Imitation in embodied agents results in self-organization of behavior

Bart Jansen[1] **Tom ten Thij**[2] **Tony Belpaeme**[1] **Bart De Vylder**[1]
Bart de Boer[2]
[1]Artificial Intelligence Lab
Vrije Universiteit Brussel
Pleinlaan 2, B-1050 Brussels, Belgium
`http://arti.vub.ac.be`

[2]Artificial Intelligence
Rijksuniversiteit Groningen
Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands

Abstract

This paper addresses the construction and learning of behavior through imitation. We explore imitation as a framework to create and sustain behavior in a population of agents. Two experiments are presented, in simulation and on actual robots. These demonstrate how the behavior self-organizes through the interplay between the imitative interactions, the environment and the embodiment of the agents.

1. Introduction

Imitation in robotic agents has received considerable attention in the last decade (e.g. Kuniyoshi et al., 1994; Gaussier et al., 1998; Schaal, 1999; Alissandrakis et al., 2001; Billard and Matarić, 2001; Breazeal and Scassellati, 2002; Demiris, 2002). Imitation manifests itself in its different guises throughout the natural world (for a review see Whiten and Ham, 1992). In particular higher-order primates and humans seem to benefit profoundly from imitative learning (Nadel and Butterworth, 1999). Because of this, imitation receives quite some attention from robotics researchers (Dautenhahn and Nehaniv, 2002), especially as imitation 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 this 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.

Through imitation a learner acquires skills and behaviors from peers or parents by observing and mimicking those skills and behaviors. While imitating, the imitator needs to observe behavior and map it onto its motor system. This is called the correspondence problem. What is remarkable is that the bodies of the parties involved in imitation can be largely dissimilar. Children, for example, have smaller bodies than the parents that they imitate; still this poses little problem. Meltzoff and Moore (1977) for example report how neonates can imitate adult facial expressions, thereby showing that they have the innate capability to solve the correspondence problem for dissimilar bodies[1]. In the experiment reported here, we likewise endow our agents with the capacity to solve the correspondence problem: all agents are capable of mapping observed actions on their own motor system.

We are particularly interested in the constraints posed by the embodiment of the agents. Imitation cannot be studied without considering 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. This is well illustrated by linguistic communication; imitation forms one of the cornerstones of the acquisition of language and is, for example, crucial for the learning of speech sounds (Kuhl

[1]Incidentally, this might be related to mirror neurons (Rizzolatti et al., 1996). These neurons have been found in monkeys and are active both when performing a certain action and when observing that action.

and Meltzoff, 1996). Speech is learned by imitating others, but the nature of speech sounds depends largely on the perception and production of speech, i.e. on the embodiment of our auditory and acoustic apparatus.

2. The imitation game

The imitation between agents is implemented as an *imitation game*. The imitation game[2] was introduced in the context of research into speech (de Boer, 2000). 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. In reality such pressure would be caused by the function of the sounds for distinguishing meaning: as different speech sounds are associated with different meanings, the need to differentiate between two different meanings would push apart these speech sounds. However, implementing these dynamics would have made the simulations needlessly complex. 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 turns out that the imitation game is also very useful when developing a repertoire of actions. As will be explained below, some of the peculiarities of the imitation game can be better understood when one takes its background in the learning of speech sounds into account.

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, meaningful imitation games can be played in a population of just two agents, but the population size can be much larger. In a larger population, convergence towards a shared repertoire of categories will be slower, but it will happen nevertheless. The basic dynamics of the game, however, can be demonstrated in a population of just two agents.

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. Although in many cases it is possible to identify teachers and students in learning of real language, there are cases, especially in the formation of pidgin- and creole languages, where such teacher-student roles cannot be identified, and a shared repertoire of speech sounds still emerges.

[2]Of course, Alan Turing in 1950 already used the term *imitation game* in the context of artificial intelligence for the experiment now commonly known as the Turing Test. This is not related to our use of *imitation game*.

Each agent develops a repertoire of categories. It is important that agents have true categorical perception (Harnad, 1987). This means that when they perceive an input from a continuous domain (either a speech sound or an action) they perceive it as one of the discrete categories in their repertoire, never as something in between. This is certainly the case in speech sounds (e.g. Cooper et al., 1976) and also appears to work well when working with actions.

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*. The initiator is the agent who initiates the imitation game. The imitator imitates and the initiator evaluates how well the imitator managed to imitate. It is important to notice that the only things the agents can observe are the expressions produced by either the other agents or themselves. This means that there is no such thing as telepathy; the agents do not have direct access to the other agents' categories.

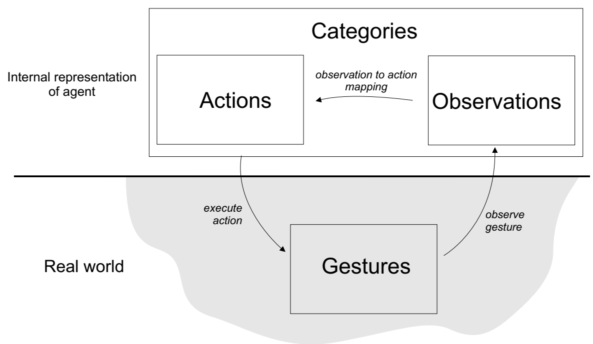

Figure 1: Agents have categories, which are made up of actions and observations. Actions can be expressed as gestures in the real world, and observing these gestures results in observations. And observation in turn can be mapped onto categories using inverse kinematics.

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). Usually, the expression will not be perfect, such that when an agent expresses the exact same category twice, the results will not be identical. The imitator perceives this expression and finds the nearest category. This effectively implements categorical perception: each expression is always perceived as a category; details of the expression are lost in categorization (nevertheless, a copy of the exact expression is retained for the duration of one game, as it might be needed to improve the categories). The imitator then expresses the category it found. This expression, too, is imperfect. 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.

Table 1 and 2 provide pseudo code for the imitation game, while figure 1 illustrates the relation between the internal representation of an agent and its actions in the world. Agents have a repertoire C of categories, each category c has an action *c.action* and an observation *c.observation* associated with it. Actions are a tuple containing start and end configuration of the robot arm[3], while observations are a series of 3D-coordinates.

initiator	imitator
if $C = \emptyset$ **new-category**(C) $c \leftarrow$ random from C $c.use \leftarrow c.use + 1$ **execute** *c.action*	

	observe O_1
	if $C = \emptyset$
	  new-category(C)
	else
	  $c_{rec} \leftarrow$ category from C so
	    $c_{rec}.observation$ closest
	    to O_1
	  execute $c_{rec}.action$.

observe O_2
$c'_{rec} \leftarrow$ category from C so
  $c'_{rec}.observation$ closest
  to O_2
if $c = c'_{rec}$
  $c.success \leftarrow c.success + 1$
  send feedback "success"
else
  send feedback "failure"	

	update($c_{rec}, O_1, feedback, C$)
do-other-updates()	**do-other-updates**()

Table 1: Pseudo code of the imitation game

3. Monitoring the performance

It is important to define good measures for monitoring the performance of the agents. Four measures were defined: the imitative success, the number of categories, the average category variance and the information flow.

3.1 Imitative success

The *imitative success* is simply the average number of successful imitation games. Imitative success is calculated as a running average over the last $n = 100$ games. Due to the nature of the game, the imitative success is already high at the start of the game, when agents have only few –not even very similar– categories. This is best explained by a small example: if the initiator has only one category, and it expresses that category as an action, anything that the imitator does is matched to that category. So the initiator judges the imitation to be successful, resulting in a high success score. This is why at the beginning of the experiment, the imitative success is already quite high. What is interesting, is looking if the imitative success stays high as agents continue to create and learn new categories.

[3]Nothing prevents the action from being a series of motor commands, instead of just the start and end configuration of the motors. In the experiments reported here, for practical considerations, actions are only a start and end position of the arm.

3.2 Number of categories

It is informative to monitor the number of categories that agents possess, as this tells us something about the noise in the environment. The number of categories that agents are able to acquire depends on what they can perceive and express. Hence, the embodiment of the agents is the limiting factor on acquiring categories.

3.3 Average category variance

Consistently successful imitation games are not possible without shared categories, i.e. categories that are similar between agents. However, it is useful to define an explicit measure of similarity of the repertoires of all agents. First we introduce a method for comparing the repertoires of two agents. *Category Variance (CV)* is the *scaled sum of minimum distances* applied to the categories of both agents and is given in equation 1. A and B are agents, a and b are categories, $a \in A$ means that a is an category in the repertoire of agent A and $|A|$ denotes the number of categories of agent A. The distance function $d(a, b)$ on categories of agents used here is Dynamic Time Warping on the observations associated to the categories.

update($c_{rec}, O, feedback, C$)
$c_{rec}.use \leftarrow c_{rec}.use + 1$
if feedback = "success"
  $c_{rec}.success \leftarrow c_{rec}.success + 1$
  shift-closer(c_{rec}, O)
else
  if $c_{rec}.success/c_{rec}.use >$ *treshold*
    $C \leftarrow C \cup$ **construct-category**(O)
  else
    shift-closer(c_{rec}, O)

shift-closer(c, O) $c.action \leftarrow c.action+$ *shift* .(**find-action**(O) $- c.action$) $c.observation \leftarrow$ **execute** ($c.action$)
find-action(O) find action that produces O using inverse kinematics
new-category(C) $c.action \leftarrow$ random $((x_1, y_1, z_1), (x_2, y_2, z_2))$ $c.observation \leftarrow$ **execute** ($c.action$) $C \leftarrow C \cup c$
construct-category(O) $c.action \leftarrow$ **find-action**(O) $c.observation \leftarrow$ **execute** ($c.action$)

do-other-updates(C)
$\forall c \in C$ **do**
  if $c.success/c.use <$ *throwawaythreshold* **and**
      $c.use >$ *minimumuses*
    $C \leftarrow C \setminus c$
with probability *addprobability* **do new-category**(C)
merge-categories(C)

merge-categories(C)
for all $c_1, c_2 \in C,\ c_1 \neq c_2$
if $\lVert c_1 - c_2 \rVert <$ *merge-threshold*
  $c_{new} \leftarrow (c_1 + c_2)/2$
  $c_{new}.success \leftarrow c_1.success + c_2.success$
  $c_{new}.use \leftarrow c_1.use + c_2.use$
  $C \leftarrow C \setminus c_1 \setminus c_2$
  $C \leftarrow C \cup c_{new}$

Table 2: Important procedures used in the imitation game

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

The category variance allows us to calculate the *Average Category Variance* ($\overline{CV}$), which is given in equation 1, of the population $P = \{A_i\}_{i=1\ldots N}$ consisting of N agents, indicating how similar the categories of all agents are on average.

$$\overline{CV}(P) = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} CV(A_i, A_j) \quad (2)$$

3.4 Information flow

When taking an information theoretical perspective on the imitation game, one could view the initiator as the source or sender of information, the physical system consisting of the robot arm and the stereo head as the noisy channel and the imitator as the destination or receiver. The source can send different symbols, corresponding to the different action categories of the initiator. The communication can succeed or fail, indicated by the success or failure of the imitation game: we assume that a symbol was correctly transmitted if its repetition by the receiver over a similar channel yields the original symbol according to the sender. If the source emits symbols randomly and independently chosen according to a fixed distribution, the information content (in bits) or entropy H of a symbol depends on its frequency p and is given by:

$$H = \log_2(\frac{1}{p}). \quad (3)$$

The initiator selects actions with equal chance out of his repertoire, which makes the frequency of each action inverse proportional to the number of categories. This results in the following measure of information flow:

$$I = \begin{cases} \log_2(|A|) & \text{if the game succeeds} \\ 0 & \text{if the game fails} \end{cases} \quad (4)$$

with $|A|$ the number of categories of the initiator. If $|A|$ is a power of 2, I gives exactly the number of bits necessary to distinguish between these categories.

The creation of a new category by an agent has a twofold effect on the information flow. As this new category needs time to propagate through the population and get adopted by other agents, in the short term the imitative success and hence the information flow will decrease. In the long term, however, if the imitative success is able to recover, the information flow will increase as the number of categories increases.

4. Experimental setup

The imitation experiments, both in simulation and on physical systems, are implemented as embodied agents, consisting of a robot arm and a stereo head. At the moment one such setup is present at our laboratory and different agents use the same physical system during an experiment (figure 2).

As an agent must be able to observe other agents' actions as well as its own, the stereo head is shared between two agents playing an imitation game. This causes all agents to have the same point of view on an action, in contrast with the situation where two agents sit opposite each other and thus have different points of view on the action. Although this circumvents a number of interesting issues, such as the recognition of self and other, and the use of deictic coordinates, it presents a necessary simplification to test the principles of the imitation game. Another simplification that is used, is that the actions an agent can perform are limited to motion trajectories of the robot arm from one point to another.

Within the limitations of this simplified setup, similar results have been achieved with simulations and the real setup, as will be shown in section 5. These show that both setups are viable and that results achieved in simulation are valid for the robotic setup.

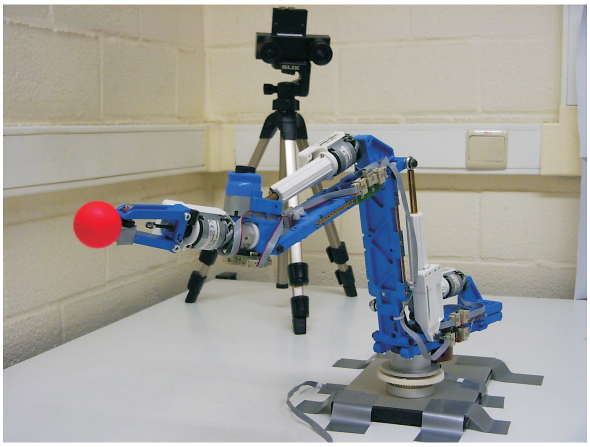

Figure 2: The setup consists of a 6 DOF arm and a stereo camera for respectively performing and perceiving actions.

The robot arm, stereo head and stereo processing software are commercially available. The arm we use is the teach-robot[4]. It has six degrees of freedom and is equipped with a gripper. The arm is position-driven; one can send goal motor positions to it but not change the motors' speed nor intervene during the execution of a command. The forward and inverse kinematics are known (De Vylder, 2002). In order to facilitate image processing and stereo vision, a colored ball is attached to the gripper of the arm. Currently only the position of this ball is focused on when observing an action. The vision system contains a MEGA-D stereo head, which delivers a series of simultaneously taken left and right images. The Small Vision System (SVS) (Konolige, 1997) is used for camera calibration and stereo vision. When an agent observes an action (both its own and that of another agent), it tracks the ball and builds a series of three-dimensional coordinates. We will describe this process in some detail.

For each pair of image frames, the left image is searched for possible ball segments using color templates, the right image is left unprocessed. A simple tracking mechanism then selects the most probable ball segment. Of all the pixels of this segment, the 3D-coordinates are calculated using SVS. Finally, the current ball position is estimated using expectation maximization (Bilmes, 1998). If the SVS calculation fails or is not reliable, the right image is searched for the most probable ball segment as well. The center of the ball segments in the right and left image are then assumed to correspond, from which we can estimate the ball's 3D-coordinates. This process runs at 10 frames per second.

The observation of an action by an agent results in a time series of 3D points. The interval between successive points is determined by the frame rate of the vision system and is independent of the duration of an action. As this duration varies over different actions, the length of such time-series varies as well. We assume that an agent knows when to observe its own actions and it is informed about the starting and stopping of actions of others.

Despite the fact that agents all use the same physical setup, it is important to notice that action execution and observation are subject to considerable amounts of noise. Among other disturbing factors are the inaccurate robot control (up to 2 cm inaccuracy in some workspace regions) and the varying illumination conditions. This noise poses limitations on the number and kinds of actions the agents can discriminate, but is inherent in working with embodied agents acting in the real world.

5. Results

5.1 Simulation results

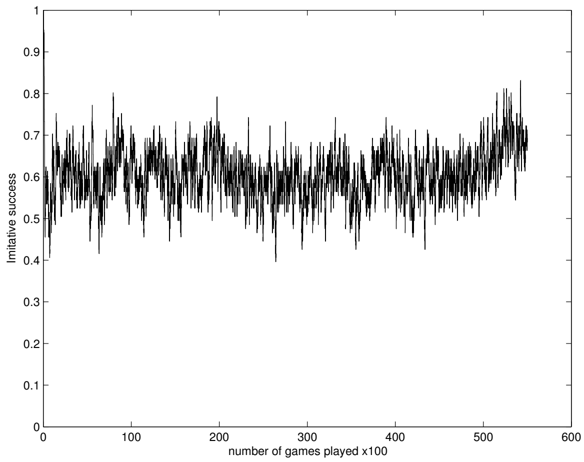

Figure 3: Average imitative success for the simulation.

Before performing experiments on real robots, we have implemented the imitation game in simulation, to verify the functionality of the imitation game. In the simulations a model of the kinematics of our physical robot was used to calculate simulated observations for actions. To approximate uncertainties found in the real system, noise is added to articulations of the robot arm and to

[4]Microelectronic Kalms

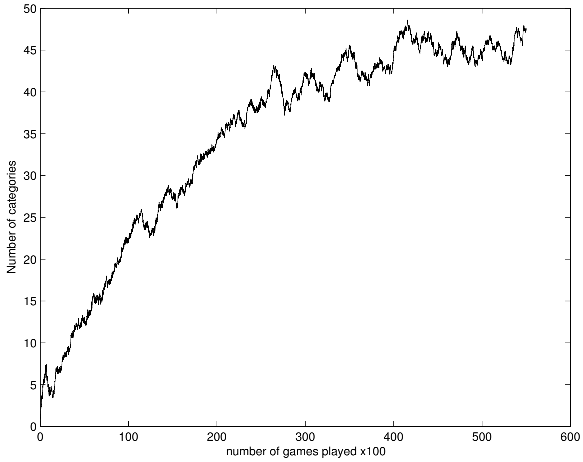

Figure 4: Average number of categories for the simulation

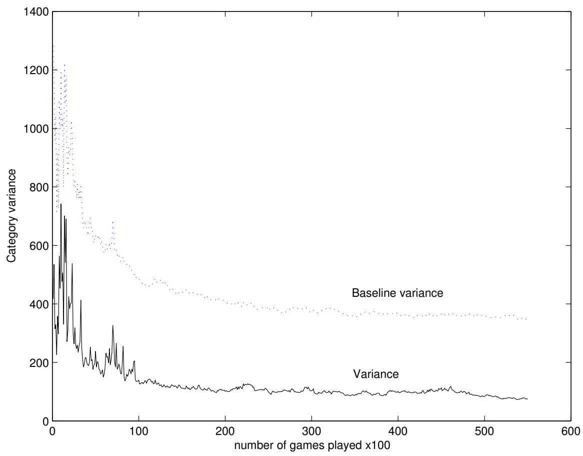

Figure 5: Average category variance for the simulation.

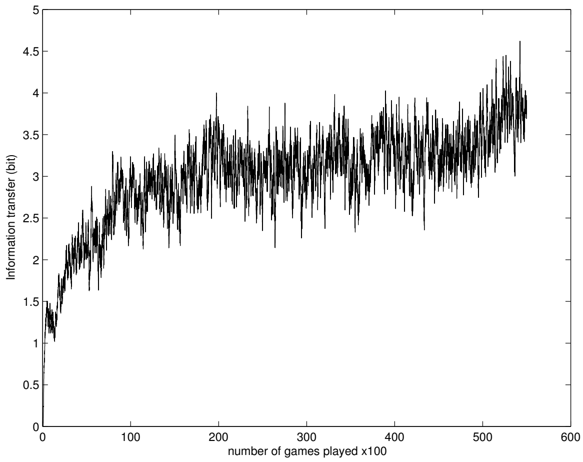

Figure 6: Information transfer for the simulation.

the perception. The model consists of a basic geometrical model of the robot arm, without simulating electrical and mechanical properties of the arm. The model takes as input 6 joint angles, and returns the position of all joints and of the end-effector (De Vylder, 2002). Noise is added to the simulation by adding a normal distributed term to these positions.

The simulation has 10 agents, and runs 55000 imitation games. Figure 3 shows the evolution of imitative success over all 55000 games, note how the success stays stable. The agents perform better than baseline imitative success of 0.53, which was determined by letting the agents first create random categories and then play imitation games. Figure 4 shows the average number of categories in the population. Figure 5 shows the average category variance $\overline{CV}$, which tells us how similar categories are between agents, $\overline{CV}$ is markedly better than random performance. Figure 6 shows how the information transfer increases as the agents continue playing imitation games, at 55000 games the agents have learned enough categories to transfer up to 4 bits of information.

These results illustrate that the agents succeed in developing a shared repertoire of categories. Similar results were obtained in larger populations and in simulations with higher amounts of noise added to the simulated observations, see (Jansen et al., 2003; Jansen, 2003). The work described here is distinguished from previous work by the fact that the imitation game now performs –and performs well– on a real-world system. This shows that the game implements dynamics that converge reliably under various conditions, both in simulation and on actual systems.

5.2 Robot experiment results

As experiments on the actual robotic setup run slower, the experiment uses only two agents. 13000 imitation games are played to obtain the results shown here.

Figure 7 through 10 show the imitative success, the number of categories, the average category variance and the information transfer. The success is again stable and slightly higher than in the simulation (due to the fact that only two agents are learning from each other, which is easier than having 10 agents learn from each other). The number of categories has not yet started to settle, meaning that both agents have not yet reached their limit of distinguishing and producing actions. The average category variance is considerably lower than baseline performance, showing that the categories resemble each other quite well. Figure 11 again illustrates this by plotting all actions of both agents at game 13000. And finally the information transfer is close to 6 bits, showing that both agents have acquired a coherent enough repertoire of categories to express 2^6 successfully distinguishable actions.

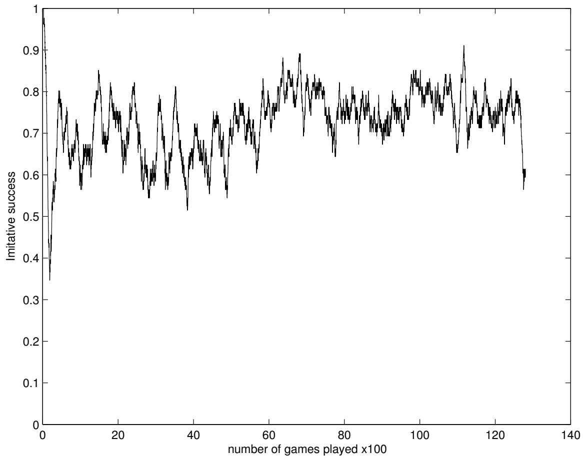

Figure 7: Average imitative success for the robotic setup.

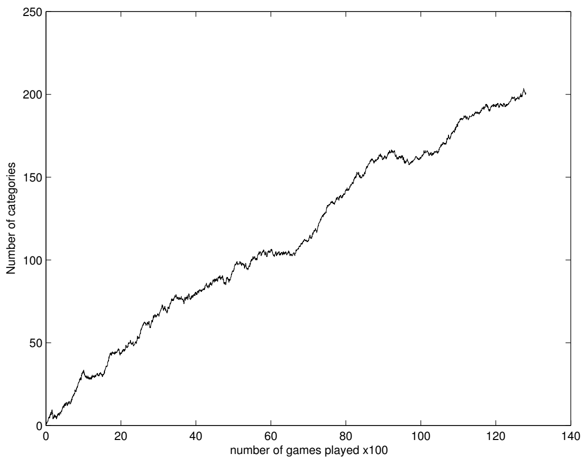

Figure 8: Average number of categories for the robotic setup

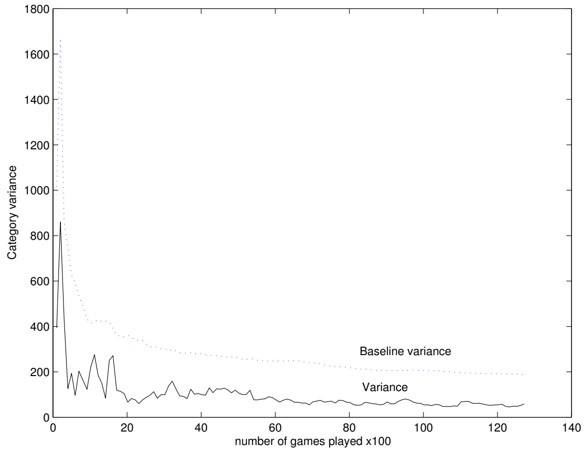

Figure 9: Average category variance for the robotic setup.

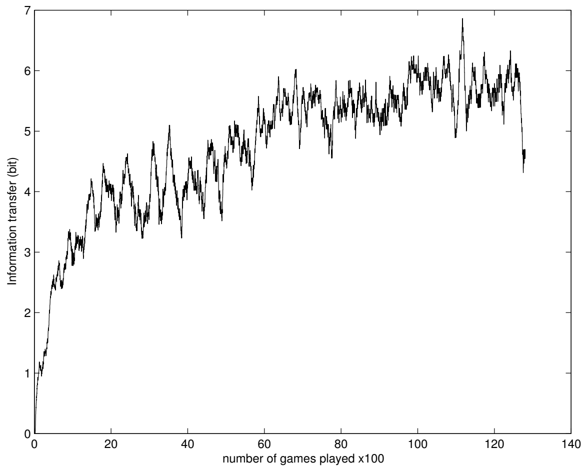

Figure 10: Information transfer for the robotic setup.

6. Conclusion

We demonstrated how actions can emerge and be sustained in a population of agents through simple imitative interactions. These interactions, called imitation games, are basic turn taking interactions between two agents, where one agent demonstrates an action and the other imitates. Successful imitation reinforces actions, while unsuccessful actions will disappear over time from the agent's repertoire. The imitation between the agents forms a self-organizing system. As with any self-organizing system it has two components: random fluctuation, implemented by the agents creating random actions, and selective feedback, implemented by the successfulness of imitations having an influence on which actions are retained in the agent's repertoire. This feedback provides pressure on the actions to be imitable, and is the driving factor for attaining a coherent and shared action repertoire in the population.

This is related to the theories of Steels (Belpaeme et al., 1998; Steels, 2000), who views human language and cognition as a self-organising system. In this view, human conceptualisation and language are the result of social and linguistic interactions between individuals in a population. Language and conceptualisation are thus open-ended systems in constant flux and evolve both synchronically and diachronically under influence of its users.

The take home message of this paper is that imitation serves to tune internal representations (in this case the internal representations are action categories, but these could also be speech sounds or sign language) of agents without the need to explicitly transfer these representations between agents. The agents' internal representations become shared, without resorting to some kind of "telepathic" transfer of representations. Moreover, the representations become shared in a robust manner, as imitation weeds out representations that are poorly observable or reproducible. These representations are also

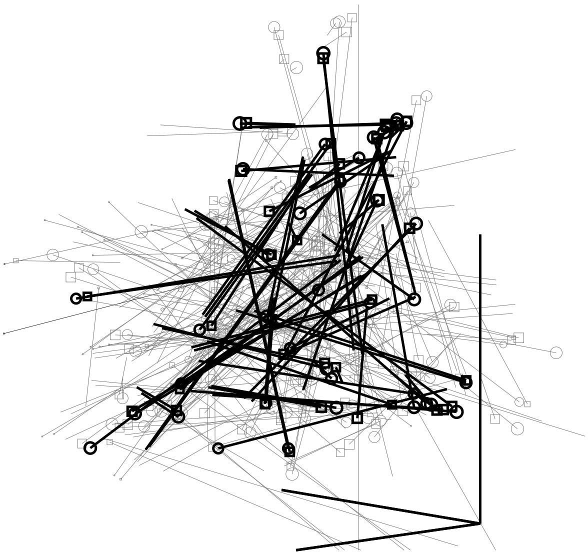

Figure 11: The actions of two agents plotted in 3D space, to demonstrate how actions resemble each other. Squares and circles denote the starting point of the action for agent 1 and 2 respectively. The most successful actions have been plotted in black, less successful actions in grey.

adaptive, as the dynamics of the imitation game serve to discard any representations that are unsuccessful and gracefully accepts changes to the parameters, embodiment or the environment.

In the experiments reported here, the agents acquire repertoires of action categories that are not situated. The actions do not serve to communicate, as they for example would in sign language. Nor do they serve to manipulate objects or perform tasks, such as stacking blocks. A logical extension to the imitation game described in this paper, would be to make the actions that agent learn "useful". For this however, the agents would need intentionality. They would need to understand what it is that the other is trying to demonstrate. Is the other agent just waving its arm? Or is the action aimed at pushing a block across the table?

To build agents that actually imitate the use of an action, we would need to construct a basic form of Theory of Mind for agents. This would enable agents to deduce the intention behind an action, and would lead agents to imitate the intent of an action, instead of the action itself. We hope to learn from developmental psychology how intentionality can bootstrap itself, without being programmed explicitly into the agents. One route, which might be promising, looks at how infants learn to coordinate their body movements; new born babies learn to coordinate their limbs and explore the impact of their actions on the environment through play (Vygostky, 1967). Perhaps creating a challenging and stimulating environment for the agent and letting them explore the impact of their actions on the environment, might bootstrap the self-organization of intentional behavior, which could then later be used to acquire new skills through imitation.

Acknowledgements

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 Flemish Fund for Scientific Research (F.W.O. Vlaanderen).

References

Alissandrakis, A., Nehaniv, C. L., and Dautenhahn, K. (2001). Through the looking-glass with ALICE – trying to imitate using correspondences. In Balkenius, C., Zlatev, J., Kozima, H., Dautenhahn, K., and Breazeal, C., (Eds.), *Proc. First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, volume 85. Lund University Cognitive Studies. September 17-18, 2001, Lund, Sweden.

Belpaeme, T., Steels, L., and van Looveren, J. (1998). The construction and acquisition of visual categories. In Birk, A. and Demiris, Y., (Eds.), *Proceedings of the 6th European Workshop on Learning Robots*, Lecture Notes on Artificial Intelligence, Berlin. Springer.

Billard, A. and Matarić, M. J. (2001). Learning human arm movements by imitation: Evaluation of a biologically inspired connectionist architecture. *Robotics and Autonomous Systems*, 37(2-3):145–160.

Bilmes, J. A. (1998). A gentle tutorial of the EM algorithm and its application to parameter estimation for gaussian mixture and Hidden Markov Models. Technical report, U.C. Berkeley technical report TR-97-021.

Breazeal, C. and Scassellati, B. (2002). Challenges in building robots that imitate people. In Dautenhahn, K. and Nehaniv, C. L., (Eds.), *Imitation in Animals and Artifacts*, pages 363–390. The MIT Press, Cambridge, MA.

Cooper, F. S., Delattre, P. C., Liberman, A. M., Borst, J. M., and Gerstman, L. J. (1976). Some experiments on the perception of synthetic speech sounds. In Fry, D. B., (Ed.), *Acoustic Phonetics*, pages 258–283. Cambridge University Press, Cambridge, UK.

Dautenhahn, K. and Nehaniv, C. L., (Eds.) (2002). *Imitation in Animals and Artifacts*. The MIT Press, Cambridge, MA.

de Boer, B. (2000). Self-organisation in vowel systems. *Journal of Phonetics*, 28(4):441–465.

De Vylder, B. (2002). Forward and inverse kinematics of the Teach-Robot. Technical report, AI-MEMO 02-02, Artificial Intelligence Laboratory – Vrije Universiteit Brussel.

Demiris, Y. (2002). Imitation as a dual-route process featuring predictive and learning components: a biologically-plausible computational model. In Dautenhahn, K. and Nehaniv, C. L., (Eds.), *Imitation in Animals and Artifacts*, pages 327–361. The MIT Press, Cambridge, MA.

Gaussier, P., Moga, S., Quoy, M., and Banquet, J.-P. (1998). From perception-action loops to imitation processes: A bottom-up approach of learning by imitation. *Applied Artificial Intelligence*, 12(7-8):701–727.

Harnad, S., (Ed.) (1987). *Categorical Perception*. Cambridge University Press, Cambridge.

Jansen, B. (2003). An imitation game for emerging action categories. In Banzhaf, W., Christaller, T., Dittrich, P., Kim, J. T., and Ziegler, J., (Eds.), *Advances in Artificial Life, 7th European Conference on Artificial Life (ECAL 2003), Dortmund, Germany*, volume 2801 of *Lecture Notes in Computer Science*, pages 800–809. Springer.

Jansen, B., De Vylder, B., Belpaeme, T., and de Boer, B. (2003). Emerging shared action categories in robotic agents through imitation. In Dautenhahn, K. and Nehaniv, C. L., (Eds.), *Second International Symposium on Imitation in Animals and Artifacts 2003*, pages 145–152. The Society for the Study of Artificial Intelligence and the Simulation of Behavior.

Konolige, K. (1997). Small vision system: Hardware and implementation. In *Proceedings of the Eighth International Symposium on Robotics Research*. http://www.ai.sri.com/konolige/svs/Papers.

Kuhl, P. K. and Meltzoff, A. N. (1996). Infant vocalizations in response to speech: Vocal imitation and developmental change. *The Journal of the Acoustical Society of America*, 100(4):2425–2438.

Kuniyoshi, Y., Inaba, M., and Inoue, H. (1994). Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Transactions on Robotics and Automation*, 10(6):799–822.

Meltzoff, A. and Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, 198:75–78.

Nadel, J. and Butterworth, G., (Eds.) (1999). *Imitation in infancy*. Cambridge studies in cognitive and perceptual development. Cambridge University Press, Cambridge, UK.

Rizzolatti, G., Fadiga, L., Gallese, V., and Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3:131–141.

Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(6):233–242.

Steels, L. (2000). Language as a complex adaptive system. In Schoenauer, M., (Ed.), *Proceedings of Parallel Problem Solving from Nature VI*, Lecture Notes in Computer Science. Springer, Berlin, Germany.

Vygostky, L. S. (1967). Play and its role in the mental development of the child. *Soviet Psychology*, 5(3):6–18.

Whiten, A. and Ham, R. (1992). On the nature and evolution of imitation in the animal kingdom: Reappraisal of a century of research. In Slater, P., Rosenblatt, J., Beer, C., and Milinski, M., (Eds.), *Advances in the study of behavior*, volume 21, pages 239–283. Academic Press, San Diego, CA.