Imitation in embodied agents results in self-organization of behavior

Bart Jansen [a]   Tom ten Thij [b]   Tony Belpaeme [a]

Bart De Vylder [a]   Bart de Boer [b]

[a]Artificial Intelligence Lab, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium
[b]Artificial Intelligence, Rijksuniversiteit Groningen, Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands

1   Introduction

Imitation in robotic agents has received considerable attention in the last decade, especially as it is a powerful means to transfer skills between agents in a way that comes natural to humans. Most robotic studies emphasize on constructing artifacts that implement a particular form of learning by demonstration. Typically, in this form of learning a pre-defined repertoire of behaviors is transferred from one agent (a teacher) to another (the student). The repertoires of behaviors and the roles of the imitating agents are fixed. Hence the focus of that kind of work lies on machine learning methodologies for imitation. Our work, on the other hand, takes a different point of view. Instead of constructing an imitation application, we study imitation as an adaptive and flexible way to acquire new behavior using a robotic multi-agent system.

The experiments have been performed on a robotic arm that uses four degrees of freedom and a stereoscopic camera. Using the arm and camera as embodiment, agents engage in games where they try to imitate the action they saw the other agent perform.

Actions are not coupled to meaning, so it is not a model of the cultural learning of real gestures. The only feedback between agents during the experiment is given by the initiator of the game as one bit of information on whether he judges the game's outcome as successful. It turns out that a population using a scheme as simple as described in this paper is able to develop a shared repertoire of action categories in a very robust manner.

We are particularly interested in the constraints posed by the embodiment of the agents. Their perception and their physical implementation poses limits to what they can observe and what actions they can perform. This principally influences the behavior that can be successfully imitated, and determines the collective behavior that will be sustained in a population of imitating agents.

2   The imitation game

The imitation between agents is implemented as an *imitation game*. The imitation game was introduced in the context of research into speech. In this context it is logical to use imitation as a learning mechanism, as children learn to speak through imitation, but imitation was also used as a means to put pressure on agents to develop sounds that are as distinctive as possible. Thus, the aim of the imitation game was to develop a repertoire that contains as large a number of clearly distinctive speech sounds as possible. At the same time, these sounds had to be distinguishable and learnable under the physical and cognitive constraints of production, perception and learning. It turned out that this imitation game is also very useful when developing a repertoire of actions.

The generic imitation game (whether for speech or actions) is based on a population of agents that each have an open repertoire of categories. In principle, imitation games can be played in a population of just two agents, but the population size can be much larger. Crucial is that the agents in the population do not have fixed roles, nor do they start with pre-defined categories in their repertoires. They start out empty and no agent can therefore be the teacher from which the other agents learn their categories. A shared repertoire emerges through repeated interactions between the agents.

In any imitation game, two agents are randomly selected from the population. One agent is randomly assigned the role of *initiator*; the other gets the role of *imitator*. To start the imitation game, the initiator selects one random category from its repertoire (unless its repertoire is empty, in this case it starts with creating a random category). It then expresses this category (by uttering speech, or in this paper, by making a gesture with a robot arm). The imitator perceives this expression and finds the nearest category.

The imitator then expresses the category it found. It is perceived by the initiator, who, like the imitator, finds its closest category. If this closest category is the same as the category it initially expressed, the game is successful. Apparently both agents' categories are equally similar and sufficiently distinguishable such that twice-repeated categorical perception of imperfect expressions does not change the category. However, when the category the initiator perceives is not the same as the category it initially selected, the game is a failure. Apparently, the agents' categories are not sufficiently similar, or perhaps they are not sufficiently distinguishable. In each case, the initiator gives feedback to the imitator. This feedback consists of a single bit of information saying whether the game succeeded or failed. Now both agents can update their repertoires in order to improve expected future performance at the imitation game. In the implementation game as implemented here, the imitator makes the most important changes.

3   Results

We have implemented the imitation game as described above on a real robot arm and with a real vision system. In order to facilitate image processing and stereo vision, a coloured ball us attached to the gripper of the arm. Currently only the position of this ball is focussed on when observing an action. Particularities of the vision system are described in more detail in the original paper.

Using this experimental setup, we investigate how a population of agents succeeds in constructing a shared repertoire of action categories, while high imitative success is maintained. A set of qualitative measures is introduced to evaluate the success of the imitative attempts under various conditions.