

Imitation or something simpler?

Modelling simple mechanisms for social information processing

Jason Noble and Peter M. Todd
Center for Adaptive Behavior and Cognition
Max Planck Institute for Human Development
Lentzeallee 94, D-14195 Berlin, Germany
noble, ptodd@mpib-berlin.mpg.de

1 Introduction

In the pine forests of Israel, black rats *Rattus rattus* have hit upon a novel feeding technique. They strip the scales from pine cones to obtain the nutritious seeds inside (Terkel, 1996). The behaviour seems to be socially learned in some way, rather than genetically inherited: Terkel has shown that rat pups will learn to strip cones if they are born to a naive mother but then fostered to an experienced one. What mechanism is implicated in the transmission of this behaviour? It could be that rat pups are genuinely imitating their mothers, but this would mean that they are solving a complex correspondence problem in translating novel visual input—from observations of cone manipulation—into appropriate motor outputs. Might a simpler mechanism be involved?

The social transmission of pine cone stripping behaviour is only one example of the general strategy of gaining information from the behaviour of one's conspecifics. Every animal is constantly facing the problem of what to do next: given the various cues that it can detect in its environment at any one time, what behaviour is most likely to further its ultimate goals of survival and reproduction? In all but the most solitary of species, the cues available from the environment will include information about the behaviour of other animals.

Investigation of the general question as to how animals might learn from or be otherwise influenced by the behaviour of their conspecifics has often been overshadowed by a focus on the specific issue of imitation. This emphasis has been around for some time: Galef (1988), in a historical review, points out that the debate over whether nonhuman animals are capable of imitation dates back at least to the dispute between Darwin and Wallace over whether the human mind is of evolved or divine origin. One product of this long controversy has been a plethora of definitions of imitation (see Zentall and Galef, 1988; Heyes and Galef, 1996), which we lack the space to adequately review. For the purposes of our argument, we will simply take true imitation to be the goal-directed copying of another's behaviour. As Tomasello (1996) notes in his discussion of "imitative learning" (p. 324), successful imitation—in the sense we are interested in—requires not only perceiving and reproducing the bodily movements of another, but understanding the changes in the environment caused by the other's be-

haviour, and finally being able to grasp the “intentional relations” between these, i.e., knowing how and why the behaviour is supposed to bring about the goal. We realize that in stipulating such a definition of imitation, we are setting the bar quite high. Our aim is not to define imitation away, but to offer an explicit definition and thereby focus attention on the issue of mechanism. We feel that “imitation,” with its connotations of intelligent, purposeful action, has been used in the past as a vague, catch-all term to cover instances of social learning for which the underlying mechanisms were very poorly understood.

Many authors have focused on imitative learning because it seems likely to be implicated in the origin of culture (e.g., Tomasello et al., 1993). The existence and the title of this book are themselves testaments to the abiding interest that psychologists, biologists and cognitive scientists have had in imitation. Nevertheless, our contention in this chapter is that the power of simpler social information processing mechanisms has been underestimated. In this we are not alone: several phenomena that were once seen as clearly imitative, such as the opening of milk bottles by birds (Fisher and Hinde, 1949) and the washing of food by monkeys (Kawamura, 1959; Kawai, 1965), have since been questioned (Galef, 1976; Sherry and Galef, 1984; Galef, 1988; Cheney and Seyfarth, 1990; Byrne, 1995). We will argue that individual-based simulation modelling, already a popular technique in the field of adaptive behaviour, is an excellent tool for exploring how animals can be successful using simple mechanisms. We will also review the literature on social learning, and identify mechanisms that provide parsimonious explanations for many apparently imitative phenomena.

1.1 A cautionary tale

Human observers of animal and robot behaviour have a propensity to invoke mechanisms that are more complex than those strictly needed to explain the observable facts. Braitenberg (1984) demonstrates this nicely in his book *Vehicles* by asking us to imagine a robot that appears to behave aggressively in the presence of light sources. This robot turns to face any light source within the range of its sensors and accelerates towards it, eventually smashing into it at high speed. If the source moves, the robot will pursue it accordingly. Braitenberg suggests that we will find it difficult to avoid attributing a *dislike* of the light source to the robot (for empirical support of this point see Bumby and Dautenhahn, 1999). An observer—let’s say a cognitive psychologist—asked to speculate about the internal workings of the robot might guess that it contains some sort of light-detection mechanism, which invokes an attack response, passing information to a central processor that in turn sets off the motor movements needed to bring the robot closer to the light. In fact, the robot is very simply constructed: Figure 1 shows that wires from two light sensors at the front of the machine are cross-connected to two motors at the back. This is all that is required to bring about the behaviour described.

Once you know how the robot works, it is clear that the given design will get the job done. But it is very difficult to correctly deduce the robot’s internal workings from its behaviour alone. Braitenberg refers to this phenomenon as the “law of uphill analysis and downhill invention.” We want to apply two lessons from Braitenberg’s fable to our review of social learning. Firstly, the mechanism that produces a given behaviour may be a great deal simpler than appearances suggest—quite possibly this has something to do with the human tendency to interpret the world from an “intentional stance” (Dennett, 1987). Secondly, the best way to understand the mechanisms necessary for a certain behaviour may be to build an artificial system that performs that behaviour (see




Figure 1: Braitenberg's (1984) Vehicle 2b—crossed connections between light sensors and motors produce apparently purposeful pursuit of a light source.

also Simon, 1995, in this regard).

1.2 Why and how questions

The idea that animals are influenced by the behaviour of others poses both why and how questions, that is, questions of evolved function and questions of mechanism. In recent years there has been some progress towards understanding the function of social learning. Models of cultural transmission (Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 1985), highly horizontal (within-generation) transmission (Laland et al., 1993, 1996), and what economists call herding behaviour (Banerjee, 1992; Bikhchandani et al., 1998) help to delineate the conditions under which it will be advantageous for individuals to learn from others rather than finding things out for themselves. However, these models tend to result in rather general conclusions. For example, Laland et al. (1996, p. 140) summarize the results of work on cultural transmission thus:

When environments change very slowly, all information should be transmitted genetically, since the modest demands for updating are easily met by the genetic system responding to selection. When environmental change is very rapid, tracking by pure individual learning is favored. At intermediate rates of change, social learning is an advantage.

Results like these are useful, but to get a complete picture of any one instance of social learning we also need to understand the mechanism involved—something which is less often attempted. Terkel's (1996) black rats, for instance, clearly gain a selective advantage if they learn to strip pine cones, because there is little alternative food available in their environment. But *how* do they do it? What process allows the rat to take in sensory information relating to pine cones and the feeding habits of conspecifics, and from this acquire the behaviour? In asking these questions we seek an explanation at the cognitive level; we want to know what behavioural program or algorithm underlies the rats' performance. Another way of considering the problem is to suppose that we wanted to construct a simulated or robotic rat. Assuming we had given our construction sensory and motor capabilities that matched those of the real

animals, what would be the minimal set of instructions we could give it that would allow it to mimic real *behaviour*?

Much of the work that professes to be about mechanisms of social learning does not really confront this challenge. For example, Bandura (1969) discusses the process of “identification,” whereby a human observer comes to adopt the attitudes and behaviours of a model. Tomasello et al. (1993) talk about “perspective taking” as being central to the ability to truly imitate another. In neither of these cases do the authors go any further than labelling the phenomenon; we are left in the dark as to just how identification or perspective taking might be achieved. This is a deficiency we believe could be redressed through the use of simulation modelling.

1.3 Modelling simple mechanisms

As observers of animal behaviour, we are in a similar position to the observer of Braitenberg’s hypothetical vehicle. We can see the outward behaviour, but the underlying mechanism is hidden from our view. What’s more, in the animal case there is no indulgent designer who will take the cover off and show us how it works. How, then, are we to avoid ascribing mechanisms of spurious complexity to the animals whose social behaviour we study? We believe that one answer lies in the “downhill invention” part of Braitenberg’s law, that is, in model building.

A number of authors (Judson, 1994; Miller, 1995; Di Paolo, 1996; Todd, 1996) have recently made a case for individual-based simulation modelling as a way of making progress in areas like social learning, in which the empirical data are often insufficient to distinguish between competing theories. In such simulations, individual organisms interact in a shared environment—although both the organisms and the environment may be very simple. Sometimes the behavioural strategies of the organisms, as well as other model parameters, are set by the designer. Alternatively, many of these simulations incorporate an evolutionary process, in which organisms that are more successful have a greater likelihood of passing on their strategy to the next generation: this kind of artificial selection can be employed either as a general tool for optimization, or as a way of looking explicitly at evolutionary dynamics.

The method is inspired by work in the field of artificial life and the simulation of adaptive behaviour—for introductory reviews, see Langton (1989) and Meyer (1994) respectively. Simulation is not new as a research tool in biology, but conventional biological simulations tend to model whole populations, abstracting away from the individual organism, and they tend to be extensions of simple game-theoretic models, thus incorporating radical simplifying assumptions such as random mating and the absence of a spatial distribution. Social learning and social information processing in general are self-evidently about interactions between individuals mediated by an environment, and we believe that individual-based simulations can be useful tools in their investigation, much as they have been for studying the evolution of individual learning (e.g., Todd and Miller, 1991; Belew and Mitchell, 1996). It has long been recognized within fields like artificial life that complex global phenomena can arise from simple local rules, and this is precisely what we will suggest is happening in many social information processing contexts: individuals follow a simple rule (e.g., “stay close to your mother”) and, in combination with some form of learning, the overall pattern of behaviour that arises makes human observers suspect complex imitative abilities.

Previous work on behavioural simulation has certainly considered social dynamics in contexts such as communication, dominance and territorial behaviour, and flocking or schooling. However, as yet relatively few models have addressed the specific topic

of social information processing. Some simulations have modelled stigmergic communication, a process by which one individual influences another via alterations in the environment (Theraulaz and Bonabeau, 1999). For example, Deneubourg et al. (1992) and Bonabeau et al. (1994) looked at the kinds of nest structures that could be built if a group of simulated wasps could not perceive each other, but only each other's construction work. These models capture an indirect form of social influence. A more clear-cut example is the work of Toquenaga et al. (1995); the authors constructed a simulation of foraging and nesting behaviour in egrets. They used their model to demonstrate that a tendency to approach others (see section 2.3 below) is more likely to evolve when food resources are patchy rather than evenly distributed; when food comes in discrete patches, approaching others often equates to approaching food. Furthermore, the overall pattern of behaviour that occurs when all of the birds exhibit this approach tendency strongly resembles flock foraging and colonial roosting.

We see simulations like that of Toquenaga et al. as suggestive of a productive feedback cycle between evolutionary simulation models and empirical research on mechanisms of social learning. The model shows that a complex pattern of foraging and roosting *could* emerge from something as simple as a tendency to approach conspecifics. This idea now stands as a kind of null hypothesis; we only have to change our ideas about egrets if someone can demonstrate that the real animals behave in a way that cannot be accounted for by this simple rule. Toquenaga et al.'s simulation could also be used to generate predictions of interest to field biologists: for example, what happens if the food distribution suddenly or periodically changes from patchy to even? How does this compare with parallel real world cases? Through this kind of interaction between the model and the world, we can build our confidence that a certain process not only could explain the observed results, but is in fact the real mechanism involved. In a field such as social learning and imitation, where the data are relatively clear but their interpretations hotly contested, this would be no small achievement.

2 Some simple mechanisms

Once we are focused on *how* social information can affect the behaviour of individuals, a variety of simple mechanisms become candidates for the production of imitative or imitation-like behaviour. In this section we present a selection of such social information processing mechanisms; we owe much to review articles by Galef (1988) and Zentall (1996). We claim that these mechanisms are *simpler* than true imitation; by this we mean roughly that the algorithms involved can be specified using fewer bits of information. We do not anticipate serious resistance to this idea. True imitation, as we have defined it in section 1, necessarily involves multiple processes dealing with behaviour recognition and matching, noticing the effects of behaviour on the environment, and imputing goals to others. Even if a clever short cut could be demonstrated for each sub-task, the overall algorithm for true imitation would remain relatively complex because of the number of steps involved.

Observers tend to suspect imitation when they see one animal perform a behaviour X , and then see another animal doing X shortly afterwards. Perhaps the simplest way for such behavioural matching to occur is for two animals to independently choose the same action because they are faced with similar environmental stimuli and have similar inherited or learned behavioural priorities. For example, Brockmann (1995) looked at mating in horseshoe crabs: unattached males congregate around certain copulating pairs but not others. This behaviour could have been due to mate choice copying by

the males (see Dugatkin, 1996) but Brockmann established that it is in fact the result of common preferences for certain characteristics in females. Tooby and Cosmides (1992, p. 115-6) describe a thought experiment that makes the same point: if all the people on Earth were replaced with jukeboxes, and if each jukebox was equipped with a clock and a location sensor that determined the song it would play, then at any one time all the jukeboxes in Rio de Janeiro would play one song, and all those in Beijing would play another. The result might be suggestive of cultural transmission or social learning, but appearances would be deceptive.

2.1 Social facilitation

One very simple kind of social influence involves behaving differently depending on the presence of conspecifics. Imagine a species that follows the behavioural rule “Don’t do anything unless others are nearby.” Suppose that two animals are hungry and alone, and then come into contact with each other: now they can begin to eat. An observer might declare that one animal was imitating the other, but of course it is doing no such thing. This is approximately the logic of social facilitation (Zajonc, 1965), although Zajonc’s original theory is more detailed, and is intended to explain the effects of an audience on human performances. Clayton (1978) suggests that social facilitation in animals could be due to the reduction of isolation-induced fear, and this makes intuitive sense: a rule like “Don’t do anything except watch for predators unless others are nearby” could clearly be adaptive, and observations that birds will spend less time on vigilance when feeding with others (e.g., Sullivan, 1984; Bekoff, 1995) support the idea. Hemelrijk’s (1999) simulation of dominance interactions among primates offers a synthetic example of social facilitation. Her simulated organisms simply turn on the spot when alone, but if a conspecific comes within detection range, their behaviour switches accordingly: the two will approach each other (and possibly fight), but neither organism is truly imitating the other.

2.2 Contagious behaviour

Contagious behaviour is exemplified by a rule such as “If others are fleeing, flee also.” The idea is that the stimuli produced by the performance of a particular behaviour serve as triggers for others to behave in the same way. For instance, consider a robot that has a direct connection from its auditory sensors to its motors, such that the characteristic sound of another robot moving rapidly causes it to do likewise. In a group of these robots, any stimulus that causes one of them to flee will lead to a chain reaction of rapid movements. Possible examples of contagious behaviour include flight responses, movement in flocks or schools, and chorusing by birds and dogs (Galef, 1988). Laughing and yawning are excellent examples of contagious behaviour in humans (Provine, 1996).

Zentall (1996) stresses that contagious behaviour must have a genetic basis, i.e., it must involve the triggering of an instinctive response. Zentall wants to prevent the term from being used to explain too broad a range of phenomena, and clearly there is a distinction to be made between mere contagious behaviour and the purposeful copying of an act novel to the observer. However, the capacity for simple associative learning is virtually universal in animals, and behavioural contagion combined with learning may result in more than the sum of the parts. Continuing with our robotic example, suppose that these robots have the ability to build associative links between stimuli, i.e., to learn, perhaps through the operation of a Hebbian neural network. Let us further suppose that

one of the robots has developed a “fear” of the colour green. When a green stimulus is presented, the focal robot flees, precipitating a wave of contagious flight by the others. But the other robots learn, and if this sequence is repeated several times they will develop an association between the previously neutral green stimulus and their own flight response, i.e., they will come to fear green as well. In other words, contagious behaviour plus learning could lead to something like cultural transmission.

2.3 Stimulus enhancement

Stimulus enhancement (also called local enhancement) is what happens when animals obey a rule like “Follow someone older than you, and then learn from whatever happens.” Galef’s (1988b, p. 15) description is apt:

A tendency on the part of naive individuals to approach conspecifics, alterations conspecifics have made in the environment, or objects they have contacted, can increase a naive individual’s probability of exposure to one set of stimuli rather than others. Enhanced exposure can lead to habituation, familiarity, perceptual learning, latent learning, increased probability of manipulation of one portion of the environment, and so forth.

For example, if you follow your parents everywhere, and your parents sometimes eat chocolate, we do not need to postulate a capacity for genuine imitation to explain the fact that you develop a liking for chocolate. Perhaps you sample chocolate pieces dropped by your parents; you then learn that chocolate-eating is good. Again, a simple behavioural tendency—in this case, following a conspecific—combines with the capacity for learning to result in the potential transmission of acquired behaviours.

Stimulus enhancement explains the acquisition of pine-cone-stripping behaviour in Terkel’s (1996) work on black rats. Terkel and his colleagues found that although naive rats never learned to strip cones unaided, the animals were capable of learning the trick if partially-stripped cones were provided, and especially so if they were exposed to cones with progressively fewer rows of scales removed. Furthermore, the investigators noticed that young rats pay close attention to whatever their mother is eating, and often manage to steal partially-eaten cones from her. Thus, we can explain the transmission of the feeding technique by supposing that the pups are programmed to “watch whatever mother is eating,” that this results in them being exposed to cones at various stages of stripping, and that individual learning then does the rest.

Cecconi et al. (1996) describe an instance of stimulus enhancement in simulation: young organisms are carried around for a time on the shoulders of their parents, and are thus able to learn effective foraging strategies before striking out on their own. The authors use this result to argue that delayed maturation in species like our own may be due to the benefits of this type of social learning while young. Billard and Dautenhahn (1998) employed stimulus enhancement in getting a robot to learn a simple vocabulary: a learning robot followed a teacher and thereby gained an adequate perceptual context for acquiring new word meanings through simple associative learning. The model of roosting and foraging in egrets (Toquenaga et al., 1995) described in section 1.3 is also an example of stimulus enhancement.

2.4 Observational learning

If we add slightly more sophisticated learning abilities to stimulus enhancement, we get observational learning. The algorithm involved is approximately “Pay attention to

what others are doing or experiencing, and if the results *for them* appear to be good or bad then learn from this.” Mineka and Cook’s (1988) work on fear acquisition in monkeys illustrates the idea: the authors took naive, lab-raised rhesus monkeys and allowed them to observe an experienced conspecific reacting fearfully to the presence of a snake. The observers, previously indifferent to snakes, rapidly acquired a persistent fear of them. Mineka and Cook argue that the sight of a fearful conspecific is therefore an unconditioned stimulus for a fearful response. It is easy to see that in the wild, this sort of learning could result in the transmission of acquired fears, and would appear to be imitative. All that needs to be assumed is that the monkeys have evolved both an innate ability to recognize the cues associated with fear on the part of a conspecific (such as grimacing and retreating), and the tendency to learn to fear a co-occurrent stimulus (i.e., the snake).

Observational learning can also exist in a simpler form: explicit evaluation of the conspecific’s experience as good or bad may be omitted. For example, Norway rats will develop a marked preference for a novel food that they smell on the breath of a conspecific (Galef, 1996). We might say that the first rat, the observer, learns that the new food is good because it observes positive consequences for the second rat, the demonstrator. That is, the observer notes that the demonstrator is still alive to tell the tale after eating a new and potentially toxic substance. It turns out, however, that the rats are not sensitive to the consequences of eating poisonous foods: they do not learn that a food is bad if the demonstrator has become ill after eating it; in fact they develop a preference as usual. So the rats’ heuristic is simply “Pay attention to what others are eating and do likewise.” Tuci et al. (1999) simulated this phenomenon, and showed that given certain assumptions about the rats’ environment (e.g., the lethality of poison and the behaviour of sick animals) their failure to evaluate the demonstrator’s health is not a mistake, but is actually an adaptive strategy.

Matarić (1994) put both kinds of observational learning to work in engineering cooperative robots that could carry out tasks as a group. Each learning robot had three sources of feedback: information about its own performance, information on the reinforcement of other group members (i.e., knowing whether the other robots were being rewarded or punished for their behaviour), and reinforcement for simply performing the same behaviour as other group members (i.e., a reward for conformity). The second and third sources of feedback correspond to observational learning with and without evaluation of the conspecific’s experience. Of course, Matarić’s work does not establish that observational learning is the equal of true imitation in fostering robot culture; we mention it only in order to show that observational learning can be effectively implemented in a real system.

2.5 Matched-dependent behaviour

We know from laboratory work on animal learning that species such as rats and pigeons can readily be trained to discriminate, e.g., to press one bar when a red light is on and to press another for a green light. Miller and Dollard (1941) showed that this sort of learning was equally possible when the behaviour of another animal served as the discriminative stimulus; they trained rats to follow a leader left or right at a maze junction. Thus, simple reinforcement learning can result in social learning if the contingencies are right. There is no implication that the follower understands the leader’s intentions, nor even that the follower is aware of the match between the leader’s behaviour and its own. The only rule we need to postulate is along the lines of “Given the perceived situation, do something that has previously resulted in a positive outcome.”

Along similar lines, Skinner (1953) suggested that a wild pigeon could learn through trial and error that scratching in a field was likely to be rewarding (i.e., likely to result in ingesting food) if other pigeons could be seen scratching there. In fact the pigeon need not even observe others feeding: learning a correspondence between hidden food and the evidence of feeding, such as scratch marks, would amount to the same thing. The general point is that contagious behaviour may sometimes be learned. Suppose that one of the hypothetical robots discussed earlier was to be deprived of its contagious flight response, perhaps by severing the direct connection between its auditory sensor and its motor output. The robot would be able to reacquire the behaviour by learning that a certain noise—the sound of others fleeing—was associated with an imminent negative experience that could only be avoided by activating a certain behaviour, that is, flight.

2.6 Cross-modal matching

Vocal mimicry on the part of birds occupies a special place in discussions of social learning. It is often argued that because the original stimulus and the animal's response are in the same sensory modality—that is, because birds can hear their own sounds in the same way that they hear sounds made by others—a relatively simple pattern-matching mechanism could account for the phenomenon. In contrast, copying the movements of another animal requires cross-modal matching; the observer must be able to translate the visual input associated with another's movements into appropriate motor outputs. Consider that there is no trivial link between the sight of watching someone else scratch their nose, and the experience of scratching your own nose (see Heyes, this volume).

None of the simple mechanisms discussed so far requires an ability to perform cross-modal matching. Even though processes like contagious behaviour or learned copying could mean that the sight of one animal doing X was a sufficient stimulus for another animal to do X , there is no suggestion of a systematic ability to copy movements. However, imagine an animal that *was* able to identify the movements of others, and map them to movements of its own muscles. If such an ability was combined with observational learning, we would get the behavioural rule “If someone else moves their head (or forelimb or tail or . . .) thus or so, make the same movement yourself.” (As with observational learning, the rule might be conditional on positive outcomes for the observed animal.) Admittedly, this is a reasonably complicated rule, and it is not easy to see how such a mapping would be achieved. Some authors would claim that this rule is indeed the ability to imitate, but Heyes (1993) makes the point that cross-modal matching can explain quite sophisticated social learning without postulating the kind of intentional perspective-taking and sensitivity to goals that Tomasello et al. (1993) find to be characteristic of human imitation. Work on “mirror neurons” in monkeys (Gallese and Goldman, 1998) and humans (Iacoboni et al., 1999) is highly suggestive that, at least in primates, direct mappings may exist between movements seen and movements performed. Meltzoff's (1996) findings on the imitative powers of very young infants also point to an innate ability to perform cross-modal matching in humans. Finally, the qualified success of the COG project in getting a humanoid robot to mimic behaviours such as head-nodding (Scassellati, 1999) shows that cross-modal matching may result in apparent imitation without requiring the ability to attribute intentions.

3 Conclusion

The implications of our emphasis on the explanatory possibilities of simple mechanisms are different depending on one's goals. For scientists interested in explaining animal behaviour, the recognition that some simpler mechanism could be responsible for apparently imitative behaviour is not the end of the story. In particular, computer modelling helps investigators to pose the right empirical questions, firstly by acting as an "intuition pump," and secondly by allowing exploration of the range of parameters for which a particular mechanism could be effective and evolutionarily stable. On the other hand, if one is an engineer, and interested in the idea that robots with the ability to imitate would be useful artifacts, then the message is that simpler mechanisms such as stimulus enhancement or contagious behaviour may give rise to the desired effects. Of course, it may be the case that adding true imitation learning to a group of artificial agents can greatly enhance their ability to adapt to a novel environment, much as theoretically taboo Lamarckian inheritance can greatly speed up technologically pragmatic evolutionary systems (see Ackley and Littman, 1994). Nevertheless, building a truly imitative system is unlikely to be easy. To the extent that we want to build artificial systems that mimic nature, or want to use artificial systems to understand nature, copying the preferences of most researchers in the adaptive behaviour field for simple mechanisms is the better place to start.

References

- D. H. Ackley and M. L. Littman. A case for Lamarckian evolution. In C. G. Langton, editor, *Artificial Life III*, pages 487–509. Addison-Wesley, Reading, MA, 1994.
- A. Bandura. Social learning theory of identificatory processes. In D. A. Goslin, editor, *Handbook of Socialization Theory and Research*, pages 213–262. Rand-McNally, Chicago, 1969.
- A. V. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107:797–817, 1992.
- M. Bekoff. Vigilance, flock size, and flock geometry: Information gathering by western evening grosbeaks (Aves, fringillidae). *Ethology*, 99:150–161, 1995.
- R. Belew and M. Mitchell, editors. *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Addison-Wesley, Reading, MA, 1996.
- S. Bikhchandani, D. Hirshleifer, and I. Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives*, 12(3):151–170, 1998.
- A. Billard and K. Dautenhahn. Grounding communication in autonomous robots: An experimental study. *Robotics and Autonomous Systems*, 24:71–79, 1998.
- E. Bonabeau, G. Theraulaz, E. Arpin, and E. Sardet. The building behavior of lattice swarms. In R. A. Brooks and P. Maes, editors, *Artificial Life IV*, pages 307–312. MIT Press, Cambridge, MA, 1994.
- R. Boyd and P. J. Richerson. *Culture and the Evolutionary Process*. The University of Chicago Press, Chicago, 1985.
- V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge, MA, 1984.
- H. J. Brockmann. Satellite male groups in horseshoe crabs, *Limulus polyphemus*. *Ethology*, 102(1):1–21, 1995.
- K. E. Bumby and K. Dautenhahn. Investigating children's attitudes towards robots: A case study. In K. Cox, B. Gorayska, and J. Marsh, editors, *Proceedings of the Third Cognitive Technology Conference, CT'99*, pages 391–410. M.I.N.D. Lab, Michigan State University, East Lansing, MI, 1999.

- R. W. Byrne. *The Thinking Ape: Evolutionary Origins of Intelligence*. Oxford University Press, Oxford, 1995.
- L. L. Cavalli-Sforza and M. W. Feldman. *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton University Press, Princeton, NJ, 1981.
- F. Cecconi, F. Menczer, and R. K. Belew. Maturation and the evolution of imitative learning in artificial organisms. *Adaptive Behavior*, 4(1):27–48, 1996.
- D. L. Cheney and R. M. Seyfarth. *How Monkeys See the World*. University of Chicago Press, Chicago, 1990.
- D. A. Clayton. Socially facilitated behavior. *Quarterly Review of Biology*, 53:373–391, 1978.
- J.-L. Deneubourg, G. Theraulaz, and R. Beckers. Swarm-made architectures. In F. J. Varela and P. Bourgin, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 123–133. MIT Press / Bradford Books, Cambridge, MA, 1992.
- D. C. Dennett. *The Intentional Stance*. MIT Press / Bradford Books, Cambridge, MA, 1987.
- E. A. Di Paolo. Some false starts in the construction of a research methodology for artificial life. In J. Noble and S. R. Parsowith, editors, *The Ninth White House Papers: Graduate Research in the Cognitive and Computing Sciences at Sussex*. Cognitive science research paper 440, School of Cognitive and Computing Sciences, University of Sussex, 1996.
- L. A. Dugatkin. Copying and mate choice. In Heyes and Galef (1996), pages 85–105.
- J. Fisher and R. A. Hinde. The opening of milk bottles by birds. *British Birds*, 42:347–357, 1949.
- B. G. Galef, Jr. Social transmission of acquired behavior: A discussion of tradition and social learning in vertebrates. *Advances in the Study of Behavior*, 5:77–100, 1976.
- B. G. Galef, Jr. Imitation in animals: History, definition, and interpretation of data from the psychological laboratory. In Zentall and Galef (1988), pages 3–28.
- B. G. Galef, Jr. Social enhancement of food preferences in Norway rats: A brief review. In Heyes and Galef (1996), pages 49–64.
- V. Gallese and A. Goldman. Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, 2(12):493–501, 1998.
- C. K. Hemelrijk. An individual-orientated model of the emergence of despotic and egalitarian societies. *Proceedings of the Royal Society of London: Biological Sciences*, 266:361–369, 1999.
- C. M. Heyes. Imitation, culture and cognition. *Animal Behaviour*, 46:999–1010, 1993.
- C. M. Heyes and B. G. Galef, Jr., editors. *Social Learning in Animals: The Roots of Culture*. Academic Press, San Diego, CA, 1996.
- M. Iacoboni, R. P. Woods, M. Brass, H. Bekkering, J. C. Mazziotta, and G. Rizzolatti. Cortical mechanisms of human imitation. *Science*, 286:2526–2528, 1999.
- O. P. Judson. The rise of the individual-based model in ecology. *Trends in Ecology and Evolution*, 9:9–14, 1994.
- M. Kawai. Newly acquired pre-cultural behavior of the natural troop of Japanese monkeys on Koshima Inlet. *Primates*, 6:1–30, 1965.
- S. Kawamura. The process of sub-culture propagation among Japanese macaques. *Primates*, 2:43–60, 1959.
- K. N. Laland, P. J. Richerson, and R. Boyd. Animal social learning: Toward a new theoretical approach. In P. P. Klopfer, P. P. Bateson, and N. Thomson, editors, *Perspectives in Ethology*, pages 249–276. Plenum Press, New York, 1993.
- K. N. Laland, P. J. Richerson, and R. Boyd. Developing a theory of animal social learning. In Heyes and Galef (1996), pages 129–154.

- C. G. Langton. Artificial life. In C. G. Langton, editor, *Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems (ALIFE '87)*, pages 1–48. Addison-Wesley, Redwood City, CA, 1989.
- M. J. Matarić. Learning to behave socially. In D. Cliff, P. Husbands, J.-A. Meyer, and S. W. Wilson, editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 453–462. MIT Press / Bradford Books, Cambridge, MA, 1994.
- A. N. Meltzoff. The human infant as imitative generalist: A 20-year progress report on infant imitation with implications for cognitive psychology. In Heyes and Galef (1996), pages 347–370.
- J.-A. Meyer. The animat approach to cognitive science. In H. Roitblat and J.-A. Meyer, editors, *Comparative Approaches to Cognitive Science*. MIT Press, Cambridge, MA, 1994.
- G. F. Miller. Artificial life as theoretical biology: How to do real science with computer simulation. Cognitive Science Research Paper 378, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK, 1995.
- N. E. Miller and J. Dollard. *Social Learning and Imitation*. Yale University Press, New Haven, CT, 1941.
- S. Mineka and M. Cook. Social learning and the acquisition of snake fear in monkeys. In Zentall and Galef (1988), pages 51–73.
- R. R. Provine. Contagious yawning and laughter. In Heyes and Galef (1996), pages 179–208.
- B. Scassellati. Imitation and mechanisms of joint attention: A developmental structure for building social skills on a humanoid robot. In C. L. Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, volume 1562 of *Lecture Notes in Artificial Intelligence*, pages 176–195. Springer, Berlin, 1999.
- D. F. Sherry and B. G. Galef, Jr. Cultural transmission without imitation: Milk bottle opening by birds. *Animal Behaviour*, 32:937–938, 1984.
- H. A. Simon. Artificial intelligence: An empirical science. *Artificial Intelligence*, 77:95–127, 1995.
- B. F. Skinner. *Science and Human Behavior*. Macmillan, New York, 1953.
- K. A. Sullivan. The advantages of social foraging in downy woodpeckers. *Animal Behaviour*, 32:16–21, 1984.
- J. Terkel. Cultural transmission of feeding behaviour in the black rat (*Rattus rattus*). In Heyes and Galef (1996), pages 17–47.
- G. Theraulaz and E. Bonabeau. A brief history of stigmergy. *Artificial Life*, 5(2):97–116, 1999.
- P. M. Todd. The causes and effects of evolutionary simulation in the behavioral sciences. In Belew and Mitchell (1996), pages 211–224.
- P. M. Todd and G. F. Miller. Exploring adaptive agency II: Simulating the evolution of associative learning. In J.-A. Meyer and S. W. Wilson, editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 306–315. MIT Press / Bradford Books, Cambridge, MA, 1991.
- M. Tomasello. Do apes ape? In Heyes and Galef (1996), pages 319–346.
- M. Tomasello, A. C. Kruger, and H. H. Ratner. Cultural learning. *Behavioral and Brain Sciences*, 16: 495–552, 1993.
- J. Tooby and L. Cosmides. The psychological foundations of culture. In J. H. Barkow, L. Cosmides, and J. Tooby, editors, *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, pages 19–136. Oxford University Press, New York, 1992.
- Y. Toquenaga, I. Kajitani, and T. Hoshino. Egrets of a feather flock together. *Artificial Life*, 1(4):391–411, 1995.

- E. Tuci, J. Noble, and P. M. Todd. I'll have what she's having: A simulation analysis of the copying of food preferences in Norway rats. In K. Dautenhahn and C. Nehaniv, editors, *Proceedings of the AISB'99 Symposium on Imitation in Animals and Artifacts*, pages 74–79. Society for the Study of Artificial Intelligence and Simulation of Behaviour, 1999. Available WWW: <http://www-abc.mpib-berlin.mpg.de/users/noble/Research/Social/> [1999, 12 August].
- R. B. Zajonc. Social facilitation. *Science*, 149:269–274, 1965.
- T. R. Zentall. An analysis of imitative learning in animals. In Heyes and Galef (1996), pages 221–243.
- T. R. Zentall and B. G. Galef, Jr., editors. *Social Learning: Psychological and Biological Perspectives*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1988.