

The Evolution of Imitation and Mirror Neurons in Adaptive Agents

Elhanan Borenstein*

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

Eytan Ruppin†

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

Abstract

Imitation is a highly complex cognitive process, involving vision, perception, representation, memory and motor control. The underlying mechanisms that give rise to imitative behavior have attracted a lot of attention in recent years and have been the subject of research in various disciplines, from neuroscience to animal behavior and human psychology. In particular, studies in monkeys and humans have discovered a neural mirror system that demonstrates an internal correlation between the representations of perceptual and motor functionalities. In contradistinction to previous engineering-based approaches, we focus on the evolutionary origins of imitation and present a novel framework for studying the evolution of imitative behavior. We successfully develop evolutionary adaptive agents that demonstrate imitative learning, facilitating a comprehensive study of the emerging underlying neural mechanisms. Interestingly, these agents are found to include a neural “mirror” device analogous to those identified in biological systems. Further analysis of these agents’ networks reveals complex dynamics, combining innate perceptual-motor coupling with acquired context-action associations, to accomplish the required task. These findings may suggest a universal and fundamental link between the ability to replicate the actions of other (imitation) and the capacity to represent and match others’ actions (mirroring).

1 Introduction

1.1 Imitation and Mirror Neurons

The past twenty years have seen a renewed interest in imitation in various fields of research (Prinz and Meltzoff, 2002) such as developmental psychology (Meltzoff, 1996), experimental studies of adult social cognition

(Bargh, 1997), and most relevant to our work, neurophysiology and neuropsychology (Rizzolatti et al., 1996, 2002). Research in this last field had led to the exciting discovery of *mirror neurons*. These neurons were originally found in the ventral premotor cortex (area F5) in monkeys, an area which is characterized by neurons that code goal-related motor acts (e.g. hand or mouth grasping). Some of the neurons in this area, which have been termed *mirror neurons*, discharge both when the monkey performs an action and when it observes another individual making a similar action (Gallese et al., 1996; Rizzolatti et al., 2002). Most mirror neurons exhibit a marked similarity in their response to action observation and execution, and in some cases this similarity is extremely strict (Rizzolatti et al., 2001). An analogous mechanism, whereby cortical motor regions are activated during movement observations was also demonstrated in humans using TMS (Fadiga et al., 1995), MEG (Hari et al., 1998), EEG (Cochin et al., 1998) and fMRI (Iacoboni et al., 1999; Buccino et al., 2001). Mirror neurons are thus the first identified neural mechanism that demonstrates a direct matching between the visual perception of an action and its execution. The ability to match the actions of self and other may have a functional role in fundamental cognitive processes, such as understanding the actions of others, language and mind reading (Rizzolatti et al., 2001). In particular, imitation of motor skills requires the capacity to match the visual perception of a demonstrator’s action to the execution of a motor command. The neural mirror system, demonstrating such an internal correlation between the representations of perceptual and motor functionalities, may form one of the underlying mechanisms of imitative ability.

1.2 Context-Based Imitation

Learning by imitation, like any cognitive process, must be considered an intrinsically embodied process, wherein the interaction between the neural system, the body and the environment cannot be ignored (Keijzer, 2002; Dautenhahn and Nehaniv, 2002a). In particular, every action, either observed or performed, occurs within a certain *context*. A context can represent the time or place in which the action is made, various properties of the environment, the state of the individual performing the action or the social interaction partners (see, for example, Dautenhahn, 1995). Clearly, there is no sense in learning a novel behavior by imitating another’s actions if you do not know the context in which these actions are made – a certain action can be extremely beneficial in one context, but have no effect (or even be deleterious) in a different context. Discussing an agent-based perspective on imitation, Dautenhahn and Nehaniv (2002a) consider the problem of imitating the right behavior in the appropriate context, i.e., “when to imitate”, as one of the five central questions (“Big Five”) in designing experiments and research on imitation. We hence use the term *context-based imitation* in the sense

of being able to reproduce another’s observed action whenever the context in which the action was originally observed, recurs.¹ For example, an infant observing his parents may learn by imitation to pick up the phone (*action*) whenever the phone is ringing (*context*).

Context-based imitation can thus be conceived as constructing a set of associations from contexts to actions, based on observations of a demonstrator performing different actions within various contexts. These associations should comply with those that govern the demonstrator’s behavior, and should be learned (memorized) so that each context stimulates the production of the proper motor action even when the demonstrator is no longer visible. It should be noted however, that “*action*” is an abstract notion, and in reality, an imitating individual (agent) should also be capable of matching a *visual perception* of the demonstrator’s action to the corresponding *motor command* that activates this action.² The key objective of this study is to gain a comprehensive understanding of the mechanisms that govern such context-based imitative learning and to examine the nature of the associations between visual perception, motor control and contexts that are being formed in the process.

1.3 Evolving Imitating Agents

Imitation is an effective and robust way to learn new traits by utilizing the knowledge already possessed by others and it has already been applied by researchers in the fields of artificial intelligence and robotics. 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. Borenstein and Ruppin (2003) employed learning by imitation to enhance the evolutionary process of autonomous agents. For an up-to-date introduction to work on imitation in both animals and artifacts see the cross-disciplinary collection (Dautenhahn and Nehaniv, 2002b). Furthermore, some researchers, motivated by the recent discovery of a neural mirror system, have implemented various models for imitative learning, employing neurophysiologically inspired mechanisms. Billard (2000) presented a model of a biologically inspired connectionist architecture for learning motor skills by imitation. The architecture was validated through a mechanical simulation of two humanoid avatars, learning several types of movements sequences. Demiris and Hayes (2002) and Demiris and Johnson (2003) developed a mirror-neuron based computational architecture of imitation inspired by Meltzoff’s Active Intermodal Matching mechanism (Meltzoff

¹Animal behavior and human psychology literature introduces a wide range of definitions of imitation, focusing on what can constitute true imitation vs. other forms of social learning (Zentall, 2001; Nehaniv and Dautenhahn, 2002). Our definition addresses the importance of the observed action’s *context* for a successful behavior.

²In this study we focus on visually based imitation. However, it should be noted that other forms of imitation, such as vocal imitation, need not involve visual modality (see, for example, Nehaniv and Dautenhahn, 2002; Herman, 2002).

and Moore, 1997) and combined it with an “active” distributed imitation architecture. They have demonstrated that this dual-route architecture is capable of imitating and acquiring a variety of movements including unknown, partially known, and fully known sequences of movements. Oztop and Arbib (2002), focusing on the grasp-related mirror system, argued that mirror neurons first evolved to provide visual feedback on one’s own “handstate” and were later generalized to understanding the actions of others. They have conducted a range of simulation experiments, based on a schema design implementation of that system, providing both a high-level view of the mirror system and interesting predictions for future neurophysiological testing. Other researchers (Marom et al., 2002; Kozima et al., 2002) claimed that the mirror system structure can be acquired during life through interaction with the physical or social environment and demonstrated models whereby perceptual and motor associations are built up from experience during a learning phase.

The studies cited above, however, assume that the agents’ basic ability and motivation to imitate are innate, explicitly introducing the underlying functionality, structure or dynamics of the imitation mechanism into the experimental system. **In contrast to this engineering-based approach, we wish to study the neuronal mechanisms and processes underlying imitation from an evolutionary standpoint, and to demonstrate how imitative learning *per se* can evolve and prevail.** Evolutionary autonomous agents form an intuitively appealing approach for modelling and studying the evolution of biological neural mechanisms (Ruppin, 2002). Using a simulated environment, wherein agents *evolve* to perform a simple imitative task, facilitates a thorough examination of the resulting mechanism in “ideal conditions”: Full control of the environment and experimental setup, and perfect knowledge of the agents’ behavior and neural dynamics. Clearly, acknowledging the evolutionary origins of imitation and examining the emerging (rather than engineered) device can shed light on the common *fundamental* principles that give rise to imitative behavior. It is important to note, however, that our key goal in this model is *not* to simulate the neural mechanism that underlie imitative behavior in the human or primate brain nor to incorporate the full range of social skills required for imitative learning (e.g. extraction of the context from the environment or coping with a different embodiment). The model described in this paper is clearly a simplified conceptual model and does not presume to encapsulate many of the well established biological and neuronal data on imitation. Rather, the aim of such evolutionary autonomous agents model is to examine *generic* and *universal* properties of complex living systems (the “*life as it could be*” paradigm (Langton, 1995)). *The key point in this study is thus to examine the emerging characteristics of the mechanism evolved to support imitation in a system where no constraints on the underlying mechanisms or representations were explicitly encoded.*

In this study, we thus set out to pursue two objectives: **We first present a novel experimental framework for evolving context-based imitative learning in evolutionary adaptive autonomous agents** (Ruppin, 2002; Floreano and Urzelai, 2000). We demonstrate the evolution of imitating agents that comprise a simple mechanism of imitative behavior. **We then turn to systematically analyze the structure and dynamics of the resulting neurocontrollers.** This analysis reveals neural devices analogous to those found in biological systems, including clear examples of internal coupling between observed and executed actions. Further analysis of the network adaptation dynamics reveals a hybrid mechanism, combining innate perceptual-motor coupling with acquired context-action associations. We conclude with a discussion of the implications of our findings for imitation theory and a description of future work.

2 The Experimental Setup

2.1 The Environment

The agents in our simulation inhabit a world that can be in one of several *world states* $\{s_1, s_2, \dots, s_n\}$. In each time step, the world state is randomly selected from $\{s_1, s_2, \dots, s_n\}$ with a uniform distribution. These states can represent, for example, the presence of certain food items or the size of an observed object and hence form the *context* in which actions are observed and performed. The world state, however, is not visible in every time step and is seen (i.e. included in the agent’s sensory input) only in 60% of the time steps. An additional set, $\{a_1, a_2, \dots, a_m\}$, represents the repertoire of motor *actions* that can be performed by the agent or by the demonstrator. A *state-action injective mapping* is also defined, assigning a certain action as the proper action for each world state s_i . Within the simulations described below, both n and m are set to 4, allowing $4! = 24$ different state-action mappings. Regularly performing the proper action assigned to the current state of the world is deemed a successful behavior and confers a positive fitness. Similarly, when the world state is not visible, a successful agent should not perform any action. It is assumed that the environment is also inhabited by a demonstrator (teacher), successfully performing the proper action in each time step. The demonstrator’s action is visible (i.e. included in the agent’s sensory input) only in 20% of the time steps. The partial visibility of the world state and demonstrator ensures that during the agent’s life it will encounter both scenarios wherein the demonstrator is not visible, forcing the agent to “memorizing” the proper state-action mapping, and scenarios wherein the world state is not visible, in which a successful agent should “observe” the demonstrator’s action but not perform any action. The specific visibility values defined above have no

significant effect on the resulting agent, but rather provide a good blend of the various visibility scenarios during the agents' life, facilitating the examination of the agents' neurocontroller in these scenarios. **Furthermore, the above mapping, from world states to actions, is randomly selected anew in the beginning of each agent's run in the world.** The motivation for this state-action mapping shuffle is twofold. First, it prevents such a mapping from becoming genetically determined. To demonstrate a successful behavior, agents must *learn* the proper mapping by observing the demonstrator, promoting an imitation based mechanism to evolve. Second, it represents a scenario of a changing environment, wherein novel world states appear over time (new food sources, other species, etc.), making prior state-action mappings obsolete.

2.2 The Agent

Figure 1 illustrates the structure of the agent's sensorimotor system and neurocontroller. The agent's sensory input in each time step comprises 8 binary values, including the current world state (if visible) and a 4-cell retinal "image" of the demonstrator's action (if visible). The retinal image is determined according to a predefined mapping from actions to retinal binary patterns which remains fixed throughout the simulation.³ In time steps wherein the world state or demonstrator are not visible, the corresponding input neurons are set to 0. Each of the agent's output neurons represents a motor action command, determining which actions (if any) will be executed by the agent. The output neurons (as well as the hidden neurons) are continuous neurons ranging from 0 to 1, and can thus be perceived as indicating the probability of activating each motor action. A successful agent should thus produce in each time step an activation level close or equal to 1 in the motor neuron that corresponds to the appropriate action, and values close or equal to 0 in the rest of the motor neurons. In time steps where the world state is not visible (and thus, no action should be performed by the agent), a successful agent should produce activity level close or equal to 0 in all motor neurons.

Considering the agent's task and the environment it inhabits, the architecture of the agent's neurocontroller should encompass several characteristics. Clearly, it should be capable of acquiring new behaviors during the agent's life to allow imitative learning. However, to support complex dynamics which may employ both a fixed component and a learned behavior, the neurocontroller should also allow a combination of innate and acquired elements. Moreover, the precise blend of innate and acquired properties should be determined through genetic evolution. An interesting architecture that satisfies these requirements has been proposed by Floreano

³The selected retinal representation is of no specific significance, however, we use the representation illustrated in Figure 1 (wherein each action is represented by a multi-bit configuration) rather than a trivial one (wherein each action is represented by a single bit) to examine the emergence of internal localized representation of complex input patterns.

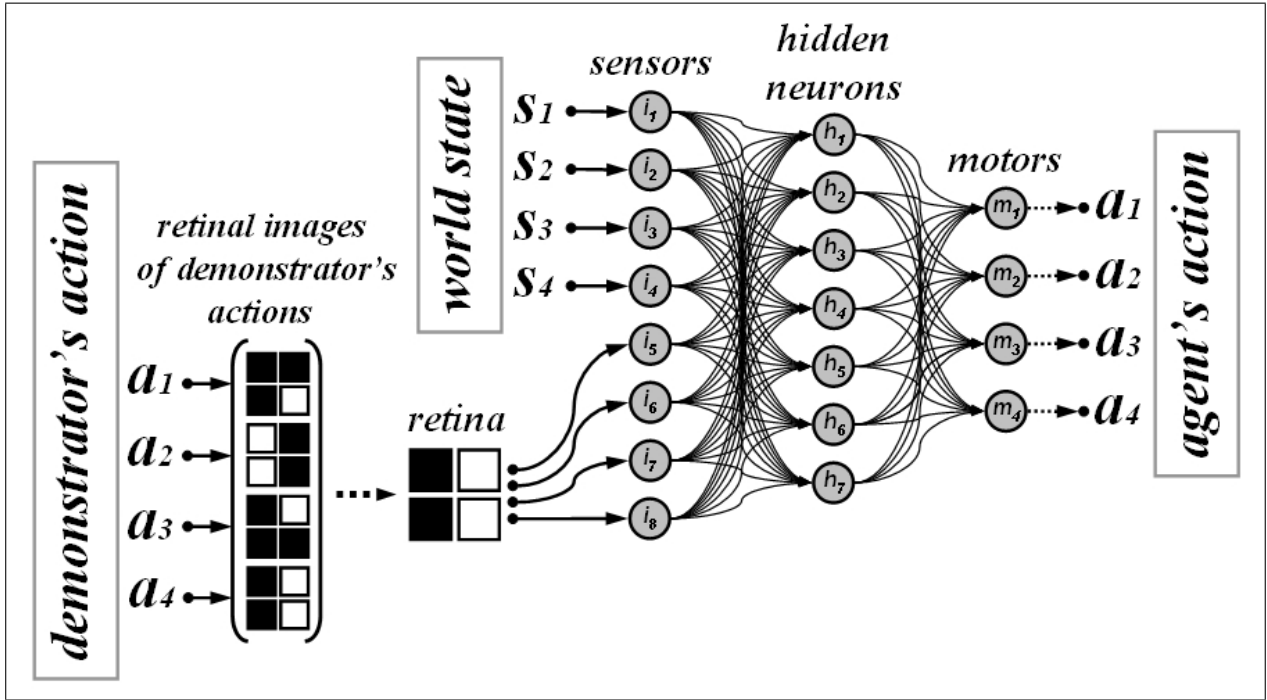


Figure 1: The agent's sensorimotor system and neurocontroller. The sensory input is binary and includes the current world state and a retinal "image" of the demonstrator's action (when visible). The retinal image for each possible demonstrator's action and a retinal input example for action a_4 are illustrated. The motor output determines which actions are executed by the agent. The network synapses are adaptive and their connection strength may change during life according to the specified learning rules.

and Urzelai (2000), and is applied with a few modifications in the model described below.

Each agent employs a simple feed-forward neural network as a neurocontroller (i.e. the agent cannot perceive its own actions). These networks however are *adaptive*, whereby the genotype of each individual encodes not only the initial synaptic weights but also a *Hebbian learning rule* and *learning rate* for each synapse (Floreano and Urzelai, 2000). In particular, each synapse in the network, (i, j) , connecting neuron j to neuron i , is encoded by 4 genes, defining the following properties:

- (i) w_{ij}^0 - the initial connection strength of the synapse (real value in the range $[0, 1]$).
- (ii) s_{ij} - the connection sign (1 or -1).
- (iii) η_{ij} - the learning rate (real value in the range $[0, 1]$).
- (iv) Δw_{ij} - the learning rule applied to this synapse.

Each synaptic weight w_{ij} is initialized with w_{ij}^0 at the beginning of the agent's life and is updated after every time step (a sensory-motor cycle) according to:

$$w_{ij}^t = w_{ij}^{t-1} + \eta_{ij} \Delta w_{ij} .$$

Δw_{ij} encodes one of five learning (modification) rules (here, o_j and o_i denote the activity of the presynaptic neuron and postsynaptic neuron respectively):

- (1) No learning: $\Delta w_{ij} = 0$.
- (2) Plain Hebb rule: $\Delta w_{ij} = (1 - w_{ij}) o_j o_i$.
- (3) Postsynaptic rule: $\Delta w_{ij} = w_{ij}(-1 + o_j) o_i + (1 - w_{ij}) o_j o_i$.
- (4) Presynaptic rule: $\Delta w_{ij} = w_{ij} o_j(-1 + o_i) + (1 - w_{ij}) o_j o_i$.
- (5) Covariance rule:

$$\Delta w_{ij} = \begin{cases} (1 - w_{ij}) F(o_j, o_i) & \text{if } F(o_j, o_i) > 0 \\ (w_{ij}) F(o_j, o_i) & \text{otherwise} \end{cases}$$

where $F(o_j, o_i) = \tanh(4(1 - |o_j - o_i|) - 2)$.

These rules have been selected based on neurophysiological findings (i.e. they encapsulate some of the common mechanisms of local synaptic adaptation found in biological nervous systems) and were modified to satisfy some computational constraints (e.g. in this adaptation process synapses cannot change sign and their strength is kept in the range $[0, 1]$). For a detailed description of these adaptation dynamics see Floreano and Urzelai (2000). The network topology is static throughout the process and for the purpose of our simulation was set to 8-7-4 (i.e., 8 input neurons, a hidden layer with 7 neurons, and 4 output neurons), with an additional threshold unit in each layer. Such *evolutionary adaptive autonomous agents*, inspired by those presented in Todd and Miller (1991) and Floreano and Urzelai (2000), demonstrate a learning process that is supervised only indirectly, through natural selection.

2.3 The Evolutionary Process

A population of the agents described above evolve to successfully behave in the environment. Each agent lives in the world for 500 time steps. Fitness is evaluated according to the agent's success in performing the proper

action assigned to the current world state (i.e. activating only the appropriate motor neuron), according to the state-action mapping, in each time step. An agent should perform an action only if the world state is visible and regardless of the demonstrator’s visibility. We use the Mean-Square Error (MSE) measure to calculate the distance between the agent’s motor output (continuous values ranging from 0 to 1) and the desired output (a value of 1 for the appropriate motor neuron and 0 for the rest), averaged over the agent’s life. A MSE value of 0 thus indicates a perfectly behaving agent. The agent performance during the first 100 time steps is not evaluated (infancy phase). Fitness value is then calculated as $(1 - MSE)$ and averaged over 20 trial runs in the world.

The initial population is composed of 200 individuals, each assigned a randomly selected haploid genome (i.e. each individual holds one copy of the genome), encoding the initial connection weights, learning rules and learning rates. Each new generation is created by randomly selecting agents from the previous generation and allowing them to reproduce. Agents are selected according to their fitness, using linear scaling and a roulette wheel selection scheme (Mitchell, 1996). During reproduction, 2% of the genes are mutated. Connection strength genes and learning rate genes are mutated by adding a randomly selected value from the interval $[-0.3, 0.3]$, connection sign genes are mutated by flipping the sign and learning rule genes are mutated by randomly selecting one of the available rules. The genomes of the top 20% of individuals are copied to the next generation without mutation. Variations in these parameter values have no significant effect on the resulting agents.

3 Results

3.1 The Evolution of Imitation

Within the settings described in the previous section the proper action assigned to each world state is randomly selected anew at the beginning of the agent’s life. The appropriate state-action associations can thus be inferred only from the demonstrator’s observed actions. Agents cannot rely on genetically coded behavior and must incorporate some sort of imitation-based learning strategy in order to demonstrate a successful behavior. Although no such learning strategy was explicitly introduced into the system, examining the fitness of the best agent in the population as a function of generation clearly demonstrates that such imitating agents have evolved (Figure 2). Evidently, after approximately 2000 generations, the evolved agents successfully master the behavioral task, regularly executing the proper action in each world state.

Having successfully evolved imitating agents, we turned to examine the structure, dynamics and neural

mechanisms that these agents employ. We have performed numerous evolutionary simulation runs, of which approximately half resulted in near-optimal imitating agents (exhibiting an evolutionary dynamics similar to those shown in Figure 2). Unsuccessful simulation runs seemed to stem from early convergence of the population to sub-optimal solutions (wherein agents did not produce a distinct motor action in each time step). In the remainder of this paper we focus on analyzing one such successful agent – the best agent in the last generation of a specific evolutionary simulation run. Other successful agents, from various evolutionary runs, were analyzed and demonstrated similar dynamics.

Direct evidence of the agent’s successful imitative behavior and the resulting learning dynamics are demonstrated in Figure 3, depicting the activity of one of the motor neurons (m_2) in different states of the world. In this specific simulation run, the state-action mapping was arbitrarily set so that a_2 is the proper action in world state s_4 and not in any other state. In the beginning of its life, the agent activates motor m_2 (i.e., performs action a_2) whenever the world state is visible. However, after only a few demonstrations of the appropriate behavior, the proper state-action mapping is learned and this motor is activated only when the world state is s_4 , as expected. In fact, as demonstrated in Figure 4, the ability to learn by imitation the appropriate state-action

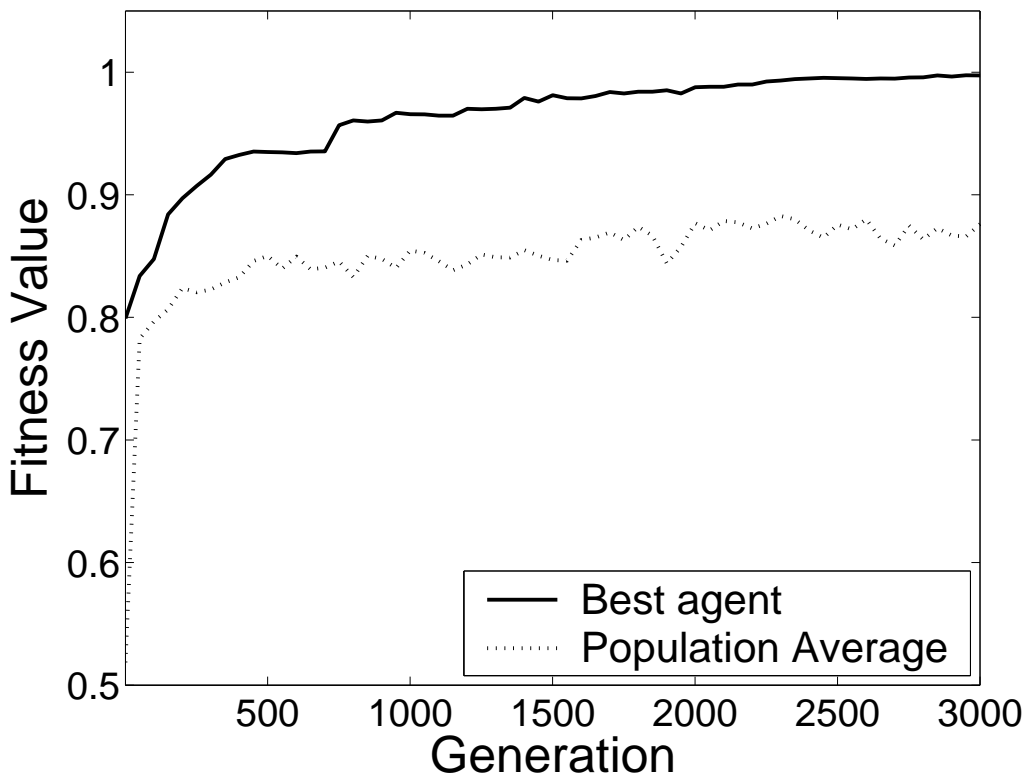


Figure 2: The fitness of the best agent in the population and the population average fitness as a function of generation.




Figure 3: The activation level of one motor neuron (m_2) during the first 150 time steps. The different shapes indicate whether the world state was s_4 and whether it was visible. The triangles at the bottom further represent time steps in which the demonstrator was visible.

mapping remains active during the agent's life, allowing the agent to learn a new mapping when necessary. In this experiment, the state-action mapping was initially set, as before, so that a_2 is the proper action in world state s_4 . However, in the middle of the agent's life (time step 250) the state-action mapping, and accordingly the demonstrator's behavior, was changed so that a_2 is the proper action in world state s_2 . Evidently, although the agent learned a certain mapping in the beginning of its life, it can quickly adapt to a new mapping after observing a few demonstrations of the new appropriate behavior.




Figure 4: The activation level of motor neuron m_2 during the agent's life, demonstrating the agent's ability to learn new behaviors. In this simulation run the state-action mapping was modified in step 250, making a_2 the proper action in world state s_2 rather than s_4 as it was initially set. The triangles at the bottom further represent time steps in which the demonstrator was visible.

3.2 The Emergence of Mirror Neurons

Examining the network hidden layer reveals an interesting phenomenon with regard to the internal representation of actions. As stated above, to support imitative learning, wherein associations from contexts to motor commands should be inferred from observations of the demonstrator’s actions, an agent should be capable of matching the visual perception of an observed action to the motor command that generates the corresponding action. Figure 5, depicting the activation level of 3 hidden neurons, attests to the emergence of such inherent perceptual-motor coupling. Apparently, various neurons in the hidden layer are active both when the agent performs a certain action and when it observes the demonstrator making a similar action, **forming internal mirror neurons analogous to those found in biological systems**. For the purpose of this study, we thus define mirror neurons as neurons that show a neural activation level significantly higher than 0 for both observation and execution of a certain action, and are not active in any other scenario. Although other definitions may be applied, the above definition forms a suitable analogy to the characteristics of biological neural mirroring. Interestingly, as seen in Figure 5, the activation level of mirror neurons during action observation is typically lower than the activation level during action execution. An analogous phenomenon can also be detected in neuronal recording data in the literature, and should be further investigated. However, in our simulation, the relatively small number of hidden neurons and mainly, the feed-forward nature of the network may account for this phenomenon, forcing mirror neurons to participate also in motor excitation.⁴ These constraining properties of the artificial network, a direct consequence of several computational limitations, may also induce some constraints on the biological implications of this model, including, for example, the lack of clear distinction between active and passive perception. Such mirror neurons were found in most of the agents that evolved in our simulation environment. However, typically, not all actions in the repertoire were associated with a corresponding mirror neuron, and there have been a few cases where successful agents did not seem to incorporate any clear neural mirroring matching our above definition. There was also no evident correlation between the initial conditions or the simulation parameters and the emergence of mirror neurons.

The functional characteristics of the emerging mirror neurons were further examined through a set of intervention experiments, wherein hidden neurons were externally activated (stimulated) or inactivated (lesioned). These experiments confirmed that the detected mirror neurons convey the required information about the action to be performed. For example, when the world state is not visible (a scenario that would usually result

⁴Furthermore, the relatively small number of hidden neurons may form a bottleneck that promotes the use of these neurons for both action perception and action execution and consequently the formation of mirror neurons. However, the fact that the same single neuron is activated in the observation and activation of the same specific action, the essence of mirroring, is surprising.

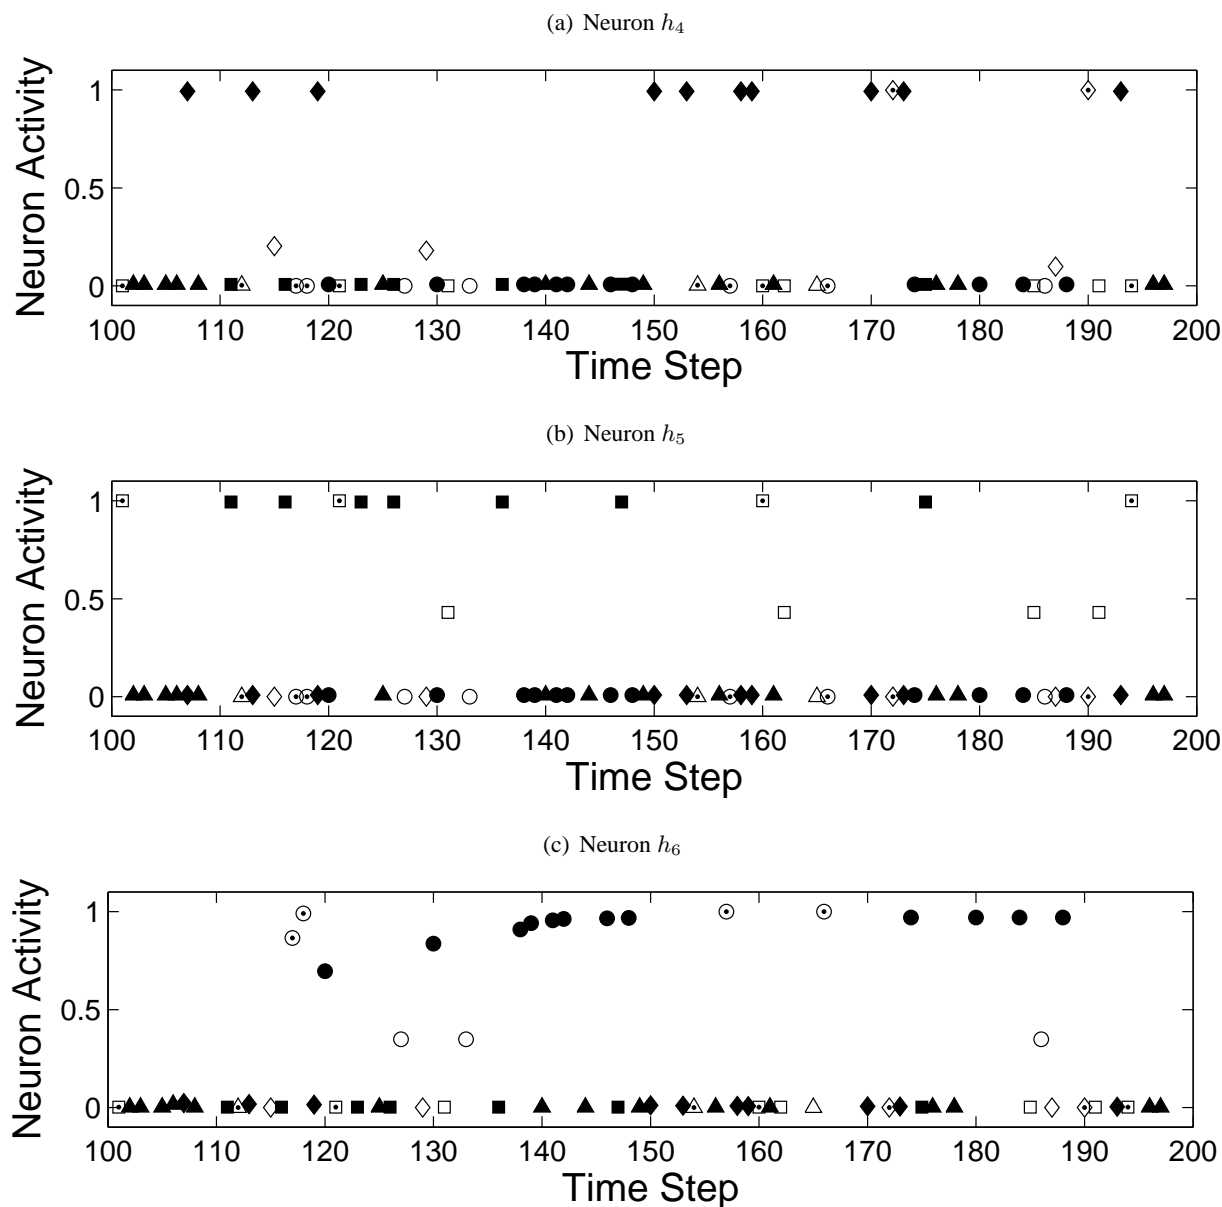


Figure 5: The activation level of 3 hidden neurons (h_4 , h_5 and h_6) during time steps 100-200 with an indication of the executed or observed action. Circles, squares, diamonds and triangles represent actions a_1 , a_2 , a_3 , a_4 respectively. An empty shape indicates that the action was only observed but not executed, a filled shape indicates that the action was executed by the agent (stimulated by a visible world state) but not observed, and a dotted shape indicates time steps in which the action was both observed and executed.

in no action being performed) an ‘artificial’ stimulus of a mirror neuron resulted in the agent’s performance of the action associated with that mirror neuron. Similarly, inactivating a mirror neuron inhibits the production of the associated action and in some cases resulted in the production of the wrong action.⁵ Furthermore, ap-

⁵Recent reversible inactivation studies (Fogassi et al., 2001) demonstrated a distinction between two sectors in area F5 in monkeys: Mirror neurons are located in sector *F5 convexity*. Canonical neurons (neurons that respond to the presentation of three-dimensional objects of different size and shape) are located in sector *F5 bank*. While inactivation of area F5 bank produced a severe deficit of the

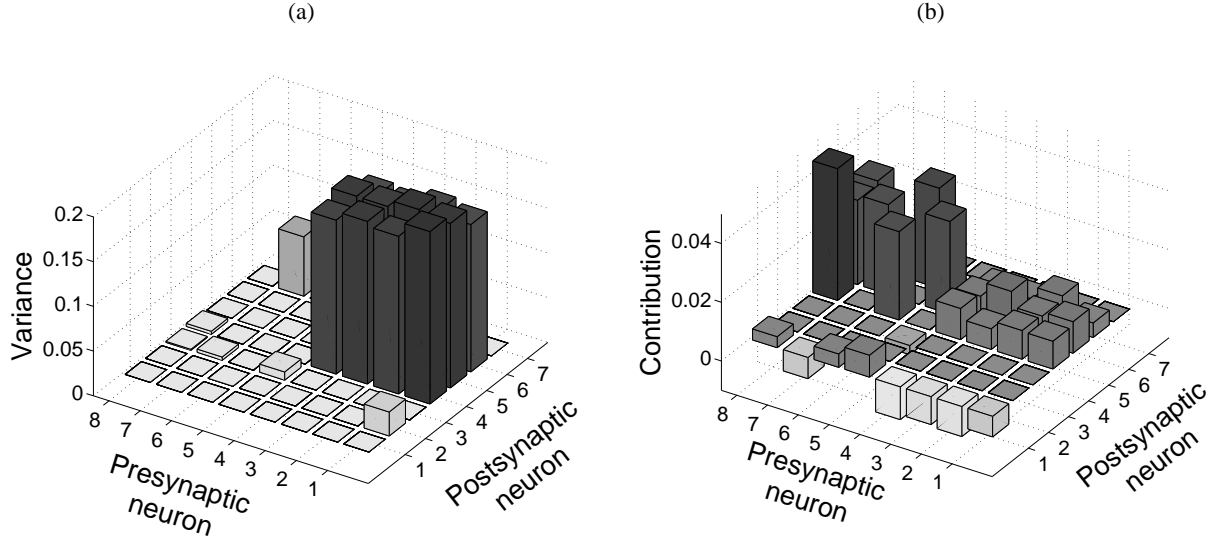


Figure 6: An illustration of the connection strength variance (a) and the overall contribution (b) of the synapses connecting the sensory input layer (presynaptic) to the hidden layer (postsynaptic). Neurons 1-4 of the presynaptic input layer represent the world state while neurons 5-8 are the retinal neurons, representing the observed demonstrator’s action. Neurons 4-6 of the hidden postsynaptic layer have been identified as mirror neurons.

plying *multiple* neurons activation/inactivation settings, it has been shown that even actions that could not be associated with a fully localized representation (i.e. a single mirror neuron) are still represented in the hidden layer through a distributed configuration of neurons. These findings also account for the cases mentioned above wherein successful agents did not seem to incorporate any clear localized mirror neurons.

3.3 The Developmental Dynamics

We finally turn to examine the ontogenic, developmental aspects of the resulting neurocontroller. Our main objective is to identify which components in the neural mechanism are innate and which are acquired during the agent’s life. We first determine which synapses play a significant role in the *learning* process. Clearly, variation in the synapse strength during life or the genetically coded learning rate are not appropriate indicators as they cannot differentiate between learning processes that genuinely adapt the agent to the world and unrelated self-organization processes. We thus measure the variance in the connection strength *at the end of the agent’s life across 1000 simulation runs* (i.e. the particular agent that was analyzed above, living 1000 different lifetimes). A low variance value indicates that the synapse dynamics are independent of the world characteristics (e.g. the state-action mapping), and thus cannot contribute to the learning process that adapt the agent to the world. As

required actions, inactivation of the cortical convexity determined only a motor slowing, preserving the appropriate action production. Clearly, within our simple model, such distinction between canonical and mirror neurons could not have developed and the mirror neurons that have emerged play a crucial role in the visuomotor pathway.

demonstrated in Figure 6a, this measure highlights the acquired nature of the synapses connecting the world state neurons (input neurons 1-4), with the mirror neurons we have identified (hidden neurons 4-6). **Clearly, the acquired state-action associations are induced by these synapses.** The markedly lower variance values in other synapses from this layer and in synapses connecting hidden layer neurons to motor neurons (not illustrated here), suggest that these synapses do not play an important part in the learning process and encompass the innate properties of the network. We then turned to determining the overall contribution of each synapses to the agent’s successful behavior, either learned or innate. Examining the effect of numerous multiple lesion configurations, we have utilized the Multi-perturbation Shapley value Analysis (MSA), a rigorous way to determine the importance of system elements (Keinan et al., 2004). In each configuration, a set of synapses are cancelled out by setting both their initial strength and learning rate to 0. The resulting contribution of each synapse connecting the input layer to the hidden layer is illustrated in Figure 6b. Evidently, the synapses that have been identified above as participating in the learning process possess a non-negligible contribution value. However, the most important synapses are among those connecting the retinal neurons (input neurons 5-8), representing the observed action, with the mirror neurons (hidden neurons 4-6). **These connections manifest the strong innate associations between the visual perception of observed actions and the internal representation of these actions, developed during the evolutionary process.**

Based on the findings described above, a simple model of the mechanism that evolved in our settings to support imitative behavior can be inferred (Figure 7). Notably, the required perceptual-motor coupling was not explicitly engineered into the agents, but rather emerged through evolution as an *innate* property. Furthermore, to support an effective mechanism of imitation, visually perceived actions are linked to the corresponding motor commands via fully localized internal elements, representing each action, in the form of mirror neurons. The acquired context-action stimuli can then be constructed through a simple mechanism of Hebbian learning without external supervision or reinforcement signals.

4 Discussion

This study presents an experimental framework for studying the evolution and dynamics of imitation in evolutionary autonomous agents. This framework provides a fully accessible, distilled model for imitation and can serve as a vehicle to study the mechanisms that underlie imitation in biological systems. As stated in Section 1.3, our experimental setup employs a simplified model that is not presumed to encapsulate many of the well estab-

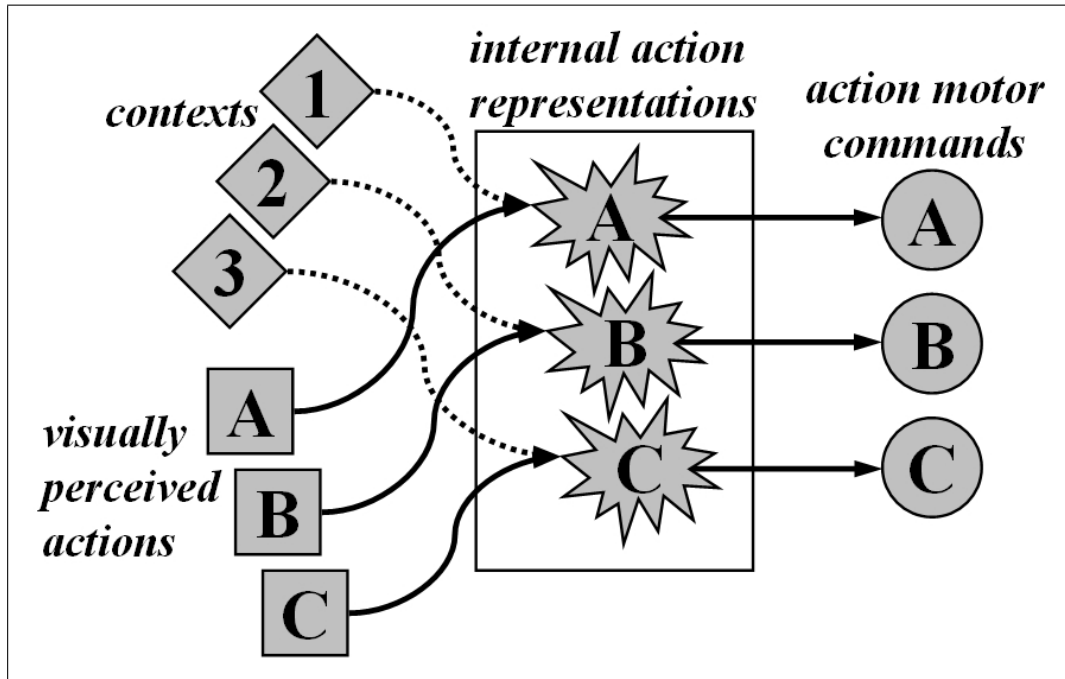


Figure 7: A simple model of context-based imitation. Solid arrows represents innate associations, while dashed arrows represents associations that are acquired during the agent’s life via Hebbian learning.

lished biological and neuronal data on imitation, nor to simulate a fully realistic social learning scenario. Rather, the aim of this model is to examine the *generic* and *universal* properties of imitative learning mechanisms. Our confidence in this framework is based on two observations: First, being an evolutionary developed mechanism, rather than an engineered one, we believe it is likely to share the same fundamental principles driving natural systems. Second, our analysis of the resulting mechanism reveals phenomena analogous to those found in biological neural mechanisms.

The key point in our findings is that while creating a system in which only the evolution of imitation is solicited, a neural mirroring system had emerged. That is, even though no constraints on the underlying mechanisms or representations were explicitly encoded into the system, such mirror neurons have been demonstrated. These findings imply a **fundamental and essential link between the ability to imitate and a mirror system.** In fact, in this regard, we believe that the simplicity of our model is one of its key assets: The emergence of neuronal mirroring to support imitation even in such a simple model, may suggests a universal and fundamental link between the ability to replicate the actions of other (imitation) and the capacity to represent and match others’ actions (mirroring).

It is also important to note, that although it has been hypothesized that mirror neurons underlie imitative learning functionality, the precise role of the mirror system remains unknown (Rizzolatti et al., 2001). The

linkage between imitation and mirroring demonstrated in our study corroborates this hypothesis and may prove to be interesting for understanding the mechanisms that give rise to social cognitive skills. Moreover, the mirror neurons that emerged in our model, being a clear instance of shared internal representation between observed and executed actions, also provide interesting insights that may be applied to artificial intelligence and robotic research. Although the use of internal representation is prevalent in engineered systems, the existence of such a representation in evolved systems has been challenged (Cliff and Noble, 1997). The model presented in this paper, promoting the use of observed actions of “*others*” for learning proper motor actions of “*self*”, provides a simple example of evolved internal representation.

Clearly, the simple model presented in this paper cannot account for the full range of imitative behaviors found in nature (e.g. recognition of novel or compound actions). However, focusing on low-level, innate imitation, this model addresses the essential questions concerning the mechanism underlying imitative behavior. It successfully demonstrates how the required associations between perceived actions, motor commands and contexts can be constructed within a hybrid adaptation process, combining evolution and lifetime learning.

The framework presented in this paper can be further enhanced to examine central issues concerning the development of imitation in animals and artifacts and the functional role of the neural mirror system. We wish to use this basic model to determine the physical and social environmental conditions that promote the emergence of mirror neurons. In particular, our framework can be enhanced to simulate a more realistic scenario of social learning. For example, we wish to examine how an extension of the agent’s sensory input, and a complex social environment inhabited by demonstrators with varying levels of success, affect the resulting imitation strategy. Questions concerning the dependencies between observed and executed actions and the formation of neural mirroring are especially of great interest: How will the representation of actions that cannot be executed by the observer (e.g. due to different embodiment) differ from those of imitated actions? How will a hierarchical repertoire of actions affect the emerging representation? Another intriguing possibility would be to utilize this framework to explore the role of mirror neurons in the evolution of communication (Rizzolatti and Arbib, 1998; Arbib, 2002) and in predicting the actions of others (Ramnani and Miall, 2004). We hope that further extensions of this basic model will allow us to obtain testable predictions regarding imitative behavior in humans and primates, and shed new light on some of the key issues concerning perception, mirroring and cognition.

Acknowledgements

E.B. is supported by the Yeshaya Horowitz Association through the Center of Complexity Science. We would like to thank the anonymous reviewers for their insightful recommendations for improving this paper.

References

- M. Arbib. The mirror system, imitation, and the evolution of language. In Kerstin Dautenhahn and Chrystopher Nehaniv, editors, *Imitation in Animals and Artifacts*. The MIT Press, 2002.
- J.A. Bargh. The automaticity of everyday life. In Jr. R.S. Wyer, editor, *The automaticity of everyday life: Advances in social cognition*, pages 1–61. Erlbaum, Mahwah, NJ, 1997.
- A. Billard. Learning motor skills by imitation: a biologically inspired robotic model. *Cybernetics & Systems*, 32:155–193, 2000.
- A. Billard and K. Dautenhahn. Experiments in learning by imitation: grounding and use of communication in robotic agents. *Adaptive Behavior*, 7(3/4):411–434, 1999.
- E. Borenstein and E. Ruppin. Enhancing autonomous agents evolution with learning by imitation. *Journal of Artificial Intelligence and the Simulation of Behavior*, 1(4):335–347, 2003.
- G. Buccino, F. Binkofski, G.R. Fink, L. Fadiga, L. Fogassi, V. Gallese, R.J. Seitz, K. Zilles, G. Rizzolatti, and H.-J. Freund. Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. *European Journal of Neuroscience*, 13:400, 2001.
- D. Cliff and J. Noble. Knowledge-based vision and simple visual machines. *Philosophical Transactions of the Royal Society: Biological Sciences*, 352:1165–1175, 1997.
- S. Cochin, C. Barthlmy, B. Lejeune, S. Roux, and J. Martineau. Perception of motion and qEEG activity in human adults. *Electroencephalography And Clinical Neurophysiology*, 107:287–295, 1998.
- K. Dautenhahn. Getting to know each other—artificial social intelligence for autonomous robots. *Robotics and Autonomous Systems*, 16:333–356, 1995.
- K. Dautenhahn and C.L. Nehaniv. The agent-based perspective on imitation. In Kerstin Dautenhahn and Chrystopher Nehaniv, editors, *Imitation in Animals and Artifacts*. The MIT Press, 2002a.

- K. Dautenhahn and C.L. Nehaniv, editors. *Imitation in Animals and Artifacts*. MIT Press, Cambridge, Mass., USA, 2002b.
- Y. Demiris and G. Hayes. Imitation as a dual-route process featuring predictive and learning components: a biologically plausible computational model. In Kerstin Dautenhahn and Chrystopher Nehaniv, editors, *Imitation in Animals and Artifacts*. The MIT Press, 2002.
- Y. Demiris and M. Johnson. Distributed, predictive perception of actions: a biologically inspired robotics architecture for imitation and learning. *Connection Science Journal*, 15(4):231–243, 2003.
- L. Fadiga, L. Fogassi, G. Pavesi, and G. Rizzolatti. Motor facilitation during action observation: a magnetic stimulation study. *Journal of Neurophysiology*, 73:2608–2611, 1995.
- D. Floreano and J. Urzelai. Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13:431–443, 2000.
- L. Fogassi, V. Gallese, G. Buccino, L. Craighero, L. Fadiga, and G. Rizzolatti. Cortical mechanism for the visual guidance of hand grasping movements in the monkey: A reversible inactivation study. *Brain*, 124: 571–586, 2001.
- V. Gallese, L. Fadiga, L. Fogassi, and G. Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119: 593–609, 1996.
- R. Hari, N. Forss, S. Avikainen, E. Kirveskari, S. Salenius, and G. Rizzolatti. Activation of human primary motor cortex during action observation: A neuromagnetic study. *Proceedings of the National Academy of Sciences USA*, 95:15061–15065, 1998.
- G. Hayes and J. Demiris. A robot controller using learning by imitation. In *Proceedings of the 2nd International Symposium on Intelligent Robotic Systems*, 1994.
- L.M. Herman. Vocal, social, and self imitation by bottlenosed dolphins. In Kerstin Dautenhahn and Chrystopher Nehaniv, editors, *Imitation in Animals and Artifacts*. The MIT Press, 2002.
- M. Iacoboni, R.P. Woods, M. Brass, H. Bekkering, J.C. Mazziotta, and G. Rizzolatti. Cortical mechanisms of human imitation. *Science*, 286:2526–2528, 1999.

- F. Keijzer. Representation in dynamical and embodied cognition. *Cognitive Systems Research*, 3:275–288, 2002.
- A. Keinan, C.C. Hilgetag, I. Meilijson, and E. Ruppin. Causal localization of neural function: The Shapley value method. *Neurocomputing*, to appear, 2004.
- H. Kozima, C. Nakagawa, and H. Yano. Emergence of imitation mediated by objects. In C.G. Prince, Y. Demiris, Y. Marom, H. Kozima, and C. Balkenius, editors, *Proceedings Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, volume 94, pages 59–61, Edinburgh, Scotland, 2002.
- C. Langton. *Artificial Life: An Introduction*. MIT Press, Boston, MA, 1995.
- Y. Marom, G. Maistros, and G. Hayes. Toward a mirror system for the development of socially-mediated skills. In C.G. Prince, Y. Demiris, Y. Marom, H. Kozima, and C. Balkenius, editors, *Proceedings Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, volume 94, Edinburgh, Scotland, 2002.
- A. Meltzoff. The human infant as imitative generalist: a 20-year progress report on infant imitation with implications for comparative psychology. In C.M. Hayes and B.G. Galef, editors, *Social Learning in Animals; The Roots of Culture*, New York Academic Press, 1996.
- A. Meltzoff and K. Moore. Explaining facial imitation: a theoretical model. *Early Development and Parenting*, 6:179–192, 1997.
- M. Mitchell. *An introduction to genetic algorithms*. MIT Press, 1996.
- C.L. Nehaniv and K. Dautenhahn. The correspondence problem. In Kerstin Dautenhahn and Chrystopher Nehaniv, editors, *Imitation in Animals and Artifacts*. The MIT Press, 2002.
- E. Oztop and M.A. Arbib. Schema design and implementation of the grasp-related mirror neuron system. *Biological Cybernetics*, 87:116–140, 2002.
- W. Prinz and A.N. Meltzoff. An introduction to the imitative mind and brain. In A.N. Meltzoff and W. Prinz, editors, *The imitative mind: Development, evolution and brain bases*, pages 1–15. Cambridge University Press, Cambridge, MA, 2002.

- N. Ramnani and R.C. Miall. A system in the human brain for predicting the actions of others. *Nature Neuroscience*, 7:85 – 90, 2004.
- G. Rizzolatti and M.A. Arbib. Language within our grasp. *Trends in Neurosciences*, 21:188–194, 1998.
- G. Rizzolatti, L. Fadiga, L. Fogassi, and V. Gallese. From mirror neurons to imitation: Facts and speculations. In A.N. Meltzoff and W. Prinz, editors, *The imitative mind: Development, evolution and brain bases*, pages 247–266. Cambridge University Press, Cambridge, MA, 2002.
- G. Rizzolatti, L. Fadiga, V. Gallese, and L. Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3:131–141, 1996.
- G. Rizzolatti, L. Fogassi, and V. Gallese. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2:661–670, 2001.
- E. Ruppin. Evolutionary autonomous agents: A neuroscience perspective. *Nature Reviews Neuroscience*, 3(2): 132–141, 2002.
- P.M. Todd and G.F. Miller. Exploring adaptive agency II: Simulating the evolution of associative learning. In J.A. Meyer and S.W. Wilson, editors, *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 306–315, Cambridge, MA., 1991. MIT Press.
- T.R. Zentall. Imitation in animals: evidence, function, and mechanisms. *Cybernetics and Systems*, 32(1-2): 53–96, 2001.