The role of population dynamics in imitation

Tony Belpaeme, Bart de Boer, Bart De Vylder, and Bart Jansen
*Artificial Intelligence Lab - Vrije Universiteit Brussel
Pleinlaan 2, B-1050 Brussels, Belgium
{tony,bartb,bartdv,bartj}@arti.vub.ac.be

Abstract

We propose an extension of the study of imitation in artefacts in which imitation takes place at the population-level. Instead of assuming a one-to-one teacher-student relation between agents, agents in our set-up take on no particular role and have to build up a repertoire of actions through imitative interaction with all other agents in the population. This allows us to study the sufficient and necessary conditions for acquiring an action repertoire through imitation.

1 Introduction

Imitation is a potent mechanism to culturally propagate and maintain knowledge and abilities. Not only humans rely heavily on imitation, also animals such as dolphins, some bird species and some primates rely on imitation to acquire gestures and articulations (see *e.g.* Dautenhahn and Nehaniv, 2002). Studies on imitation in cognitive science have concentrated on mimicry, joint attention, the relationship between imitator and imitatee, theory of mind, intentionality, speech (Kuhl and Meltzoff, 1996) and learning affordances of objects and tools (Tomasello, 1999). All these issues have been picked up by artificial intelligence in constructing artefacts which learn from imitation, be it in simulation (*e.g.* Alissandrakis et al., 2002) or on robotic platforms (*e.g.* Kuniyoshi et al., 1994; Billard and Hayes, 1997). However, little attention has been paid to the role of the social medium in which imitation takes place. Often, imitation is only considered between two agents: one acting as teacher and one as student. The teacher is in possession of a full repertoire of gestures, actions or articulations, which the student has to acquire through imitation learning. This is of course a valid approach when one is interested in the paradigm of imitation to program artefacts by demonstration (Bakker and Kuniyoshi, 1996). However, when studying imitation from a broader perspective, including aspects such as the propagation of learned behaviour, it is also necessary to take into account the social setting, or rather the population, in which imitation occurs. This paper attempts to lay bare the issues that are concerned with imitation in agent populations, and presents the theoretical framework of our experiments.

2 Towards imitation in populations

Often imitation in artefacts is studied as a one-way relationship between teacher and student. The teacher, be it a human or an artificial system, possesses a repertoire of behaviour which the student has to acquire (Kuniyoshi et al., 1994; Billard and Matarić, 2001). However, imitation in a teacher-student relationship addresses only a subset of the issues surrounding imitation. Teacher-student imitation is a potent approach for building intelligent systems, and it is true that a lot of imitation occurring in human behaviour (social behaviour, tool use, language etc.) as well as behaviour observed in higher primates (tool use in chimpanzees for example). However, it does not tackle some issues that are fundamental to more basic imitation, such as the emergence of roles, the transition from non-imitative behaviour to imitative behaviour, the construction of a repertoire of behaviours through imitation, the propagation of imitated behaviour in a population, and the stability of imitated behaviour over time. We would like to re-evaluate the issues related to imitation between two agents and extend them to study imitation on a population-level.

Dautenhahn and Nehaniv (2002, chapter 1) summarised the issues for imitating agents in five central questions: *Whom* to imitate? *When* to imitate? *What* to imitate? *How to map* observed behaviour onto an internal representation and eventually onto imitated behaviour? And, *how to evaluate* the quality of the imitation? In the simplest teacher-student model, only the question of how to map observed behaviour onto an internal representation and onto imitated behaviour is addressed. The other questions can be resolved by carefully defining the roles of the agents and the rules of their interactions. When one starts out with a homogeneous population, all the above questions have to be addressed. This is the reason why we work with populations of agents where no roles or target repertoires of behaviours are pre-defined. In order to facilitate the work, some rules of interaction are pre-defined, but one of the aims of the research is to determine which rules have to be defined beforehand, and which can be emergent.

2.1 Agents

In our view, a population consists of agents with each agent having the following properties.

- An agent needs to be able to solve the *correspondence problem*: perceived behaviour should be mapped onto deictic coordinates. For example, lifting the right arm should result in the imitating agent lifting its right, not its left, arm. The correspondence problem might be a challenge when the population is heterogeneous. If agents have different embodiments, for example when the population consists of both robots and humans, the mapping between robot behaviour and human behaviour, and vice versa, is far from trivial[1].

- Each agent has to implement a *concept of identity*. When behaviour is observed it should be matched with representations of previously observed behaviour. When an agent has a concept of identity, it also has a concept of novelty, which is needed to acquire new behaviour.

- The agents should engage in *imitative interactions*, i.e. the agent should have some behaviour which facilitates imitation. This might be very minimal. Quinn (2001) for example demonstrates how robots engage in imitative following behaviour without this being explicitly programmed. In this context the neural correlates for imitative interactions are also relevant. Mirror neurons, for example, are both active when a monkey observes a specific action executed by another monkey and when the monkey itself engages in the same action (Gallese et al., 1996). It is also possible to investigate the recruitment of innate functionality for imitation in this context. Arbib (2002) for example uses the mirror neuron system to propose an evolutionary account of natural language which relies on the imitation of motor actions. However, in our work we do not wish to elaborate on the origins of imitative behaviour, therefore our agents are already endowed with mechanisms for observing actions, categorising actions (using techniques such as dynamic time warping), and producing actions. For implementation details see (Jansen et al., 2003).

- The action space in which actions are observed and performed is *continuous*. Though this is not a requirement to study imitation in artificial systems (Alissandrakis et al., 2001), a continuous action space is interesting as actions are temporal sequences with noise on the perception and execution. The noise level will influence the successful acquisition of new actions. An important problem that has to be solved in a system with a continuous space of possible actions is how to make the transition from the continuous space of possibilities to a discrete number of action categories (see *e.g.* Oudeyer, 2002).

- The agents have *no specific role* in the population. Each agent can be either the model or the imitator; instead turn taking—not necessarily between the same agents—is used.

- All agents start out as *tabulae rasae*. The action repertoire is completely bootstrapped through interactions with other agents.

- A *source of innovation* is needed for the agents to explore new actions. Agents explore the action space by performing a random action. If other agents pick up such an action, it might propagate through the population. This implies that other agents need to assess whether an action is to be learned or not. This depends on the concept of identity mentioned above: if an observed action is similar to an action category already in the repertoire of the agent, the observed action is used to refine that category. However, if the observed action is new to the imitator, it will learn that action. The choice between either refining an existing category or learning a new category depends on the implementation of the identity measure.

2.2 Imitation games

The agents engage in simple imitative interactions dubbed *imitation games*. An imitation game is played between two agents: one acting as the model, the other as the imitator. The model executes an action which is observed by the imitator. The imitator maps this onto an internal representation and tries to repeat the observed action, not by mimicking the observed action, but by executing the action by mapping the internal representation again onto motor actions. The model then gives feedback on how well the imitation succeeded. The feedback is then used by the imitator to refine its internal representation so that later imitations might be more successful.

2.3 Self-organisation

Through the population dynamics the emerging action repertoire will become shared throughout the population. Yet, the repertoire will only consist of actions that are well imitated. Actions that are not transferred to other agents will, under influence of the imitation game, die out. This is a self-organising system, and just as any other self-organising system it has two components: random fluctuation (implemented in the exploration of agents creating random actions) and positive feedback (given by the model to the imitator) (Steels, 1998). This feedback provides pressure on the action repertoire to by imitable. This self-organising effect arrives at a coherent and shared repertoire of actions throughout the population.

---

[1] The experiments at the VUB AI-lab for the moment do not consider heterogeneous populations.

2.4 Embodiment

It is important that experiments on imitation in artefacts be implemented on robots and not just in simulation. Theories of learning and imitation, especially when imitation takes place in a population, are of necessity very complex. Testing these theories is therefore not always possible in the traditional way with “pen-and-paper” analysis. This is where computer simulations prove to be extremely useful. However, in computer simulations it is often possible to ignore detail that is considered irrelevant for the problem at hand, and implement a much-simplified model. Because of the complex dynamics of learning and imitation, this might result in throwing away the baby with the bath water. Therefore we argue for doing experiments with robots that operate in the real world as well. In a real world situation it is not possible to make oversimplifications. This is not to say that simulations are not useful. Both methods complement each other in important ways. However, eventually models of learning and imitation should be shown to work in the real world. Especially with the recent availability of off-the-shelf actuators and perception systems this has become feasible, as researchers can concentrate now on the cognitive aspects rather than on the details of robotics.

3 Imitation and communication

Experience with imitation in a different domain, namely the emergence of vowel systems, has led us to believe that cultural learning of repertoires of actions through imitation is feasible. The aim of the experiments concerned (de Boer, 2000, 2001) was to investigate the role of self-organisation in determining the universal tendencies of sound systems as used by human languages. For this, a computer model was used in which a repertoire of vowel signals emerged in a population of agents. The main behaviour responsible was imitation, coupled with a drive to imitate the other agents as well as possible with a repertoire of sounds that was as large as possible. As vowel sounds are both easiest to implement in a computer model and best investigated by linguists, these were the speech signals of choice. Each agent in the population was *a priori* able to produce and perceive vowel sounds in a human like way, but each agent also started out *tabula rasa*. As all agents were identical, the correspondence problem was avoided. Through repeated interactions aimed at imitation and occasional random insertions, shared vowel systems emerged. An agent could extend its repertoire of vowels when imitation failed (this was therefore used as a means of novelty detection). It was found that the emerged vowel systems were realistic in the sense that they were very much like the vowel systems that are found in human languages.

Individual agents discover vowel categories through learning of articulatory and acoustic prototypes. Acoustic as well as articulatory prototypes are represented by low-dimensional vectors, and with each vowel category both an acoustic and an articulatory prototype are associated. In response to either success or failure of imitation, prototypes can be moved, merged, removed or added. It must be said that the learning procedure was optimised for the task of vowel learning and is rather *ad hoc*. It remains therefore to be discovered how the procedure must be updated for it to work with more complex actions, such as movements of a robot arm.

These experiments showed that it is indeed possible to have a repertoire of shared actions (in this case, vowel articulations) emerge in a population of agents that imitate each other. No roles or action repertoires were predefined in the population. This shows that successful imitation can emerge without explicit teacher-student relations defined beforehand. The fact that not just random systems emerged, but systems that show remarkable similarities to human vowel systems shows that even though action categories are not defined beforehand, only a limited set of possible sets of categories can stably emerge. The stability of potential sets is determined by the constraints of perception and production, but also by how well the learned systems conform to what is already present in the population, and by the relative positions of the action categories within a system. Such constraints cannot necessarily be determined beforehand and sometimes the constraints that do emerge are unexpected. It is therefore quite possible that given a certain experimental set-up, sets of action categories that have been defined *a priori* by a human operator are not stably transferable from a teacher to a student. In the case of emergent action categories, this problem is avoided.

The experiments with vowels showed that imitation is a feasible way of learning shared repertoire of action categories in a population. A number of corners have been cut in these experiments in order to make them work. Most notably, production and perception were rather simplified with respect to the real thing. The experiments on robots, proposed in the companion paper (Jansen et al., 2003) try to implement a very similar imitation game on a real robot arm. This will clarify which problems will have to be solved in order to implement imitation of complex movements in the real world.

4 Conclusion

We have agued for experiments that take a broader perspective on imitation than is generally assumed in experiments on imitation in artefacts, or even in much research on imitation in animals and humans. Often fixed roles of student and teacher are assumed, as well as a fixed repertoire that has to be transferred. Although many cases of imitation occurring in humans or higher primates are of this kind, it is just a subset of the whole problem of imitation. When imitating other agents, in general an agent also has to decide what to imitate and when to imitate, and possibly also which other agents to imitate. When doing this, it has to solve the problems of the correspondence between observed actions and actions it can perform itself, of the identity between different actions, of role taking and when imitation fails, of how to invent new actions itself. Having a repertoire of actions emerge in a population of agents circumvents the problem of implementing behaviours in a teaching agent that cannot be learned reliably by a student agent. It is very hard to determine beforehand, given the actuators, sensors and learning mechanism of an agent, what behaviours it can and cannot learn. It is therefore quite possible to design a teacher-and-student experiment that is bound to fail. When behaviours are emergent, only imitable action repertoires will emerge. These can always be used later in a teacher-student setting. In this respect emergence in a population through imitation is not just a superset of teacher-student imitation, but also a necessary precursor.

The architecture that is proposed for investigating this is the *imitation game* in a population of agents. All agents in the population are initially without a repertoire of actions (*tabula rasa*) and have to develop a repertoire of actions that they can imitate from each other (it has been shown in simulation that such a set-up works for agents that have to develop a set of shared speech sounds).

In order to avoid inadvertently cutting corners, it is necessary to show that the set-up also works with a population of robots. Simulations are definitely useful for developing the theory, but eventually one has to show that the theory also works in the real world. For this, we have started to work on experiments that employ off-the-shelf robotic components to implement an imitation game for gestures (see the companion paper Jansen et al. 2003). In such a robotic setting, problems of among others noise, calibration, handling time-series data, detecting novelty and establishing identity make the problem of imitation much more challenging. But such a setting also confronts the researcher with a better assessment of where the real difficulties are with imitation. Also, the constraints of perception and articulation limit what repertoires of actions can successfully be learned and transferred. Thus robotic work automatically takes into account that an agent's abilities are not only limited by its cognitive system, but also by its embodiment.

We hope that this paper has provided material for discussing the issues surrounding imitation in populations, and the use of robots and simulations for investigating them.

Acknowledgements

Tony Belpaeme is a postdoctoral fellow of the Flemish Fund for Scientific Research (FWO - Vlaanderen). Bart Jansen is sponsored by a grant from the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT).

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., editors, *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.

Alissandrakis, A., Nehaniv, C. L., and Dautenhahn, K. (2002). Do as i do: Correspondences across different robotic embodiments. In Polani, D., Kim, J., and Martinetz, T., editors, *Proceedings of the Fifth German Workshop on Artificial Life (GWAL5)*, pages 143–152. 18-20 March 2002, Lbeck, Germany.

Arbib, M. A. (2002). Grounding the mirror system hypothesis for the evolution of the language-ready brain. In Cangelosi, A. and Parisi, D., editors, *Simulating the evolution of language*. Springer, London.

Bakker, P. and Kuniyoshi, Y. (1996). Robot see, robot do: An overview of robot imitation. In *Proceedings of the AISB Workshop on Learning in Robots and Animals*, pages 3–11. Brighton, UK, April 1996.

Billard, A. and Hayes, G. (1997). Transmitting communication skills through imitation in autonomous robots. In Birk, A. and Demiris, Y., editors, *Proceedings of EWLR97, Sixth European Workshop on Learning Robots*. Brighton, UK, July 1997.

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.

Dautenhahn, K. and Nehaniv, C. L., editors (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 Boer, B. (2001). *The origins of vowel systems*. Oxford University Press, Oxford, UK.

Gallese, V., Fadiga, L., Fogassi, L., and Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, 119(2):593–609.

Jansen, B., De Vylder, B., Belpaeme, T., and de Boer, B. (2003). Emerging shared action categories in robotic agents through imitation. Submitted to the Second International Symposium on Imitation in Animals & Artifacts 2003.

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.

Oudeyer, P.-Y. (2002). Phonemic coding might be a result of sensory-motor coupling dynamics. In Hallam, J., editor, *Proceedings of the 7th International Conference on the Simulation of Adaptive Behavior*, pages 406–416, Cambridge, MA. The MIT Press.

Quinn, M. (2001). Evolving communication without dedicated communication channels. In Kelemen, J. and Sosík, P., editors, *Proceedings of the European Conference on Artificial Life 2001*, Lectures Notes in Computer Science, pages 357–366, Berlin. Springer.

Steels, L. (1998). Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation. In Hurford, J. R., Studdert-Kennedy, M., and Knight, C., editors, *Approaches to the Evolution of Language: social and cognitive bases*. Cambridge University Press, Cambridge, UK.

Tomasello, M. (1999). *The cultural origins of human cognition*. Harvard University Press, Cambridge, MA.