

Social learning and cultural evolution in embodied and situated agents

Alberto Acerbi, Stefano Nolfi
Institute of Cognitive Sciences and Technologies, CNR
Viale San Martino della Battaglia 44, 00185, Roma, Italy

Abstract – In this paper we investigate whether social learning or the combination of individual and social learning can provide an adaptive advantage for artificial embodied agents that have to develop behavioural abilities that are too difficult or too costly to be acquired through individual learning. Obtained results demonstrate that social learning provides an adaptive advantage when individuals are allowed to learn socially from experienced individuals and individually. Moreover, the obtained results indicate how the iteration of a social and individual learning through out generations lead to a cumulative cultural evolutionary process in which novelties are integrated with previously developed skills and are successfully transmitted in further generations.

I. INTRODUCTION

Social learning refers to the process in which agents learn new skills by interacting with other agents (for a discussion on the definition of the term and of the relation between social learning and other adaptive processes, see [1, 2, 3, 4]). Many species have evolved a capacity to use information provided by other individuals to enhance their individual learning capabilities [5]. By exploiting the interaction with other experienced individuals (e.g. parents or adults), in fact, learning individuals might shortcut the time required to develop the skills already developed by these individuals and/or might increase the chances to acquire skills that are hard to be developed through individual learning. Indeed social learning played an important role in the evolution of complex behaviours in vertebrates [5] and in humans [6, 7].

Using social learning algorithms to evolve artificial embodied agents poses a series of interesting questions. To begin with, social learning in nature is supposed to be realized by numerous different mechanisms at individual level, from cognitive high-demanding intentional imitation learning [8] to simple heuristics in which individual learning is “guided” by social interactions [9, 10, 5]. In addition, whether or not social learning can provide an adaptive advantage in specific conditions, in respect to individual learning, depends on several factors such as the characteristics of the specific behavioural skill to be acquired, the rate at which environmental conditions varies, the availability of the necessary cognitive and behavioural pre-requisites, the existence of a variety of strategies in the population etc. For example, as demonstrated through analytic models developed by the anthropologist Alan Rogers [1] and by Richerson and Boyd [7, 2], imitating the behaviour of other individuals is

adaptive only when the imitation is relatively unusual and most individuals acquire the same behavioural skills through individual learning.

In this paper we investigate through a simulation study the factors that affect the relative costs and benefits of developing a skill through individual versus social learning. As we will clarify in the next sections, by individual and social learning we mean a learning process in which variations of free parameters are retained or discarded on the basis of whether they lead or not to an improvement of individual performance or on the basis of whether they reduce or not the differences between the sensory-motor mapping produced by an inexperienced and an experienced individual, respectively.

In particular we show that social learning can successfully lead to the acquisition of skills that are difficult to acquire through individual learning if the process of social learning is accompanied by a phase in which individuals adjust on their own what they have learn socially. Indeed, at least in the case in which social learning consists of an attempt to imitate the fine-grained behaviour displayed by experienced individuals, minor differences between the sensory-motor mapping produced by experienced and inexperienced individuals at the end of learning might lead, in the case of sequential behaviours, to different outcomes that significantly impact on the individuals' performance.

Finally, we show how an iterated form of social learning, in which inexperienced individuals of a given generation learn socially, by interacting with the best individuals of the previous generation, and individually lead to a cumulative cultural evolution process.

In section 2 we describe our experimental set-up. In section 3 we describe a model in which inexperienced individuals learn on the basis of a social learning process or of a combination of social and individual learning. In section 4, we present the results obtained by iterating the social learning process through out generations. Finally, in section 5 we discuss the implications of the obtained results from a theoretical point of view (i.e. from the point of view of understanding in which conditions social learning can have an adaptive value) and from an engineering point of view (i.e. from the point of view of developing artificial agents displaying complex behaviours).

II. EXPERIMENTAL SETUP

In this section we describe the experimental set-up used to perform our experiments and the learning algorithm used to develop experienced agents, i.e. agents able to correctly perform a foraging task. These agents will be used, in the experiments described in the following sections, to study whether inexperienced agents can acquire the skills of experienced agents through a combination of social and individual learning.

The foraging task involves two qualitatively different abilities: (a) an ability to categorize objects, and (b) an ability to approach or avoid objects depending on whether they belong to a food or poisonous element category. The latter skill involves the ability to produce a sequential behaviour (i.e. sequence of motor movements that, in interaction with the environment, allow the robot to reach or avoid the perceived element).

The agents, the environment, and the task

Consider the case of agents that "live" in an environment containing edible and poisonous food elements (α and β , respectively) that should be able to reach ("eat") α and avoid β . Both α and β objects occur in two different varieties that have different perceptual properties.

Agents are allowed to "live" for 200 trials, each lasting 50 time steps. During each trial the agent is placed in a bounded squared arena with a side of 1 unit, containing a single object located in the centre randomly selected from the four object types. At the beginning of each trial the agent is placed in a location randomly selected around the object and at a minimum distance from the object of 0.5 units. The performance of the agent is measured by computing a score of 0.01 and -0.01 for each trial in which the agent reach an object α or β , respectively. Since each agent face during its life one hundred α objects and one hundred β objects, agents' performance might range from -1.0 to +1.0.

Agents are provided with a neural network controller that includes eight sensory neurons, five internal neurons, and two motor neurons (Fig. 2). The architecture of the neural controller include two modules: (1) a module A that receive as input the perceptual property of the current food element (i.e. "11" or "00" in the case of α objects and "10" or "01" in the case of β objects) and that should produce as output an appropriate categorization, (2) a module B that receive as input the data from a simulated camera, and the category of the current food element (produced by module A) and that should produce as output the appropriate motor action (Fig. 2). The simulated camera is a linear camera with a visual degree of 60° provided with 60 photoreceptors with a view angle of 1° each. Each photoreceptor is activated to 1.0 or to 0.0 depending on whether the food object is located in its view field or not. Each of the 6 sensory neurons encodes the average activation level of 10 corresponding photoreceptors. Noise is simulated by adding a random value in the range [0.0,

0.05] to the value of each sensory neuron. The two motors neurons control two simulated wheels. The activation state of the motor neuron is normalized in the range [-0.03, 0.03] corresponding to the maximum distance covered by each wheel during 100ms.




Fig. 1. The environment and a typical behavior exhibited by an agent able to solve the problem in the case of an edible food element.




Fig. 2. The agents' neural controller.

Creating experienced agents through an individual learning procedure

To synthesize experienced agents to be used in the experiments on social learning we decided to train agents provided with the neural architecture described in the previous section through a simulated annealing algorithm [11]. Simulated annealing was chosen for its simplicity and for the fact that it does not pose any constraint on the learning procedure or on the architecture of the learning agent.

More precisely we used the following procedure: (1) The free parameters of the agents neural controller (i.e. the connection weights and biases) are set randomly with an uniform distribution within the range [-5.0, 5.0]; (2) The agent is allowed to live for 200 trials and its performance are evaluated; (3) The free parameters are varied; (4) The performance of the agent are re-evaluated and the variation of the free parameters are retained or discarded depending on whether they produce an increase or a decrease of the agent's performance. To reduce the computational cost, maladaptive variations are discarded after 20 trials while possibly adaptive variations are retained only after a test period of 200 trials; (5)

The third and fourth step are repeated for 100000 learning cycles.

Variations are introduced by replacing each free parameter with a new value, randomly selected with a uniform distribution on the same range, with a probability of 5 %.




Fig. 3. Performance of 10 agents trained with simulated annealing at the end of the training phase. Test performance is obtained by testing the agent for 1000 trials.

Variations that lead to equal or better performance are always retained. Variations that lead to worse performance, instead, are retained only with a probability computed according to the following equation:

$$p = \exp\left(\frac{f(t) - f(t-1)}{T}\right) \quad (1)$$

Where f is the evaluation function, t is the learning cycle, and T is a temperature that decreases linearly from 0.02 to 0.0 during the 100000 learning cycles. This means that maladaptive variations are retained only with a probability that is inversely proportional to their negative effects and that is progressively reduced as the learning process proceeds.




Fig. 4. Output of the categorization module produced for the four objects. Each column represents the categorization output produced

by the individual of a different replication of the experiment at the end of the learning phase. Circles and squares indicate the output produced for α and β objects, respectively.

The experiment was replicated for 10 individuals starting with different randomly assigned free parameters (Fig. 3). As can be seen, performance varies in different replications from rather good performance (in the case of replications 7 and 10) to very poor performance (in the case of replications 1, 2 and 4).

Significant differences in performance within the 10 replications of the experiment are also observed with respect to agent's categorization abilities (Fig. 4). In fact, an effective categorization, in which objects corresponding to the same category are represented similarly and objects of α and β categories are well separated in the representation space, only occurs in the case of replication 7 and 10 that corresponds to the best performing individuals.

These results indicate that the chosen problem is difficult to learn individually. In the following section we will analyze whether social learning and the combination of individual and social learning can enhance agents' learning capabilities.

III. SOCIAL LEARNING

In this section we describe the results obtained by training agents with the architecture described in section II in two different conditions: (1) on the basis of a social learning, or (2) on the basis of a combination of social and individual learning. Social learning consists in the attempt to approximate the behaviour produced by an experienced agent (i.e. the best agent obtained through the individual learning procedure described in the previous section). Individual learning consists in the attempt to improve agent's foraging abilities (i.e. the same process described in section II).

Since the problem to be solved require two abilities, i.e. an ability to categorize different food elements and an ability to display an approaching or avoidance behaviour, for both conditions we run three sets of experiments in which social learning involved the approximation of: (1) the motor actions produced by the experienced agent only, (2) the motor and categorization actions produced by the experienced agent, (3) the categorization output produced by the experienced agent (see Fig. 5).

It should be noticed that in a realistic setting the motor and categorization actions produced by the experienced agents are not directly accessible to inexperienced agents. They are accessible only indirectly through sensory states that provide indirect and often incomplete information on what another agent is doing and on how it is categorizing different environmental conditions (e.g. by producing different vocalizations in different environmental circumstances). This aspect implies that agents that learn socially need to solve the so-called "correspondence problem" [12] that consists in the need to translate the sensory information that provide indirect information on what another agent is doing into corresponding motor outputs to be produced by the inexperienced agent. This important aspect is simplified in our experiments by assuming that the learning and the experienced agents are placed in the

same environmental condition (i.e. that the learning agent is “on the shoulder” of the experienced agent) and that the learning agent has direct access to the motor commands executed by the experienced agent. The rationale for this simplification is that in this paper we want to study whether and in which conditions social learning can produce an adaptive advantage independently from the extent to which this correspondence problem is solved.




Fig. 5. A schematisation of the social learning process. The learning robot try to approximate the motor action produced by the experienced robot (full arrow) and/or the categorization performed by the experienced robot (dotted arrow).

Our hypothesis in performing these experiments is that the availability of a an explicit good model, constituted by the best experienced individual, would ease the task of developing the requested capability with respect to an individual learning scenario in which the solution of the problem should be discovered from scratch.

Experimental setting and results

During social learning the inexperienced agent is placed “on the shoulder” of the experienced agent (i.e. the inexperienced agent receives the same sensory state of the experienced agent and moves on the basis of the motor actions produced by the experienced agent, for other related models see [13-19]). Social learning is realised on the basis of the simulated annealing algorithm described in section II. The evaluation function, however, is based on the sum of the discrepancies of the motor and/or categorization outputs of the two agents: variations are retained or discarded on the basis of their impact on the ability to approximate the experienced agent’s output.

In the experiments in which agents are allowed to learn socially and individually, the learning process involves two phases. During the first phase, lasting 50000 learning cycles, variations are retained or discarded on the basis of their impact on the approximation learning task only, During the second phase, lasting 50000 learning cycles, variations are retained or discarded on the basis of the weighted sum of agent’s ability to approximate the sensory-motor mapping produced by the experienced individual and to forage. Such weighted sum of (Δw) is computed according to the following equations:

$$Dw = f(I)W + f(S)(1 - W) \quad (2)$$

$$W = \frac{C}{C_{max}} \quad (3)$$




Fig. 6. Average performance and standard deviation for 10 replications in the following experimental conditions: IL: agents that learning individually only (see section 2.2); SIL: agents that learn socially and individually; SL: agents that learn socially only. Social learning has been realized in three different conditions: (1) by approximating only the motor output of the experienced agent; (2) by approximating the motor and the categorization output of the experienced agent; (3) by approximating only the categorization output of the experienced agent.

Where $f(I)$ indicates the performance with respect to the individual learning task and $f(S)$ indicates the performance with respect to the social learning task, C is the learning cycle during the second phase (ranging between 0 and 50000), and C_{max} is a constant equal to 50000.

As shown in Fig. 6, the possibility to learn socially only from the best experienced agent do indeed lead to better performance with respect to the situation in which agents learn to forage individually only, with respect to the average performance of the group. Notice that condition SL3, i.e. the experimental condition in which agents learn socially by approximating the categorization of the experienced agents, can’t be adaptive by definition since the performance of the learning agents are never evaluated in respect to the foraging ability.

We nonetheless can observe a significant difference between the condition in which agents learn socially and individually and the condition in which they learn only socially --- the performance in the former condition are much better both from the point of view of the average performance and from the point of view of the performance of the best learning individual.

The fact that the social learning alone lead to worse performance with respect to the combination of individual and social learning can be explained by considering that the reduction of the discrepancy between the outputs produced by the inexperienced and the experienced agents does not necessary produce an improvement on the foraging

performance (Fig. 7). Indeed, small differences between the outputs of the two agents (e.g. summed square error < 0.05) might correspond to huge differences in term of performance. This fact, in turn, can be explained by considering that sequential behavioural abilities, such as for example reaching a food element through a sequence of movements, are dynamical properties resulting from a sequence of non-linear interactions between the agent and the environment where small differences at the level of the interactions might lead to significant differences at the level of the resulting behaviour [20, 21].

Other differences between the cases in which agents learn socially or socially and individually will be discussed in the following sections.




Fig. 7. Summed square difference (error) between the motor actions produced by the experienced and the learning agent (bottom) and performance of the learning agent on the foraging task (top) during a typical experiment in which an inexperienced agent learns to approximate the sensory-motor mapping produced by the best experienced agent. Performance on the foraging task has been measured by stopping the learning process and by testing the learning individual after each learning trial.

IV. ITERATED CULTURAL TRANSMISSION

In this section we describe a set of experiments in which we tested whether the abilities acquired through the social or social and individual process described in the previous section can be further transmitted to other individuals so to realize an iterated cultural transmission process in which learning agents later become and act as experienced agents.

This iterated cultural transmission process was realized by repeating the social learning process described in the previous section for a certain number of generations, by selecting as experienced agent for generation $g+1$ the agent of generation g with the best performance with respect to the foraging task, and by using as inexperienced agents individuals with randomly assigned free parameters. Notice how this experimental setting involves a vertical transmission process only (for related experiments involving both vertical and

horizontal transmission process, see [22, 15]). In the experiment reported in this section the populations of learning agents consisted on 10 individuals, with randomly initialised free parameters. Only a single learning agent was selected as experienced agent for the succeeding generation. The best individual obtained through individual learning (see section II) was used as experienced agent for generation 1. Iterated learning was realized by using the simulated annealing algorithm described in the previous section and by allowing the inexperienced agent to approximate the motor and categorization outputs of the experienced agent (i.e. we selected the variation of the experiments performed in the previous section that produced the best results). All other parameters were kept the same as those described in the previous section.




Fig. 8. Best and average performance through out generations in an iterated cultural transmission process in which agents learn through a combination of individual and social learning. The plain line represents the moving average of the data.




Fig. 9. Best and average performance through out generations in an iterated cultural transmission process in which agents learn through social learning only. The plain line represents the moving average of the data.

The obtained results indicate that, in the condition in which individuals learn socially and individually, the skills developed by the best individuals of previous generations are successfully transmitted to individuals of succeeding generations (Fig. 8) while, in the case in which individuals learn socially only, performance deteriorate through out generations (Fig. 9). The increase in the average performance observed in the former case indicates that social learning in combination with individual learning enhances agents' ability to develop effective foraging skills. The increase in the performance of the best individual through out generations indicate that the iterated social learning process lead to a cumulative cultural evolutionary process in which the skills and the characteristics of the skills that are acquired through social and individual learning evolve through out generations (for a related result in the context of the emergence of a communication system, see [23]).

These conclusions are further demonstrated by comparing the performance agents that have been subjected to: (1) an individual learning process only, (2) a social and individual learning process (i.e. individuals of generation 2), and (3) an iterated social and individual learning process (i.e. individuals of generation 10). As shown in Fig. 10, in fact, agents subjected to the iterated learning process outperform both agents that have been subjected to individual and social learning (but not to the iterated learning process) and agents that have be subjected to individual learning only. The difference between the three conditions is significant according to a non-parametric analysis of variance (Kruskall-Wallis test).

To identify the nature of the innovations occurred through out generations and successfully transmitted to individuals of successive generations in the experiments in which agents learned individually and socially we measure the percentage of α and β elements reached (Fig. 11) and the average time needed to reach α elements (Fig. 12) for agents of generation 2 and 10 (i.e. for agents that were not or where exposed to iterated individual and social learning processes, respectively). Obtained results indicate that from generation 2 to generation 10 performance increase both with respect to the ability to produce correct behaviors (i.e. by reaching α elements and avoiding β elements) and with respect to the ability to reach α elements quickly so to reduce the risk to not reach the food elements within the given time. The fact that agents of generation 10 reach food elements more quickly than agents of generation 2 indicate that progresses through out generations concern agents' motor skills. The fact that the best agent of generation 2 erroneously reach β_2 (i.e. one of the two poisonous elements) in about 10% of the cases while the best agent of generation 10 never produce this type of error demonstrates how progresses through out generation concern also agents' categorization abilities (i.e. agent ability to treat β_1 and β_2 objects in the same manner).




Fig. 10. Performance of 100 agents trained in the three experimental conditions: (IL) individual learning, (ICT-gen 1) individual and social learning, (ICT - gen 10) iterated individual and social learning. Boxes represent the inter-quartile range of the data. The horizontal lines inside the boxes indicate the median values. The horizontal lines outside the boxes indicate the minimum and maximum values.




Fig. 11. Percentage of objects reached in 1000 trials. α_1 , α_2 , β_1 , and β_2 indicate the performance for the two food and two poisonous elements, respectively. The black and white histograms indicate the performance for agents of generation 2 and 10, respectively.




Fig. 12. Average time (in simulation step) in which agents reach the two “edible” elements. The black and white histograms indicate the performance for agents of generation 2 and 10, respectively.

V. DISCUSSION AND CONCLUSIONS

In this paper we investigated whether social learning or the combination of individual and social learning can provide an adaptive advantage for artificial embodied agents that have to develop behavioral abilities that are difficult or too costly to acquire through individual learning only.

To achieve this objective we devised a simple experimental scenario that requires the acquisition of an ability to categorize different environmental conditions (i.e. edible and poisonous food elements) and to produce simple sequential behaviors (i.e. approaching and avoiding a food element). We developed a new algorithm that allows embodied agents to adapt on the basis of the combination of a social and individual learning processes and a cultural evolutionary algorithm in which individuals of succeeding generations learn socially and individually by interacting with the best individuals of previous generations. We showed how agents that learn on the basis of social and individual learning outperform agents that learn on the basis of an individual learning only or social learning only. Finally, we showed how the iteration of this process for several generations can lead to the emergence of an ‘artificial culture’ in which new skills or new characteristics of agents’ skills discovered by specific individuals through learning are culturally transmitted throughout generations together with other pre-existing skills.

The adaptive advantages obtained by combining social and individual learning can be explained by considering that: (a) the possibility to develop through social learning strategies that resemble that of successful individuals reduces the risk that individual learning is stacked in local minima, and (b) the possibility to refine the strategy acquired through social learning through individual learning overcomes the problem that small differences between the sensory-motor mapping

produced by the experienced and inexperienced agent might cause significant differences in performance.

The fact that the iteration of individual and social learning processes combined with the selection of the best individual and the end of the learning process as experienced individual leads to an increase in performance through out generations can be explained by considering that: (a) the combination of individual and social learning leads to a reliable transmission of skills acquired in previous generations, and (b) the possibility to develop skills that are similar but not identical to those of experienced individuals allow agents to explore new types of solutions that can be further optimized through individual learning so to lead to solutions that are better than those of current experienced individuals. Overall these two factors lead to a cumulative cultural evolutionary process in which skills developed in previous generations are preserved in successive generations and in which novelties occurring as the result of social and individual learning are integrated with previously developed skills.

The obtained results also indicate how social and individual learning have different characteristics that make these two forms of learning effective in certain cases and weak in other cases. In particular, we showed how social learning, if realized by reducing the differences in the sensory-motor mapping produced by an inexperienced agents and an experienced agents at the level of fine-grained behavior, might produce rather poor performance when the skills to be acquired consist of sequential behaviors.

The complementary characteristics of the two learning processes confirm the importance of combining them and indicate that the two processes might play partially different roles. In particular, these considerations suggest that the main role of imitational learning might be in the acquisition of macro-behaviors and behavioral arbitration mechanisms (i.e. the acquisition of an ability to identify the type and the number of behaviors to be executed in different circumstances) while the main role of individual learning might be in the acquisition of elementary behavioral abilities (i.e. the ability to produce a sequence of fine-grained interaction with the environment that lead to the production of a given elementary behavior). For a similar view, see [16, 24].

These considerations are coherent with the observation that social learning in real animals is always paired and intertwined with individual learning [3, 5].

In future work we plan to apply the model proposed in this paper in more realistic robotic scenarios and to investigate the possibility to implement in artificial agents simpler forms of social learning [3, 9, 10] that do not require an explicit imitation of the behavior exhibited by experienced agents and in which the social transmission of skills occur more indirectly.

ACKNOWLEDGMENT

This work was supported by the ECAGents project, a project funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant 001940.

The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

REFERENCES

- [1] A. R. Rogers, "Does biology constrain culture?", *American Anthropologist*, 90, pp. 819-831, 1989.
- [2] R. Boyd and P. J. Richerson, "Why does culture increase human adaptability?" *Ethology and Sociobiology*, 16, pp. 125-143, 1995.
- [3] B. G. Galef, "Introduction: social learning and imitation", in *Social learning and imitation: The roots of culture*, C. M. Heyes and B. G. Galef, Eds., New York: Academic Press, 1996, pp. 3-16.
- [4] T. Kameda and D. Nakanishi, "Cost-benefit analysis of social/cultural learning in a non-stationary uncertain environment: An evolutionary simulation and an experiment with human subjects", *Evolution and Human Behavior*, 23, pp. 373-393, 2002.
- [5] B. G. Galef and K. N. Laland, "Social learning in animals: Empirical studies and theoretical models", *BioScience*, 55(6), pp. 489-499, 2005.
- [6] P. J. Richerson and R. Boyd, *Not by genes alone: How culture transformed human evolution*, Chicago: Chicago University Press, 2004.
- [7] R. Boyd and P. J. Richerson, *Culture and the Evolutionary Process*, Chicago: University of Chicago Press, 1985.
- [8] M. Tomasello, "Do apes ape?", in *Social learning in animals: The roots of culture*, C. M. Hayes and B. G. Galef, Eds., San Diego: Academic Press, 1996, pp. 319-346.
- [9] E. Visalberghi and D. Frigaszy, "Do monkeys ape?" in *Language and intelligence in Monkeys and Apes*, S. T. Parker and K. R. Gibson, Eds., Cambridge: Cambridge University Press, 1990, pp. 247-273.
- [10] J. Noble and P. M. Todd, "Imitation or something simpler? Modelling simple mechanisms for social information processing", in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds., Cambridge: MIT Press, 2002, pp. 423-440.
- [11] S. Kirkpatrick, C. D. Gelatt Jr. and M. P. Vecchi, "Optimization by simulated annealing", *Science*, 220(4598), pp. 671-680, 1983.
- [12] K. Dautenhahn and C. L. Nehaniv, "The correspondence problem", in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds., Cambridge: MIT Press, 2002, pp. 41-61.
- [13] D. Denaro and D. Parisi, "Cultural evolution in a population of neural networks", in *Neural Nets.WIRN-96*, M. Marinaro and R. Tagliaferri, Eds., Berlin: Springer Verlag, 1996, pp. 100-111.
- [14] D. Parisi, "Cultural evolution in neural networks", *IEEE Experts*, 12(4), pp. 9-11, 1997.
- [15] A. Acerbi and D. Parisi, "Cultural transmission between and within generations" *Journal of Artificial Societies and Social Simulations*, 9(1), electronic publication, 2006.
- [16] S. Schaal, "Is imitation learning the route to humanoid robots?", *Trends in Cognitive Sciences*, 3(6), pp. 233-242, 1999.
- [17] G. Hayes and J. Demiris, "A robot controller using learning by imitation", in *Proceedings of the 2nd International Symposium on Intelligent Robotic System*, A. Borowsky and J. L. Crowley, Eds., Grenoble: LIFTA-IMAG, 1994, pp. 198-204.
- [18] K. Dautenhahn, "Getting to know each other: artificial social intelligence for autonomous robots", *Robotics and Autonomous System*, 16, pp. 333-356, 1995.
- [19] M. Ito and J. Tani, "On-line imitative interaction with a humanoid robots using a dynamic neural network model of a mirror system", *Adaptive Behavior*, 12(2), pp. 93-115, 2004.
- [20] S. Nolfi, "Behaviour as a complex adaptive system: On the role of self-organization in the development of individual and collective behaviour", *ComplexUs*, 3(3/4), pp. 195-205, 2006.
- [21] H. J. Chiel and R. D. Beer, "The brain as a body: Adaptive behaviour emerges from interactions of nervous system, body and environment", *Trends in Neurosciences*, 20(12), pp. 553-557, 1997.
- [22] D. Livingstone and C. Fyfe, "Modelling the evolution of linguistic diversity" in *Advances in Artificial Life. Proceedings of the Fifth European Conference, ECAL99*, D. Floreano, J. Nicoud and F. Mondada, Eds., Berlin: Springer Verlag, 1999, pp. 704-708.
- [23] P. Vogt, "Cumulative cultural evolution: can we ever learn more?", in *From Animals to Animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour*, S. Nolfi et al., Eds., Berlin: Springer Verlag, 2006, pp. 738-749.
- [24] S. Schaal and D. Sternad, "Programmable pattern generators" in *Proceedings of 3rd International Conference on Computational Intelligence in Neuroscience*, Research Triangle Park, NC, 1999, pp. 48-51.