

John Benjamins Publishing Company



This is a contribution from *Interaction Studies 10:1*
© 2009. John Benjamins Publishing Company

This electronic file may not be altered in any way.

The author(s) of this article is/are permitted to use this PDF file to generate printed copies to be used by way of offprints, for their personal use only.

Permission is granted by the publishers to post this file on a closed server which is accessible to members (students and staff) only of the author's/s' institute.

For any other use of this material prior written permission should be obtained from the publishers or through the Copyright Clearance Center (for USA: www.copyright.com).

Please contact rights@benjamins.nl or consult our website: www.benjamins.com

Tables of Contents, abstracts and guidelines are available at www.benjamins.com

Representations underlying social learning and cultural evolution

Joanna J. Bryson

University of Bath / Konrad Lorenz Institute for Evolution and Cognition Research

Social learning is a source of behaviour for many species, but few use it as extensively as they seemingly could. In this article, I attempt to clarify our understanding of why this might be. I discuss the potential computational properties of social learning, then examine the phenomenon in nature through creating a taxonomy of the representations that might underly it. This is achieved by first producing a simplified taxonomy of the established forms of social learning, then describing the primitive capacities necessary to support them, and finally considering which of these capacities we actually have evidence for. I then discuss theoretical limits on cultural evolution, which include having sufficient information transmitted to support robust representations capable of supporting variation for evolution, and the need for limiting the extent of social conformity to avoid ecological fragility. Finally, I show how these arguments can inform several key scientific questions, including the uniqueness of human culture, the long lifespans of cultural species, and the propensity of animals to seemingly have knowledge about a phenomenon well before they will act upon it.

Keywords: social learning; cultural evolution; imitation; cognition; memetics

1. Introduction

Why are humans the only species exhibiting exponentially accumulative culture? Language obviously currently facilitates this process, but language is also an *example* of the sort of accumulated cultural artifact which needs to be explained. Researchers have demonstrated that a wide variety of species — from apes through fish and ants — regularly exploit socially-transmitted behaviour (de Waal and Johanowicz, 1993; Laland and Williams, 1997; Whiten et al., 1999; Van Schaik et al., 2003; Perry et al., 2003; Galef Jr. and Laland, 2005; Franks and Richardson, 2006).

Thus the basic capacity of social learning is present in these species, and has proven adaptive at least in limited forms.

One way to approach this question is to postulate a large set of tradeoffs between different mechanisms for acquiring behaviour. If we extend the definition of *learning* beyond its ordinary-language sense to mean *any mechanism of behaviour acquisition*, then we can define a continuum of representations extending from genetic coding to the current very-short-term perceptual memory an individual organism has for whatever image happens to be on its retina (Bryson, 1996). Along the way — for at least some species — there are intermediate repositories such as long-term individual memory, and the versions of these memories distributed through a population that make up culture. Thinking about learning in this way, we can ask why for some species learning is focussed at the purely genetic end of the spectrum, while for others individual experience and culture matter more.

In this article I attempt to answer this question by examining it from a computational perspective. Theoretical computer science studies specifically what can be learned or computed in what sort of systems, and how much time it might take. The two key determinants for types of computation are representations and processes. *Representations* are simply the way in which information is encoded when it is inside an agent, and *processing* amounts to converting information from one form of encoding to another. For example, information about an environment might be transmitted to an agent in via light, and might initially be represented on the agent's retinae as a relatively direct (though two-dimensional) mapping of brightness and location. Further cognitive processing results in the information also being represented in various parts of the agent's brain as sets of edges, colours, objects, and egocentric locations. Under appropriate conditions of internal motivation, this information might be further processed into representations in the motor cortex that ultimately result in muscular contractions and thus action.

For the purpose of this article I define *culture* as any behaviour routinely acquired from conspecifics by non-genetic means. The qualifier "routinely" implies that this is behaviour shared by some significant subset of an adult population, not just chance behaviour-patterns expressed by one or two individuals due to shared experience of a unique event (Bloch and Sperber, 2002).

Individual long-term memory does not exist in all species, and even where present, the extent to which it is general-purpose is often overestimated (Gallistel et al., 1991; Roper, 1983; Sherry and Schacter, 1987). Even simple stimulus-response conditioning does not work for all stimuli to all responses. Pigeons can learn to peck for food, but cannot learn to peck to avoid a shock. They can, however, learn to flap their wings to avoid a shock, but not for food (Hineline and Rachlin, 1969). Similarly, rats presented with 'bad' water learn different cues for its badness depending on the consequences of drinking it. If drinking leads to shocks,

they condition to visual or auditory cues, but if drinking leads to poisoning they learn taste or smell cues (Garcia and Koelling, 1966). Such examples indicate that even forms of learning routinely considered to be associative in experimental psychology are subject to genetic biases and constraints. It is precisely such biases and constraints towards the capacity to represent and process particular types of information that I hypothesise vary between species, and thus form proximate explanations of what behaviour acquisition capacities species possess.

In this article I review the representations underlying social learning. I do so by starting from a simplified taxonomy of social learning types observed in nature. I use this to suggest the representations that necessarily support these forms of learning, and to examine how much more complex or costly social learning might be than simple individual learning. Next I consider whether the learning necessary for the culture we see in animal species is sufficient to produce the sort of cultural accumulation we see in humans. To begin though, I provide background and motivation by demonstrating the potential utility of social learning and clarifying when this mechanism is useful for acquiring behaviour.

2. Motivation: The computational complexity of learning

Learning is easy. Assuming that all you mean by 'learning' is *changing values inside a representation*. Constraining learning so that it does something useful is the hard problem. This was the conclusion of Marler (1991) after he examined the surprising diversity of mechanisms that have evolved to satisfy one relatively simple problem: the transmission of birdsongs between individuals of a species. This has also been the experience of artificial intelligence (Bishop, 2006). The lesson from machine learning is that the only way to get a learning system to do its work in a tractable amount of time is by appropriately restricting the conceptual space in which it will search for appropriate values (Wolpert, 1996). In Bayesian terms, the problem is selecting an appropriate class of models for the learning domain (Chater et al., 2006). These models are a necessary precondition for stating the learning-system's prior expectations, which are then modified by evidence (perception) to form the agent's current theory of the world. These sorts of techniques and results derive from the formal analysis of the computational complexity of learning (Sipser, 1997). They also provide an explanation for the fact that the vast majority of individual learning in nature tends to be specialised to task (Gallistel et al., 1991; Roper, 1983).

Results concerning the computational tractability of learning in general have strong implications for social learning in particular. Even in species that do possess some general learning capacity, the probability of an individual stumbling across

a useful piece of knowledge within its lifetime might tend to be low. Where an individual agent, either animal or artifact, has some capacity for general learning, it may very well be in its interest to acquire knowledge that has already proven useful to other similar agents. This line of reasoning has led to the recent surge in interest in culture and social learning (Schaal, 1999; Hurley and Chater, 2005; Wood and Bryson, 2007).

Although results concerning tractability, learning and concurrent processing are well established in computer science, to date they have not fully penetrated the social learning literature. Even leading researchers who strongly advocate the importance of social-learning phenomena like cultural evolution can still underestimate the computational power of a society of communicating learners. For example, Boyd and Richerson (1995, 1985) propose an explanation for limits on social learning which presumes that social learning is less adaptive relative to individual learning the faster the learners' environment changes. This is because information not gained directly from the environment might be out-dated. An assumption of this model is that the only advantage of social learning is avoiding the cost of exploratory behaviour, which may be dangerous. They assume that in the limit case, if no agent is learning from its own experience, then all agents' knowledge will become out-of-date as the environment changes. They propose therefore that the faster the environment changes, the more agents should rely on individual learning.

This model of the tradeoffs does not properly take account of the full role of time in behaviour (including learning), nor of the impact of concurrency. If individual adaptation has been well-circumscribed by genetic evolution, then the main problem with individual learning is not how risky it is in terms of physical danger, but how *improbable* it is to learn at all. For many behaviours cognitive species such as higher primates engage in (e.g. nut cracking with stones, speaking English), discovery of a technique might be so unlikely that the *temporal* cost to an individual of learning it might be intractable. That is, without social learning only a very small percentage of individuals might ever learn the skill in their own lifetimes. In contrast, learning socially might take relatively little time, perhaps a few exposures to the correct technique being performed by a suitable model, plus some experimentation and refinement by the individual. Thus a rapidly changing environment might in fact favour agents that learn quickly socially.

Note that the seed 'innovation' that is passed through a society by this sort of process might not even be the result of deliberate individual exploration, but rather a chance event observed or even misperceived by conspecifics. Consider a herd of animals forced to cross a river — once the first few are successful in reaching the other side, the following animals can attempt to follow the path of the new leaders. Where the number of behaviour options is large and their consequences

not evident, innovation by chance events might be nearly as likely as innovation through individual insight. From the perspective of a process of cultural evolution, whether a new behaviour was discovered by chance or insight is irrelevant. All that matters is whether the behaviour is prone to replication (Blackmore, 1999).

Evolution itself could be construed as a form of social learning under the definition proposed in this article's introduction. A species accumulates a set of behaviours which are 'remembered' genetically. Of course, biological evolution is necessarily a slow process in that useful traits spread through a species only by transmission from parent to child. Conventional social learning introduces the possibility for a rapid spread of innovative behaviours between every member of a society. However, as I discuss below, there may be good reason for learning to be damped and slowed, and for limits on the pressure for behavioural conformity. But this discussion is best taken after we are more informed about the types of social learning currently observed in nature. For this I turn to examining its types and representations, with an eye to analysing their computational attributes.

3. Types of social learning

This section presents a simplified taxonomy of the forms of generic social learning generally considered to be exhibited in nature. This taxonomy is derived from somewhat more elaborate ones which have become well-established over the last decade or so following earlier efforts (Whiten and Ham, 1992). For more complete descriptions, see Zentall (2001) or Whiten (2006). After presenting this taxonomy, I will analyse its implications for representational substrates, then critique its plausibility.

In the below descriptions, the term *model* refers to another agent that already holds and expresses a behaviour being socially learned. Modelling does not need to be deliberate.

- *Social facilitation*: Also known as contagious behaviour. The increased propensity to express an already known behaviour when others express it. The classic example is yawning. This mechanism may not in itself be seen as true learning since there is not necessarily a long-term change in behaviour. However, it can be part of a mechanism for learning to express a behaviour in a particular context.
- *Local enhancement*: An agent acquires a propensity to be in a particular area where it sees its model. This in turn (and in combination with other species-specific biases) leads to the agent displaying a similar behaviour as its model. For example, an agent that follows a model into a patch of unfamiliar plants may discover subsequently that some of the plants are edible just through

random exploration. This is an example of social learning where a new behaviour is learned, but not directly from the model.¹ Rather, a small amount of information from the model facilitates *individual* learning by the agent.

- *Stimulus enhancement*: An agent becomes interested in an object another agent has acted upon, and in the course of exploring that object discovers affordances known to the previous agent. It thus now expresses a new behaviour, again similar to its model.
- *Goal emulation*: An observing agent notices the model has accomplished something interesting, and acquires the goal of accomplishing the same thing. Again, with enough species-specific and/or environmental constraints, the end behaviour itself may be quite similar, or the agent may find quite a different way of achieving the same goal. But the agent's new behaviour would have been very unlikely to be expressed without the observation of the model's original achievement.
- *Program-level imitation*: Postulated originally by Byrne (1995) and elaborated with further evidence by Byrne & Russon (1998), program-level imitation is the acquisition of sequential or even hierarchical 'plans' organising actions into complex behaviours. Byrne & Russon (1998) give the example of an orangutan living near a camp that begins doing laundry. This is also sometimes referred to as 'staged emulation', because the individual actions are not necessarily learned new, but rather the combination of the actions are associated with each other and with a set of stimuli.
- *Gestural- or vocal-script imitation*: Temporally-precise replication of behaviour over a large number of degrees of freedom of continuous manual or vocal gestures. This element is the closest of this taxonomy to the ordinary-language meaning of imitation. An example would be copying everything, including a speaker's accent and timing in repeating an exact verbal phrase, or imitating the timing and posture associated with an individual's gesture or gait.

Notice that I have deliberately *not* included simple *emulation* or *imitation* here. I discuss this further in the Analysis section below, but briefly these terms have been subsumed by the more precise taxonomy elements I do list.

I also have not addressed highly specialised, species-specific, evolutionarily 'ritualised' forms of learning, such as tandem running in ants (Franks and Richardson, 2006) or imprinting in some hatchling birds (Lorenz, 1937). As I discussed in the introduction of this article, the capacity for general-purpose learning in animals is probably overestimated. However, Lorenz (1937) discriminates imprinting from associative learning in several ways:

1. it occurs in a particular period in life,
2. it is irreversible and unrepeatably, and

3. acquisition of the stimulus is completed long before the response behaviour is expressed.²

Thus in this case the behaviour is nearly entirely specified genetically, with only small details that cannot be learned by biological evolution (like individual variation in mothers' appearance or the location of a nest) left to individual adaptation.

4. Primitive elements of social learning

What does this taxonomy imply about the computational properties of animal social learning? To answer this question, we first need to attempt to analyse how these forms of social learning might be implemented. In this section, I consider what elements must constitute the capacities described above.

Computer scientists — including those who build Artificial Life (ALife) and Artificial Intelligence (AI) — often speak of 'primitives'. *Primitives* are the fundamental components, the building blocks of behaviour. Like atoms, these primitives are themselves constructions, e.g. in animals of some combination of neural coding, sensory organs and structural components, etc. But any particular description of behaviour requires discussion to be grounded in basic units at some level of abstraction. In ALife or AI the primitives are built in conventional computer code. The intelligence of the system then must express these primitives in reasonable contexts and reasonable orders. There is some evidence that brains also work this way. Complex gestures and stimuli can be represented by single nerve cells as witness by the cells' response to these stimuli (Rizzolatti et al., 2000; Perrett et al., 1987). Further, stimulating single cells can result in complex behaviour expression such as taking a hand from the ground to the face while simultaneously opening the mouth (Graziano et al., 2002). While clearly entire systems are required to analyse the stimuli or construct the response, the fact that these single-cell representations exist indicates the brain is indeed hierarchically organised, with executive systems ordering the primitive actions (Prescott, 2007) and the cerebellar system providing overall fluency and coherence (Altman & Bayer, 1996).

The first stage in understanding computation is to understand the 'data structures' or representations that are being manipulated (Abelson et al., 1996). By *representation* I mean the stored form derived from experience which is used to generate action.

In an attempt to understand the representations underlying social learning, I will begin by defining a few primitive cognitive capacities. Each of these is a perceptual process — a function that transforms sensory input into a useful representation. Notice that not all of these primitives will necessarily appear in the final theo-

ry. Rather, I am starting with a set of primitives I believe underly common theories of social learning. This article will not in the end necessarily support all of them.

- *Context Identification* The learning necessary to recognise a particular stimulus or, more likely, stimulating situation. This is a form of perceptual memory. It cannot be a simple retinotopic map (e.g. to remember an image), since exact visual context matches are exceedingly rare. Rather, it must be sufficiently abstract to generalise across differences such as relative distance or changes in light conditions.
- *Goal Mapping* is attribution to another agent of a particular aim, desire or intent. It is generally believed that such goals can only be identified through being mapped to a similar sort of aim, desire or intent of the observing agent. For example, “Maybe she did that because she was hungry [as I am].”
- *Action Mapping* is the association between a behaviour or element of a behaviour from the observed animal to a similar behaviour within the repertoire of the observer. To keep things simple, we take ‘behaviour’ in a very general sense here, including perceptual acts such as focusing attention necessary to a task, as well as gross motor movement.
- *Body Mapping* is the identification of a particular body part of an observed agent to a corresponding body part of the observer. Although once thought a cognitive task (Piaget, 1964), this is now seen as some combination of an evolved and/or a developed perceptual capacity (Meltzoff and Moore, 1983; Rizzolatti et al., 2000; Hurley, 2005)
- *Coordinate Mapping* is the identification of a particular location in space with respect to the observed agent to the equivalent egocentric-space coordinate of the observer.

The above primitives create semantic-like content necessary for intelligent action. Notice that their definitions also necessarily imply another, more basic set of representational primitives: *contexts*, *goals*, *actions*, *body parts*, and *coordinates*.

In addition, constructing intelligent behaviour from the above may require more syntax-like process primitives — capacities such as:

- the ability to *associate* two primitives, for example a context to a goal,
- the ability to *chain* two items, for example two sequential steps in a procedure,
- the ability to *heighten attention* to a particular context or aspect of context, and
- the ability to *desire* (acquire) a new goal.

Again, I am not proposing that all these capacities are available in all (or even any) agents capable of social learning. I *am* claiming that these capacities are needed

in order to display all the forms of social learning mentioned in the original taxonomy of social learning.

5. Analysing the taxonomy

There can be no social learning without individual learning, since individuals perform social learning. In fact, social learning might be seen as one or more special forms of individual learning. Constraints to facilitate learning (as were described in the first two sections of this article) may have evolved in order to take best advantage of the information available from conspecifics (Bryson and Wood, 2005; Bargh and Chartrand, 1999; Wood, 2008). Examples of such useful constraints range from predispositions to visually attend to conspecifics through giving more weight to imitating individuals with higher prestige.

In this section I look for the simplest, most parsimonious computational explanations for the types of social learning found in the established social-learning taxonomies, which I described two sections ago. This analysis is done in terms of the primitives just described. There are three questions I wish to address with this analysis:

1. Can the observed capacities for social learning tell us anything about our *representational* capacities?
2. Can we simplify or eliminate any of the assumptions currently existing in the literature?
3. Is there anything about human social learning that might explain (rather than *be explained by*) the great extent of our cultural accumulation?

In the 1990s there was considerable discussion about whether there was any evidence in nature (excluding humans) for social learning in general and imitation in particular (Whiten and Ham, 1992; Tomasello and Call, 1997; Dautenhahn and Nehaniv, 1999). Imitation came to be defined in some circles as occurring only when a “truly novel action” was acquired. Social influence on the production of known actions became known as *emulation*. In terms of the taxonomy I’ve presented, one of the significant debates at the time concerned whether something akin to program-level imitation (which one might also call “staged emulation”) could be used to compose a novel act and create true imitation, or if only a process more like script imitation qualified. Since the focus of this article is on mechanisms and their computational properties, I eschew the semantic debates, although the arguments that came from them have certainly informed the taxonomy and this analysis. I have avoided using either *imitation* or *emulation* except when qualified to specify my meaning.

5.1 Simple social learning mechanisms

Social facilitation, location enhancement and stimulus enhancement are very little more than ordinary individual learning, at least in terms of the mechanics of the learning. Location and stimulus enhancement assume context identification, plus either an association with an established behaviour or individual learning of new behaviour, either of which occurs as a consequence of being attentive to the location or stimulus. These forms of social learning in no way assume goal, action or body mapping. In fact, social facilitation requires no learning at all. However in the case where the social facilitation keeps happening in the same context, it may result in learning by increasing the probability of associating that context with a known action (Chartrand and Bargh, 1999; Eenshuistra et al., 2004). A possible example of learning resulting from social facilitation might be the gradual tuning of the context in which innate warning cries are expressed by vervet monkeys (Seyfarth et al., 1980).

Goal acquisition through goal emulation might seem as simple as stimulus enhancement, since it might also require the acquisition of a single primitive element, the goal. However, motivations are fundamental to an agent's intelligence. It is not easy to see how a totally new goal — with its attendant drives and emotions — would be incorporated into an agent. Goal emulation may be more like operant conditioning. An action or a perceptual context might become identified with a pre-existing drive, and thus become desirable itself. This reduction can be applied to simplify or eliminate goal mapping as a primitive. Goal emulation could be accounted for through action mapping, with the additional recognition or association of the observer's own desire to its perception of the target's action. If the two animals are in a similar state, whether due to shared history (e.g. a troop hasn't eaten yet today) or shared responsiveness to a perceptual context (e.g. an innate fear of loud noises), then the probability of sharing a drive may be high enough for reasonably accurate learning to occur.

At its simplest then, goal emulation might be viewed as the association of a behaviour to a context, where that context is some combination of a perceptual context and an internal drive. Put even more simply, it is socially acquired stimulus and response.

5.2 Program-level imitation

Program-level imitation is essentially an ordered set of goal emulations — a structured association of contexts to actions. Note though that this is the first type of learning specified which requires more than simple associative links between generically useful context and action representations.

The extent of the “structure” in the structured association is much debated. It is tempting to take what appears to be the simplest explanation, and assume that simply associating sufficient perceptual context (perhaps including recent memory of prior events) to action responses will allow an otherwise undifferentiated set of stimulus-response pairs to form the representation for learning a new task. However, this is not what humans or animals appear to do. In extensive experimentation with modelling human learning, Anderson et al. (1997) determined that intelligence driven by sense-action pairs requires specification of a subset of pairs to be active in a particular task context. Even within the task-specific subset, they also require each pair to be associated with a probability for being useful, referred to as a *utility value*. In my own research, I have found evidence that even this amount of information is not sufficient. Rather than probabilities of success, accurate representations of *priority* of one task-element over another are needed to guarantee task consummation. We have evidence that this better describes the behaviour of monkeys at least (Wood et al., 2004; Bryson and Leong, 2007), as well as being a useful representation for organising artificial intelligence (Bryson and Stein, 2001; Bryson, 2003).

From the computational perspective, it is easy to see why animals would need this restriction of scope. First, if search is sequential — which for some cognitive tasks seems to be the case (Horowitz and Wolfe, 1998) — then searching more options takes a longer time. Secondly, even if the search can be done concurrently, it might be difficult to discriminate between a large number of alternatives reliably (von der Malsburg, 1995; Cowan, 2001). Thus forming hierarchical associations between a higher-level goal context with a prioritised set of (sub-)goal emulation pairs might be a sensible structure. However, we have known since Lashley (1951) that some action chains are stored as simple sequences with no intervening perceptual-context checks. There is simply no other way to account for their speed of execution (Davelaar, 2007).

Forming and ordering the goal-emulation pairings may be two functions of two different regions of the hippocampal complex (Bryson and Leong, 2007). There is also evidence of neural representations for meta-level task information such as sequence ordering or identifying the present task (Tanji, 1996; Fogassi et al., 2005; Iriki, 2006).

Whiten (1998) has reported that chimpanzees not only imitate sequential³ behaviour, but do so more accurately on subsequent trials if the demonstration is repeated. This increase of fidelity — essentially moving from goal emulation to program-level imitation — might result from individual learning of the affordances of the task during the course of early imitation attempts. Once affordances are known, they can provide lower-cost (more abstract) representations and thus facilitate more detailed learning. Or there may be a social drive to emulate with more care when prompted

by a repeated demonstration. However, these increasing-fidelity results have not yet been well-supported through replication, although the hierarchical structure of social task learning has, at least in young children (Whiten et al., 2006).

5.3 Vocal- and gesture-script imitation

There is one further type of social learning and underlying representation that I would argue is not well described in the currently popular taxonomies (Zentall, 2001; Whiten, 2006). Humans, songbirds, and a few other species are capable of temporally-precise imitation of vocal and physical gestures on a large number of axes, or degrees of freedom. Examples of these axes in the case of vocal imitation are: volume, pitch, timbre, and formants; all of which can be replicated with high precision and in precise and temporally-accurate sequence. In the case of manual gestures, the same precision and temporal accuracy hold for a large number of different joint angles. Such a performance requires encoding a high-information-content temporal script which are fundamentally different from the sorts of representations described above, not least because they have precise duration.

One might imagine that gestures could be extended sequences of many body or coordinate mappings — a sort of mega-program imitation. However, as explained in the previous section, it is well-established that there is no neurological means by which some of our most-rapid sequences of action expression can be launched by individual context triggers. Rapid sequences must be set in motion as a unit with special neurological mechanisms providing timed release control (Lashley, 1951; Henson and Burgess, 1997; Davelaar, 2007). Presumably, script imitation would also require such special mechanisms.

The special case of vocal imitation in songbirds (and some other birds) has been the subject of extensive neuroscience research (Leonardo, 2004). The upshot seems to be that male songbirds exploit a specialised neurological substrate for learning songs, and that this substrate is not capable of learning and production simultaneously. Similar constraints have been found in humans, indicating humans act using at least some of the the same representational substrate they use to learn (Müsseler and Hommel, 1997).

While the neuroscience of human vocal and gesture imitation is not yet as well-understood as in birds, it may well have similarly dedicated mechanisms (Doupe and Kuhl, 1999). Pöppel (1994) documents a privileged representation of “phrases”, within which humans are capable of precise temporal memory and replication. These have a maximum duration of two to three seconds — the exact duration seems to be under intelligent (though not deliberate) control and situation-appropriate. That is, we tend to remember salient phrases of speech, music or gesture with appropriately-lengthed memory. The maximum possible duration of

such episodes is presumably a cognitive constraint. Pöppel (1994) draws attention to the fact that most poetry and music consists of phrases of this length.

I believe the capability for high-fidelity, temporally-accurate gesture imitation may be the key to the puzzle posed in the introduction to this article — why human culture is different, at least from other primates? I will discuss these implications further in the Culture section below. But at the level of representations, the most relevant data comes down to a lack of evidence that primates other than humans are capable of temporally-precise gesture or vocal imitation (Fitch, 2000; Bispham, 2006; Fitch, 2007). There is no documentation of vocal imitation, entrainment or the like in primate species despite extensive observation and in the case of some language experiments active coaching. Vocal imitation does exist in some marine mammals, bats, birds and elephants where it seems to play an important role in social signalling and bonding (Poole et al., 2005; Beecher and Campbell, 2005).

Of course, in non-human animals it is difficult to distinguish for certain whether negative results indicate a lack of the target capacity in the animal or a lack of experimental ability to elicit the behaviour. However, humans have a strong tendency for entrainment — the tendency to follow a rhythm being produced by another. The strength of the tendency is partially indicated by the English idiom of “falling into step”. For a human, resisting entrainment is like resisting gravity, it takes conscious will. However, there are no reported cases of other primate species entraining to the same rhythm, even where there has been effort to coax it (Bispham, 2006; Fitch, 2007).

If apes are not capable of what I am calling here *gesture-script imitation*, how can they perform “do as I do” tasks? These involve imitating the gestures of a demonstrator (normally human) such as clapping one’s self, or jumping up and down (Custance et al., 1995). These sorts of imitation certainly do require sophisticated body mapping, and some process of action sequencing. But because chimpanzee and human bodies are similar, it may be that a temporally-low-resolution representation — snap shots of the body configuration at the start and end points — is sufficient to generate comparable actions within the tolerance required by those coding this research (Custance et al., 1995). In fact, in less-closely related species, less-careful body mapping is sometimes demonstrated (Custance et al., 1999; Topál et al., 2006). Even in human children, precise body mapping is only followed when the children assess it to be an important part of the demonstration (Gergely et al., 2002). Thus while “do as I do” can be quite sensibly categorised as gesture imitation in ordinary language, in the context of this paper it does not fall into this category because it does not require a rich, temporally-precise script. Rather, it is either directed social facilitation if the end-point positions are familiar and easily recognised or goal emulation otherwise.

5.4 Conclusions

The evidence of the previous section leads to the following conclusions about the representation of social learning:

- The majority of social learning observed in nature does not require computationally complex information to be transmitted between the individuals. There is no need for precise locations, postures, temporal measurement, or a count of iterations involved in steps with distinct cycles. Rather it can mostly be summarised as learning salient contexts — a capacity any intelligent species already requires. The difference in social learning is that salient contexts are flagged by the presence or attention of a model. In some cases, actions appropriate to those contexts also need to be learned.
- Internalising the entire motivational content of novel goals is probably not necessary. Rather goal acquisition probably consists of mapping of desires to new contexts for satiating existing emotions or drives.
- There is some evidence that some species may learn structures (such as sequences or hierarchies) of context / goal pairings. This has mostly been suggested in primates to date, but even here the evidence is generally considered equivocal.
- Species that perform temporally-precise vocal (or other gestural) imitation require a different, specialist representation to encode temporal scripts containing far more information than ordinary social transmissions.

6. What makes humans different?

I now return to the question of why humans are unique in having exponentially accumulating culture. My explanation hinges on a difference in representational capacities, described in the previous section — in particular, that for storing and recalling short temporally-precise scripts.

6.1 Accelerating cultural evolution

That humans have this extra capacity other primates do not is probably an accident of sexual selection. For example, our ancestors may have pair-bonded through duets, and this may have supplied selective pressure for vocal imitation (Vanechoutte and Skoyles, 1998). My theory is that this accident (whatever its precise nature) provided us with a representation suitable for a memetic cultural

evolution explosion, similar to the sudden increases in diversity of life we see in paleontological history.

There has been a lot of attention to innovation in recent literature regarding evolution (both cultural and biological), but this obsession can lead to neglecting the difficulty of *preserving* “good tricks” (Dennett, 1995) once discovered or stumbled upon (Dawkins, 1976). Random variation cannot be entirely guarded against. Variation is key to evolution — there must be something to select between for directed change to occur. On the other hand, the most important information *can* be well-defended. One of the key mechanisms for doing this is redundancy. If one copy of a key instruction is corrupt, there might be another copy available. For example, if a core system of development fails, other systems may be able to compensate (Winslow et al., 2007).

Much more information is stored in the two-to-three second phrases of detailed transcription that I have postulated underly our capacity for gesture imitation than in the simple context-action pairs underlying programme-level imitation. It is like the difference between a black-and-white photograph and a film with sound and colour. Thus knowledge represented in this information can be highly redundant, as we know language is (Buzo et al., 1980). My suggestion is that this redundancy provides the robustness necessary for important information to be stored, but ignorant, unsupervised processes of variation still operate, occasionally introducing new features of greater utility. Thus the difference between our culture and the culture of our fellow apes might be as vast as the difference between simple replicating chemicals in the “primordial soup” and the complex genetic code of eukaryotes.

I’ve discussed support and implications of this idea at greater length elsewhere (Bryson, 2008). Briefly these amount to known results in genetic algorithms concerning redundancy and complexity (Burke et al., 1998), and a set of theories on the origins of language that postulate that evolution of language from crossover-like operations on standard primate long calls, which imply semantics-laden phrases preceded the evolution of a lexicon (Wray, 1998; Arbib and Bickerton, 2008)

6.2 Representations of Self and Other

My theory stands in contrast to some other recent but well-developed theories of human uniqueness, which have to do with the capacity for representing and reasoning about the self and others (Stiller and Dunbar, 2007; Moll and Tomasello, 2007). Essentially, cooperative behaviour like language and culture is considered more evolvable if the agent can represent both its own needs and the needs of the other. This is because global optima can be conceived of, recognised and then

made into goals. The situation improves further if agent A can reason about agent B reasoning about agent A.

I am quite convinced that these sorts of cognitive capacities, now present, accelerate the development of culture and human society. However, we know that they are not necessary for the evolution of altruistic behaviour like cultural communication (Nowak, 2006; Čače and Bryson, 2007). My own suspicion is that these capacities are themselves, like language, another example of the cognitively-useful culturally-provided mechanisms we are now good at accumulating. Although altruistic giving and goal assistance is displayed in strikingly young children (Moll and Tomasello, 2007), these children have all been raised from birth in strongly social contexts. While some social behaviour has no doubt been canalised in the million years since we have been using advanced stone tools (Ambrose, 2001), this doesn't indicate that such capacities were necessarily the initial catalyst for rapid cultural accumulation. Given their late appearance in ontology (Piaget, 1964), actual cognitive reasoning about others' notions of oneself would probably not arise without tutelage, even if there is some level of genetic predisposition. Parents are strongly motivated to provide such tutelage since they want their children to predict and accommodate their own behaviour.

6.3 What limits cultural evolution?

Even if the above hypothesis about why human capacities for social learning are unique is correct, it does not in itself explain why more species don't exploit social learning to a greater extent than they currently do. As I argued before, memetic evolution has the potential to be a powerful means for discovering and disseminating new and more optimal behaviour. However, I believe that even where such cultural evolution exists, it must co-evolve with a set of constraints that damp its effects on the society and its ecosystem.

In the most basic case, this is obvious if you think about it — if you are in a room with other people, look around yourself. Would it be a good idea if all of you converged on identical behaviour right now? Would you all even fit in the same chair? Whatever the 'crossover' equivalent is in a type of social learning — the mechanism of recombining good tricks from other conspecifics — it must maintain diversity in order to support individual survival.

Some people worry that this kind of argument is group-selectionist — why should an individual take on greater risk by assuming a more variant position? Firstly, in terms of *biological* variation it is not the individual that determines its variation, but its parents. If the environment is volatile, it may be in the interest of the parents' genes to produce a variety of offspring. This genetic variation in turn varies the optimal behaviour for the individual, since genetically different

individuals may find themselves ‘fittest’ at different survival strategies. Secondly, there may be no obvious evidence that one slight variant in behaviour is better than the next. Animals tend to be good at distributing themselves over available foraging resources in proportion to the resources’ value (Seth, 2007). But finally, altruistic behaviour is actually fairly easy to evolve in vertebrate circumstances. These include living in a world where your current location is determined by your previous one and having a lifespan that overlaps with your offspring and parents’ (Mitteldorf and Wilson, 2000; Nowak, 2006; Taylor and Irwin, 2000; Čače and Bryson, 2007).

In this case, what pressures do damp the utilisation of social learning? There may be selection against excessive conformity due to longer-term consequences, such as population booms and busts. When a new, rich food source is discovered a population might rapidly expand while exploiting it, but if the resources is over-exploited to extinction then the subsequent stress to the population might lead to a demographic crash risking extinction.

Perhaps social learning is not only a valuable source of new behaviour, but also a risky strategy prone not only to some probability of spreading negative traits (Bechlivanidis, 2006), but to positive feedback cycles such as I have just described. In that case, we may find that where social learning exists it has been stabilised with co-evolved limits and damping mechanisms. For example, an action may be ‘learned’ in some sense quickly, but its expression inhibited in most circumstances until sufficient social cuing releases the behaviour. This is one way to control for noisy input, and may be useful for individual learning in general, but particularly for the cases concerning social learning.

This perspective has implications for a number of open research questions.

- *Why do cognitive species have such extended periods of development?* A long development period is necessary for any species that learns novel behaviour (and so is a candidate for cultural evolution) because individual experiences must be aggregated and integrated (Barrickman et al., 2008). Development, with its different phases of specialised learning, may hold key biological constraints for facilitating cultural evolution.
- *Why do primates learn to recognise behavioural patterns more quickly than they learn to express them?* This phenomenon, also described as “looking vs. knowing” or “perceiving vs. acting”, has been well-documented in infants (Spelke et al, 1992; Hood et al., 2000) and monkeys (Santos and Hauser, 2002; Hauser, 2003). As per the above arguments, it may not be a good idea to act on knowledge too early. However, relatively uncertain knowledge may safely be used to inform choices in *observation*. This increases the probability that the scarce

resources associated with learning (see the Introduction) will be focussed on something that needs their attention.

- *Why are humans the only species that have rapidly accumulating cultural evolution?* Being the only species capable of transmitting precise temporal scripts (possibly, but not necessarily, through vocal imitation) may give us two advantages. First, we may be better computational platforms for cultural evolution because our base representation for the cultural ‘genotype’ carries enough information to protect critical information while allowing unsupervised variation. Second, also because we can transmit more information faster, we can transmit culturally not only actions but principles for applying them. Thus the limits and damping systems can be customised and transmitted memetically *along with the behaviour itself*, rather than having to be entirely biological or genetic.

7. Summary

I began this article by discussing the computational difficulties of all learning (broadly defined), and described the potential computational advantages of social learning as a special case. In this I critiqued the Boyd and Richerson (1995, 1985) account of the tradeoff between individual and social learning, showing that in fact the power of concurrency — of many agents trying to learn in parallel — can lead to society members learning about environmental change much more rapidly than if they had each been searching on their own.

I next described a simplified social-learning taxonomy drawn from extant proposals, then provided a novel account for the representational and process requirements necessary for each element of the taxonomy. I used this to explain that most social learning is relatively simple for a species capable of individual learning, with the exception being a special capacity for learning temporally-precise ‘scripts’ on a large number of axes. Of the primates, only humans appear to have this special representation, and I suggest some combination of this capacity and our other cognitive capacities derived from being an ape is what made humans capable of cultural accumulation.

Finally I discussed why, given that social learning is powerful and accessible, it is not used more widely. Here I suggest that cultural evolution may require damping so that sufficient variation exists that a society is not competing within itself too extremely, nor subject to other negative consequences of overwhelming positive feedback.

Acknowledgements

Thank you to Andy Whiten and Mark Wood for many discussions of social learning and its representations. An earlier, much shorter version of this paper entitled “Representational Requirements for Evolving Cultural Evolution” appeared on <http://www.interdisciplines.org/adaptation> in April 2007, thanks to Adrianna Wozniak for her editing and comments. Many of the arguments, evidence and implications in the present article were developed during my sabbatical at the Konrad Lorenz Institute for Evolution and Cognition Research. Thanks to their board and the University of Bath for allowing me to pursue this research, and thanks to the KLI fellows particularly Christophe Heintz and Davide Vecchi for discussions. Thanks also to the reviewers of the present article for extensive and extensively helpful suggestions.

Notes

1. In fact my example is deliberately ambiguous as to whether the model was familiar with the food or not, though the canonical example is a mother duck leading ducklings.
2. “To explain the process in question [imprinting] as one of associative learning, one would have to assume that the reaction is, in some rudimentary stage, already present at the time when its object is irreversibly determined, an assumption which psychoanalysts would doubtless welcome, but about which I have doubts.” (Lorenz, 1937)
3. Possibly hierarchical, see Whiten, 1998 p. 280.

References

- Abelson, H., Sussman, G. J., and Sussman, J. (1996). *Structure and Interpretation of Computer Programs*. MIT Press, Cambridge, MA, 2nd edition.
- Altman, J. and Bayer, S. A. (1996). *Development of the Cerebellar System: In Relation to Its Evolution, Structure, and Functions*. CRC Press, Boca Raton, FL.
- Ambrose, S. H. (2001). Paleolithic technology and human evolution. *Science*, 291:1748–1753.
- Anderson, J. R., Matessa, M. P., and Lebiere, C. (1997). ACT-R: A theory of higher level cognition and its relation to visual attention. *Human Computer Interaction*, 12(4):439–462.
- Arbib, M. A. and Bickerton, D. (2008). Holophrasis versus compositionality in the emergence of protolanguage. *Interaction Studies*, 9(1). Special Issue.
- Bargh, J. A. and Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, 54(7):462–479.
- Barrickman, N. L., Bastian, M. L., Isler, K., and van Schaik, C. P. (2008). Life history costs and benefits of encephalization: A comparative test using data from long-term studies of primates in the wild. *Journal of Human Evolution*, 54(5):568–590.
- Bechlivanidis, C. (2006). An examination of the role of prestige in cultural evolution with the aid of agent-based modelling. Technical Report CSBU-2006-21, Department of Computer Science, University of Bath, UK.

- Beecher, M. D. and Campbell, S. E. (2005). The role of unshared songs in singing interactions between neighbouring song sparrows. *Animal Behaviour*, 70:1297–1304.
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer, London.
- Bispham, J. (2006). Rhythm in music: What is it? Who has it? And why? *Music Perception*, 24(2):125–134.
- Blackmore, S. (1999). *The Meme Machine*. Oxford University Press.
- Bloch, M. and Sperber, D. (2002). Kinship and evolved psychological dispositions. *Current Anthropology*, 43(5):723–748.
- Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press.
- Boyd, R. and Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16(2):125–143.
- Bryson, J. J. (1996). The design of learning for an artifact. In Sharkey, N., Hayes, G., Heemskerk, J., and Prescott, T. J., editors, *Proceedings of the AISB Workshop on Learning in Robots and Animals*, pages 12–15, Brighton, UK.
- Bryson, J. J. (2003). Action selection and individuation in agent based modelling. In Sallach, D. L. and Macal, C., editors, *Proceedings of Agent 2003: Challenges in Social Simulation*, pages 317–330, Argonne, IL. Argonne National Laboratory.
- Bryson, J. J. (2008). Embodiment versus memetics. *Mind & Society*, 7(1):77–94.
- Bryson, J. J. and Leong, J. C. S. (2007). Primate errors in transitive ‘inference’: A two-tier learning model. *Animal Cognition*, 10(1):1–15.
- Bryson, J. J. and Stein, L. A. (2001). Architectures and idioms: Making progress in agent design. In Castelfranchi, C. and Lespérance, Y., editors, *The Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL2000)*, pages 73–88. Springer.
- Bryson, J. J. and Wood, M. A. (2005). Learning discretely: Behaviour and organisation in social learning. In Demiris, Y., editor, *Third International Symposium on Imitation in Animals and Artifacts*, pages 30–37, Hatfield, UK. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour.
- Burke, D. S., De Jong, K. A., Grefenstette, J. J., Ramsey, C. L., and Wu, A. S. (1998). Putting more genetics into genetic algorithms. *Evolutionary Computation*, 6(4):387–410.
- Buzo, A., Gray Jr., A. H., Gray, R. M., and Markel, J. D. (1980). Speech coding based upon vector quantization. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 28(5):562–574.
- Byrne, R. W. (1995). *The Thinking Ape*. Oxford University Press.
- Byrne, R. W. and Russon, A. E. (1998). Learning by imitation: A hierarchical approach. *Brain and Behavioral Sciences*, 21(5):667–721.
- Chartrand, T. L. and Bargh, J. A. (1999). The chameleon effect: The perception-behavior link and social interaction. *The Journal of Personality and Social Psychology*, 76(6):893–910.
- Chater, N., Tenenbaum, J. B., and Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in the Cognitive Sciences*, 10(7):287–291.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Brain and Behavioral Sciences*, 24(1):87–114.
- Custance, D. M., Whiten, A., and Bard, K. A. (1995). Can young chimpanzees (*Pan troglodytes*) imitate arbitrary actions? Hayes and Hayes (1952) revisited. *Behaviour*, 132:837–859.
- Custance, D. M., Whiten, A., and Fredman, T. (1999). Social learning of an artificial fruit task in capuchin monkeys (*Cebus apella*). *Journal of Comparative Psychology*, 113(1):13–23.

- Čače, I. and Bryson, J. J. (2007). Agent based modelling of communication costs: Why information can be free. In Lyon, C., Nehaniv, C. L., and Cangelosi, A., editors, *Emergence and Evolution of Linguistic Communication*, pages 305–322. Springer, London.
- Dautenhahn, K. and Nehaniv, C., editors (1999). *Proceedings of the AISB'99 Symposium on Imitation in Animals and Artifacts*, Edinburgh. AISB.
- Davelaar, E. J. (2007). Sequential retrieval and inhibition of parallel (re)activated representations: A neurocomputational comparison of competitive queuing and resampling models. *Adaptive Behavior*, 15(1):51–71.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.
- Dennett, D. C. (1995). *Darwin's Dangerous Idea*. Penguin.
- Doupe, A. J. and Kuhl, P. K. (1999). Birdsong and human speech: Common themes and mechanisms. *Annual Reviews in Neuroscience*, 22(1):567–631.
- Eenshuistra, R. M., Weidema, M. A., and Hommel, B. (2004). Development of the acquisition and control of action–effect associations. *Acta Psychologica*, 115(2–3):185–209.
- Fitch, W. T. (2000). The evolution of speech: A comparative review. *Trends in Cognitive Sciences*, 4(7):258–267.
- Fitch, W. T. (2007). The biology and evolution of music: A comparative perspective. *Cognition*, 100(1):173–215.
- Fogassi, L., Ferrari, P. F., Gesierich, B., Rozzi, S., Chersi, F., and Rizzolatti, G. (2005). Parietal lobe: From action organization to intention understanding. *Science*, 308(5722):662–667.
- Franks, N. R. and Richardson, T. (2006). Teaching in tandem-running ants. *Nature*, 439(7073):153.
- Galef Jr., B. G. and Laland, K. N. (2005). Social learning in animals: Empirical studies and theoretical models. *BioScience*, 55(6):489–499.
- Gallistel, C., Brown, A. L., Carey, S., Gelman, R., and Keil, F. C. (1991). Lessons from animal learning for the study of cognitive development. In Carey, S. and Gelman, R., editors, *The Epigenesis of Mind*, pages 3–36. Lawrence Erlbaum, Hillsdale, NJ.
- Garcia, J. and Koelling, R. A. (1966). The relation of cue to consequence in avoidance learning. *Psychonomic Science*, 4:123–124.
- Gergely, G., Bekkering, H., and Király, I. (2002). Rational imitation in preverbal infants. *Nature*, 415:755.
- Graziano, M. S. A., Taylor, C. S. R., Moore, T., and Cooke, D. F. (2002). The cortical control of movement revisited. *Neuron*, 36:349–362.
- Hauser, M. D. (2003). Knowing about knowing: Dissociations between perception and action systems over evolution and during development. *Annals of the New York Academy of Sciences*, 1001(1):79–103. special issue on The Self: From Soul to Brain.
- Henson, R. N. A. and Burgess, N. (1997). Representations of serial order. In Bullinaria, J. A., Glasspool, D. W., and Houghton, G., editors, *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist Representations*. London. Springer.
- Hineline, P. N. and Rachlin, H. (1969). Escape and avoidance of shock by pigeons pecking a key. *Journal of Experimental Analysis of Behavior*, 12:533–538.
- Hood, B., Carey, S., and Prasada, S. (2000). Predicting the outcomes of physical events: Two-year-olds fail to reveal knowledge of solidity and support. *Child Development*, 71(6):1540–1554.
- Horowitz, T. S. and Wolfe, J. M. (1998). Visual search has no memory. *Nature*, 357:575–577.

- Hurley, S. (2005). Active perception and perceiving action: The shared circuits model. In Gendler, T. and Hawthorne, J., editors, *Perceptual Experience*, chapter 6. Oxford University Press.
- Hurley, S. and Chater, N., editors (2005). *Perspectives on Imitation: From Neuroscience to Social Science*. MIT Press, Cambridge, MA.
- Iriki, A. (2006). The neural origins and implications of imitation, mirror neurons and tool use. *Current Opinion in Neurobiology*, 16(6):660–667.
- Laland, K. N. and Williams, K. (1997). Shoaling generates social learning of foraging information in guppies. *Animal Behaviour*, 53(6):1161–1169.
- Lashley, K. S. (1951). The problem of serial order in behavior. In Jeffress, L. A., editor, *Cerebral mechanisms in behavior*. John Wiley & Sons, New York.
- Leonardo, A. (2004). Experimental test of the birdsong error-correction model. *Proceedings of the National Academy of Science*, 101(48):16935–16940.
- Lorenz, K. Z. (1937). The companion in the bird's world. *The Auk*, 54(3):245–273.
- von der Malsburg, C. (1995). Binding in models of perception and brain function. *Current Opinion in Neurobiology*, 5:520–526.
- Marler, P. (1991). The instinct to learn. In Carey, S. and Gelman, R., editors, *The Epigenesis of Mind*, pages 37–66. Lawrence Erlbaum, Hillsdale, NJ.
- Meltzoff, A. N. and Moore, M. K. (1983). Newborn infants imitate adult facial gestures. *Child Development*, 54:702–709.
- Mitteldorf, J. and Wilson, D. S. (2000). Population viscosity and the evolution of altruism. *Journal of Theoretical Biology*, 204(4):481–496.
- Moll, H. and Tomasello, M. (2007). Cooperation and human cognition: The Vygotskian intelligence hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480):639–648.
- Müsseler, J. and Hommel, B. (1997). Blindness to response-compatible stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, 23(3):861–872.
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563.
- Perrett, D. I., Mistlin, A. J., and Chitty, A. J. (1987). Visual neurones responsive to faces. *Trends in Neuroscience*, 9:358–364.
- Perry, S., Baker, M., Fedigan, L., Gros-Louis, J., Jack, K., MacKinnon, K., Manson, J., Panger, M., Pyle, K., and Rose, L. (2003). Social conventions in wild white-faced capuchin monkeys: Evidence for traditions in a neotropical primate. *Current Anthropology*, 44(2):241–268.
- Piaget, J. (1964). Part I: Cognitive development in children: Piaget — development and learning. *Journal of Research in Science Teaching*, 2(3):176–186.
- Poole, J. H., Tyack, P. L., Stoeger-Horwath, A. S., and Watwood, S. (2005). Elephants are capable of vocal learning. *Nature*, 434:455–456.
- Pöppel, E. (1994). Temporal mechanisms in perception. *International Review of Neurobiology*, 37:185–202.
- Prescott, T. J. (2007). Forced moves or good tricks in design space? Landmarks in the evolution of neural mechanisms for action selection. *Adaptive Behavior*, 15(1):9–31.
- Rizzolatti, G., Fogassi, L., and Gallese, V. (2000). Cortical mechanisms subserving object grasping and action recognition: A new view on the cortical motor functions. In Gazzaniga, M. S., editor, *The New Cognitive Neurosciences*, chapter 38, pages 538–552. MIT Press, Cambridge, MA, second edition.

- Roper, T. J. (1983). Learning as a biological phenomenon. In Halliday, T. R. and Slater, P. J. B., editors, *Genes, Development and Learning*, volume 3 of *Animal Behaviour*, chapter 6, pages 178–212. Blackwell Scientific Publications, Oxford.
- Santos, L. R. and Hauser, M. D. (2002). A non-human primate's understanding of solidity: Dissociations between seeing and acting. *Developmental Science*, 5, in press.
- Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(6):233–242.
- van Schaik, C. P., Ancrenaz, M., Borgen, G., Galdikas, B., Knott, C. D., Singleton, I., Suzuki, A., Utami, S. S., and Merrill, M. (2003). Orangutan cultures and the evolution of material culture. *Science*, 299(5603):102–105.
- Seth, A. K. (2007). The ecology of action selection: Insights from artificial life. *Philosophical Transactions of the Royal Society, B — Biology*, 362(1485):1545–1558.
- Seyfarth, R. M., Cheney, D. L., and Marler, P. (1980). Monkey responses to three different alarm calls: Evidence of predator classification and semantic communication. *Science*, 14:801–803.
- Sherry, D. F. and Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, 94(4):439–454.
- Sipser, M. (2005). *Introduction to the Theory of Computation*. PWS. Boston, MA. 2nd edition.
- Spelke, E. S., Breinlinger, K., Macomber, J., and Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, 99:605–632.
- Stiller, J. and Dunbar, R. I. M. (2007). Perspective-taking and memory capacity predict social network size. *Social Networks*, 29(1):93–104.
- Tanji, J. (1996). Involvement of motor areas in the medial frontal cortex of primates in temporal sequencing of multiple movements. In Caminiti, R., Hoffmann, K., Lacquaniti, F., and Altman, J., editors, *Vision and Movement: Mechanisms in the Cerebral Cortex*, volume 2, pages 126–133. Human Frontier Science Program, Strasbourg.
- Taylor, P. D. and Irwin, A. J. (2000). Overlapping generations can promote altruistic behavior. *Evolution*, 54(4):1135–1141.
- Tomasello, M. and Call, J. (1997). *Primate Cognition*. Oxford University Press.
- Topál, J., Byrne, R. W., Ádám Misklósi, and Csányi, V. (2006). Reproducing human actions and action sequences: “Do as I do!” in a dog. *Animal Cognition*, 9(4):355–367.
- Vanechoutte, M. and Skoyles, J. (1998). The memetic origin of language: Modern humans as musical primates. *Journal of Memetics — Evolutionary Models of Information Transmission*, 2(2).
- de Waal, F. B. M. and Johanowicz, D. L. (1993). Modification of reconciliation behavior through social experience: An experiment with two macaque species. *Child Development*, 64:897–908.
- Whiten, A. (1998). Imitation of the sequential structure of actions by chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, 112:270–281.
- Whiten, A. (2006). The dissection of imitation and its ‘cognitive kin’ in comparative and developmental psychology. In Rogers, S. and Williams, J. H. G., editors, *Imitation and the development of the social mind: Lessons from typical development and autism*, pages 227–250. Guilford Press, New York.
- Whiten, A., Flynn, E., Brown, K., and Lee, T. (2006). Imitation of hierarchical action structure by young children. *Developmental Science*, 9(6):574–582.
- Whiten, A., Goodall, J., McGew, W. C., Nishida, T., Reynolds, V., Sugiyama, Y., Tutin, C. E. G., Wrangham, R. W., and Boesch, C. (1999). Cultures in chimpanzees. *Nature*, 399:682–685.

- Whiten, A. and Ham, R. (1992). On the nature and evolution of imitation in the animal kingdom: Reappraisal of a century of research. *Advances in the Study of Behaviour*, 21:239–83.
- Winslow, B. B., Takimoto-Kimura, R., and Burke, A. C. (2007). Global patterning of the vertebrate mesoderm. *Developmental Dynamics*, 236(9):2371–81.
- Wolpert, D. H. (1996). The lack of *a priori* distinctions between learning algorithms. *Neural Computation*, 8(7):1341–1390.
- Wood, M. A. (2008). *An Agent-Independent Task Learning Framework*. PhD thesis, Department of Computer Science, University of Bath, United Kingdom.
- Wood, M. A. and Bryson, J. J. (2007). Skill acquisition through program-level imitation in a real-time domain. *IEEE Transactions on Systems, Man and Cybernetics Part B — Cybernetics*, 37(2):272–285.
- Wood, M. A., Leong, J. C. S., and Bryson, J. J. (2004). ACT-R is *almost* a model of primate task learning: Experiments in modelling transitive inference. In *The 26th Annual Meeting of the Cognitive Science Society (CogSci 2004)*, pages 1470–1475, Chicago. Lawrence Erlbaum Associates.
- Wray, A. (1998). Protolanguage as a holistic system for social interaction. *Language and Communication*, 18(1):47–67.
- Zentall, T. R. (2001). Imitation in animals: Evidence, function, and mechanisms. *Cybernetics & Systems*, 32(1):53–96.

About the author

Joanna Bryson holds a first degree in Behavioural Science from Chicago, Masters degrees in Artificial Intelligence and Psychology from Edinburgh, and a PhD in Computer Science from MIT, where she worked on action selection for cognitive systems in the Artificial Intelligence Laboratory. She has held research positions at the Primate Cognitive Neuroscience Laboratory at Harvard and the Human Communication Research Centre in Edinburgh. Since 2002 she has been at Bath, where she founded Artificial Models of Natural Intelligence. She is currently on sabbatical studying the biological evolution of cultural evolution as a resident fellow of the Konrad Lorenz Institute for Evolution and Cognition Research in Altenberg, Austria.

Author's address

J. Bryson
Artificial models of natural Intelligence
University of Bath
Bath BA2 7AY, United Kingdom
Tel.: +44-1225-383934
Fax: +44-1225-383493
J.J.Bryson@bath.ac.uk

Present address

Konrad Lorenz Institute for Evolution and Cognition Research
Adolf Lorenz Gasse 2
A-3422, Altenberg, AUSTRIA