

Detecting, Understanding, and Explaining Animal Imitation¹

Richard W. Byrne

[Chapter 9 in S Hurley & N Chater (eds) *Perspectives on imitation: from mirror neurons to memes*, Vol.1, pp.255-282. MIT Press, Cambridge MA, 2005.]

Abstract: Imitation has fascinated psychologists because of two main qualities, both of which appear near to magic at first sight: the ability to deal with opaque correspondences between actions of self and other, and the ability to acquire novel, complex skills by observation alone. Comparative psychologists have converged on the ‘two-action method’ for detecting animal imitation, but what the results mean varies widely from experiment to experiment. Some test opaque correspondence, none test complex skill learning. Other experiments are aimed directly at the ability to learn complex novel actions, but they too have problems. The ‘do as I do’ method produces results vulnerable to other explanations, and few kinds of animal copy the repertoire of other species, limiting what can be learnt from this. Observational evidence of several kinds also suggests imitative ability; and, like that from experiments, none is perfect. Site-to-site variation in repertoires is always hard to link directly to imitation, because environmental variation is usually great; rapid learning of improbably complex behavior is a matter of relative plausibility; and the natural experiment of manual disablement may also be explained in more than one way. One reason for the extreme skepticism applied to animal imitation is a lack of a plausible, mechanistic model of the process. One such is outlined, and whether it best matches ‘imitation’, or some other category of behavior, is discussed. Regardless, the ability to pick up complex, new behaviors from observation may have been important in ape and human evolution, and a possible evolutionary scenario is set out.

9.1. Introduction. Imitation is magical. That element of magic is what so fascinated comparative and developmental psychologists for much of the 20th century, and explains the academic popularity of imitation today—in the face of a persisting lay view that imitation is a mere sham of intelligent behavior, a cheap trick. The magic is twofold. Firstly, how can the imitator recognize that an action was “the same” when performed by itself as when observed done by another? The perspectives may be very different, with little visual or auditory similarity: it is magic. Secondly, how can a package of skill be “transferred” from one individual’s repertoire to that of another? If this trick can be done, it offers a rich method for the transmission of useful procedures between individuals, and even across the generations. But how might a rich complex of goals, rules, processes and schemata be got from simply seeing an act done by another: it is magic.

Confronted with magic, the scientist has only two options. Either the observations must be mistaken, or a non-magical way must be found to accommodate the data: perhaps taking a different viewpoint, perhaps by extending or changing theories. Imitation has seen both options applied in recent years. However, a major pitfall in bringing science to this particular magic has been the shifting use of definitions, as I will try to illustrate. There is a real danger, if “imitation” means different things in different contexts and to different researchers, that an *impossible* set of characteristics will be attached to it—making the poor scientist’s task quite impossible. For this reason, no apology will be made for sometimes fussy-seeming care in definition.

9.2. Imitation in the comparative psychology tradition. Since Thorndike first drew psychologists’ attention to the problems posed by imitation, and provided one of the simplest useful definitions, “learning an act by seeing it done” (Thorndike, 1898), comparative psychologists have sought an acid test of whether a species can or cannot imitate. Now, they are confident they have found it: the two-action method. For comparative psychologists this paradigm now forms the operational definition of imitation (Galef, pers. comm., 25th May 2002). In the two-action method, separate groups of individuals are each allowed to watch a model, often a conspecific, solving the same task (Dawson & Foss, 1965; Heyes, 1993). Each different model applies a different method, but all methods succeed. Then, individuals are tested on their own. If they each tend to match the method that was shown to them, the species is deemed to have shown imitation. By this criterion, a growing list of species are now claimed to show imitation, or “true imitation” as it is sometimes known by these researchers.

Although various workers define imitation in different ways, in reality all are using an operational definition (“what the 2-action method measures”), so it is to the *method* that we must look. When we do, we find that what has been found varies from case to case: “true imitation” is not the same in different experiments. There are three main dimensions of variation:

9.2.1. Opacity. Consider the now-famous phenomenon of neonatal imitation, in which a baby, even when only minutes old, is shown a distinctive facial gesture such as tongue protrusion to one side, or mouth opening and shutting, and copies the act seen (Meltzoff, 1996; Meltzoff & Moore, 1977). Nothing new is being learned (the act may be “novel” in a trivial sense, since a neonate may not have yet done it, but these simple facial gestures are

within the latent repertoire of new-born humans, and were indeed chosen for that reason). No problem is being solved: the biological function of neonatal imitation is rather considered to be a social one. What is remarkable is the highly *opaque* nature of the translation, from visual inspection of an adult's face to muscular movements needed to match the gesture: how does the infant do it? These data dramatically illustrate what has been called the "correspondence problem", of knowing that two actions are the same when they have been experienced in very different ways. Explanations for this magic have varied from disbelief (Anisfeld, 1991, 2004), through re-construal of the behavior as less magical and explicable on existing theory (Heyes, 1993, this volume; Mitchell, 1987), to development of a new theory, involving an innate, amodal system for 1:1 mapping among sense modalities (Meltzoff, 1990a, 2004).

But in other 2-action experiments, the translation is not always so opaque. Closest to the neonatal human data are perhaps studies in which Japanese quail match the simple action, pecking or stepping up, seen to achieve a food reward (Dorrance & Zentall, 2001). A bird's eyes are close to its beak, so the visual and proprioceptive experience of pecking are very different to the sight of another bird's peck, and a quail cannot even see its own feet when it steps up, so visual-visual matching is problematic. Budgerigars can easily see their feet, but a case can still be made that matching actions of the beak must involve rather opaque translation (Dawson & Foss, 1965; Galef et al., 1986, Heyes, this volume). Not so, when a primate learns to tackle a task, with its hands. A number of primate species have been shown to match the hand action they have been shown by a human (poking, twisting, etc) applied to a box clasp which affords several different manipulations (Bugnyar & Huber, 1997; Custance et al., 1999; Stoinski et al., 2001; Voelkl & Huber, 2000; Whiten et al., 1996): but to recognize an action as the same when done by oneself and another is a good deal less magical when it's done by a hand. Hands are uniquely shaped and highly familiar, so there is little ambiguity in how two hands are to be matched. The correspondence problem remains, but it is now a tractable one that an artificial intelligence engineer would feel confident in tackling by visual-visual mapping. Not magic.

9.2.2. Complexity. When I watch a competent mechanic deal with a blockage in the carburetor of my car, I may be able to learn a new skill by observation, and apply it myself when the car next fails. I will have acquired a novel and rather complex behavioral routine: this illustrates the second magical aspect of imitation, as a rich and powerful method of social learning. In the two-action method, there is usually no attempt to demonstrate

learning of any *new* skills: the actions tested are familiar ones in the species' existing repertoire. The "true imitation" of most comparative psychologists, operationally defined by the two-action method, has little to do with skill learning. Only one experiment of this genre has aimed to test learning of a sequence of actions (Whiten, 1998). Each of 4 chimpanzees, all of whom had rich experience of human living in a "language acquisition" program, learnt the particular three-action sequence it was shown, although all sequences were equally effective in opening the box that contained food. However, it took several cycles of observe-and-test before the sequence learning reached statistical significance. Zoo gorillas, similarly tested, failed to show any such learning (Stoinski et al., 2001).

9.2.3. Learning A useful distinction, introduced by Janik and Slater (2000) in a review of vocal learning in animals, is that between production learning, where a new behavioral complex is acquired, and contextual learning, where a familiar behavior is applied in novel circumstances. Contextual learning is illustrated by the predator-specific alarm calls of vervet monkeys (Struhsaker, 1967). Developing monkeys start to produce these distinct calls without any particular experience: call form seems to be tightly defined genetically, developing reliably in many environments. However, the referents to which the alarm calls are given do change during development: an infant may give an "eagle alarm" to a falling leaf, but later it will restrict the call to broad-winged soaring birds, and finally only to a serious predator, the martial eagle *Stephanoaetus coronatus* (Seyfarth & Cheney, 1986).

Applying this distinction to imitation, it is clear that *production imitation* is generally neither sought nor found in the two-action method. However, *contextual imitation* occurs when the appropriate site or occasion on which to apply a familiar action is learnt by observation. Contextual imitation may be a useful cognitive capacity for transmitting habits, through populations or over generations, even if any learning of new behavior patterns relies on simpler means than production imitation, such as trial and error learning.

However, in many cases, when two-action experiments have been done, no effort has been made to distinguish an alternative possibility. This is a simple idea, though it has been given several names (Byrne, 1994; Byrne & Russon, 1998; Thorpe, 1956, 1963): response facilitation, priming, perceptual-motor resonance, and perhaps also contagion (although it would seem that Thorpe intended this to apply only to innate perceptual-motor linkages, rather than more generally).

Imagine a species with a highly restricted repertoire of, say, 17 discrete behaviors, and no capacity to build new ones. Confronted with any foraging problem, an individual

applies actions from this repertoire, and those that are followed by advantageous consequences are more likely to be applied again in similar circumstances (trial-and-error learning). The actions are generated in an order that is determined by the activity level of whatever brain unit, or record, controls the corresponding action. But suppose now that when one of the 17 actions is seen to be done by another individual, the activity level of the corresponding brain unit is temporarily increased: the action is “primed”, that particular response has been “facilitated”. (This may or may not involve a serious correspondence problem, as discussed above.) Tested with the two-action method, because such observation changes the relative probabilities of trying one action rather than another, the species would automatically show the behavior held to demonstrate “true imitation.” Yet, to the extent that anything is learned, it would be by trial-and-error learning—of when to re-apply a successful action—rather than observational learning.

What is needed to demonstrate observational learning in the two-action paradigm is a *delay before testing*, preferably filled with other activities. If the delay were long enough, one might presume that the activity levels of brain units would return to their resting levels, although nothing is known of the time-course of this process so it is not clear how long would be long enough. Better, if the critical actions were used for purposes irrelevant to the experimental task in the time before learning was tested, one could be confident that genuine contextual learning by imitation had occurred. Zentall’s quails showed learning when first tested 30 min after their observation of a model, and Heyes’ budgerigars up to 22 hours, periods which were inevitably filled with other activities (although whether these activities employed the motor actions later tested is not certain), so contextual imitation is highly likely. In most experiments of the two-action paradigm, contextual imitation has yet to be demonstrated.

9.3. Skill learning by imitation in the everyday lives of animals. This chapter is about skill learning, so we will now set aside the fascinating issues attendant on the “hard problem” of perceptual-motor correspondence. The two-action method has successfully shown that young (and even new-born) human babies, Japanese quails and budgerigars are all able to make such correspondences—or at least, the majority of researchers believe that is what the evidence means. Problems for the future include the generality of this capacity across species; and whether it has any role in skill learning, or, as suggested by Byrne and Russon

(1998), instead functions because mimicry is important in social recognition and social bonding.

We now focus on the second magical aspect of imitation, its potential for allowing the observational learning of complex skills. Here, the comparative psychology tradition, searching for “true imitation” as operationally defined by the two-action method, has little to say because in all the two-action experiments carried out with birds and new-born humans, the behaviors examined have been simple, familiar ones. This is no criticism: the issue has been perceptual opacity not skill acquisition. Unfortunately, the picture is no different in most two-action experiments with non-human primates, where the fidelity with which skills can be passed on has been the ostensible focus. Here, the actions have been manual—twisting a handle, pushing out a bolt, pulling a lever—so the correspondence between the act as done by self and other has been rather transparent. Moreover, no effort has been made to ensure that these actions were not part of the individuals’ pre-existing repertoire before the experiment; usually, since simple actions are chosen for testing, they most likely were.

Acquisition of novelty has been examined only in two experimental paradigms, both with great apes. Firstly, as noted above, learning of novel, three-action sequences each made up of simple manual actions was examined in 4 chimpanzees and 6 gorillas (Stoinski et al., 2001; Whiten, 1998). The chimpanzees were kept in enriched conditions, and might be argued to have been “enculturated” by humans, and they gradually learnt the sequence; the gorillas were part of a zoo collection, and they did not. Secondly, imitation of novel actions has been reported in experiments where great apes are first trained to “do as I do” for a reward, then tested with other manual and bodily gestures chosen to be unfamiliar to the subjects (Custance et al., 1995). Observers blind to the purpose of the experiment are shown video clips of the apes’ behavior just after a gesture has been demonstrated to it, and can reliably detect which particular action had been shown to the subject, out of a small set of possible gestures. However, the match between actions of ape and human demonstrator was reported to be quite poor. This opens the door to an alternative explanation than imitation. Great apes, unlike most animals, have very large repertoires of manual action (e.g. the only published chimpanzee ethogram includes 515 separate behaviors, yet this categorizes the actions only at a very broad level rather than specifying the exact form of action, as is typical of ethograms for simpler species: Nishida et al., 1999). It is difficult to be sure, then, that the gestural copying reported in experiments

reflects observational learning of new behavior patterns, as the researchers believe, or whether instead rare actions in the pre-existent repertoire are *primed* (made more salient) by seeing them demonstrated, so the individual is then likely to perform an action it otherwise seldom would.

Experimental evidence for observational learning of novel, complex actions by animals is therefore weak at present. Instead, we must turn to observational data for significant evidence of non-vocal production imitation in animals. Inevitably, observational evidence is more easily challenged than experimental evidence, and the skeptical approach of dismissing the observational data is still quite feasible. If animals do not acquire new, complex skills by observational learning, there is no need for a new theory. However, as my blocked carburetor example reminds us, humans *can* augment their skill repertoire by observation, even if this has become less common in a modern, western environment where verbal instruction is given pride of place. In a non-western or medieval European world, crafts are typically acquired by a combination of uninstructed observation and individual trial and error (Parker, 1996). It is hard to deny that humans have the capacity to acquire novel organizations of behavior by observational learning. Given this fact, it may be productive to examine non-human primates' behavior for signs that they too have some such capacity. Moreover, a cognitive theory is needed to explain the capacity, whichever species are finally agreed to learn skills by imitation.

Observational evidence is most sturdy when animals acquire complex behaviors that are not part of their natural repertoire (Pepperberg, this volume, makes the same point for vocal learning). Great apes brought up or kept as pets by human carers seemingly readily acquire "human" behaviors. In careful analyses of rehabilitant orangutans, this has been shown to imply acquisition by imitation (Russon, 1996, 1999; Russon & Galdikas, 1993, 1995). With natural behaviors, alternative explanations are, however, harder to exclude. In particular, a complex performance might have been acquired by gradual, trial-and-error learning because it was the most *efficient* way to do the job, given the physical constraints of the task. On this view, while learning may be helped by many social influences—*social facilitation*, increasing the motivational disposition for tackling that sort of task; *stimulus enhancement*, focusing attention on particular places and objects; *response facilitation*, changing the probability with which actions are tried out; and various sorts of *emulation* (Byrne, 1998a), allowing the individual to learn aspects of the task structure—production imitation is not one of them.

Attempts to avoid this critique have pointed to various characteristics as unlikely to be products of socially-guided, trial-and-error learning alone:

9.3.1. Site to site variation in behavioral style. Social traditions may exist as a result of mechanisms other than imitation, and are known in many species (Roper, 1983). But if precisely the *same* task is done differently in different populations, it would seem that the persistence of the less efficient variants must reflect imitation (or, more improbably, teaching), since trial and error should operate as a “hill climbing” algorithm and force convergence on the most efficient method.

Just this has been argued to apply to wild chimpanzees, with respect to their style of eating insects with tools. *Dorylus* ants are captured by dipping sticks into dense masses of insects. All the individuals in some populations use long sticks in a bimanual process (McGrew, 1974), but in others they rely on the apparently inefficient monomanual use of short sticks (Boesch & Boesch, 1990; Sugiyama et al., 1988). *Macrotermes* termites are extracted with fishing probes. In some populations, it is routine to re-sharpen probes or use both ends, both apparently efficient methods; but in other populations blunted tools are discarded and wholly new ones made (McGrew et al., 1979).

Unfortunately, if the environment differs in some critical way between populations, a simpler explanation may be that the behavioral styles reflect adaptive solutions to slightly different environmental problems (Tomasello, 1990). It is nearly impossible to establish identity of environments over the functionally relevant aspects of a complex task, most of which are unknown.

Moreover, it might be expected that occasional innovations would, over extended time, produce just the same hill-climbing process as trial-and-error learning, even with a social habit that was acquired ontogenetically by imitation. This means that the idea of population differences in important skills must be regarded with considerable suspicion. Are we supposed to accept that these habits, believed to have important survival value to chimpanzees, are so *new* that there has been too little time for an optimal tradition to have been reached?

9.3.2. Improbable complexity. Most everyday actions of animals are organizationally simple, and readily understood as a result of innate dispositions and normal learning processes. If, in contrast, highly complex skills could be shown to be learnt with great reliability and considerable standardization of form, then imitation (or teaching) would be suspected. Just this has been noted for gorilla and chimpanzee plant feeding techniques

(Byrne & Byrne, 1993; Corp & Byrne, 2002), for chimpanzee insect catching and nut-cracking with tools (Boesch & Boesch, 1983; Goodall, 1986)—and for an Israeli population of black rats that has learned a special technique for eating the cones of non-native pine trees (Terkel, 1994). It is not known in detail how these skills develop, but experimental work with the rats has shown that trial and error with social facilitation and stimulus enhancement is quite insufficient: only by presenting a graded series of part-processed cones, working back from the solution, was it possible to teach naïve rats to strip cones efficiently. Given the partly nocturnal, arboreal behavior of black rats, it has not been possible to determine whether infant rats experience such graded learning sequences, either accidentally or as a consequence of deliberate maternal behavior, or whether they are able to observe adults processing.

Gorilla infants certainly have ample opportunity to see what their mothers are doing to process difficult plants, at close range and over many months. Several different techniques are acquired, apt for dealing with different plant challenges: spines, tiny hooks, stings, hard casing. Each technique is multi-stage, often involving bimanual coordination with the two hands used in different roles to achieve a single result, and different digits in a single hand are used independently to allow part-processed material to be held while further plants are processed with both hands (Byrne et al., 2001a). Very strong individual hand-preferences are found, often considered an adaptive response to complexity, and several tasks share population-level right handedness (Byrne & Byrne, 1991). “Complexity” is notoriously hard to measure in any valid way, but consistent ordering of sequences of 5 or more separate actions, chosen from a repertoire of over 70 functionally different manual processes, suggests a level of complexity at least equivalent to that of many everyday human manual tasks (Byrne et al., 2001b). Idiosyncratic variation, both from trial to trial and from animal to animal, even from mother to offspring, is found at lower hierarchical levels: grip types, which digits are involved, the extent of movement, the direction of hand preference. However, the organization of the processing sequence, including bimanual coordination, hierarchical embedding of sub-processes as subroutines, and the main sequence, is remarkably standardized (Byrne & Byrne, 1993).

Reliable acquisition, of the elaborate and highly structured organization of several different processes, by the age of 3 years in gorillas, has been argued to strongly suggest “program level imitation”: copying the structural organization of a complex process by observation of another’s behavior, regardless of how the exact details of actions are

acquired (Byrne, 1993; Byrne & Russon, 1998). This argument depends on the fact that there are many alternative methods, some much more obvious to human observers. (Skeptics are recommended to try this experiment: set naïve human subjects the problem of eating one of the gorilla foods, explaining the difficulties to them. Stinging nettle, for instance, has the most painful stings on main stem and petiole, least on the leaf underside; the leaf blades are the most nutritious part. I predict that the unhappy subjects will not readily discover the method of mountain gorillas, which minimizes the pain of consumption while maximizing intake.) The alternative methods are almost certainly less efficient than the precise technique of gorillas, so in principle trial and error learning might converge upon that, eventually. The case for imitative learning depends on the improbability of this happening, reliably for every mountain gorilla, in the 3 year period of skill development, and especially with plants (like nettles and thistle) that discourage playful exploration by the pain they induce.

9.3.3. Insensitivity to affordances. Reliable individual acquisition of complex skills would be much less improbable if the physical constraints of the task and the animal's anatomy were severe, so tightly channeling the process of learning. In the case of great ape manual skills, this seems at first sight implausible: ape hands are highly flexible grasping effectors, and plants can be processed in many ways. But a stronger challenge to imitative learning comes from the idea that animals—or at least, non-human primates—may be exquisitely sensitive to quite subtle properties of physical objects, and that this property knowledge can guide relatively rapid learning. This is the idea of affordance learning (Tomasello, 1996, 1998). While there is at present little positive evidence that it gives a realistic account, the possibility is hard to refute without experimental manipulation.

One approximation to an experimental test is provided by the unfortunate circumstance that some populations of wild great apes have experienced high levels of human-set snares, which cause severe maiming of the hands of curious infants and young juveniles that explore them (Stokes et al., 1999). Nevertheless, some of these young victims survive to adulthood, learning to deal effectively with hard-to-process plants that normally require elaborate manual skills. Since maimed hands are radically different to those of able-bodied apes, very different affordances are presented to these animals. Thus, if the techniques are a consequence of the affordance structure of hands and plants, the disabled apes should acquire idiosyncratic techniques, adjusted to efficient utilization of the remaining capacities of their hands. They do not (Byrne & Stokes, 2002; Stokes & Byrne,

2001). Instead, they use the *same* techniques as able-bodied animals, working around each of the points at which their hands do not allow actions to be done in the normal way. Their modifications affect just those aspects of the skill that vary among the able-bodied: methods of gripping, the digits or hands used for a purpose, and so on. Since the apes that sustain these hand injuries are still at an age when many skills are developing, it has been argued that this insensitivity to affordances must reflect program-level imitation from the only model available, the able-bodied mother.

None of these arguments is completely watertight, and it may still be argued that non-human apes simply *cannot* learn complex novel skills by imitation: the human capacity to do so may, after all, be derived in some unspecified way from linguistic competence. But what if we take the challenge of accepting that non-linguistic apes can use program level imitation in skill development: are we facing magic? Or can a plausible mechanistic theory be sketched, capable of explaining this facility?

9.4. A mechanistic theory of skill-learning by observation. Acquiring a complex, novel skill can be viewed as a process of building a new organization out of old components—which may themselves be quite elaborate processes, or very elementary actions (Byrne & Russon, 1998; Newell & Simon, 1972). Depending on the existing repertoire of actions, putting together the same skill may therefore be a simple or highly complicated matter. Moreover, the ability to do so will usually be strongly influenced by knowledge of the mechanical properties of the objects involved—their affordances—and thus dependent on previous, perhaps “playful” exploration. And several social influences will inevitably affect the power of trial and error to build new organizations of skill, including stimulus enhancement (Spence, 1937), response facilitation (Byrne, 1994), and indirect affordance learning by emulation (Tomasello, 1990; Whiten & Ham, 1992). All these factors will potentially affect even “imitative” learning of a complex skill.

For behavioral organization to be copied from another’s behavior, it must first be detected. The process I propose to account for this is one of *behavior parsing*. Parsing cannot directly apply to skilled action, since this does not consist of visible units: it is fluid movement, without natural joints between separate elements. The first thing we need is a system that can recognize elements in this fluid movement, and so reduce it to a string of components. For these components to be any use as building blocks for assembling a new

skill, each must correspond to an action already in the repertoire of the observer. Element detection is a non-trivial process in itself, but a neural system exists that already meets these requirements: the mirror neurons of area F5 in the pre-motor area of macaque monkeys (Gallese et al., 1996; Gallese & Goldman, 1998; Rizzolatti et al., 2002; Rizzolatti, Fadiga, Matelli et al., 1996). Monkeys are, *pace* conventional wisdom, not great imitators (Visalberghi & Fragaszy, 1990); mirror neurons are unlikely to have evolved originally as part of an imitative learning system, but rather as part of social comprehension, allowing subtle dispositions and objectives of social companions to be discerned. But mirror neurons can in principle convert a stream of fluid action into a string of elements, each of which is already part of the repertoire of the observer. (Note that the observer's prior knowledge will determine the length of this string: if the observer already can perform major parts of the process, matching will be high-level and the string short.) If this string can be learned or remembered in some way, then it can be imitated (see Heyes & Ray, 2000; Heyes, this volume).

The resulting, *action-level* imitation would treat the whole string as a unit, a linear, chain-like structure. Action-level imitation inevitably lacks the power, efficiency, and flexibility that we associate with human planning, in which hierarchical structures can be manipulated (see Dawkins, 1976; Lashley, 1951). Suppose that, rather than direct copying of linear strings of actions, the strings themselves are input to a parsing process. Is there sufficient information in the surface form of skilled behavior to betray the underlying organization? If only one "look" were allowed, a single pass through an elaborate process, the answer is certainly no. Organization is not there for the seeing, it can only be extracted from the statistical properties of many observations. How, then, might the underlying organization, the "deep structure" of action, be detected merely from observing the surface form? Are any clues available?

According to the Behavior Parsing model (BP), there are numerous clues (Byrne, 1999, 2002a). Crucially, parsing must focus on essential elements, rather than coincidental or idiosyncratic variants that occur from time to time. Every execution of a motor action is slightly different, but some aspects must remain constant across successful attempts at the same goal. ("Goal" is a mentalistic description, but in fact for the feeding activities of great apes mentioned above, and for many complex manual skills acquired by young children, consummatory activity makes the end state of observed action readily discernible.). If numerous sequences with the same end state can be compared, the elements they have in

common define the *necessary steps* of the process, and reveal their characteristic *ordering* if there is one. (Whether these elements are themselves actions, or states of the environment produced by actions, is an empirical question that does not affect the power of the model: see Byrne, 1999. Common sense suggests that it is often easier to see transformations of objects than the manual movements that effected them, suggesting that actions upon objects should be particularly easy to learn by imitation.) As well as the sequence of important steps in the course of a successful execution, the same hypothetical process of string comparison can reveal bimanual coordination, as correlated activity of the two hands, where this is so crucial that it recurs on nearly every sequence. The BP model predicts that great apes, including humans, should be acute at mentally “summing” numerous sequences of action so that correlations and exceptions are rendered obvious. Such an ability would suggest a dedicated neural system, that operates unconsciously and automatically.

Most skilled action, of non-human or human apes, seems to have a modular structure, in which some strings of actions are tightly bound together, whereas at other points the linkage is weak. This organization is highly functional, since it allows the tightly-bound modules to function as units: to be recruited as a whole into other tasks, or substituted by other modules that achieve the same local ends. But can *modularity* be “seen” by a mechanistic process? The BP model includes several ways in which this could be done. Firstly, *interruption* between modules should be relatively harmless, allowing smooth resumption of processing afterwards; in contrast, interruption within a module would force “beginning again” at the start of the module. Thus, by paying attention to the points of smooth resumption, module breaks are revealed. More generally, the distribution of *pauses* in execution should show clustering at module boundaries rather than within modules: if young apes are sensitive to these most classic data of cognitive psychology, then they will be able to parse action into modules.

As the whole process becomes less opaque and mysterious, as a result both of parsing and independent trial and error exploration, other sources of information about modular organization become available. These include *substitutions* of one module for another (for instance, if an animal occasionally processes an arboreal plant, making bimanual processing infeasible); *omissions* of a module when it is unnecessary (for instance, if a handful of leaf material is unusually free of dead material, no cleaning is needed before consumption); and *local repetitions* of a module that vary in number, revealing the unity of

the elements of in the repeated sub-string (for instance, if stripping leaves generates insufficient material for a full handful, stripping may be repeated once or several times, provided the animal has the individual digit control to allow the already-acquired material to be retained in one hand).

Moreover, the same variation in repeated sub-strings is a potent clue to *hierarchical* organization: when one or more modules is used as a subroutine within the main process, repetition inevitably occurs and thus indicates that there must be some criterion for making the decision, whether to iterate or not. (Similarly, *substitution* of one module for another under different circumstances, and recognition of *already familiar* sub-strings within an unfamiliar sequence, are indications of the flexible and efficient processing that use of a hierarchical organization can give: Dawkins, 1976.)

By the standards of artificial intelligence, the BP model is under-specified—exactly how strings of elements are dissected out for parsing in the first place, and whether the whole process could operate efficiently with no information on object affordances, need to be determined. It seems highly likely that real young apes, human or otherwise, would bring to the task from their prior exploration of the world considerable task-specific knowledge—of the effects of simple manual actions, and the opportunities for change to objects—making the learning process more efficient. Nevertheless, the BP model provides clear methods by which the structural organization of a complex, novel process can be learnt partly by program-level imitation. When this happens, the finer details of action—hand movements, grips, laterality, and so forth—may be acquired in other ways, such as individual exploration, or also by program-level imitation at lower levels in the hierarchy of behavior.

9.5. But is it imitation? In recent years, several new terms have been added to the lexicon of those grappling with the conceptual and practical issues raised by imitation. Program-level imitation clearly goes beyond processes by which trial and error learning can be aided by social circumstances (stimulus enhancement, response facilitation, etc), which can be modeled quite simply as priming of brain records. Equally, however, the BP model is no explanation of the opaque correspondence problem. To those who restrict the term “imitation” to processes that can overcome the opaque correspondence problem, program-level imitation is not really imitation. Indeed, although novel *organizations* of behavior can be imitated by BP mechanisms, it may be that the movement patterns used in the copy are

rather different ones to those of the individual copied, since these fine details may be most efficiently acquired by individual exploration. Is this really “imitation” at all? If not, what else might it be?

To capture the difference between imitation-for-a-purpose, in goal-directed copying, versus mere parody that cannot be put to use, the term *mimicry* has been introduced—meaning “imitation for its own sake” (Tomasello, Kruger & Ratner, 1993). The BP model does not require the intentions of the imitated individual to be first discerned by the copier, so at first sight it might seem appropriate to label the outcome mimicry. But intentions attributed to the individual copied are *second-order* intentions (Dennett, 1983), whereas Tomasello’s critique concerns *first-order* intentions, such as desires and beliefs of the imitator. Since we cannot tell that an infant human, chimpanzee or gorilla is absent of first-order intentions when it engages in such activity, insisting on the term mimicry is unhelpful. (Attempts to dismiss as mimicry the copying of rather complex everyday human actions, shown by many human-raised great apes, because the copying appears pointless to us, have likewise been met with the comment that human researchers cannot know what rewards or amuses a human-reared ape, and thus cannot legitimately assume purposelessness: Russon, 1996.) Moreover, it is highly functional in biological terms for an infant to copy any action of its mother or other close associate that they perform repeatedly, especially if signs of satisfaction or any consummatory activity are associated.

In any process of learning a complex skill, prior knowledge of the physical properties of the environment, and to what extent they can be modified, is likely to be crucial. Tomasello (1994) has suggested that great apes may sometimes be able to learn these things by observing others, “emulation”, but cannot learn their actions by observation, “imitation”. In early versions (see Byrne, 1998a), what was learnt by emulation was seen simply as the final result of action (called *goal emulation* by Whiten and Ham, 1992), and this certainly cannot apply to program-level imitation, where the goal of eating edible items needs no learning. Understood as *affordance learning by observation*, emulation potentially has much greater power. Since the BP model is agnostic as to whether the input consists of strings of actions, or strings of object states (each produced by an action, which itself may be hard to see clearly), it might be claimed to be a sort of repeated emulation process. However, the cardinal point of emulation seems to be that cause-effect learning enables enhanced problem-solving without copying behavior. In the BP model, no reliance whatsoever is made of understanding cause and effect, and indeed program-level imitation

may copy actions whose means of effect remain obscure; in contrast, an important aspect of behavior *is* copied, its organization, without necessarily including slavish replication of details. In some ways, then, program-level imitation is the reverse of emulation. Also note that the plausibility of emulation as a null hypothesis for young great apes has been questioned (Byrne, 2002b), since young children prove able to imitate long before they can emulate (Want & Harris, 2002).

Rather than using any of these categories of quasi-imitative process, it may be more helpful to view program level imitation as a spin off from a much more fundamental process—a part of high-level *perception*. Just as possession of a mirror-neuron system enables simple elements of goal-directed action to be “seen” directly, so complex actions, treated as strings of such elements, can be parsed to extract organization from the inevitable statistical regularities present in any behavior produced by a planning process. This chapter has focused on one consequence of that, the ability to use the observed organization to build complex skills new to the observer more efficiently than pure trial and error would ever allow. However, the BP process itself may have wider uses than program-level imitation. To “see” through the surface form of behavior, to the *deep structure* of the observed individual’s plans, is an important part of interpreting that individual’s action. Program-level imitation may therefore be seen as a handy index of an individual’s ability to *represent* the complex actions of others.

A mechanical process that can extract, from the statistics of observable behavior (and probably a good background knowledge of the local environment of action), the high-level organization of effective, organized action is but a short step from being able to “see” intentions and cause and effect. Both in everyday life, and in the underlying physics of forces, the abstract idea of *causation* is closely related to the visible facts of contact and correlation. Likewise, the efficient result of complex, directed action is very close to its *intention*. Behavior parsing may be a necessary precursor to “seeing” the plans and purposes of others, and how their behavior works to achieve these purposes.

9.6. A possible evolutionary scenario. Viewing imitation in this way, as an outcome of behavior parsing, leads to a possible explanation of the evolutionary origin of more human-like comprehension of objects and agents in great apes (Byrne, 1997). Unlike monkeys, great ape individuals have shown in a number of ways that they are able to represent and take advantage of the knowledge, ignorance, or intentions of others. Yet great apes do not

live in more complex societies than monkeys (except in virtue of the fact that their greater understanding must inevitably make their societies richer in perceived complexity), so that theories that link all intellectual advance in the primate lineage to social pressures have no explanation.

Great apes do, however, confront a much greater feeding challenge than monkeys. Their large size, combined with locomotor adaptations for arboreal feeding while hanging below branches, mean that their long-distance travel is significantly less efficient than that of monkeys. Moreover, compared with Old World monkeys, apes' guts are less well adapted to cope with coarse leaves and unripe fruit, despite their absolutely greater size. Yet great apes occur sympatrically with Old World monkeys in almost all forests where they live, and they largely eat the same range of foods: fruit, leaves, and a wide range of minor items such as fungus, insects and bark. If monkeys are better adapted for exploiting this diet in every way, they it seems mysterious that any ape species survive at all.

Suppose, however, that the response to the severe ecological challenge of feeding competition from monkeys was an adaptation in apes that allowed *effective acquisition* and *social transmission* of complex feeding skills—skills that permitted more efficient feeding than the straightforward picking, peeling and stripping of monkeys, and hence access to otherwise unavailable food sources. Suppose further that this adaptation included behavior parsing, as a basis for quick and reliable learning of any skills which painstaking trial-and-error had caused to enter a local population, allowing development of traditions that transcended one generation.

On this conjecture, behavior parsing is not only a *necessary* part of the process of discerning others' purposes and how their behavior can go about achieving them in our everyday behavior, but was an evolutionary stepping stone on the way to the modern human ability in attributing cause and intention. If so, much that is “special” about human thinking has its evolutionary origins in dealing efficiently with tricky plant problems.

¹ I would like to thank Susan Hurley and Nick Chater for their considerable work in organizing the stimulating and fruitful meeting at Royaumont from which this chapter stemmed. Ideas in the chapter have benefited from formal commentary at the meeting, by Birgit Elsner and Tom Zentall, and the latter also kindly read and commented on the first draft of the chapter. In addition, my ideas have profited from informal discussions with many people, whose (probably good!) advice I have not always taken.