Adaptation and Social Facilitation in a Population of Autonomous Robots

Davide Marocco[*]   **Alberto Acerbi**[*]

[*]Institute of Cognitive Science and Technologies, National Research Council
Via San Martino della Battaglia 44, Rome, 00185, Italy
{davide.marocco, alberto.acerbi}@istc.cnr.it

Abstract

In this paper we present an evolutionary simulation in which a population of 10 mobile robots has to develop a simple behavior consisting in the discrimination of two different foraging areas present in the environment. We show that, given a minimal selective pressure, a combination of individual learning, social facilitation, and selection at the population level can lead to effective results. Consistently with ethological findings, we argue that such dynamics in which simple mechanisms interact at different levels can account for instances of social transmission of behaviors usually interpreted as intentional acts of copying.

1.   Introduction

Human beings show a behavioral diversity that cannot be compared to the variability present in other species. This variety of behavioral patterns is not explained by genetic variation (humans, contrasted to phylogenetically related species, are greatly homogeneous from the genetic point of view) and, apparently, can be explained only in part by the environmental diversity (Pagel and Mace, 2004). Processes of social learning are, in general, presumed to play the most important role in shaping this diversity (Richerson and Boyd, 2004).

However, this correct argument is at times used to derive some implications that could not be always accurate. In the first place, it is supposed that social learning is limited to human beings, or, at least, that it is restricted to some "cognitively gifted" species like great apes or primates in general. Secondly - and closely related to the previous one - there is the conviction that genuine social learning is permitted only by sophisticated cognitive machineries and that only high-level forms of social learning, like imitation, can support a reliable transmission of behaviors between individuals. Finally, social and individual learning are considered as two processes that differ in a substantial way, with the former considered basically as a process of explicit copy of a conspecific's behavior. For reviews on

social learning in ethology see (Heyes and Galef, 1996, Zentall and Galef, 1988), also related to artificial agents and robots (Dautenhahn and Nehaniv, 2002, Nehaniv and Dautenhahn, 2007), for a general review (Hurley and Chater, 2005).

In general, it is probably true that the role that social learning plays among human beings is greater than the role that it plays among other species and it is not clear, up to now, what kind of mechanisms can permit the relative stability of the cultural patterns in human groups and support cumulative cultural evolution. However, especially in the last twenty years, ethologists have shown how social learning can be important for the development of complex behavioral skills in primates (Tomasello, 1996), but also in other vertebrates like, just to cite the more influential studies, rats (Laland and Plotkin, 1990, Galef, 1996), birds (Sherry and Galef, 1984), fish (Dugtakin, 1996). In addition to this, the social transmission of behavior in these species is often realized without the need of complex cognitive abilities but by simple processes that exploit the dynamics between learning at the individual level, the characteristics of the environment, and the genetic evolution at the population level, without an easy-to-trace distinction between the processes.

The human case of course could be different, however is not impossible that also in the human development simple mechanisms of social influence can guide and enhance individual learning and eventually play some role in the development of more complex social learning skills or act together with them. We prefer to use a minimal definition of social learning, by which we do not refer to a specific modality of learning but to a general set of adaptive modifications that take place in a single agent during its lifetime and that are influenced *in some way* by the interaction with conspecifics.

Individual based modeling or, as in our case, the development of artificial organisms (robots or simulated agents) through artificial life techniques can be an excellent way to explore this kind of dynamics (Noble et al., 1999, Noble and Todd, 2002). In particular, artificial life techniques are particularly suitable to

take into account different levels of analysis, like processes of learning at the individual level, interactions between individuals and between individuals and the environment, and selective pressures at the population level, that are often very difficult to consider resting only on empirical observations or laboratory experiments. Mathematical models, on the other hand, had proven to be extremely important and insightful tool for understanding social learning and cultural evolution dynamics (Richerson and Boyd, 2004) but, in most cases, they rely on the simplifying assumption that a behavior is simply passed on from an individual to another. Even if this simplification can be correct in most of the cases, findings by ethologists and psychologists show that a behavior can be socially acquired in many possible ways (Call and Carpenter, 2002, Galef and Laland, 2005) and it is worthwhile to argue that the specific way in which the behavior is acquired can be an important factor for the resulting dynamic at the population level.

In our contribution we present a simple experimental scenario in which a team of 10 simulated robotics agents has to develop an adaptive behavior consisting in the discrimination of two different foraging areas present in the environment. The team is subjected to a selective pressure and an on-line genetic algorithm operates on the agents' learning modalities. In particular, as we will explain in the next section, two characteristics can be inherited: the strength of the learning and its modality (individual or social). Individual learning is a simple blind process that reinforces the "preference" of an individual for one of the two areas, independently from the features of the specific area. In the same fashion, social learning just consists in the fact that the strength of this process is modulated by the number of individuals located nearby. We show, in the third section, that, given a minimal selective pressure, social learning evolves and that the combination of a blind but socially influenced individual learning with selection at the population level can lead to an adaptive social transmission of the behaviors.

2. The experimental set-up

In order to investigate the adaptive role played by the social facilitation in acquiring new behavior, we ran a robotic simulation in which a population of robots have to survive in an unknown environment. The simulation consists of multiple replications of the same basic task, in which robots have been tested in different environmental conditions. In particular, each condition provides a different selective pressure under which robots have to adapt in order to survive.

2.1 The robotic model

The environment consists in a square arena of 200 x 200 cm surrounded by walls. The floor of the arena is grey and contains two circular areas with a diameter of 60 cm colored in black and white, respectively (Fig. 1). A population consists of 10 simulated e-puck robots (Fig. 2(a)) which are provided with simple sensory-motor capabilities that allow them to move, gather information from the environment and to produce sound signals with a fixed intensity. The robots have a circular body with a radius of 37 mm, 8 infrared sensors placed around the body, 1 ground sensor placed on the bottom of the robot, 1 microphone and 1 speaker.
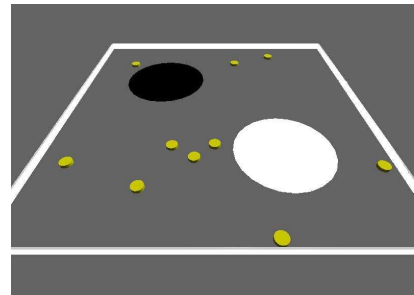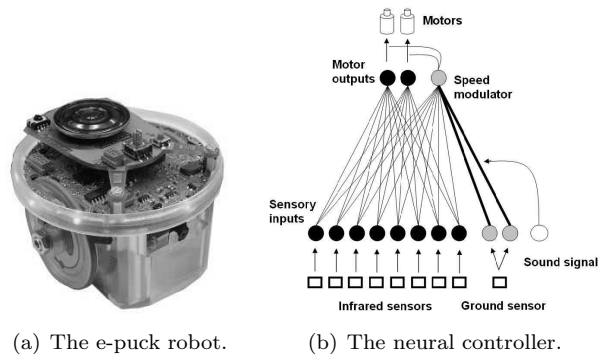

Figure 1: The environment of the experiment. Black and white circles represent the two foraging areas. Small gray circles represent the 10 robots.

(a) The e-puck robot.   (b) The neural controller.

Figure 2: The e-puck robot and the neural controller. Thick lines represent the weights that are modified in the simulation. As it is shown, the speed modulator does not operate on the outputs but it acts directly on the motors' activations. The intensity of the sound signal is used to modify the learning rate in case of *social individuals*. See the text for details

Robots' neural controllers consist of artificial neural networks with 10 sensory neurons and 3 motor neurons (Fig. 2(b)). Eight sensory neurons encode the activation states of the 8 infrared sensors (which detect obstacles up to a distance of 4 cm) and two neurons measure the activation of the ground sensors that encode the color

of the floor below the robot (gray [0 0], white [1 0], and black [0 1]). The first two motor neurons encode the desired speed of the two corresponding wheels normalized between -MaxSpeed and +MaxSpeed. The third motor neuron instead encodes the maximum speed of both wheels (i.e. the MaxSpeed parameter is normalized in the range [0;10]). The third motor neuron thus acts as a modulator that regulates the speed with which the robot moves or turns at each time step. The 8 sensory neurons that encode the state of the infrared sensors are connected to all the three motor neurons. The two sensory neurons that encode the color of the ground, instead, are connected to the third motor neuron only (i.e. to the motor neuron that regulates the maximum speed of the two wheels).

The robots continually emit a sound signal of fixed intensity ((0.1 I) - See section 2.3) through their speakers when they are located in a foraging area, independently from its color. Moreover, the robots can detect through their microphones the signals produced by other robots located nearby.

As in the case of (Acerbi et al., 2007), in order to provide the robots of the minimal sensory-motor capabilities that allow them to move around the environment and avoid walls, prior the starting of the simulation, robots are pre-trained through an evolutionary method (Nolfi and Floreano, 2002) for developing an obstacles avoidance behavior and for exploring the environment in an arena that does not contain colored areas. This pre-training process only concerns the strength of the connections between the 8 infrared sensors and the three motor neurons. See (Acerbi et al., 2007) for further details. Because of this procedure, the behavior of the robots outside the areas is already determined and cannot be modified during the simulation, therefore robots can only modify their behavior inside the areas.

In the simulation the population is placed in the environment and the task of the robots is to adapt and survive as long as possible in an environment that contains the white and the black target areas.

Each robot of the population starts its life with a fixed amount of energy (i.e. 2000 units) and the environment provide two different source of selective pressure. In all experimental conditions every time step spent outside the areas causes the loss of 1 unit of energy, while every time step spent in the black area causes a variable loss of energy units, depending on the experimental condition. The white target area does not affect the energy level of the robots. In the simulation we test the robots in 16 different conditions in which the loss of energy in the black area varies from 0 to 15 energy units per time step. In this way, the manipulation of the amount of energy lost in the black area determines the selective pressure under which robots have to survive.

Therefore, with the exception of the first experimen-

tal condition (energy loss equal to 0), the simulation task consists in locating and spending time in the white target area, while avoiding the black one. In practice, given the pre-evolved exploratory behavior, the free parameters that encode the strength of the connection between the two ground sensor neurons and the third motor neuron should be set so to allow the robot to slow down and, eventually, to stop in the white target area and to quickly move out from the black target area. It is worthy to note, however, that the robots do not have any direct feedbacks about the *quality* of a target area or about their own energy. In fact, they only have information about the colors of an area when they are in it, but they do not know how remaining in a particular area affects their energy state. This means that the experience of a particular area does not provide any information about the correct behavior that should be performed in that area.

In all experimental conditions, the population is evaluated for 1 trial lasting 20000 time steps (of 100ms each) and each robot its allowed to live for about 2000 time step. When a robot finishes up its energy it dies and when a robot lives for 2000 time step it has a small probability of dying for every further time step of its life (1‰). When a robot dies, it is replaced with a new robot placed in the environment in a randomly chosen position and orientation, with full initial energy. The dynamic of the adaptation to the environment consists of two interplaying processes: one that acts at a genetic level and another that acts at a social level.

2.2 Genetic inheritance

Each robot of the population has a genetic code consisting of two parameters (i.e. 2 genes) that determine two different aspects of its social reliance. The first parameter ω, which takes values -1 or 1, determines if a robot learns only on the basis of its own experience or if it also relies on social cues provided by the other robots. The second parameter ϕ (which varies in the range [0;1]) determines the strength of the learning. If ω is equal to -1, it determines the value of the learning rate of the robot. If ω is equal to 1, this gene determines the extent to which the robot relies on social cues. More details about the role of ω and ϕ are provided in the next section.

At the beginning of the simulation the two parameters are randomly set, so to generate a mixed population. Moreover, when a robot dies during the simulation, an on-line genetic algorithm acts on the population, by replacing the died robot with a newborn that inherits the two genes from a randomly individual selected among the robots that live from at least 2000 time steps, i.e. the oldest robots of the population. Notice that the duration of a robot's life is an indirect indication of the adaptiveness of its behavior. Indeed, the fact that a robot lives more than 2000 time steps means that it acquired the

capacity to correctly discriminate between the two areas during its ontogenetic development. During the genetic transmission process gene ω can switch its original value with a probability of 10% and gene ϕ is mutated by adding a random value in the range [-0.1;0.1].

2.3 Learning

This process affects the way in which robots modify their behavior during the life. In particular, the two synaptic weights that connect the two ground sensor neurons and the third motor neuron are trained by means a non-directional learning algorithm. In practice, the weights are initialized, at the beginning of each trial, with a null value, that is, the robots tend to go away from both the areas. When a robot happens to stay inside a target area, a learning algorithm acts to reinforce with a positive feedback the individual *preference* for that area, independently from the color of the area and consequently, from the effect on its energy. This positive feedback mechanism is realized by modifying the two weights every time an individual is inside one area through the addition of a randomly chosen value in the interval [-lrate;+lrate] to their current value, and by retaining only the variations that lead to an increase of the time spent by the individual in the corresponding target area. The learning rate is modulated differently depending on the value of ω:

$$ lrate = \begin{cases} \phi & \text{if} \quad \omega = -1 \\ \phi \cdot Q_I & \text{if} \quad \omega = 1 \end{cases} \qquad (1) $$

where Q is the intensity of the signal perceived by the robot inside the area expressed in I units and it varies linearly in the range [0.1;1.0] according to the number of individuals located in the same foraging area: $0.1I$ corresponds to one robot in the area (the learner robot itself), and $1.0I$ to all the 10 robots in the same area. In practice, in the social case the learning rate is function both of ϕ and of the intensity of the signal detected.

3. The emergence of social exploitation

As we said above, the described simulation has been run considering sixteen different values of energy loss in the black foraging area, i.e. from 0 to 15. In particular, for each replication, 100 different populations of ten robots have been randomly generated and tested in the same environmental conditions.

From the results of the simulation we can observe an interesting relationship between the exploitation of social cues by the robots and the amount of selective pressure provided by the environment.

In order to have a measure of this effect with respect to the different environmental conditions, for every replication we count the number of robots that rely on social cues in each population. To do this, we count how many robots express the parameter ω equal to 1 for each of the 100 population in all the conditions and we consider a population as *social* only if the majority of the the robots have an ω value equal to 1. In particular, in the following sections we we will call *social population* a population in which more than five robots use social cues to orient their individual learning and, similarly, *social individual* a robot that expresses that particular behavior. On the basis of this measure, the graph in figure 3 clearly shows that the number of *social populations* grows with the increase of the selective pressure, due to the increase of the energy loss in the black target area. It is interesting to note, however, that only a minimal adaptive pressure leads the population towards the exploitation of the sociality. In fact, only in the case of no adaptive pressure (energy loss equal to 0) the amount of *social populations* is less than the 50% on the total of 100 replications.
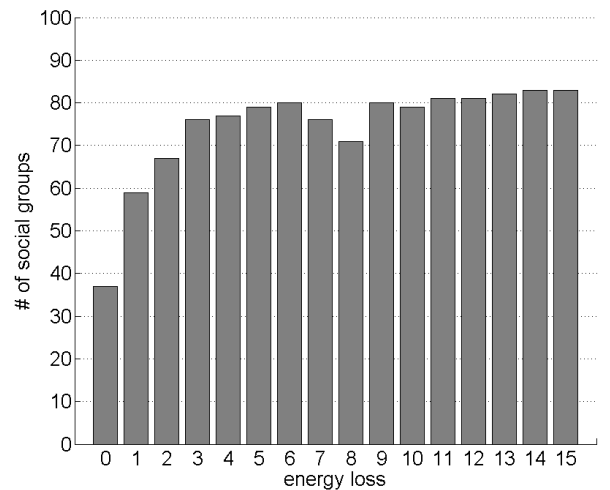

Figure 3: Number of *social populations* for each experimental condition.

The adaptiveness of the social exploitation is confirmed by the differences of the mortality (i.e. the number of robots that "die" during each trial) between the *social populations* with respect to the *non-social* ones. Indeed, by looking at Fig. 4 we can see that the mortality rate of the former, with the exception of cases 0 and 1 of energy loss, is always lower of the mortality of the latter and, after 4 units of energy loss, it remains stable despite the increment of the selective pressure provided by the environment.

We discussed above the fact that, given a minimal selective pressure, most of the replications produce *social populations*. More in detail, Fig. 5 shows the final compositions of the populations for all the replications in the different experimental conditions. We can remark that the number of populations with all *social individuals* (10)
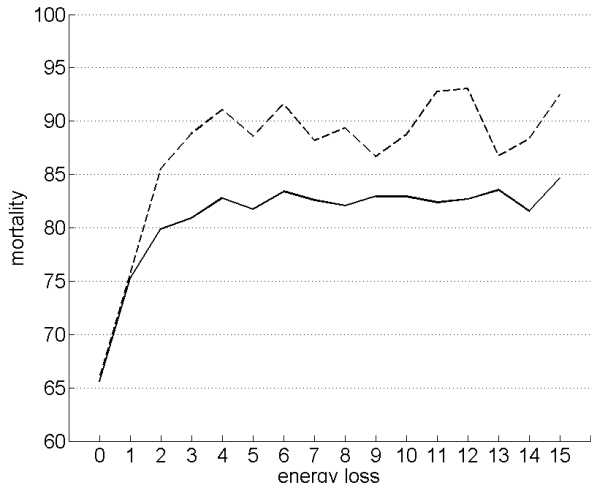
Figure 4: Average mortality for each experimental condition. Dashed line: *non-social populations*. Continuous line: *social populations*.
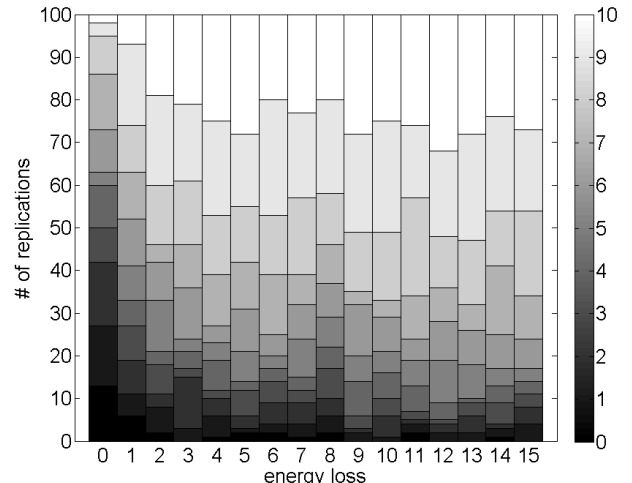

Figure 5: Final composition of the populations for all the replications in the different experimental conditions. Depending on the number of *social individuals* populations go to black (no *social individuals*) to white (10 *social individuals*). See colorbar on the right.

slightly increase during the first experiments, parallel to the increasing of the selective pressure, till it reaches a more or less stable value. Moreover, is interesting to note that the increasing of *social populations* is mainly due to the increasing of groups in which all or most individuals are social (groups made by 8, 9, or 10 *social individuals* account for approximately the 60-70% of the replications after energy loss equal to 9). These two facts, considered together, suggest a more subtle dynamic in the evolution of social exploitation.

In particular, the latter finding can indicate that, in the majority of the cases in which we observe social exploitation, we are in presence of a positive feedback that guide the adaptive search at both phylogenetic and ontogenetic level. This hypothesis contribute to interpret the overall results of the simulation.

The relation between adaptive pressure and social exploitation derives by the interactions among the three forces that, at different levels, guide the adaptation process of the populations: the individual learning algorithm, the social interaction, and the genetic transmission. In fact, given that the learning algorithm provided to the robots is totally "blind" with respect to the adaptiveness of the behavior, the only way of orienting the learning process, at the ontogenetic level, is related to the fact that individuals with adaptive behavior tend to survive more than others. Therefore, under a selective pressure provided by the environment, a *social individual* has more chances of acquiring the adaptive behavior, because robots with maladaptive behavior tend to die with more frequency.

At the same time, at the genetic level, thanks to the the transmission mechanism, individuals with the adap-

tive behavior are favored in propagating their genes. This fact produces an increase in the number of *social individuals* within a given population. In turn, this makes more advantageous (with respect to the early stages of the evolution) for a newborn to be social.

4. Conclusion

We have demonstrated on a simple experimental scenario how social facilitation of individual learning combined with a selection process operating at the population level can be useful in developing adaptive behaviors in small populations of autonomous robots. It is important to remark how this adaptive mechanism, that involves a form of social transmission, or diffusion, of behaviors, can be considered a form of social learning and that social learning, here, is not the evolved mechanism that modulate the strength of the individual learning, but the combination between this mechanism, the individual learning, and the selective pressure. Note that the effectiveness of the this process is not simply related to the magnitude of the learning rate *per se*: it is precisely the fact that social influence can modulate this magnitude during the time that makes the whole process adaptive. In fact, we have to consider that a learning rate not socially modulated will affect the behavior of a robot in the same way in both areas, making impossible for newborn robots to develop selectively a preference for only one of the two areas (see also (Acerbi et al., 2007) in which also the behaviors of populations with fixed learning rates have been analysed). As we pointed out in the Introduction, we believe that it is plausible that at least

some forms of what is named "social learning" are the outcomes of this kind of dynamics and that artificial life simulations are an important tool to test this hypothesis.

Of course, further works are needed to develop this project. For example, in this study we do not considered deeply the role of the parameters ϕ (that determined the strength of learning): how different strengths can impact the adaptivity of the process? Is it always the best solution to learn faster? We can imagine that, in relation with selective pressures on learned traits, social learners can adaptively vary the strength of learning so to learn, perhaps, only when most of the individuals are in an area. In future, we also plan to investigate the scalability of the mechanism in different environmental conditions and with different selective pressures, namely by increasing the number of foraging areas, by increasing (or decreasing) the number of robots, by increasing the complexity of the perceptual categorization process, or by changing the environmental conditions through time. Our hope is that this study could be useful to draw a preliminary sketch of a bigger picture that we consider deserving attention, possibly also in relation to human ontogenetic development.

Acknowledgments

References

Acerbi, A., Marocco, D., and Nolfi, S. (2007). Social facilitation on the development of foraging behaviors in a population of autonomous robots. In Almeida e Costa, F., Rocha, L., Costa, E., Harvey, I., and Coutinho, A., (Eds.), *Advances in Artificial Life. Proceedings of ECAL 2007*, volume 4648 of *Lecture Notes in Computer Science*, pages 625–634. Springer.

Call, J. and Carpenter, M. (2002). Three sources of information in social learning. In Dautenhahn, K. and Nehaniv, C. L., (Eds.), *Imitation in Animals and Artifacts*, chapter 9, pages 211–227. MIT Press, Cambridge.

Dautenhahn, K. and Nehaniv, C. L., (Eds.) (2002). *Imitation in animals and artificats*. MIT Press, Cambridge.

Dugtakin, L. A. (1996). Copying and mate choice. In Heyes, C. M. and Galef, B. G., (Eds.), *Social learning in animals: the roots of culture*, pages 85–105. Academic Press, San Diego.

Galef, B. G. (1996). Social enhancement of food preferences in norway rats: a brief review. In Heyes, C. M. and Galef, B. G., (Eds.), *Social learning in animals: the roots of culture*, pages 49–64. Academic Press, San Diego.

Galef, B. G. J. and Laland, K. N. (2005). Social learning in animals: Empirical studies and theoretical models. *BioSceince*, 55(6):489–499.

Heyes, C. M. and Galef, B. G., (Eds.) (1996). *Social learning in animals: the roots of culture*. Academic Press, San Diego.

Hurley, S. and Chater, N., (Eds.) (2005). *Perspectives on imitation: from neuroscience to social science*. MIT Press, Cambridge.

Laland, K. N. and Plotkin, H. (1990). Social learning and social transmission of foraging information in norway rats (rattus norvegicus). *Animal Learning and Behavior*, 18:246–251.

Nehaniv, C. L. and Dautenhahn, K., (Eds.) (2007). *Imitation and Social Learning in Robots, Human and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press, Cambridge.

Noble, J. and Todd, P. M. (2002). Imitation or something simpler? modelling simple mechanisms for social information processing. In Dautenhahn, K. and Nehaniv, C. L., (Eds.), *Imitation in animals and artifacts*, pages 423–439. MIT Press, Cambridge, MA, USA.

Noble, J., Tuci, E., and Todd, P. M. (1999). Social learning and information sharing: an evolutionary simulation model of foraging in norway rats. In Floreano, D., Mondada, F., and Nicoud, J. D., (Eds.), *Proceedings of the Fifth European Conference on Artificial Life (ECAL' 99)*, pages 514–523, Berlin. Springer-Verlag.

Nolfi, S. and Floreano, D. (2002). Synthesis of autonomous robots through evolution. *Trends in Cognitive Sciences*, 6(1):31–37.

Pagel, M. and Mace, R. (2004). The cultural wealth of nations. *Nature*, 428:275–278.

Richerson, P. J. and Boyd, R. (2004). *Not by genes alone: how culture transformed human evolution*. Chicago University Press, Chicago.

Sherry, D. F. and Galef, B. G. (1984). Cultural transmission without imitation: Milk bottle opening by birds. *Animal Behavior*, 32:937–938.

Tomasello, M. (1996). Do apes ape? In Heyes, C. M. and Galef, B. G., (Eds.), *Social learning in animals: the roots of culture*, pages 319–346. Academic Press, San Diego.

Zentall, T. R. and Galef, B. G., (Eds.) (1988). *Social learning: psychological and biological perspectives*. Lawrence Erlbaum, Hillsdale.