You did it on purpose! Towards intentional embodied agents

Bart Jansen, Bart de Boer, Tony Belpaeme

Artificial Intelligence Lab
Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium
{bartj,bartb,tony}@arti.vub.ac.be
http://arti.vub.ac.be

Abstract. The paper describes a road-map towards intentional behavior in artificial systems. We catch the developmental path in two dimensions, a social and an intentional dimension. Starting out with a babbling phase, development continues over an exploratory phase without social interactions and a phase in which action-level imitation is used. The pinnacle of development is the intentional imitation of goals. An experiment, together with preliminary results, is presented for each developmental phase.

1 Introduction

In recent years a lot of attention has gone to the construction of intentional agents, *i.e.* agents that purposefully interact with their environment and with other agents. The field has witnessed several *ad hoc* solutions towards implementing intention in agents. Especially in the field of software agents several formalisms have been defined to enable agents to maintain beliefs about actions and their possible outcomes [1]. However, little attention has been paid to the developmental path that agents should take before arriving at fully fledged intentional behavior, and the importance of embodiment and grounding in this all.

In this paper we outline a possible developmental path for mastering the difficult task of intentional imitation. We discern four steps: motor babbling, imitation of gestures, individual exploration and intentional exploration [2]. We argue that embodiment is an important factor in all four steps. To substantiate the theory, we present computer simulations and results for the first two steps and outline possible computer simulations for the last two.

2 Social and intentional dimensions of behavioral development

When studying behavioral development and learning, it is easy to be baffled by the complexity of the subject. Many different types of behavior are encountered.

One encounters such apparently disparate phenomena as motor babbling, mimesis, exploration of the world and imitation of object use. All of these phenomena are associated with different stages of development, but at the same time many of them occur in parallel. Although at first sight these phenomena may appear to have little in common, we argue that in fact they can be analyzed in terms of two developmental parameters: a social dimension and an intentional dimension.

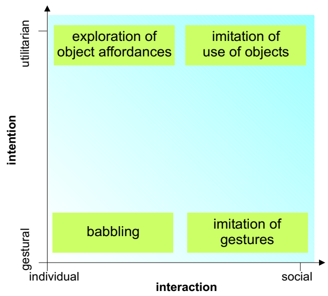

Fig. 1. social and intentional dimensions of behavioral development.

The social and intentional dimensions can be imagined as forming an abstract two-dimensional space (see figure 1). At the zero ends of both dimensions there is behavior without intention and without social content. As one moves away from zero, behavior becomes more intentional and more social, respectively. Although it is clear that this proposed space is continuous, in the sense that behavior can be more or less social or intentional, for the sake of argument it is useful to define four corners in this space. We argue that the above-mentioned four examples of developmental behavior correspond to the four corners of this space. Babbling and motor babbling fall at the zero-point. These behaviors are neither social nor intentional. They are used for exploring the possible space of movements and articulations and possibly for finding a mapping between visual perception, proprioception and action. Imitation of gestures is social in that it requires at least two agents and in that it requires the imitating agent to map the other agent's actions onto its own. However, by definition, imitation of gestures is not concerned with the intent of the other agent's actions. Agents only mime each other, this is why this type of imitation is also called mimesis [3]. Exploration of the world, on the other hand is intentional, in that the agent doing the exploration wants to figure out how its actions change the state of the world. As the agent is perfectly able to do this on its own, such behavior is not social. During the exploration phase in which affordances of objects are learned, the environment is of crucial importance. It provides situational constraints to which every action is bound [4]. The agent now does not only learn the properties of its own effectors, it learns the effect of its actions on the environment. In doing this, it acquires a representation of the physical properties of the environment. Finally, the combination of intentional and social behavior can result in imitation of object use, where by definition the agent's aim is to copy the other agent's intentions. Other social and intentional behaviors, such as constructing a theory of mind or trying to manipulate another agent are also possible. The success of an imitative attempt can be measured by observing the behavior of the robots. Measuring the success in constructing a theory of mind for instance requires inspection of the agents' internal memories. Imitation on robots was therefore heavily investigated, although no systems capable of intentional imitation were built.

3 Embodiment and behavioral development

It is perhaps obvious that a developmental component is crucial for embodied systems; what is perhaps less obvious is that taking into account embodiment is also crucial when studying development.

The importance of development for embodied systems is familiar to all researchers in robotics. It is almost always impossible to program a robot (a prototypical example of an embodied computer system) exactly to perform a certain behavior. It turns out robots perform much better when they are able to update their behaviors to at least some extent. A capability to development is therefore a crucial aspect of an embodied system. However, as we will argue below, it turns out that when studying development it is also crucial to take into account the embodiment of the developing agent. Examples of the kinds of development we are referring to are the four different kinds of behavior mentioned above. Studying such behaviors without taking into account that they take place in an agent that has specific sensors and actuators and that has to operate in a specific environment becomes such an abstract exercise that it is almost meaningless. The kinds of behaviors an agent can perform, the events it can detect and the ease with which it can manipulate its environment, all depend on the embodiment of the agent. Therefore the mechanisms necessary for development as well as the developmental trajectories an agent can follow depend on the embodiment. But from a more practical point of view, embodiment is important when studying development on robotic agents. If one wants to train robots through imitation, for example, it is not a priori clear whether the particular embodiment selected is able to imitate a pre-defined set of behaviors. When studying imitation in robotic agents, it is therefore perhaps better to have the agents develop a repertoire of basic behaviors beforehand, for example through imitation games [5–7]. Such a repertoire of basic behaviors could then in principle be used to construct more intentional compound behaviors. In any case, the embodiment selected for an agent limits the behaviors it can acquire and the ways in which it can manipulate its environment, and therefore has a strong influence on the agent's development.

4 Experiments

For the four examples of developmental learning concrete computer simulations of populations of embodied agents will be proposed. Concrete results are shown for motor babbling experiments and experiments on the imitation of gestures.

Fig. 2. The robotarm and stereovision system simulated in the experiments.

Motor Babbling An individual agent learns a model of the forward and inverse kinematics of its effector by motor babbling. The agent executes random motor commands and observes how its effector behaves. We've experimented with Locally Weighted Learning [8], combined with selection and organization of stored points. In figure 1 on the left, the average error in the model of the inverse kinematics and the number of points stored are shown. It can clearly be seen that with very few samples good predictions can be made.

Imitation game We have investigated how agents can construct a shared repertoire of actions. We propose a multi agent system in which the interactions consist of imitation games. In an imitation game two agents are randomly selected from the population. Both agents are randomly assigned the roles of *initiator* or *imitator*. A single imitation game causes the repertoires of both agents to become more similar. By repeating the game over multiple combinations of randomly selected agents, the repertoires of all agents become shared. In each game, the imitator tries to imitate the action the imitator executed. The initiator decides on the success of the game and sends binary feedback to the imitator, such that the imitator can adapt its repertoire.

The imitation game was first proposed by de Boer [5] in the context of vowel systems. He showed how a shared and human-like vowel system can emerge in a population of agents. In later experiments, the imitation game was used to show how a population of robotic agents develops a shared repertoire of actions. In this experiment, agents are equipped with a (simulated) 6-DOF robot arm with a gripper and a (simulated) stereo vision system, see figure 2. The actions performed by the agents are simple gestures with the robot arm. In figure 2 on the right results are shown for a population of 2 agents

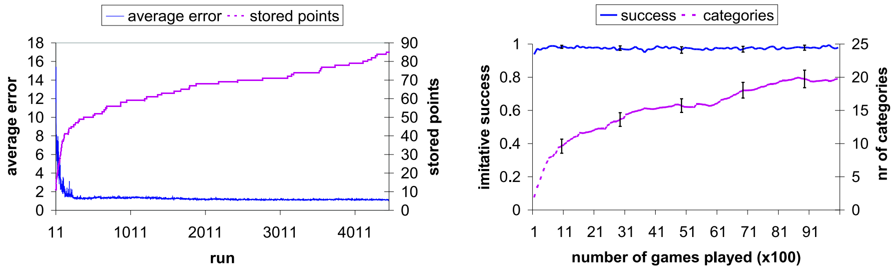

Fig. 3. On the left: the average error in the learnt model of the inverse kinematics and the number of associations stored. On the right: imitative success and number of categories for 10000 imitation games played by two agents.

playing 10000 imitation games. While the imitative success remains high, the agents succeed in constructing a shared repertoire of actions. Detailed results are available in [6, 7].

Individual exploration

We plan to set up a concrete experiment in which a robotic agent learns how its movements can cause changes in its environment. Starting from the assumption that the agent already has a model of the kinematics of its actuator, we want to investigate how the agent can learn to modify its world. In future experiments, agents will explore their environment by performing random actions. Agents will assess the effect their actions have on the environment. Actions are rewarded proportionally to the observed external effect, such that agents learn to perform actions that require the least movement to cause a maximal effect.

Intentional-social exploration

In individual exploration, agents learnt that some of their actions cause changes in the environment and learn to prefer actions that have a maximal effect. This new skill can bootstrap the difficult task of intentional imitation. As the initiator performs a sequence of actions in order to obtain a certain goal, the imitator can deduce the first agent's goal, by comparing the observed effect of the initiator's actions with the learnt associations between own actions and environmental effects. In this paradigm of goal-level imitation, both agents will thus learn to pursue the same goal. However, they might be performing different actions, because actions were learnt on an individual basis and because agents can have different embodiments and thus require other actions for accomplishing the same goal.

5 Discussion

When studying the developmental path towards intentional behavior, striking parallels are noted between motor behavior and vocal communication. Both start out with an initial "babbling" phase in which the articulatory devices—effectors or vocal tract—are explored and its inverse-kinematics are learned. Following this phase, agents explore the interaction with the environment. In motor behavior this is expressed as an exploration of the physical properties of objects and their behavior when being manipulated. In vocal communication this is less explicit, but it could exhibit itself for example in the fascination of children with echoing sounds. At the same time, through gestural imitation, the agent can achieve a representation which is influenced by observing other agents. In motor behavior, this is the phase where agents mimic the actions without picking up the purpose of the action. While in vocal communication, the agents acquire a repertoire of sounds from their peers yet without learning the meaning associated with each sound. The last phase is marked by learning to use behavior intentionally. Objects are manipulated with a clear goal in mind, and vocalizations are uttered with goal-like purpose. This marks the phase where actions are connected with meaning.

Acknowledgments

Bart Jansen is sponsored by a grant from the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT). Tony Belpaeme is a postdoctoral fellow of the Fund for Scientific Research - Flanders (Belgium) (FWO Vlaanderen). Experiments on motor babbling were performed by Steven Teerlinck.

References

1. Wooldridge, M., Jennings, N.: Intelligent agents: Theory and practice. Knowledge Engineering Review **2** (1995)
2. Meltzoff, A.M.: Elements of a developmental theory of imitation. In:. Cambridge University Press, Cambridge (2002)
3. Butterworth, G.: Neonatal imitation: existence, mechanisms and motives. In Nadel, J., Butterworth, G., eds.: Imitation in Infancy. Cambridge University Press, Cambridge, UK (1999)
4. Vygostky, L.S.: Play and its role in the mental development of the child. Soviet Psychology **5** (1967) 6–18
5. de Boer, B.: Self organization in vowel systems. Journal of Phonetics **28** (2000) 441–465
6. Jansen, B., de Vylder, B., de Boer, B., Belpaeme, T.: Emerging shared action categories in robotic agents through imitation. In Dautenhahn, K., Nehaniv., C.L., eds.: Proceedings of the Second International Symposium on Imitation in Animals and Artifacts., University of Wales, Aberystwyth (2003) 145–152
7. Jansen, B.: An imitation game for emerging action categories. In Banzhaf, W.e.a., ed.: Proceedings of the 7th European Conference on Artificial Life, Lecture Notes in Artificial Intelligence., Berlin, Springer (2003) 800–809
8. Atkeson, C.G., Moore, A.W., Schaal, S.: Locally weighted learning. Artificial Intelligence Review **11** (1997) 11–73