

To be published in Ron Sun (ed.),
Cognition and Multi-Agent Interaction:
From Cognitive Modeling to Social Simulation

Sociality in embodied neural agents

Domenico Parisi

Stefano Nolfi

1. Introduction

This chapter addresses the topic of how embodied neural agents coordinate together to exhibit interesting social behaviors. Embodied neural agents are defined in this Introduction. Sections 2 through 5 describe simulations of collective phenomena emerging from the interactions among embodied neural agents living in the same environment. Section 2 discusses spatial aggregation and proto-social behavior, Section 3 communication, and Section 4 cultural evolution. Section 5 summarizes the chapter and draws some conclusions.

Neural agents are agents whose behavior is controlled by neural networks, that is, by control systems that reproduce in simplified ways the physical structure and the physical way of functioning of the nervous system. A neural network is a set of units (neurons) linked by unidirectional connections (synapses between neurons). Connections have a quantitative weight (number of synaptic sites between pairs of neurons) and a plus or minus sign (excitatory and inhibitory synapses). At any given time every unit has an activation level (firing rate of neurons) which depends on either physico/chemical events outside the network (input units) or the sum of excitations and inhibitions arriving to the unit from connected units (internal and output units). Activation propagates from the input units to the output units through one or more intermediate layers of internal units. The pattern of activation of the output units determines some effect outside the network.

At the level of the individual agent the network's architecture of connections and the weights of the individual connections can change as a consequence of the agent's interactions with the external environment, and these changes translate into changes in behavior (learning). At the population

level an agent is a member of an evolving population of individually different agents and the architecture of connections and/or the connection weights of the agent's neural network are encoded in the agent's inherited genotype. Individual genotypes reproduce selectively and with the constant addition of new variants (genetic mutations) and this results in neural/behavioral changes in successive generations of agents (evolution).

Neural networks are simulation models, that is, they are theoretical models which are expressed as computer programs. Neural networks can be viewed as part of Artificial Life, which is an attempt at studying all phenomena of life by reproducing them in artificial systems, either simulated in a computer or physically realized in robots and other physical artifacts. When neural networks are seen in the framework of Artificial Life, research using neural networks tends to be different from classical neural network research in a number of respects (Parisi, Cecconi & Nolfi, 1990; Cliff, 1991; Cliff, Harvey and Husband, 1993; Nolfi & Parisi, 1997; Nolfi & Floreano, 2000; Parisi, 2001). While classical neural networks do not have a body, do not interact with a physical environment (their only "environment" is the researcher), are viewed as isolated individuals, and change only because of individual learning, neural networks in an Artificial Life perspective:

- have a body
- live in and interact with a physical environment
- are members of biologically and, possibly, culturally evolving populations of networks
- have a genotype which is the result of biological evolution and which determines important aspects of the network's structure and development and therefore of the individual's behavior.

Embodied neural agents adopt the same conceptual and explanatory apparatus of the natural sciences and they try to fully integrate the study of behavior in the study of nature. Everything which takes place inside a neural network and in the network's interactions with the outside environment and with the rest of the organism's body (Parisi, in press) are physical causes producing physical effects, and everything is ultimately quantitative in nature.

Embodied neural agents are part of a new research paradigm which has recently challenged the traditional view according to which intelligence is an abstract process that can be studied without taking into consideration the physical aspects of natural systems (Pfeifer & Scheier, 1999). The new paradigm tends to stress *situatedness*, i.e., the study of systems that are situated in and interact with

an environment (Brooks, 1991; Clark, 1997), *embodiment*, i.e., the assumption that systems have bodies, receive input from physically situated sensors, and produce motor actions as output (Brooks, 1991; Clark, 1997), and *emergence*, i.e., the view of behavior and intelligence as the emergent result of the fine-grained interactions between the control system of the agent, its body structure, and the external environment. An important consequence of this view is that the agent and the environment constitute a single system, i.e., the two aspects are so intimately connected that a description of each of them in isolation does not make much sense (Maturana & Varela, 1980, 1988; Beer, 1995).

Although embodied neural agents tend to be simple and to live in simple environments, if one places many agents together in the same environment interesting collective behaviors tend to emerge from their interactions. In the next three sections various aspects of sociality are explored using collections of simple embodied neural agents that live in the same environment: spatial aggregation, simple coordination, communication, and the emergence and evolution of culture. The agents do not initially possess any ability or any interesting behavior since their behavior results from the connection weights of their neural network and these connections weights initially are random. The connection weights change in the course of the simulation until the appropriate behaviors underlying interesting social phenomena emerge. In other words, the system that controls the behavior of the agents is not designed by the researcher but is evolved or learned, and the researcher only creates the conditions in which evolution or learning take place.

Evolution can be either biological or cultural and in both cases it can be simulated using a genetic algorithm. In biological evolution an agent inherits from its parent(s) a genotype encoding the connection weights of the agent's neural network. Reproduction is selective in that not all individuals have the same number of offspring. Furthermore, reproduction is accompanied by the constant addition of new variants to the pool of genotypes because random errors may occur when copies of genotypes are made and because portions of one parent's genotype may be recombined with portions of the other parent's genotype if sexual reproduction is adopted. In cultural evolution information is transmitted not via copied genotypes but through learning from others. One individual, the "student", learns to behave in the same way as another individual, the "teacher", by being exposed to the same input to which is exposed the other individual and by using the output of the other individual as teaching input to change its own connection weights as part of the backpropagation procedure. Also in this case, reproduction is selective in that not all teachers have the same number of students, and it is accompanied by the constant addition of new variability to

the cultural pool because random errors occur when a behavior is transmitted from teacher to student and because new behavioral variants which are recombinations of parts of existing variants may be invented. If the individuals that reproduce or the teachers that have students are individuals that are better able at exhibiting some particular behaviour than the individuals that do not reproduce or do not have students, and if the new variants happen to be better than the existing variants, what is observed in both cases is the evolutionary emergence of initially nonexistent behaviours.

2. Spatial aggregation and coordination

The environment in which embodied neural agents live can contain nonliving objects, organisms belonging to other species, conspecifics, and artefacts created by conspecifics. Since neural agents have a body and they live in a physical environment, all interactions of neural agents with their environment are physical interactions. If the environment contains many embodied agents, all interactions among embodied neural agents consist in alterations of the external environment which are caused by the behavior of one agent and which affect other agents.

Consider the following simulation. A collection of agents live in the same environment which contains randomly distributed food elements. The neural network which controls an agent's behavior has input units encoding the position of the single food element which is currently nearest to the agent and output units that allow the agent to move in the environment. The neural networks of all agents have the same architecture but at the beginning of the simulation each individual agent is assigned a genotype which encodes a different random set of connection weights for the agent's neural network. Each individual lives for a total number of time units (input/output cycles of its neural network) which is identical for all individuals. At birth each individual has zero energy but its energy is incremented by one unit each time the individual by moving in the environment reaches (eats) a food element. When the energy of the individual reaches a threshold, the individual generates a new individual (offspring) which inherits the same genotype of its (single) parent, with the addition of some random changes to the quantitative value of some of the weights. The offspring is placed near its parent and the parent's energy returns to zero.

While at the beginning of the simulation the agents are not very good at reaching food because of the random connection weights, the selective reproduction of the individuals that are better able to reach food and to increase their energy, and the constant addition of new variability to the pool of

genotypes because of the random variations in the inherited connection weights, lead to an improvement in the average ability to reach food in the population with each successive generation. After an initial transient phase, population size stabilizes at a value which reflects the quantity of food present in the environment (carrying capacity). Food is periodically re-introduced to compensate for the food eaten, and the carrying capacity of the environment, and therefore, population size, is a function of the length of the interval between successive food re-introductions.

The results of the simulation show that if food is re-introduced sufficiently frequently, the population distributes itself homogeneously in the environment. However, if food is reintroduced less frequently, an interesting collective phenomenon emerges with respect to the spatial distribution of the population: one observes oscillatory migratory waves of the agents in the environment. The entire population of agents may concentrate in a particular zone of the environment but, after a while, the population leaves the zone and disperses in the environment, with different individuals going in different directions. When the agents reach the wall that limits the environment, they remain near the wall for a while and then they slowly return to the initial zone in which they concentrate again. This oscillatory movement of the population repeats itself periodically until the end of the simulation (Figure 1).




Figure 1. Oscillatory migratory waves of agents. In (1) the agents concentrate in a zone of the environment which happens to have more food than other zones. In (2), after having depleted of food the originary zone, the agents migrate toward the periphery of the environment where food has accumulated in the meantime. In (3) the agents have reached the periphery. In (4), after having depleted the periphery, they return to the originary zone of concentration where food has returned since they left the zone.

How can one explain this collective phenomenon of periodic oscillatory migratory waves? For reasons that are purely based on chance some zone of the environment may contain more food than other zones and therefore the agents looking for food tend to move toward the zone with more food and concentrate there. However, as more and more agents concentrate in the same zone and eat the food which is found there, the zone tends to become progressively deprived of food. As the zone empties, the agents leave the zone in different directions to migrate in more peripheral zones which, having been without agents for some time, contain much food. The same phenomenon repeats itself in the zones located peripherally with respect to the initial zone. As the peripheral zones are

emptied of food because of the many agents that have reached those zones, an opposite wave of migration towards the initial concentration zone takes place. Therefore, the population ends up periodically migrating from a more centrally located zone to the periphery and back to the central zone (Parisi, Piazzalunga, Cecconi & Denaro, 1994).

The simulation demonstrates that interesting collective phenomena may emerge in populations of very simple neural agents even if the agents cannot be said to possess social behaviors or social abilities. The agents in this simulation do not even perceive each other. They only perceive the food elements. Furthermore, food is randomly and, therefore, at the appropriate scale, equally distributed in the entire environment. This notwithstanding, an interesting collective spatial pattern emerges from the simulation. As already mentioned, the agents do not perceive each other and they respond to input from the nonsocial environment (food) with behavior which is uniquely directed to the nonsocial environment (eating the food). However, if a population of agents lives together in the same physical environment, by altering the physical environment with their behavior (eating the food) individual agents can have an indirect influence on other individuals since each agent responds to an environment which has been altered by the behavior of other agents. This can produce emerging collective phenomena in the spatial distribution of the population such as the oscillatory migratory waves observed in the simulation.

The agents described periodically aggregate and disaggregate (disperse) spatially as a result of the changes that their behavior causes in a nonsocial environment in which the resources by themselves are randomly distributed. In real populations, both animal and human, social aggregation can result from the particular spatial distribution of resources in the environment. Many individuals can end up near to each other simply because they tend to approach the same localized resource such as a food patch or a water source or a lecture in a classroom. In these circumstances too, the agents' behavior which results in social aggregation has not evolved for that function. Each individual approaches food or water or the classroom for eating or drinking or learning, not for social purposes. However, even if it is a simple by-product of nonsocial behaviors social aggregation can be a favourable pre-condition for the emergence of social behaviors such as communication and economic exchange among individuals that happen to find themselves near each other.

In other circumstances, however, social aggregation may not be simply a by-product of behavior which has emerged for other purposes but is the result of behavior which has emerged exactly because it produces spatial aggregation. One can distinguish between two types of social behavior

that results in social aggregation and, more generally, social interaction. In Type 1 social behavior, one individual alters the environment of another individual but it does so for its own, nonsocial, reasons, while the second individual responds to the alteration of the environment by the first individual with a behavior which has emerged with the function of producing social aggregation or interaction. In Type 2 social behavior, both the behavior of the individual which alters the environment of another individual and the behavior of the individual which responds to this alteration of the environment on the part of the first individual emerge with the function of producing social aggregation or other social phenomena.

Let us consider Type 1 social behavior first. Imagine a population of agents very similar to the agents of our previous simulation with the only difference that an individual's life is made up of two successive stages. In the first life stage the individual is a 'child', which means that the input units of its neural network encode the current position of the individual's parent, not the position of the nearest food. In other words, a child sees its parent but does not see the food. This means that a child cannot find food by itself and would starve and die unless its parent gives some of its food to the child. In the second stage of an individual's life the individual becomes an 'adult' and is exactly identical to the agents of our previous simulation. An adult's neural network encodes the position of the nearest food and the individual responds by approaching and capturing the food. However, some portion of the food which is captured by an adult individual is not eaten by the individual but is given by the adult individual to its children provided the individual has children.

But children are not passive receivers of food. In order to obtain food from their parents it is their responsibility to remain in close proximity to their parents. This is why the input units of a child's neural network encode the current location of the child's parent in the environment. The child must be able to respond to this input by approaching its parent. Since a child's parent moves in the environment looking for food, this means that children should follow their parents so that a child's distance from its parent never exceeds a certain threshold. This in fact is the children's behavior which is observed after a certain number of generations in the simulation (Parisi, Cecconi & Cerini, 1995) (Figure 2).




Figure 2. A “child” (small square) evolves the behavior of following its “parent” (large rectangle) which is looking for food (circles) because this allows the “child” to obtain food from its “parent”.



In the simulation adults and their children tend to form small social aggregations of kin-related individuals (families) that move together in the environment. These social aggregations are exclusively due to the evolved behavior of children which respond to visual input originating from their parents by approaching their parents and remaining in their vicinity. Parents do not contribute with their behavior to these social aggregations since, as adults, they simply look for food. This, then, is a Type 1 situation in which agents (children) evolve a behavior in response to the input provided by conspecifics (parents) but the parents’ behavior that provides this input for their children evolves for independent, nonsocial, reasons (looking for food).

One moves toward a Type 2 situation if one assigns also to the parents a role in maintaining the spatial aggregation of their family. In the simulation that has been described the behaviour of parents that give some of their food to their children is hardwired and not evolved but also this behavior can be evolved. Children would starve to death unless their parents provide them with food, and in these circumstances their parents’ genes would not be transmitted to the next generation. Therefore, it is in the parents’ genetic interest to give some of their food to their children. In the new simulation an agent’s genotype includes not only the genes that encode the connection weights of the agent’s neural network but also an additional gene which encodes in a simple quantitative way the agent’s tendency, when it becomes an adult, to give some of its food to its children. The value of the gene varies among the agents and at the beginning of the simulation is randomly generated for each agent. Offspring inherit the same value of the gene of their parents with random mutations that may slightly increase or decrease the gene’s value. The individuals that

tend to reproduce are individuals that not only are good at finding food but also have a propensity to give some of their food to their children. Even if giving food to one's children reduces an individual's chances to generate additional children, the results of the simulation show that after a certain number of generations the "give food to your children" gene stabilizes at an intermediate quantitative value which takes into consideration both an individual's need to generate additional children and its need to keep alive and bring to sexual maturity the children which the individual has already generated. By their evolved tendency to give food to their children parents contribute to social aggregation because it is this tendency of parents that motivate children to remain near their parents.

Many real life collective behaviors, e.g., the behavior exhibited by schools of fish or flocks of birds, are in between Type 1 and Type 2. Many fish and birds move together when they look for food or, in the case of birds, migrate to distant places. If being in the vicinity of conspecifics confers some adaptive advantage in terms of increasing the probability of finding food or avoiding predation, agents that are able to perceive their conspecifics will evolve a tendency to approach their conspecifics so as to maintain proximity even as the group of agents collectively moves in the environment. As in the preceding simulation in which children actively maintain proximity to their parents, the behavior of responding to the input originating in a conspecific evolves because it produces proximity to conspecifics but the behavior of the conspecific that generates this input has not necessarily evolved for this reason. However, the spatial aggregation that is maintained in a collection of agents is different from the behavior of the pair of agents constituted by a parent and its child (or children). A parent which moves in the environment looking for food is not influenced by the behavior of its children and, in our simulations, the parent does not even perceive its children. In contrast, in a group of agents moving together in the environment each individual both perceives and is perceived by its conspecifics and each individual, either directly or indirectly, both influences and is influenced by the other individuals in the group. In fact, in a group of agents maintaining spatial proximity and moving together in the environment local causes become global causes. A local cause is an event which takes place in one particular agent and which has some direct influence on another agent. A global cause is an event or state at the level of the entire group of agents which has some influence on each individual agent belonging to the group. In our simulations a parent's behavior is a local cause of its children's behavior. In a school of fish or in a flock of birds each agent is influenced by the other agents and therefore when an individual agent influences another individual agent this influence reflects the state of the entire group. The behavior of each agent is both a local and a global cause of the behavior of the other agents.

The collective behavior of a group of agents that move together in the environment has been simulated by various researchers. For example, Reynolds (1993) evolved the control system of a group of agents (flock of birds) placed in an environment containing obstacles for the ability to collectively avoid the obstacles. The group of agents splits before an obstacle and re-unites after passing the obstacle. Baldassarre, Nolfi & Parisi (2003) (cf. also Baldassarre, Parisi and Nolfi, in press) have simulated various collective behaviors with groups of robots (“swarmbots”). A group of, say, four robots each with its own control system (neural network) are physically linked in various spatial configurations. For example, the four robots can form a line with each robot physically linked to the next robot in the line. The neural network of each robot has input units encoding the strength with which the individual robot is pushed or pulled by the other robots and the direction in which the robot is pushed or pulled by the other robots. The network’s output units control two wheels that allow the robot to move in the environment. These “swarms” of robots cannot be said to form spontaneously because the robots are already united through the physical links, but they evolve various interesting collective behaviors: they quickly line up their wheels in order to move coherently, i.e., in the same direction, and they are able to negotiate obstacles, to reach light targets if each robot is provided with additional input units encoding the location of the target, and to help single members that happen to fall in holes (Figure 3). These collective behaviors appear to be very robust in that they are exhibited even when the robots become members of new “swarms” made up of different numbers of robots and with different spatial configurations with respect to the originary “swarm” in which the robots have evolved. In all these simulations each robot causes inputs for the other robots and at the same time is influenced by the inputs caused by the other robots. Therefore, a local influence of one robot on another robot is at the same time a global influence of the entire “swarm” of robots on each individual robot.






Figure 3. **Top-Left:** The hardware prototype of an individual robot. **Top-Right:** Four simulated robots linked up to form a linear structure. **Bottom:** The trajectory followed by a star-shaped swarmbot made up of eight individual robots in an environment with obstacles, furrows, and holes. The swarmbot is depicted in its final position near the light target represented by the white sphere. The black irregular lines indicate the trajectories followed by the eight robots forming the swarmbot. While isolated robots (indicated by arrows) get stuck in furrows, the swarmbot passes over the furrows, succeeds to free its component robots that fall in holes, and searches and finds the light that was not visible from the starting position (center of the graph).

Coordinated behavior in embodied agents spontaneously emerges also with other types of tasks such as herding in response to predators and the collective building of structures. In an attempt to study the evolutionary origin of herding, Werner & Dyer (1993) co-evolved two populations of predators and prey agents which were selected for the ability to catch prey and for the ability to find food and escape predators, respectively. The author observed that, after some generations during which predators evolved an ability to catch prey, prey agents converged into small herds which were constantly splitting up and re-forming. More recently, Ward, Gobet & Kendall (2001) evolved groups of artificial fish able to display schooling behaviour. Two populations of predator and prey fish, respectively, were evolved in an environment containing randomly distributed food elements. The neural network controlling a prey's behavior included sensory neurons encoding distance and direction of nearest prey, predator, and food, and the amount of changes in water pressure in proximity to the agent, and two motor neurons encoding speed and direction of motion of the agent. An analysis of distances between prey and food and between prey and predator suggests that

schooling behaviour is correlated with an increased probability to find food clumps and a better protection from predation. Finally, Theraulaz & Bonabeau (1995) evolved a population of constructor agents who collectively build a nest structure by depositing bricks according to their perception of the local environment and to a set of behavioral rules.

An interesting phenomenon that can be studied with collective tasks is the emergence of specialization, with different individual agents spontaneously assuming different roles in the execution of the task. Specialization emerges both when the agents involved in a collective task are genetically different individuals (Yong & Miikkulainen, 2001) and when they are clones (Quinn, Smith, Mayley & Husband, in press). In Baldassarre, Nolfi & Parisi's (2003) simulations, clone agents first have to aggregate and then move together towards a target. The most effective strategy includes primitive forms of "situated" specialization in which identical individuals play different roles according to the circumstances such as leading or following the group (see next section). These forms of functional specialization seem to be due to the need to reduce interference between potentially conflicting sub-goals such as moving toward the rest of the group to maintain aggregation and moving toward the target.

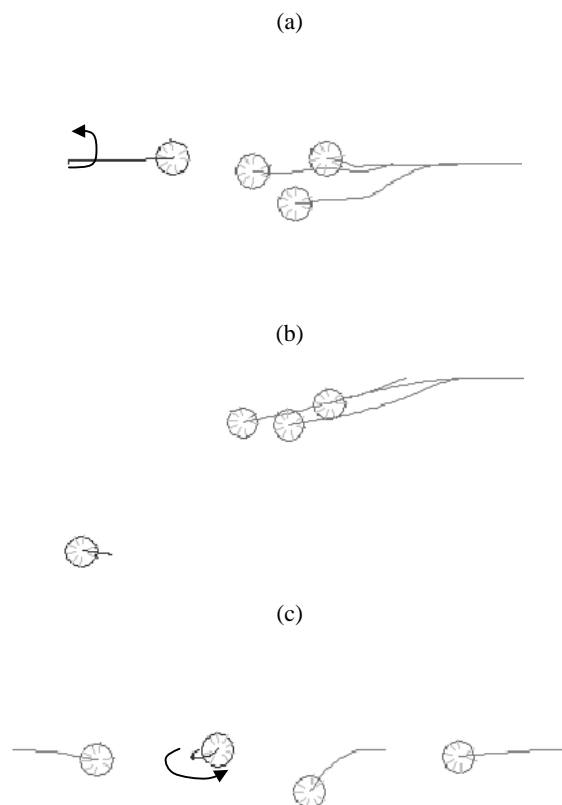
3. Communication

"Swarm" simulations still have to do with Type 1 social behavior or, perhaps, with behavior which is intermediate between Type 1 and Type 2. But of course behavior can evolve in agents for the explicit function of providing inputs to conspecifics. This behavior is called communication and the inputs that are provided to conspecifics are called signals. Communication clearly involves Type 2 situations.

Imagine a group of agents that has to reach a target in the environment but to be rewarded they must approach the target by maintaining reciprocal proximity. If the agents are initially dispersed in the environment, they may be unable to perceive each other and therefore they may be unable to aggregate and then move together towards the target. The solution is to evolve some signaling behavior that by providing an input to conspecifics allows the group to aggregate. The neural network that controls the behavior of an agent has both input units that visually encode the position of the target and input units that encode acoustic input originating from the behaviour of conspecifics (signals). The output units encode both behavior that allows the agent to move in the environment and behavior that produces a sound that can be heard by conspecifics. The sounds that

are produced by individual conspecifics sum together and result in a louder compound sound. The agents evolve an ability to recognize the direction from which the loudest sound arrives to their sensors, and therefore the direction in which the conspecifics are aggregating spatially, and to respond by moving in that direction. The results of the simulation show that the agents first respond to the sounds that they hear by aggregating together and ignoring the input from the target, and then they respond to both the sounds and the visual input from the target by moving toward the target while maintaining spatial aggregation (Baldassarre, Nolfi & Parisi, 2003).

As already mentioned, evolved agents show a form of situated specialization. Individuals that are located on the frontal side of the group with respect to the light target (“leaders”) do not turn toward the rest of the group but keep their orientation toward the light, sometimes moving backward to avoid losing contact with the rest of the group. On the contrary, individuals located behind (“followers”) turn and move toward the other members of the group. Moreover, once a compact group has formed and the group starts to move toward the light, each individual tries to maintain its current role. The final result is that the “leader” drives the whole group toward the light while the “followers” only try to remain in proximity to the “leader” and to each other so that the whole group continues to be compact (see Figure 4).






Figure 4. Behavior displayed by four agents initially located in four different starting positions and orientations. In all cases the light target is located on the left side. The lines represent the trajectories of the four agents and the circles represent the final position of the agents after a given amount of time. The arrows indicate quick changes in the orientation of individual agents.

Figure 4 shows how the agents play different functions in different circumstances. Fig. 4a shows how the individual which is closer to the light target assumes and maintains the function of “leader”. The individual turns toward the light and waits for the rest of the group before driving the entire group toward the light target. It may move backward to speed up the formation of a compact group but, as soon as the rest of the group gets closer, it starts to move toward the light target thus keeping the frontal position with respect to the rest of the group. Figure 4b shows another situation in which the individual which is closer to the light target does not turn toward the rest of the group but keeps its relative position by waiting for the rest of the group and by starting to move toward the light as soon as the rest of the group approaches. Figure 4c shows that individuals that are shadowed by other individuals and cannot see the light target (in this case the second robot from the left) turn and move toward the rest of the group. Finally, Figure 4d shows that a couple of individuals located in similar conditions with respect to the light target and to the rest of the group can assume and maintain the role of both leaders or followers. The overall result of being able to display and maintain “situated” specializations is that agents can quickly form a compact group and then move straight toward the light target.

In the simulation that has been described the behavior of producing a sound that can be heard by conspecifics is hardwired. One could do another simulation in which the behavior evolves spontaneously since it is clearly advantageous for an individual to produce such a sound. Unless an individual produces the sound, the conspecifics may not be able to know where the individual is located and to approach the individual. However, the behavior of producing signals raises an interesting problem. Communication is a Type 2 situation. It requires the evolution at the same time of the behavior of emitting the appropriate signals in the appropriate circumstances (altering the

external environment in the appropriate way) and the behavior of responding to the signal appropriately (responding to the alteration in the environment in the appropriate way). Furthermore, both behaviors must be exhibited by each individual. From an evolutionary point of view, a behavior tends to emerge only if it is advantageous for the individual that exhibits the behavior. Therefore, for communication to emerge it is necessary that both the behavior of emitting the appropriate signal is advantageous for the emitter of the signal and the behavior of responding appropriately to the signal is advantageous for the receiver of the signal. In some circumstances the two conditions may not be both satisfied and this may prevent communication to emerge. If emitters do not emit the appropriate signals in the appropriate circumstances, there are no useful signals for receivers to respond to appropriately. If receivers do not respond to signals appropriately, it makes no sense for emitters to emit the appropriate signals in the appropriate circumstances. Hence, emitting and receiving signals cannot evolve separately but they need to co-evolve.

One can imagine situations in which both the emitters of signals benefit if they emit the appropriate signals in the appropriate circumstances and the receivers of signals benefit if they respond appropriately to the signals. Consider the following simulation. A population lives in an environment in which there are large prey that can only be captured and killed if a sufficiently large group of individuals are present and hunt together the prey. The agents initially disperse in the environment and when an individual finds the prey it communicates to the other individuals where the prey is located so that the other individuals can use this information and converge to the prey's location. There are two solutions to this problem. One is a more primitive and limited solution. The individual which has found the prey immediately emits some sound (or some similar signal that can be perceived from a distance) and the other individuals perceive the sound and its direction and they immediately approach the source of the signal. This solution is primitive and limited because it only works if a number of conditions are satisfied: (a) the signal is produced by the emitter as soon as it finds the prey, (b) the signal can be perceived at sufficiently large distances, (c) the conspecifics respond immediately and, of course, (d) producing the signal does not cause the prey to escape.

A more sophisticated solution is the emergence of a true language in which different signals describe the particular location in which the prey has been found. For example, in the simulation the environment may contain various landmarks and different signals are emitted by the individual which has found the prey which co-vary with, i.e., designate, the specific landmark near which the prey has been found (e.g., "(near the) mountain", "(near the) river", etc.). This more sophisticated solution does not have the limitations of the former, simpler, one. A signal can be produced and

responded to at any time, it does not have to be strong to be perceived at large distances, and it needs not cause the prey to escape.

Notice that in the situation that has been described a communication system, whether simple or more complex, emerges because it is advantageous for both the emitters and the receivers of signals. Since the prey is too large to be hunted individually, the individual which finds the prey and emits the signal is advantaged because its signalling behavior causes other individuals to come where the prey is located so that the prey can be hunted collectively. At the same time, the individuals that receive the signal are advantaged in responding appropriately to the signal by going to where the prey is located because this allows them, again, to hunt collectively the prey. One can imagine also other situations in which a signalling system can evolve because it is advantageous to both emitters and receivers of signals. An agent can emit a signal asking another individual to do something which seems to be useful only to the emitter of the signal, but in fact the receiver of the signal responds as required because this allows the receiver to get some advantage such as avoiding being punished by the emitter of the signal or exchanging roles with the current emitter of the signal in some future occasion.

However, there may be other conditions in which the receivers of signals may be advantaged by being able to respond appropriately to the signals but the emitters of the signals have no advantages in emitting the appropriate signals in the appropriate circumstances. If this is the case, a signalling system may fail to emerge.

This has been studied in the following simulation (Mirolli & Parisi, 2004; in press). A population of agents lives in an environment that contains both edible and poisonous mushrooms. Edible and poisonous mushrooms are perceptually different but, in order to recognize them and eat the edible mushrooms while avoiding the poisonous ones, an individual must be sufficiently close to an encountered mushroom to see the mushroom appropriately. If the agent is alone and it encounters a mushroom, the only available strategy is in all cases to approach the mushroom until the agent is sufficiently close to the mushroom and is able to recognize whether the mushroom is edible or poisonous. This is not a very efficient strategy, however, since it involves a waste of time and energy if the mushroom turns out to be poisonous. If a second individual is also present and is closer to the mushroom, the second individual can send a signal to the first individual telling the first individual whether the mushroom is edible or poisonous. This behavior of the second individual, the emitter of the signal, is clearly advantageous for the first individual, the receiver of

the signal, but is it advantageous for the emitter of the signal? Why should the behavior of emitting the appropriate signal in the appropriate circumstances emerge evolutionarily if it provides no advantages for the individual exhibiting the behavior?

As a matter of fact, the results of the simulation show that if the emitter and the receiver in any given encounter are randomly selected from the entire population, a useful signalling system fails to emerge. Emitters fail to produce the appropriate signals in the appropriate circumstances (one particular signal for edible mushrooms and another, different, signal for poisonous mushrooms) and, therefore, in the absence of useful signals, receivers of signals cannot evolve the behavior of responding appropriately to the received signals. An appropriate signalling system would only benefit receivers of signals but not emitters of signals and this prevents such a signalling system from emerging.

However, it is possible to create conditions in which, given the same simulation scenario, the correct signalling system will emerge. In the simulation that has been described the signalling system is genetically transmitted in that it results from the neural network's connection weights which are encoded in the agents' genotypes. These connection weights determine both which signals are produced by emitters and how receivers respond to signals. If emitters and receivers of signals in social encounters are randomly chosen from the population, the "egoism of the gene" prevents the signalling system from emerging because in any particular encounter the emitter of the signal and the receiver of the signal tend to have different genes, i.e., they are not kin-related individuals. An emitter that produces the appropriate signals increases the reproductive chances of the receiver of the signal which will tend to evolve an ability to respond appropriately to the signal because this ability is in its own interest. However, the receiver of the signal may not produce the appropriate signals when its role changes and it becomes an emitter of signals. Hence, by increasing the reproductive chances of the receiver of the signal, a good emitter of signals may increase the reproductive chances of a bad emitter of signals. In these conditions individuals that are at the same time good emitters and good receivers of signals tend not to emerge.

But if one changes the simulation scenario and introduces the condition that in any particular encounter the emitter and the receiver of the signal are kin-related individuals, i.e., they have the same (or similar) genes because they are the offspring of the same parent, then a good signalling system does emerge as predicted by kin-selection theory. Good signalers, i.e., individuals that emit the appropriate signals in the appropriate circumstances, provide advantages to the conspecifics that

receive their signals and not to themselves but, since the receivers of the signals have the same genes of the emitters of signals, good signalling genes are maintained in the population.

This simulation seems to imply that language, at least for the particular use of language which is considered in the simulation, can emerge only within small groups of kin-related individuals. However, language is more useful if it can be used in larger groups of non-kin-related individuals. How can language emerge in such larger groups? Furthermore, while in the simulation language is genetically transmitted and it evolves biologically, human language, unlike most animal signalling systems, is learned from others and culturally rather than genetically transmitted. Can the cultural emergence and cultural transmission of language be simulated?

One way in which language can emerge in groups of non-kin-related individuals is if language is used not only to speak to others but also to speak to oneself, i.e., to think. In the simulation that has been described it is assumed that the receiver of a signal is able to keep in memory the signal heard from the emitter while the receiver is approaching the mushroom. In a variant of the same simulation, memory is not assumed but the receiver of the signal must evolve an ability to repeat the signal to itself in order to remember the signal of the emitter. This implies that good receivers of signals must also be good emitters of signals if they must benefit from the signals that they receive. In these conditions a good signalling system emerges in the population even if the emitter and the receiver of signals in any particular social encounter are not kin-related individuals.

Another condition in which an appropriate and useful signalling system does emerge is a condition in which the signalling behavior is culturally rather than genetically transmitted and there is a genetically inherited tendency to learn from others. This genetically inherited tendency to learn from others has been called “docility” by Herbert Simon (Simon, 1990). Human beings appear to possess docility more than other animals. Docility has become part of the human genotype because of the great advantages it bestows on an individual which can directly learn from others many useful abilities and behaviors without going through long, tiresome, and sometimes dangerous individual experiences. Docility implies “blind” learning. Young individuals learn anything which adults care to teach them and, in particular, without first determining if what they learn is advantageous for themselves or for others. This may explain the emergence of language as a learned and culturally transmitted ability. When an individual is learning language, the individual is learning to both emit and understand signals that in some of their uses can be advantageous for the emitter of the signal and in other uses can be advantageous for the receiver of the signals.

In the new simulation an agent's connection weights that are responsible for emitting linguistic signals and for responding to received signals are not encoded in the agent's genotype but are culturally learned by the individual at the beginning of its life. Cultural learning, i.e., learning from others, is simulated by using the behavior of another individual, the teacher, as teaching input for the learner as the learner is learning language on the basis of the backpropagation learning algorithm. The learner's connection weights are randomly assigned at birth and the learner's parent functions as its teacher. In any given learning trial both the learner and the teacher are exposed to the same input and both respond with some output on the basis of their respective connection weights. When the learner is learning to emit linguistic signals, both its neural network and the neural network of its teacher encode the perceptual properties of an encountered mushroom and the output units of both networks encode a signal that classifies (names) the mushroom as either edible or poisonous. In the early stages of learning the learner tends to emit inappropriate signals but, by comparing its own signal with the signal emitted by the teacher in response to the same mushroom and changing its connection weights to reduce the discrepancy between the two signals, the learner progressively learns to emit the same signals as the teacher. When the learner is learning to understand the signals, the input is a signal and the output is the behavior of either approaching or avoiding the mushroom. Again, in the early stages of learning the learner responds differently from its teacher, but after a certain number of trials it learns to respond in the same way as its teacher.

Considering that in the simulation only individuals that are parents function as teachers, this means that teachers are individuals that have been selected for reproduction and therefore tend to have a better language than the individuals which have not been selected for reproduction. Furthermore, the teaching input from a teacher is slightly and randomly changed when it is used by the learner for learning language, which means that, analogously to what happens with random genetic mutations, learners can in some (rare) circumstances develop a better language than their teachers' language. At the beginning of the simulation language is very bad since the teachers belonging to the first generation that teach language to the members of the second generation have random connection weights like their learners. But language gradually emerges culturally. As in biological evolution, the selection of the best individuals as teachers and the constant addition of new variability by adding some random noise to the teaching input progressively lead to the emergence of a useful language in the population - a culturally rather than biologically evolved language. Notice also that docility is not hardwired in our agents but it evolves biologically. Docility is encoded in a special gene which initially has a random value and is biologically inherited from parents to offspring with

the usual random mutations. This value determines how many language learning trials a newborn individual will have and therefore how much language it will learn. Since docility is useful to the individual, the average value of the gene tends to increase in the population and when the simulation stabilizes all individuals tend to be born with a genetically inherited tendency to learn language from their parents.

4. Culture

Culture is behavior (and beliefs, attitudes, values) which are learned from others. Behavior can be learned from others either directly, by imitating another individual or by being taught by another individual, or indirectly, by interacting with technological artefacts made by other individuals. Interactions among agents may result in learning from others. Therefore, agents that live in the same environment and interact together may learn from each other. The individuals of one generation may learn from the individuals of the preceding generation and in this way behavior can be transmitted from one generation to the next. Cultural transmission, like genetic transmission, is accompanied by cultural change or cultural evolution. Individual agents tend to exhibit different variants of the same behavior and these different variants are differentially transmitted to the next generation, with some variants generating more “copies” of themselves than other variants. Furthermore, new variants of behaviors are constantly introduced because of random errors in the “copying” process, invention of new variants, and “copying” of variants existing in other cultures. Groups of agents that interact together more than with members of other groups tend to develop different cultures because of progressive divergence and random drift.

As illustrated in the preceding section, neural agents can be used to study cultural transmission by having agents learn by using the output of other agents as teaching input, on the basis of the backpropagation procedure. In any given trial both the learner and the teacher are exposed to the same input and they both generate an output in response to this input which depends on the connection weights of their respective neural networks. The output of the learner is compared with the output of the teacher and the learner’s connection weights are changed in such a way that, after a certain number of trials, they tend to produce an output similar to the teacher’s output in response to the same input. Hence, any behavior or ability which is initially possessed by the teacher but not by the learner, is transferred to the learner.

If one assumes that the teacher already knows how to find food in the environment, i.e., to respond to visual input encoding the food's position with some motor output which allows the agent to approach the food, a learner with random connection weights at birth and therefore no initial ability to approach food will progressively learn to approach food by imitating the teacher. If one adds some random noise to the teaching input, i.e., in how the learner perceives the teacher's behavior, in some (rare) cases learners can end up being better able than their teachers to approach food (Denaro & Parisi, 1996; Parisi, 1997).

For cultural evolution to take place, two conditions must be satisfied: learners must be spatially near to teachers in order to be able to observe and imitate the teachers' behavior and, furthermore, the best individuals of the previous generation must be selected as teachers. If these two conditions are hardwired in the simulation, a population of neural agents which at birth have random connection weights will progressively acquire, in successive generations, the appropriate connection weights that allow them to approach food efficiently. The connection weights are not genetically inherited but they are culturally acquired by each individual by imitating one or more individuals of the preceding generation. In the early generations teachers do not have much to teach but this gradually changes and the ability to find food builds up through selective cultural transmission and the addition of random novelties (noise) to teaching inputs.

Both the learners' tendency to remain in proximity to teachers in order to learn from them and their ability to select as teachers the best individuals of the preceding generation may evolve genetically, with a process of co-evolution of both biology and culture. In one simulation the agent's genotype encodes the connection weights that cause the agent to approach a teacher and therefore to be in the position to learn from the teacher how to approach food. These connection weights are randomly assigned to the members of the first generation and therefore the agents are initially unable to approach teachers and learn from them. However, the connection weights that cause agents to stay close to each other in order to learn from each other evolve because they are selectively transmitted with the addition of random genetic mutations from one generation to the next (Figure 5). For biological evolution to produce better connection weights, i.e., connection weights that encode the behavior of approaching teachers, the individuals that inherit these weights must also be individuals that are able to learn from teachers the behavior of approaching food. Hence, neither biological evolution (approaching teachers) can take place without cultural evolution (approaching food) nor cultural evolution without biological evolution. The two must co-evolve (Parisi, Piazzalunga, Cecconi & Denaro, 1994).




Figure 5. At the beginning of the simulation agents are randomly distributed in the environment (left). At the end of the simulation agents have evolved a tendency to stay close to each other in order to learn from each other (right).

The ability to identify good teachers, i.e., to select as teachers the individuals of the preceding generation that are best able to approach food, can also evolve. This ability is encoded in a gene which is represented by a single number. The individuals that inherit genes with higher values are better able to select the best teachers. The gene is genetically inherited with some random noise, i.e., a randomly selected quantity is added to or subtracted from the gene's current value. Gene values are randomly assigned in the initial population of agents but, with successive generations, the average value of the gene tends to increase since selecting as teachers the best individuals of the preceding generation is a pre-requisite for the cultural emergence of the ability to approach food.

Cultural transmission can be direct or mediated by technological artefacts. While direct transmission requires face-to-face interaction, indirect cultural transmission only requires that an individual interacts with an artefact made by another individual. Particular technological artefacts tend to induce specific behaviors in the agents that use them and therefore different individuals can learn to behave in similar ways because they use the same technological artefacts. But technological artefacts are not only mediators of cultural transmission. Technological artefacts themselves can evolve. They can be transmitted from one generation to the next and, if technological transmission is accompanied by the selective reproduction of the best artefacts and the constant addition of new variants of the artefacts, what is obtained is technological evolution.

Imagine that the agents that have to look for food in the environment in order to survive and reproduce inherit not only a genotype which encodes the connection weights underlying their food

searching behavior but also some technological artefacts, e.g., vases for storing, transporting, or cooking food. These artefacts allow them to extract more energy from the food they find in the environment and therefore to increase their survival and reproductive chances. The inherited artefacts cannot be directly used by the agents that inherit them, however, but they can only function as models to be copied in order to make new artefacts. Each agent has two neural networks: a network for looking for food and a network for copying artefacts. While the connection weights of the network for looking for food are genetically inherited, the connection weights of the network for copying existing artefacts are randomly assigned at birth and they are learned using the backpropagation procedure. The observed properties of a model artefact which has to be copied function both as input to the artefact-copying network and as teaching input for learning. The network learns to produce an artefact which has the same properties of the model artefact (technically, an auto-association task). This is the artefact that the agent uses.

The individual artefacts are not all identical and some artefacts are better than others, i.e., they allow their users to extract more energy from food. If the artefacts of the preceding generation which are used as models to be copied by the individuals of the next generation are the best artefacts and if some random noise is added to the teaching input so that in some (rare) cases copies of artefacts turn out to be better than their models, what is observed is technological evolution. At the beginning of the simulation artefacts have random properties and therefore their average quality is very low. But the selective reproduction of the best artefacts and the constant addition of new artefacts due to the random noise progressively improve the average quality of the artefacts.

A number of interesting phenomena can be explored using this simulation scenario. For example, how is the selection of the best artefacts effected? The best results are obtained if the actual best artefacts are directly selected for reproduction. However, it may be closer to reality to select for reproduction those artefacts which are used by the most successful agents. In other words, when an agent has to decide which artefacts to select for reproduction, the agent does not directly judge the quality of the artefacts (which may be something too complicated and tiresome to do) but it judges how successful their users are. This inevitably gives a less good evaluation of the quality of the artefacts since the success of an individual, i.e., the total quantity of energy that the individual is able to collect, depends both on its personal ability to find food and on the quality of the artefacts the agent uses. Hence, selecting artefacts for reproduction on the basis of the success of their users tends to be less efficient than selecting them in terms of their directly assessed quality. An agent can collect much energy because the agent is very good at finding food while the artefacts which the

agent uses may not be of very high quality. However, even in these conditions one observes technological evolution, i.e., a progressive improvement in the quality of artefacts.

Another interesting question is the size of the group within which artefacts evolve. If the group is the family and an agent simply uses as models to be reproduced the artefacts used by one's parent, evolution is very slow since the artefacts used by one's parent may not be very good (Figure 6). Technological evolution is faster if the group is larger. In a simulation two populations of agents are contrasted. One population lives as a single integrated community. The other population is segmented into a number of separate communities (enclaves). If the population of artefact users is divided up into communities of agents that do not interact together, artefacts are selected for reproduction among the best artefacts of the local community, not the absolutely best artefacts at the level of the entire population. This has the consequence that technological improvement is slowed down (Figure 6). (For some real historical cases, see Diamond, 1997).




Figure 6. Evolutionary increase in the quality of artefacts when model artefacts to be reproduced are selected from among all the artefacts of the population (top), from among the artefacts of the local community (middle), and are those used by an agent's parent (bottom).

A final result which emerges from the simulations is that the presence of artefacts tends to increase the average energy (wealth) of a population of agents, which is inevitable since artefacts augment the quantity of energy extracted from food, but also to increase the economic stratification of the population. In other words, in a population with artefacts there is a greater difference in energy (wealth) between the average individual and the best individual than in a population without artefacts. This may have occurred in historical reality, for example with the introduction of farming

technologies in populations that previously obtained their food from hunting and gathering (Haas, 2001).

5. Summary and conclusions

In this chapter some computer simulations have been described that show how interactions among simple embodied neural agents living together in the same environment can produce interesting social phenomena related to spatial aggregation, the performance of tasks that require social coordination, communication, and cultural and technological transmission and evolution.

Agents may aggregate spatially because they modify the external environment for other reasons, and other agents respond to these modifications in ways that produce spatial aggregation but do not have this function, or they may aggregate spatially because they develop behaviors that have the function to keep them in proximity to other agents. Agents that are near to each other can coordinate their respective behaviors to accomplish tasks that no individual agent would be able to accomplish by itself alone. Communicating agents develop behaviors that cause specific inputs for other agents and the other agents develop an ability to respond appropriately to these inputs. Communication may be difficult to develop because it requires agents that are able to both emit and understand signals and both behaviors must be advantageous for the individual that exhibits them. Finally, agents that interact may learn by imitating other agents and new behaviors and new technological artefacts can emerge if behaviors and artefacts are selectively transmitted from one generation to the next and with the constant addition of new variants.

Embodied neural agents tend to be simple in the sense that the neural networks which control their behavior contain a small number of units connected together in simple architectures and result in simple behaviors and abilities. Human beings have larger and more structurally complex neural networks which result in much more complex behaviours. Furthermore, the neural networks of human beings have lots of recurrent connections that produce the kind of self-generated inputs that underlie what is called “mental life”: mental images, rememberings, thoughts, predictions, evaluations of courses of action, decisions. However, there appear to be no obstacles in principle to progressively moving from simple to more complex neural networks and behaviors for embodied agents.

The reason why neural agents tend to be so simple is that neural networks cannot be designed or programmed but they must evolve or learn whatever abilities or behaviors they possess. The behavior of a neurally controlled agent depends on the particular architecture and connection weights of its neural network, and the architecture and connection weights that result in some desired behavior cannot be identified a priori and programmed by the researcher. This is why research using neural agents tends to be concerned with simple behaviors: any complex behavior must start as simple and must become progressively more complex by a spontaneous process of learning or evolution. This, however, might be seen as an asset rather than a liability if one assumes that in order to really understand how human agents behave individually and socially one should be able to reconstruct how their behavior has become what it is.

Acknowledgements

This research has been supported by the ECAgents project funded by the Future and Emerging Technologies program (IST-FET) of the European Commission under the EU RD contract IST-1940. The information provided is the sole responsibility of the authors and does not reflect the Commission's opinion. The Commission is not responsible for any use that may be made of data appearing in this publication.

References

Baldassarre G., Nolfi S. & Parisi D. (2003). Evolving mobile robots able to display collective behavior. *Artificial Life*, 9: 255-267.

Baldassarre G., Nolfi S. & Parisi D. (2003). Evolution of collective behaviour in a team of physically linked robots. In Raidl G., Guillot A., Meyer J.-A. (eds.), *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics*. New York: Springer, 581-592.

Baldassarre G., Parisi D. & Nolfi S. (in press). Coordination and behavior integration in cooperating simulated robots. In *From Animals to Animats 8: Proceedings of the VIII International Conference on Simulation of Adaptive Behavior*.

- Beer R.D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173-215.
- Brooks R.A. (1991) Intelligence without reason. In J. Mylopoulos & R. Reiter (eds.), *Proceedings of 12th International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann.
- Clark A. (1997) *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Cliff D.T. (1991) Computational neuroethology: A provisional manifesto. In J-A. Meyer & S.W. Wilson (Eds.), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. Cambridge MA: MIT Press/Bradford Books.
- Cliff D., Harvey I. & Husbands P. (1993) Explorations in evolutionary robotics. *Adaptive Behavior*, 2:73-110.
- Denaro, D. & Parisi, D. (1996) Cultural evolution in a population of neural networks. In M. Marinaro & R. Tagliaferri (eds.) *Neural Nets - WIRN '96*. New York: Springer, 3-19.
- Diamond, J. (1997) *Guns, Germs, and Steel. The Fate of Human Societies*. New York: Norton.
- Haas, J. (ed.) (2001) *From Leaders to Rulers*. New York: Plenum.
- Maturana H.R. & Varela F.J. (1980). *Autopoiesis and cognition: the realization of the living*. Dordrecht: Reidel.
- Maturana H.R. & Varela F.J. (1988). *The tree of knowledge: the biological roots of human understanding*. Boston: New Science Library.
- Mirolli, M. & Parisi, D. (2004) *Language emerges only in kin-related groups or if it used to talk to oneself*. Paper presented at the 5th International Conference on the Evolution of Language. Leipzig.

- Mirolli, M. & Parisi, D. (in press) Language, altruism and docility. In J. Pollack et al. (eds.) *Proceedings of Artificial Life IX*.
- Nolfi, S. & Floreano, D. (2000) *Evolutionary Robotics*. Cambridge, MA: MIT Press.
- Nolfi, S., & Parisi, D. (1997) Neural networks in an Artificial Life perspective. In W. Gerstner, A. Germond, M. Hasler & J.-D. Nicoud (eds.) *Artificial Neural Networks - ICANN '97*. New York: Springer, 733-737.
- Parisi, D. (1997) Cultural evolution in neural networks. *IEEE Expert*, 12, 9-11.
- Parisi, D. (2001) Neural networks and Artificial Life. In G. Parisi & D. Amit (eds.) *Frontiers of Life*. San Diego: Academic Press.
- Parisi, D., Cecconi, F. & Cerini, A. (1995) Kin-directed altruism and attachment behavior in an evolving population of neural networks. In N. Gilbert & R. Conte (eds.) *Artificial Societies*. London: UCL Press.
- Parisi, D., Cecconi, F. & Nolfi, S. (1990). Econets: Neural networks that learn in an environment. *Network*, 1:149-168.
- Parisi, D., Piazzalunga, U., Cecconi, F. & Denaro, D. (1994) Social aggregation in evolving neural networks. In C. Castelfranchi & E. Werner (eds.) *Proceedings of MAAMAW '92*. New York: Springer.
- Parisi, D. & Ugolini, M. (2002) Living in enclaves. *Complexity*, 7, 21-27.
- Pfeifer R. and Scheier C. (1999) *Understanding Intelligence*. Cambridge, MA: MIT Press.
- Quinn, M., Smith, L., Mayley, G. & Husband, P. (in press) Evolving teamwork and role allocation with real robots. In *Proceedings of the 8th International Conference on the Simulation and Synthesis of Living Systems (Artificial Life VIII)*.

Reynolds C.W. (1993). An evolved, vision-based behavioral model of coordinated group motion. In Meyer J.A., Roitblat H.L., Wilson S.W. (eds.), *From Animals to Animats 2: Proceedings of the second international conference on the simulation of adaptive behaviour*. Cambridge, MA: MIT Press.

Simon, H. A. (1990) A mechanism for social selection and successful altruism. *Science*, 250: 1665-1668.

Theraulaz G. & Bonabeau E. (1995). Coordination in Distributed Building. *Science*, 269: 686-688.

Ugolini, M., Parisi, D. Simulating the evolution of artifacts. In D. Floreano & J.-Nicoud (eds.) *Advances in Artificial Life*. New York: Springer, 489-498.

Yong, C.H. & Miikkulainen, R. (2001). *Cooperative Coevolution of Multi-Agent Systems*. Technical Report AI01-287, Department of Computer Sciences, University of Texas at Austin.

Ward C.R., Gobet F. & Kendall G. (2001). Evolving collective behavior in an artificial ecology. *Artificial Life*. Vol. 7, 191-209.

Werner G.M. & Dyer M.G. (1993). Evolution of herding behavior in artificial animals. In Meyer J.A., Roitblat H.L., Wilson S.W. (eds.), *From Animals to Animats 2: Proceedings of the second international conference on the simulation of adaptive behaviour*. Cambridge, MA: MIT Press.