AISB Journal **1**(4), © SSAISB, 2003
http://www.aisb.org.uk

Enhancing Autonomous Agents Evolution with Learning by Imitation

Elhanan Borenstein* and Eytan Ruppin†

* School of Computer Science, Tel Aviv University, Tel-Aviv 69978, Israel, *borens@post.tau.ac.il*
† School of Computer Science and School of Medicine, Tel Aviv University, Tel-Aviv 69978, Israel, *ruppin@post.tau.ac.il*

Abstract

This paper presents a new mechanism to enhance the evolutionary process of autonomous agents through lifetime adaptation by imitation. Imitation is an effective method for learning new traits and is naturally applicable within the evolutionary paradigm. We describe a set of simulations where a population of agents evolve to solve a certain task. In each generation, individuals can select other agents from the population as models (teachers) and imitate their behavior. In contradistinction to previous studies, we focus on the interaction between imitation and evolution when imitation takes place only across members of the same generation, and does not percolate across generations via vertical (cultural) transmission. We show how this mechanism can be applied to successfully enhance the evolution of autonomous agents, when other forms of learning are not applicable.

1 Introduction

A large body of work in recent years has studied *the interaction between lifetime learning and genetic evolution* when lifetime adaptations, acquired by learning, are not inherited. Hinton and Nowlan (1987) introduced a simple model that demonstrates how learning can guide and accelerate evolution. Nolfi et al. (1994) presented experimental results supporting this view, even when the learning task differs from the evolutionary task. Other researchers (Nolfi and Parisi, 1997; Floreano and Mondada, 1996) studied the interaction between learning and evolution in robots and artificial agents systems. These studies employed various sources of training data such as external oracles, regularities in the environment or "self-generated" teaching data. There is, however, an additional source of training data; one which is naturally available within the evolutionary paradigm - the knowledge possessed by other members of the population. This knowledge can be harnessed to improve the evolutionary process in the form of *learning by imitation*.

The motivation for using learning by imitation to enhance evolution is twofold. First, imitation is an effective and robust way to learn new traits by utilizing the knowledge already possessed by others. The existence of true imitative behavior in the animal kingdom is still in debate, however, social learning can be found in a variety of species providing clear benefits over other forms of learning (Kawamura, 1963; Whiten and Ham, 1992; Zentall, 2001). Second, while oracles or other forms of supervised training data are scarce in agent environments, learning by imitation is still a valid option, using other members of the population as teachers.

Extending these studies further, *our goal is to put forward a novel framework for merging these two approaches and study learning by imitation within the scope of the interaction between learning and evolution.* We wish to explore *learning by imitation* as an alternative to conventional supervised learning and to apply it as a tool to enhance genetic evolution. We will label this framework as *imitation enhanced evolution (IEE).*

Learning by imitation has already been applied by researchers in the fields of artificial intelligence and robotics in various experiments. Hayes and Demiris (1994) presented a model of imitative learning to develop a robot controller. Billard and Dautenhahn (1999) studied the benefits of social interactions and imitative behavior for grounding and use of communication in autonomous robotic agents. For an up-to-date introduction to work on imitation in both animals and artifacts see the cross-disciplinary collection (Dautenhahn and Nehaniv, 2002). Furthermore, various frameworks that study the interaction between cultural transmission and evolution have already been well established (e.g. Boyd and Richerson, 1985; Cavalli-Sforza and Feldman, 1981; Laland, 1992). Gene-culture co-evolution accounts for many adaptive traits (Feldman and Laland, 1996). Studies and simulations of the evolution of language (Ackley and Littman, 1994; Kirby and Hurford, 1997; Arbib, 2002) assume, by definition, some sort of cultural transmission.

It is important to realize though, that in contradistinction to these studies, our framework does not employ cultural evolution. In fact, we preclude culture from evolving in the first place. *Following in the footsteps of the studies of the interaction between learning and evolution cited above, we thus avoid any direct form of acquired-knowledge transfer between generations either genetically or culturally.* We work in a strict Darwinian framework, where lifetime adaptations are not directly inherited (although, as demonstrated in some of the studies cited above, they may be genetically assimilated through the Baldwin effect, 1896) and may affect the evolutionary process only by changing the individual's fitness, and thus the number of its offsprings. In terms of cultural transmission (see Boyd and Richerson, 1985, for a detailed definition), we allow *horizontal* transmission alone (where individuals of the same generation imitate each other) and exclude any form of *vertical* transmission (where members of the current generation transmit their knowledge to members of the next generation). Numerous field studies suggest that at least in non-human societies, horizontal transmission is far more common than vertical transmission (Laland, 1992). Furthermore, to prevent any form of cultural evolution from taking place, within each generation, only innate behaviors are imitated; that is, we prevent behaviors acquired by imitation from being imitated again by another member.

A simple model that fits this framework has been studied before by Best (1999). He demonstrated an extension of the computational model presented in Hinton and Nowlan (1987), introducing social learning (namely *imitation*) as an additional adaptive mechanism. The reported results exemplify how horizontal cultural transmission can guide and accelerate the evolutionary process in this simplified model. Best has also demonstrated how social learning may be superior to conventional learning and yield faster convergence of the evolutionary process. However, Best's model has several limitations. The evolutionary fitness function (which is the one used in Hinton and Nowlan, 1987) represents a worst-case scenario where only the exact solution has a positive fitness value. There is no probable path that a pure evolutionary search can take to discover this solution. Additionally, there is no distinction between genotypes and phenotypes and thus no real *phenotypic* adaptation process. Imitation is carried out simply by copying certain *genes* from the teacher's genome to the student.

We wish to generalize this framework and study the effects of learning by imitation in a more realistic scenario of *autonomous agents evolution* (see Ruppin, 2002, for a general review). The definition of imitation in the literature varies considerably (Billard and Dautenhahn, 1999), but for the purpose of this paper we use imitation (or learning by imitation) in the sense of having an individual (student) being able to match its behavior to that of a demonstrator (teacher). In particular, using autonomous agents to model the population members, this form of imitation is implemented by using the teacher's output for each sensory input, as the target output in a back-propagation training algorithm. We focus on the effects that imitation may have on the genetic evolutionary process, starting with the most basic question: *can imitation enhance the evolution of autonomous agents (in the absence of vertical transmission), in an analogous manner to the results previously shown for supervised learning, and how?* Although it was shown that *learning* can guide the evolutionary process (e.g., via the Baldwin effect), the contribution of *imitation* to evolution is not obvious; while in late stages of the evolutionary process the best agents may already possess sufficient knowledge to approximate a successful teacher, in early stages of the process it may be the case of "the blind leading the blind", resulting in a decrease of the population's average fitness.

This paper presents a set of simulations, where lifetime learning by imitation was used to adapt individuals that go through an evolutionary process. The results are compared with those of a simple evolutionary process, where no lifetime learning is employed, and with those of an evolutionary process that employs conventional supervised learning.

The remainder of this paper is organized as follows. We begin in Section 2 with a brief overview of the effect of lifetime adaptation on the evolutionary process. In Section 3 we present the *IEE* model in details. To validate the effectiveness of our model we introduce in Section 4 a set of tasks which were used to test our model and the experimental results in Section 5. The paper concludes with a discussion of future work and a short summary.

2 The Effects of Lifetime Adaptation on Genetic Evolution

Studies of the interaction between lifetime learning and evolution (Hinton and Nowlan, 1987; Nolfi et al., 1994; Nolfi and Parisi, 1997; Floreano and Mondada, 1996) have shown that learning can accelerate and guide the genetic evolutionary process. These studies demonstrated (through both theoretical analysis and simulations) how the *dynamics* of the lifetime adaptation process can account for this positive effect. The phenotypic modifications that take place in an individual subject to lifetime adaptation (e.g. learning), significantly depend upon its innate configuration. Individuals which initially have a low fitness value, may attain higher fitness through learning. The expected fitness gain though, will be higher for individuals which are initially closer to the optimum configuration. As illustrated in Figure 1, learning can thus help to reveal the innate potential of each individual in the population. One may consider lifetime adaptation as a local search process that can enhance the global search (evolution) by determining which configurations lie in the vicinity of the global optimum solution and are thus worthwhile retaining in the population (as they have a better chance to produce successful offsprings). From a mathematical standpoint, lifetime adaptation can be conceived as a *functional* that can potentially transform an initially ragged fitness function into a smoother function, making the evolutionary process more effective.

Our hypothesis is that learning by imitation, that is, using the best individuals in the population as teachers, may be sufficient to reveal the innate *potential* of the population members. The results reported in the following sections clearly validate this assumption.

In this study we focus on the simple case where the learning (imitation) task is similar to the evolutionary task. This case most probably does not closely represent the imitation processes found in nature. Lifetime adaptation in humans and other cultural organisms operates on high-level traits which are not coded directly in their genome. However, we believe that this simple scenario can provide valuable insights into the roots of imitative behavior. We further discuss this topic in Section 6.

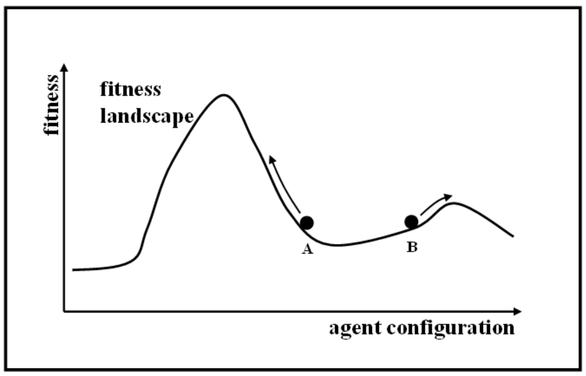

Figure 1: An illustration of the effect that lifetime adaptation may have on the genetic evolutionary process. Both agents start with the same innate fitness value (indicated by the black dots). Applying lifetime adaptation (illustrated as a simple hill climbing process) will result in the selection of agent A which is closer to the optimal solution. Inspired by Nolfi and Floreano (1999)

3 The Model

A haploid population of agents evolve to solve various tasks. Each agent's neurocontrollers is a simple feed-forward (FF) neural network (Hertz et al., 1991). The initial weights of the network synapses are coded directly into the agent's genome (the network topology is static throughout the process). The initial population is composed of 100 individuals, each assigned randomly selected connection weights from the interval [-1,1]. The *innate fitness* of each individual is determined by its ability to solve the specific task upon birth. Within the pure evolutionary process, the innate fitness will determine the reproductive probability of this individual. Each new generation is created by randomly selecting the best agents from the previous generation according to their innate fitness, and allowing them to reproduce (Mitchell, 1996). During reproduction, 10% of the weights are mutated by adding a randomly selected value from the interval [-0.35,0.35]. The genomes of the best 20 individuals are copied to the next generation without mutation.

When conventional supervised learning is applicable (i.e., an explicit oracle can be found) we also examined the effect of supervised learning on the evolutionary process. Each individual in the population goes through a lifetime learning phase where the agent employs a back-propagation algorithm (Hertz et al., 1991), using the explicit oracle as a teacher. Its fitness is then reevaluated to determine its *acquired fitness* (i.e., its fitness level after learning takes place). In order to simulate the delay in fitness acquisition associated with acquired knowledge, we use the average of the innate and acquired fitness values as the agent's *final fitness* value. This fitness value is then used to select the agents that will produce the next generation.

In the IEE paradigm, agents do not use conventional supervised learning, but rather employ learning by imitation. In every new generation of agents, created by the evolutionary process, each agent in the population selects one of the other members of the population as an imitation model (teacher). Teachers are selected stochastically, where the probability of selecting a certain agent as a teacher is proportional to its *innate* fitness value (i.e., its initial fitness levels before learning takes place). The agent employs a back-propagation algorithm, using the teacher's output for each input pattern as the target output, mimicking a supervised learning mode. The imitation phase in each generation can be conceived as happening simultaneously for all agents, preventing behaviors acquired by imitation from being imitated. Only the *innate* behavior of the teacher is imitated by the student. The *acquired fitness* and *final fitness* are evaluated in the same method that was described in the case of conventional learning.

As stated above, acquired knowledge does not percolate across generations. Each time a new generation is produced, all lifetime adaptations possessed by the members of the previous generation are lost. Newborn agents inherit only the genome of their parents which does not encode the acquired network adaptations that took place during the parent's lifetime. Successful individuals that were copied from the previous generation also go through a new genotype-to-phenotype ontogenetic development process and thus lose all adaptations acquired during the previous generation.

To summarize, learning by imitation in a population of evolving agents (IEE) works as follows:

1. *Create the initial population. Assign the network weights of each individual with randomly selected values.*
2. *Repeat:*
   (a) *For each individual in the population:*
      i. *Evaluate the innate fitness* F_i.
   (b) *For each individual* S *in the population:*
      i. *Set* S *to be the student.*
      ii. *Select a teacher* T *from the population. The probability of selecting a certain individual as a teacher is proportional to its innate fitness value* F_i.
      iii. *Train* S *with back-propagation algorithm. Use the output of* T *as the desired output (when computing the output of* T*, use the innate configuration of* T*).*
      iv. *Evaluate the acquired fitness* F_a *of* S.
   (c) *For each individual in the population:*
      i. *Evaluate the final fitness* $F_f = \frac{F_i + F_a}{2}$.
   (d) *Create the next generation by selecting the best individuals according to* F_f *and allow them to reproduce as described above.*

4 The Tasks

The model described in the previous section was tested on three different tasks. The first two are standard classification benchmark problems. The third is an agent-related task used in previous studies of the interaction between learning and evolution.

4.1 The Parity Problem

The agents evolved to solve the five bit parity problem. A network topology of 5-6-2-1 was used (i.e., 5 input neurons, two hidden layers, the first with 6 neurons and the second with 2, and 1 output neuron), with an additional threshold unit in each layer. All 32 possible input patterns were used both for evaluating the network performance and for training.

4.2 The Triangle Classification Problem

A simple two-dimensional geometrical classification problem was used in this task. The network receives as input a point from the unit square and should determine whether it falls within the boundaries of a predefined triangle. A network topology of 2-5-1 was used (with an additional threshold unit in each layer). The test set and training set consisted of 100 points randomly selected from the unit square.

4.3 Foraging

The task in this simulation is similar to the one described by Nolfi et al. (1994). An agent is placed on a two-dimensional grid-world (Figure 2). A number of food objects are randomly distributed in the environment. As its sensory input the agent receives the angle (relative to its current orientation) and distance to the nearest food object. The agent's output determines one of four possible actions: turn 90 degrees left, turn 90 degrees right, move forward one cell, or do nothing (stay). If the agent encounters a food object while navigating the environment, it consumes the food object. The agent's fitness is the number of food objects that were consumed during its lifetime. Each agent lives for 100 time steps

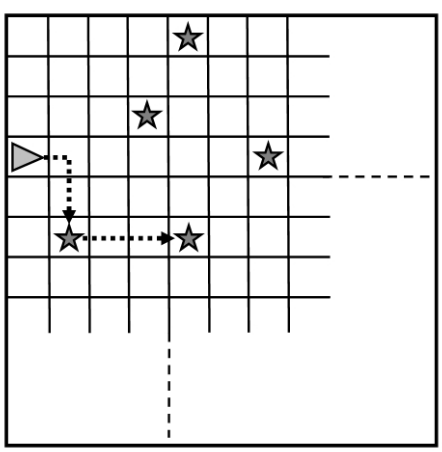

Figure 2: The foraging task: The agent (triangle) navigates in a 2D grid-world. Food objects (stars) are randomly distributed in the world. The agent can turn 90 degrees left, turn 90 degrees right, move one cell forward, or stay. Each time the agent encounters a food object, it consumes the food object and gains one fitness unit. Inspired by Nolfi and Floreano (1999)

in a 30x30 cells world which initially contains 30 food objects. A network topology of 2-6-2 was used (with an additional threshold unit in each layer).

In this task, unlike the previous ones, there is no explicit oracle we can use to train the agent. Nolfi et al. (1994) used available data to train the agent on the task of predicting the next sensory input, which differs, but is in some sense still "correlated" with that of finding food (the evolutionary task). In our model, we can still use the same mechanism of learning by imitation to train the agent on the original evolutionary task, using the best individuals in the population as teachers.

There are several strategies we can apply to determine which sensory input patterns should be used for training. Randomly selecting arbitrary input patterns, as we did in previous tasks, is not a suitable strategy here as the real input distribution that an agent encounters while navigating the environment may differ considerably from a uniform distribution. However, two behaviorally motivated strategies may be considered: a *query* model and an *observational* model. In the query model, the student agent navigates in the environment and for each sensory input pattern it encounters, the student queries the teacher to obtain the teacher's output for this pattern. The teacher's output is than used as the target output in back-propagation training of that pattern. In the observational model, the student "observes" the teacher agent as the teacher navigates in the environment and uses the sensory input patterns encountered by the teacher as training patterns (again, using the teacher's output for the back-propagation algorithm). Using this model we can further limit the observed patterns to those which occur during time steps that precede the event of finding food. This constraint will allow the student to imitate only useful behavioral patterns. We will label this strategy as *reinforced agent imitation (RAIL).*

5 Results

We first studied IEE in the two classification tasks described in Sections 4.1 and 4.2, where conventional supervised learning can still be applied. In these tasks we were able to compare the effects that both lifetime adaptation mechanisms (i.e., learning and imitation) have on the evolutionary process. The results clearly validate that the IEE model

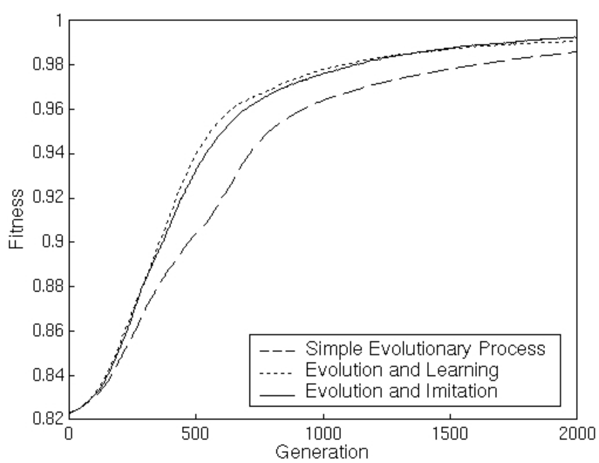

Figure 3: The triangle classification task: the innate fitness of the best individual in the population as a function of generation.

consistently yields an improved evolutionary process. The *innate fitness* of the best individuals in populations generated by applying learning by imitation is significantly higher than that produced by standard evolution.

Figure 3 illustrates the *innate* performances of the best agent as a function of generation, in populations evolved to solve the triangle classification problem (Section 4.2). To evaluate the agent's classification accuracy we use the Mean-Square Error (MSE) measure to calculate the distance between the network predicted classification and the true classification, averaged over all the patterns in the test set. Fitness is defined as $(1-Error)$. The results of a simple evolutionary process (dashed line) and of an evolutionary process that employs conventional supervised learning (dotted line) are compared with those of an evolutionary process that employs learning by imitation (solid line). Each curve represents the average result of 4 different simulation runs with different, randomly assigned, initial connection weights. The results presented in Figure 3 demonstrate how applying either of the learning paradigms yields better performing agents than those generated by a simple evolutionary process. In fact, applying learning by imitation produces practically the same improvement throughout the process as does conventional supervised learning.

When facing the 5-bit parity task, the effect of applying lifetime adaptation is even more surprising. Figure 4 illustrates the *innate* performances of the best agent as a function of generation, in populations evolved to solve the 5-bit parity problem. Each curve represents the average result of 10 different simulation runs with different, randomly assigned, initial connection weights. While simulations applying the IEE model still outperform the simple evolutionary process, using conventional supervised learning actually results with a significant decrease in performances. The problematic nature of this specific task may account for these poor results. The parity problem, although often used as a benchmark, is considered to be a difficult and atypical classification problem (Fahlman, 1989). Learning algorithms facing this task tend to get trapped in local minima. However, learning from an imperfect teacher, as is the case in learning by imitation, induces a certain level of noise into the learning process and may thus help to prevent the process from getting stuck.

Evidently, learning by imitation has a similar (if not superior) effect on the evolution-

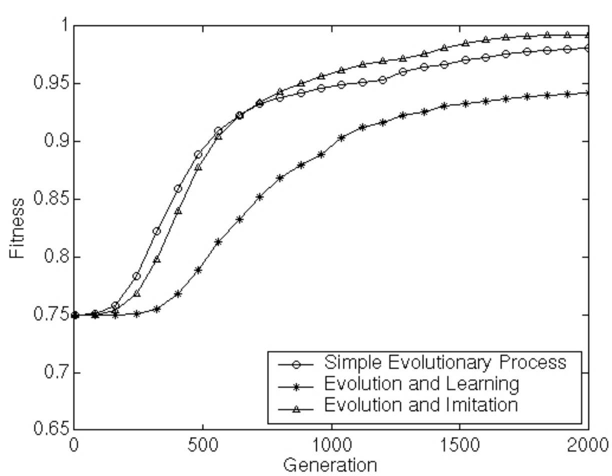

Figure 4: The 5-bit parity task: the innate fitness of the best individual in the population as a function of generation.

ary process to the one that was previously shown for conventional supervised learning. The knowledge possessed by the best members of the population can be used as alternative training data for other members, even in the early stages of the evolutionary process. We then turned to use IEE to enhance evolution where explicit training data is not available. This is the case in the foraging task described in Section 4.3.

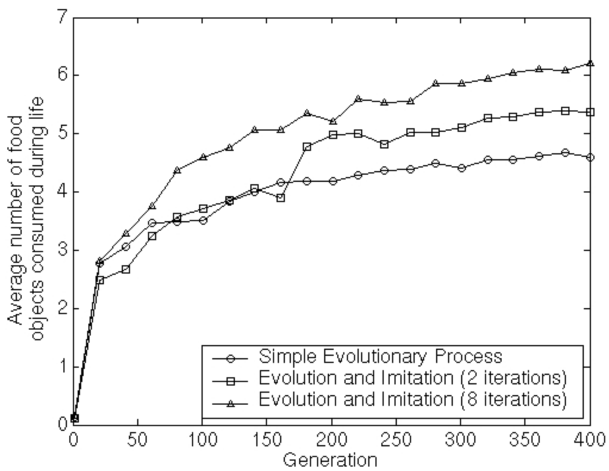

Figure 5: The foraging task: the average *innate fitness* of the population as a function of generation. The results of a simple evolutionary process are compared with those of simulations that employed lifetime imitation with two distinct adaptation forces (2 and 8 learning iterations).

Figure 5 illustrates the results of the simulations in which the agents faced the foraging task. The average *innate fitness* of the population in a simple evolutionary process is compared with the average *innate fitness* of populations that applied learning by imitation. The agents in this simulation employed the *RAIL* strategy of imitation. Fitness is measured as the number of food objects an agent consumes during its lifetime. Each curve represents the average result of 10 different simulation runs with different, randomly assigned, initial connection weights. As can be seen in Figure 5, autonomous agents produced by our model demonstrate better performances than those generated by the simple evolutionary process; that is, their *innate* capacity to find food in the environment is superior.

We also examined the effect of employing different *adaptation forces*. In our experiments, the adaptation force is implemented simply as the number of learning (backpropagation) iterations we apply in each lifetime adaptation phase. The results illustrated in Figure 5 also demonstrate that a higher adaptation force (i.e., a higher number of iterations in each imitation phase) further improves the performance of the resulting agents. This effect coincides with an analogous effect reported by Best (1999) where higher transmission force resulted with faster convergence of the evolutionary process.

To further explore the effects of lifetime imitation on evolution, we examined the improvement in fitness during lifetime as a function of generation. The improvement can be evaluated by calculating the difference between the *acquired fitness* and the *innate fitness* (i.e., $F_a - F_i$) in every generation. The results illustrated in Figure 6 clearly demonstrate that in very early stages of the evolutionary process, the best agents in the population already possess enough knowledge to improve the fitness of agents that imitate them. In fact, the contribution of imitative learning decreases as the evolutionary process

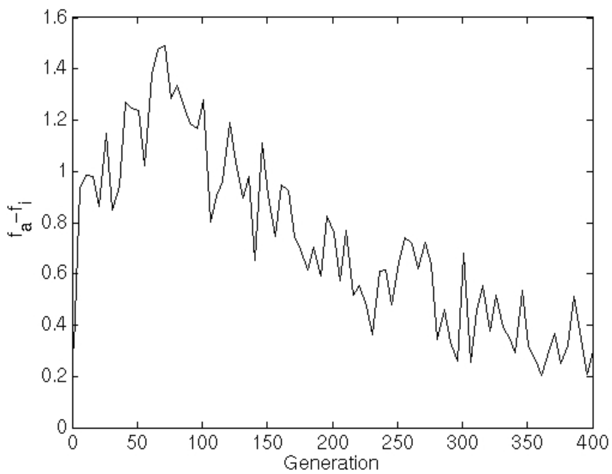

Figure 6: The foraging task: the improvement of the population average fitness gained by lifetime imitation as a function of generation.

proceeds, probably due to population convergence to high performance solutions.

An additional observation on the interaction between lifetime adaptation and evolution can be obtained from examining the diversity of the population throughout the evolutionary process. The average genome variance of the population, i.e., the variance among the population members, in the value of each gene (encoding a certain network weight) averaged over all genes, can serve as a measure of the population's diversity. As demonstrated in Figure 7, during the first few generations, the population's initial diversity decreases rapidly due to the selection pressure of the evolutionary process. However, throughout most of the following generations, the diversity found in populations subject to lifetime adaptation by imitation is higher than the diversity of populations undergoing a

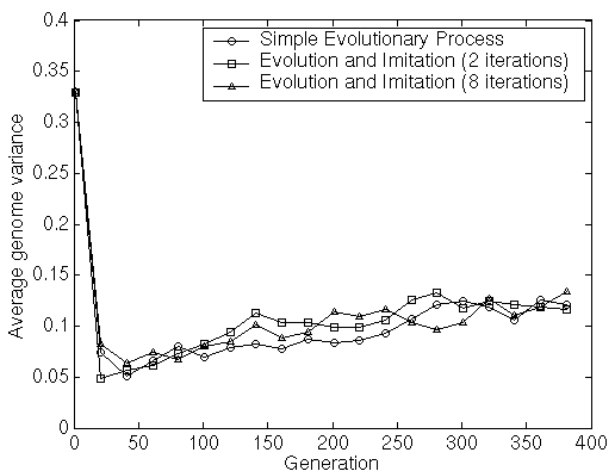

Figure 7: The foraging task: the average genome variance as a function of generation with and without imitation. Populations that employ lifetime adaptation, maintain a higher diversity throughout the evolutionary process.

simple evolutionary process. Allowing members of the population to improve their fitness through lifetime adaptation before natural selection takes place facilitates the survival of suboptimal individuals and helps to maintain a diversified population. This feature can partly account for the benefit gained by applying lifetime adaptation to agents evolution.

6 Discussion

This paper demonstrates how learning by imitation can be applied to an evolutionary process of a population of agents, utilizing the knowledge possessed by members of the population. Our IEE model proves to be a powerful tool that can successfully enhance evolutionary computation simulations in agents.

In our model, the agents' ability and incentive to imitate is assumed to be instinctive. Quoting Billard and Dautenhahn (1999), "our experiments address learning by imitation instead of learning to imitate". The imitation paradigm presented in this paper additionally assumes that the agents can estimate the fitness of their peers (i.e., more successful agents are larger and look healthier, etc.). More specifically, the RAIL strategy, where agents imitate only successful behavior, assumes that agents can detect significant changes in the fitness of their peers during their lifetime or identify specific activities that may contribute to their fitness. The model presented in Section 3 can provide a framework to explore ways in which these assumptions can be relaxed. Coding the imitative behavior patterns themselves into the genome might result in the spontaneous emergence of imitative behavior in a population of agents. Behavior patterns that can be coded may include attributes such as the imitation model selection scheme, imitation strategy, imitation period, etc. Our model can also be extended to study the incentive that should be provided to an agent to make it assume the role of a teacher. Teaching, or even allowing someone else to imitate one's actions is, by definition, an altruistic behavior, and might have various costs associated with it. We wish to explore the conditions which may lead to the emergence of active teaching even in the presence of a fitness penalty for such a behavior. Such favorable teaching conditions may arise when the fitness associated with various actions is correlated with the frequency of these actions in the population (see also Boyd and Richerson, 1985, for a discussion of frequency-dependent bias). A good example of this case can be found in the emergence of normative behaviors (Axelrod, 1986; Flentge et al., 2001). Since the IEE model presented here entails a relatively simple form of cultural transmission, confined to horizontal transmission of innate behaviors, it can serve as a solid testbed for future studies of the emergence, evolution and prevalence of imitation.

7 Summary

Our study focuses on the effects of imitation on the evolution of agents in the absence of cultural evolution. We show that introducing the adaptive mechanism of lifetime learning by imitation can significantly enhance the evolutionary processes, resulting in better performing agents. This paradigm is particulary useful in evolutionary simulations of autonomous agents, when conventional supervised learning is not possible. Our model can serve as a theoretical and experimental framework to further explore central issues concerning the interaction between imitation, learning and evolution.

Acknowledgements

We wish to thanks the anonymous reviewers for their helpful comments. We are grateful to Daniel Polani for his valuable insights.

References

Ackley, D. and Littman, M. (1994). Altruism in the evolution of communication. In Brooks, R. and Maes, P., editors, *Artificial Life IV: Proceedings of the International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press.

Arbib, M. (2002). The mirror system, imitation, and the evolution of language. In Dautenhahn, K. and Nehaniv, C., editors, *Imitation in Animals and Artifacts*. The MIT Press.

Axelrod, R. (1986). An evolutionary approach to norms. *American Political Science Review*, 80(4):1095–1111.

Baldwin, J. (1896). A new factor in evolution. *American Naturalist*, (30):441–451.

Best, M. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 7(3/4):289–306.

Billard, A. and Dautenhahn, K. (1999). Experiments in learning by imitation: grounding and use of communication in robotic agents. *Adaptive Behavior*, 7(3/4):411–434.

Boyd, R. and Richerson, P. (1985). *Culture and the evolutionary process*. The University of Chicago Press, Chicago.

Cavalli-Sforza, L. and Feldman, M. (1981). *Cultural transmission and evolution: a quantitative approach*. Princeton University Press.

Dautenhahn, K. and Nehaniv, C. (2002). *Imitation in Animals and Artifacts*. MIT Press, Cambridge, Mass., USA.

Fahlman, S. E. (1989). Faster-learning variations on back-propagation: An empirical study. In *Proceedings of the 1988 Connectionist Models Summer School*, Los Altos, CA. Morgan-Kaufmann.

Feldman, M. and Laland, K. (1996). Gene-culture coevolutionary theory. *Trends in Ecology and Evolution*, 11(11):453–457.

Flentge, F., Polani, D., and Uthmann, T. (2001). Modelling the emergence of possession norms using memes. *Journal of Artificial Societies and Social Simulation*, 4(4).

Floreano, D. and Mondada, F. (1996). Evolution of plastic neurocontrollers for situated agents. In Maes, P., Mataric, M., Mayer, J., Pollack, J., and Wilson, S., editors, *From Animals to Animates*, volume IV, Cambridge, MA. MIT Press.

Hayes, G. and Demiris, J. (1994). A robot controller using learning by imitation. In *Proceedings of the 2nd International Symposium on Intelligent Robotic Systems*.

Hertz, J., Krogh, A., and Palmer, R. (1991). *Introduction to the theory of neural computation*. Santa Fe Institute.

Hinton, G. and Nowlan, S. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.

Kawamura, S. (1963). The process of sub-culture propagation among japanese macaques. In Southwick, C. and Nostrand, V., editors, *Primates Social Behaviour*, pages 82–90, New York.

Kirby, S. and Hurford, J. (1997). Learning, culture and evolution in the origin of linguistic constraints. In Husbands, P. and Harvey, I., editors, *4th European Conference on Artificial Life*, volume IV, pages 493–502, Cambridge, MA. MIT Press.

Laland, K. (1992). A theoretical investigation of the role of social transmission in evolution. *Ethology and Sociobiology*, 13(2):87–113.

Mitchell, M. (1996). *An introduction to genetic algorithms*. MIT Press.

Nolfi, S., Elman, J., and Parisi, D. (1994). Learning and evolution in neural networks. *Adaptive Behavior*, 1(3):495–502.

Nolfi, S. and Floreano, D. (1999). Learning and evolution. *Autonomous Robots*, 7(1):89–113.

Nolfi, S. and Parisi, D. (1997). Learning to adapt to changing environment in evolving neural networks. *Adaptive Behavior*, 1:99–105.

Ruppin, E. (2002). Evolutionary autonomous agents: A neuroscience perspective. *Nature Reviews Neuroscience*, 3(2):132–141.

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., editors, *Advances in the Study of Behavior*, pages 239–283, San Diego, CA. Academic Press.

Zentall, T. (2001). Imitation in animals: evidence, function, and mechanisms. *Cybernetics and Systems*, 32(1-2):53–96.