Jansen is currently working on Brucare, a project funded by the Brussels Capital Region

REFERENCES

[1] P. Bakker and Y. Kuniyoshi, "Robot see, robot do: An overview of robot imitation," in *Proceedings of the AISB Workshop on Learning in Robots and Animals*, 1996, pp. 3–11, brighton, UK, April 1996.

[2] S. Schaal, A. Ijspeert, and A. Billard, "Computational approaches to motor learning by imitation," *Philosophical Transactions of the Royal Society of London: Series B*, vol. 358, pp. 537–547, 2003.

[3] A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn, "Imitating with ALICE: Learning to imitate corresponding actions across dissimilar embodiments," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 32, no. 4, pp. 482–496, 2002.

[4] K. Dautenhahn and C. L. Nehaniv, "The agent-based perspective on imitation," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds. Cambridge, MA: The MIT Press, 2002, pp. 1–40.

[5] C. Schmidt, N. Sridharan, and J. Goodson, "The plan recognition problem: An intersection of psychology and artificial intelligence," *Artificial Intelligence*, vol. 11, no. 1-2, pp. 45–83, 1978.

[6] S. Carberry, "Techniques for plan recognition," *User Modeling and User-Adapted Interaction*, vol. 11, no. 1-2, pp. 31–48, 2001.

[7] J. Nadel and G. Butterworth, Eds., *Imitation in infancy*, ser. Cambridge studies in cognitive and perceptual development. Cambridge, UK: Cambridge University Press, 1999.

[8] R. Krauss and S. Fussell, "Perspective-taking in communication: Representations of others' knowledge in reference," *Social cognition*, vol. 9, pp. 2–24, 1991.

[9] B. Keysar, D. Barr, J. Balin, and J. Brauner, "Taking perspective in conversation: The role of mutual knowledge in comprehension," *Psychological Science*, vol. 11, pp. 32–38, 2000.

[10] A. Thomaz, M. Berlin, and C. Breazeal, "An embodied computational model of social referencing," in *Proceedings of the Fourteenth IEEE International Workshop on Robot and Human Interactive Communication (Ro-Man 2005)*, 2005.

[11] B. de Boer, "Self-organisation in vowel systems," *Journal of Phonetics*, vol. 28, no. 4, pp. 441–465, 2000.

[12] B. Jansen, T. ten Thij, T. Belpaeme, B. De Vylder, and B. de Boer, "Imitation in embodied agents results in self-organization of behavior," in *From Animals to Animats: Proceedings of the 8th Conference on the Simulation of Adaptive Behavior*, S. Schaal, A. J. Ijspeert, A. Billard, and S. Vijayakumar, Eds. Cambridge, MA: The MIT Press, 2004, pp. 455–466.

[13] B. Jansen, "Simultaneous learning of basic and compound actions by imitating compound actions," in *Evolvability and Interaction: Evolutionary substrates of communication, signalling and perception in the Dynamics of Social Complexity*, ser. Technical Report No:393, P. McOwan, K. Dautenhahn, and C. Nehaniv, Eds. University of Hertfordshire, 2003.

[14] B. Jansen and T. Belpaeme, "A computational model of intention reading in imitation," *Robotics and Autonomous Systems*, vol. 54, no. 5, 2006, in press.

[15] B. Jansen, "Robot imitation - the emergence of shared behaviour in a population of agents," Ph.D. dissertation, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium, 2005.

[16] L. Fogassi, P. F. Ferrari, B. Gesierich, S. Rozzi, F. Chersi, and G. Rizzolatti, "Parietal lobe: From action organization to intention understanding," *Science*, vol. 308, no. 5722, pp. 662 – 667, 2005.

[17] G. Rizzolatti and M. Arbib, "Language within our grasp," *Trends in Neurosciences*, vol. 21, pp. 188–194, 1998.

[18] M. Johnson and Y. Demiris, "Perceptual perspective taking and action recognition," *International Journal of Advanced Robotic Systems*, vol. 2, no. 4, pp. 301–308, 2005.