Acting on the world: understanding how agents use information to guide their action

Jackie Chappell

School of Biosciences, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK

Abstract

How do animals work out which parts of their environment are the most important or interesting to them, and gather information on those parts to guide their action later? In this essay, I briefly outline what we already know about how animals gather and represent information about the world. I then discuss a few of the unsolved problems relating to how animals collect information, before suggesting some approaches which might be useful in unravelling these problems.

Keywords: animal cognition, exploration, behavioural flexibility, evolution, development

1. Introduction

"I am only interested in everything." — Les Murray, poet

In many respects, writing a reflective essay on a topic within Aaron Sloman's research interests is an easy task: after all, he has so many interests! However, the real difficulty lies in narrowing the selection down to focus on a particular part. Appropriately, this problem of selection or attention is analogous to that faced by animals living in complex environments. They need to work out which aspects of their complex environments are the most interesting or important[1], store and structure information about the important parts, and then use that information to guide their actions. This is the problem I will discuss in this essay.

Here I will attempt briefly to outline the current state of the art in animal cognition, specifically referring to how animals gather and represent information

---

Email address: `j.m.chappell@bham.ac.uk` (Jackie Chappell)

[1] Here, I assume that the importance of something to the animal is governed by whether or not it affects the animal's evolutionary fitness.

about the world in order to take action on it. Throughout I will use the informal meaning of information as referring to semantic content that is *about* something that actually exists or could exist (Sloman, 2011). Since this is such a vast topic, I will only be able to touch briefly on a few of the most active areas of research. I will then outline a few of the many unsolved problems, followed by some suggestions about the most promising approaches for investigating these problems.

I refer in the title to ‘agents’ because most of what I say could apply equally to artificial agents that have to negotiate and explore complex environments and yet still carry out the actions and tasks that they were programmed to complete. Indeed, as will become clear throughout this essay, I think that there are benefits in considering the problem from the perspective of the challenges posed by the environment (which is a more familiar perspective for AI researchers), rather than thinking only about the capabilities of the animal. However, since my own expertise is in animal cognition and not AI, my examples will be drawn from the animal literature, including work on humans.

2. What we know about animal cognition

When we observe animals, what often impresses us most is how easily they adapt their behaviour appropriately to rapidly changing contexts, how they deal with complex environments, and how they solve problems. For example, imagine a Sumatran orangutan (*Pongo abelii*) moving around in the canopy to forage. Orangutans are the largest arboreal animals, and as such, they face severe problems in navigating the discontinuous rainforest canopy. They need to cross gaps between trees, but to do so, they have to pass through the zone of thin twigs and branches at the periphery of the tree’s canopy, rather than paying the high energetic cost of descending to the ground, crossing and ascending the next tree (Thorpe et al., 2007). These twigs are insufficient to support their weight (males weigh up to 100kg and females up to 40kg), so they need to find a way of traversing the gap safely. Orangutans are able to use their body mass to deform compliant branches in order to bridge gaps, including swaying compliant trunks of trees like an inverted pendulum (Thorpe and Crompton, 2006, 2005).

Many questions are raised by this kind of behaviour: do orangutans follow habitual routes, or do they plan their route two or more steps in advance to avoid moving into a tree from which it is difficult to cross to another without backtracking? How do they decide what would make a suitable support for a particular crossing? Many different variables are important in determining whether a particular support is suitable, including the size of the gap, the nature of the supports on the far side of the gap, the orangutan’s body weight and limb span, the type of locomotion it chooses to use, the type of tree or liana used and its diameter. How does the orangutan collect, represent, organise and then use these kinds of information to guide its action? These are real problems which have an important effect on the orangutan's evolutionary fitness: it needs to move through the canopy to feed, and if it falls from a height it is likely to be seriously injured.

In the remainder of this section, I will outline what we currently know about these kinds of processes in animals in a selection of domains, before moving on to discuss the areas in which we still lack understanding in Section 3.

2.1. Behavioural flexibility

One of the features that makes the behaviour of some animals so impressive is the behavioural flexibility they display. While some animals show stereotyped sequences of behaviours irrespective of context, most animals — even some invertebrates — are able to alter their actions depending on the context, or if a sequence of actions is interrupted. In all but the simplest of environments, this kind of capability ensures that the animal is able to deal with environments that are highly variable, and that may alter dynamically during execution of behaviours.

Web construction by spiders appears to be a behaviour under fairly tight genetic control. The sequence of actions required in order to build a web are highly constrained, and vary between species. Most orb-weaving spiders show variation in their behaviours while they are attaching the 'proto-hub' to the surrounding supports in the environment. However, once they have built this scaffolding, they can move on their own threads exclusively, and as a consequence the sequence of behaviours used to construct the frame, radii and add the sticky spiral is highly stereotyped (Benjamin and Zschokke, 2004). In contrast two species of spiders in the genus *Linyphia* that construct sheet webs show highly variable web construction sequences, frequently switching between building the supporting structure and laying down sticky threads (Benjamin and Zschokke, 2004). Furthermore, many species are able to adjust the parameters of their web designs depending upon the prevailing conditions. For example, the spacing between the spirals of orb webs is crucial in determining the size of prey that can be trapped (Sandoval, 1994), and spiders can alter the location, orientation and strength of their web depending upon the direction and strength of the wind or other environmental factors.

Male satin bower birds (*Ptilonorhynchus violaceus*) construct elaborate 'bowers' to attract females. The bowers are consist of two parallel walls of sticks planted in the ground and graduated in height, and are decorated with objects of a particular colour (blue in the case of satin bower birds). Females choose males partially on the basis of their bower, and the symmetry of its walls. It has recently been shown that if one wall is experimentally destroyed (something which apparently happens vary rarely in nature), males preferentially replace the sticks in the destroyed wall, not the intact one (Keagy et al., 2011). Furthermore, they show a ‘templating’ behaviour in which they hold a stick vertically in their bill, moving until they match the length of the held stick with its counterpart in the intact wall. They then turn to the same position on opposite side and place the stick in the corresponding position in the destroyed wall (Keagy et al., 2011). This ability to re-order or re-organise actions depending on the current situation allows the animal to begin to ‘debug’ problems which occur during execution. Without this ability, animals (and robots) can become trapped in a loop when a condition which would allow them to proceed to the next action is not present.

Another aspect of flexibility is the ability to group features in the environment that share some property or functional aspect, and respond to features in the that group in the same way. For example, animals are able to group items into ‘food’ and ‘non-food’ categories, and assign novel items to the correct groupings by using shared characteristics. Furthermore, many animals can use relational rules in order to learn, rather than having to learn specific pairings of stimuli. Bees can learn to choose the stimulus that is above a reference mark on a given trial, even when the stimuli themselves are novel and change on each trial (Avarguès-Weber et al., 2011; Chittka and Jensen, 2011). Similarly, a grey parrot (*Psittacus erithacus*) who learned vocal labels for many properties of objects (colour, shape, material, number etc.) was able to answer questions about which property was shared or differed among a set of presented objects (Pepperberg, 1987).

2.2. Physical cognition

The physical world provides another source of complexity and variability. The animal encounters objects, materials and surfaces which differ markedly in their properties and the affordances (Gibson, 1977) they provide. In some circumstances, when the variability is low in the features of the environment that affect the animal’s fitness, mechanisms of associative learning may suffice. That is, the animal may learn to associate the perceptual qualities of a particular object with a particular affordance. In AI terms, this would be equivalent to only recognising that a cup can be grasped if the cup is the same size, shape and colour as the one used during training[2]. However, in more complex situations where the number of configurations of variables relevant to the animal is high (perhaps like the situation described for the orangutan at the beginning of this section), there

---

[2] Although animals can generalize somewhat between learned stimuli.

is a potential for combinatorial explosion (Perlovsky, 1998; Bellman, 1961) and other forms of learning and cognition will be required.

Since the physical world has certain regularities (the principles of continuity, connectedness, solidity and gravity are a few examples), it is possible that animals are born with basic knowledge about these predictable properties, but later add to and modify these rules after experience with the complexity of the world. This appears to be the case for human children, who have certain expectations about the domain of the physical world (as well as number and agency among others; see Spelke and Kinzler 2007). There has been much less research in this area in non-human animals, but recent research suggests that at least apes (Cacchione et al., 2009; Cacchione and Call, 2010) and domestic dogs (Kundey et al., 2009) may have basic core concepts in the area of solidity.

Physical cognition also involves understanding the causal properties of objects, and the effects that the individual's own actions will have on the world. In many cases this also requires an understanding of the physical properties of objects (for example, whether an object is rigid enough to act as a tool, or whether an object's shape will allow it to roll along a surface), but it extends to understanding the causal role of the actor and each of the objects involved in achieving the desired outcome. This has been extensively investigated in the domain of tool use, but tool use is just one example of a suite of behaviours involving the kinds of information processing described here. We know that New Caledonian crows (*Corvus moneduloides*) — a tool manufacturing species — can select tools with appropriate properties depending on the context of the task (Chappell and Kacelnik, 2002, 2004; Bluff et al., 2010), and can even manufacture an appropriate tool (a hook) using a novel material (metal wire) (Weir et al., 2002). Rooks (*Corvus frugilegus*) — a non-tool using species in the wild — were also able to spontaneously modify and use materials as tools, as well as solving other problems involving physical cognition (Bird and Emery, 2009b,a). Thus, the kinds of cognitive abilities required to use and manufacture tools do not appear to be adaptive specialisations of tool use itself. However, using tools seems to impose an additional cognitive load on some species, making physical cognition problems more difficult to solve. When chimpanzees were tested on a trap tube problem (in which the subject has to remove a food item from the apparatus while avoiding moving it over a trap), with and without the requirement to use a tool, the chimpanzees performed significantly better when no tool was required (Seed et al., 2009).

One issue which frequently arises in the literature is whether any non-human animals can solve problems involving "unobservable causes", defined as those "based on structural or functional relationships between objects rather than on perceptually based exemplars" (Penn and Povinelli, 2007, p. 107). Gravity is the unobservable cause most frequently discussed, but support is another. If we return to the orangutan example mentioned earlier, the compliance (or flexibility) of supports is not apparent to the animal until it applies its weight to the support, observes another individual doing so, or perhaps observes the effect of the wind on the trees. Diameter of the support might give the orangutan a rough estimate of its likely compliance, but this is likely to be complicated by the material (supports formed from tree branches and lianas will differ in compliance for a given diameter), by the length of the support and whether it is connected to other supports in a mesh of branches. In these circumstances it is possible that animals have certain 'expectations' about the causal effects of their own actions, which they may 'test' by performing actions, much as human children have been shown to do (e.g. Cook et al., 2011).

2.3. Planning

In highly complex environments with a great deal of variability, there may be many potential actions to take in order to achieve a goal, not all of which will be equally effective. In these situations, planning might help the animal to select among the potential actions. Planning is an immensely complicated topic, and in the field of animal cognition, we are only just starting to investigate how (or indeed, *whether*) non-human animals might plan. In the simplest case, the animal might select one action or a sequence of actions in a pre-determined order, which will immediately result in obtaining the goal. There is a decision among options to be made, but only one step is involved (many people would argue about whether this should really be called planning at all: see Sloman 1999 for a discussion). More complex kinds of planning might involve situations in which each option, if chosen, leads to a selection of further options, and so on, along a branching decision tree. Thus, after even a few steps, the number of possible combinations explodes, making it very difficult to exhaustively evaluate all the options.

Planning for most animals (if it occurs at all) probably falls somewhere between those two extremes. One important issue is how animals select among the alternatives, particularly when there are multiple steps, each with their own consequences. Do animals mentally simulate the consequences of alternative actions? That would imply that they have some kind of internal model of the world (either stored from previous experience, or extrapolated on the basis of the kinds of understanding outlined in Section 2.2) with which to generate such simulations. The problem for behavioural biologists is that this activity mostly internal, so we are faced with difficulties in determining whether it is occurring at all.

One experimental strategy is to pose a problem and allow subjects a period in which they can observe, but not interact with, the problem. The rationale is that if subjects formulate a plan for solving the problem during the preview period using mental simulation, they should be faster at completing the task than subjects denied a preview period. Dunbar et al. (2005) employed this design, presenting four types of puzzle box (which differed in difficulty) to chimpanzees, orangutans and human children. Subjects were presented with the puzzle boxes to open in order to obtain a reward, and were either allowed to view the boxes before attempting to open them ('prior view condition') or not ('no prior view condition'). In all species, there was an effect of the prior view condition such that subjects solved the task faster when given a prior view. However, there are problems with this approach: differences in motivation or behavioural style may affect the time taken to solve the task, but have no implications for mental simulation or its absence.

Another approach is to examine the choices made by subjects at each step, while altering the difficulty of the task, and including some options that are superficially attractive, but eventually lead to failure when chosen over less attractive, but ultimately successful options. For example Fragaszy et al. (2003) tested capuchin monkeys and chimpanzees on a 2D computerized maze task, in which subjects had to move a cursor through a maze towards a goal. Both species were more successful at completing the mazes than would be expected by chance, suggesting that they could 'look ahead' and plan their path appropriately. However, capuchins were less successful when they had to take a path that initially lead away from the goal direction (Fragaszy et al., 2009). This experimental approach is a useful one, but has the disadvantage that it requires extensive training of subjects to use the apparatus. In an attempt to design a task probing similar processes, but which did not require training or the use of tools, we tested orangutans on a horizontal puzzle tube, in which two sets of obstacles on each side of a centrally-placed reward could hinder the subject retrieving the food (Tecwyn et al., 2011). Subjects therefore had to look ahead to determine which of the two alternative paths would allow them to obtain the food. Two of the three subjects solved the task, but there were interesting differences in strategy between them, suggesting the use of different combinations of rules in order to solve the task.

3. Unsolved problems in animal cognition

I have presented the research in Section 2 as the ‘state of the art’, but in reality we are still far from having a full picture of how animals collect, store, represent and use information about the world. When working with fully working organisms (in contrast to robots or simulations that you have to construct from scratch), it is easy to take certain aspects of their behaviour for granted. In particular, the more we probe in certain areas, the less we realise we understand about the mechanisms involved. In this section, I will summarize a few of the areas where our knowledge is lacking, focussing particularly on the ways in which animals gather information.

3.1. Perception

In order to collect information about the world, animals first have to perceive it: that is, they need to receive and translate signals from the environment, then store them as neural representations. It may seem a relatively trivial task to work out what the perceptual content is for animals, particularly as we understand the mechanisms of many animal sensory systems quite well. However, we encounter problems when studying species that have rather different sensory systems from our own. To give one example, many species of birds have laterally-placed eyes, in contrast to our own which are frontally-placed. Thus it is difficult to determine what the bird is attending to without having detailed knowledge about the arrangement and extent of both the monocular and binocular visual fields. Since the arrangement of birds’ visual fields are dependent upon their ecological niche (Martin, 2007), these can differ substantially between species, and their arrangement and coverage can often be surprising.

For example, we have recently shown that while Senegal parrots (*Poicephalus senegalus*) have a relatively broad frontal binocular field and good visual coverage above and somewhat behind the head, they cannot see below their bill tip (Demery et al., 2011). This means that they cannot see what is held in their bill, which has important implications for understanding what information they can collect from their environment during exploration (see Sections 3.2 and 3.3).

3.2. Attention

Animals have limited capacity for processing information at any one time, which tends to be less than the rate at which the environment provides information. Thus, they have to direct their attention selectively. While we know a great deal about the physiological and neurobiological mechanisms of attention (particularly visual attention) in humans (Desimone and Duncan, 1995), most of the research in non-human animals has been conducted on other primates or rats. However, it is clear that attention has a major role in shaping the information processed by animals in important areas of their lives (Dukas, 2002). For example, when preying on simulated cryptic (camouflaged) prey, located in their central visual field, blue jays were significantly less likely to detect a stimulus in their peripheral visual field (simulating detection of a predator) when the prey detection task was more difficult (Dukas and Kamil, 2000), even though their success rate for detecting prey did not differ between the easy and difficult task.

Attention therefore has both perceptual and cognitive aspects: the sense organs need to be directed towards a stimulus, but processing resources also need to be focussed on particular parts of the environment, effectively filtering it. Since attentional mechanisms filter and restrict the information that is available for an animal to learn, attention can act as a kind of feedback loop: learning and genetic pre-dispositions can influence what the animal attends to, and what the animal learns is constrained by what it attends to. It is therefore very important for us to attempt to understand attentional mechanisms in animals to get a fuller picture of their whole information processing system.

How can behavioural biologists go about studying this? In some cases there may be ‘behavioural markers’ of attention for sensory modalities such as vision (but see discussion in Section 3.1), but in others, attention may be difficult or impossible to observe in naturally-behaving animals. We may need assistance from physiologists and neuroscientists to identify such behavioural markers, or we could use objects equipped with instruments in order to detect, for example, the location of touch and force generated by the animal during haptic exploration of an object. Collaborations with computer scientists may also help us to model the putative effects of attention and thus identify markers of attention in animals.

3.3. Exploration

Exploration[3] has been studied extensively in non-human animals in relation to navigation and spatial behaviour, but there has been less research on more general aspects of exploration and play (Held and Špinka, 2011; Power, 2000; Chappell et al., In Review). As with attention (see Section 3.2), we need behavioural markers of exploration, which can be difficult to define. In an excellent collaboration between biologists and roboticists, Grant et al. (2009) made detailed behavioural observations of the way in which rats used their vibrissae to explore their environment, and showed that they alter the speed and pattern

---

[3] By which I mean interacting with the environment without an immediate goal other than the collection of information

of their ‘whisking’ behaviour in order to learn about the shape of objects and their surface texture. They then confirmed that the pattern of this behaviour optimizes the efficiency with which this information is collected by modelling the behaviour in a robot (Pearson et al., 2007).

Since information collected during exploration provides some of the raw material for animals’ information processing, we need to study the process in more detail (Chappell et al., In Review). For example, do animals start out with perceptual or attentional biases towards exploring certain kinds of objects or collecting certain kinds of information? This could both limit and shape the kinds of information processing they are later capable of (as discussed in Section 3.2). Is exploration directed primarily towards novel stimuli, in order to collect new information? Do any animals ‘test’ the properties of objects when their expectations are violated, or re-initiate further bouts of exploration, in the way that we know happens in human children (Cook et al., 2011)?

4. Where do we go from here?

In the preceding sections I have attempted to provide a brief summary of the state of our knowledge about animal cognition as it relates to the physical domain, and have highlighted areas where our knowledge is currently lacking for various reasons. This is, of course, a huge and complex topic, and since most of the interesting stuff goes on ‘invisibly’ inside the animal’s brain, it can be a challenge to investigate. What, therefore, is the best way to proceed?

We still have a long way to go before we have established what mechanisms are involved in the kinds of processes I have discussed above. For example, despite some elegant experiments we have little idea for many species about the different kinds of processes (in the world) that the animal can distinguish between. We also know little about how animals organise the information they collect, whether they form abstractions of some kind, and how they reuse information at a later time. Well-designed experiments can probe some of these issues by examining the details of the ways in which animals choose between different options, how they approach problems and so on. We can also present animals with objects which have unusual properties or behave deceptively, and use animals’ initial actions and subsequent reactions to those objects to investigate whether they have prior expectations, using techniques similar to those used with pre-verbal human infants.

However, this kind of experimental approach would be much more fruitful if combined with a requirements analysis (Chappell and Thorpe, 2010; Sloman, 2005). If we understand what the requirements of the environment are (requirements will differ substantially between species with differing niches), we are in a better position to focus our own attention on the relevant questions, and avoid overlooking important possibilities because the behaviour of the animal is too subtle to be noticed. This brings me back again to the importance of understanding the sensory and attentional world of the animal being studied: as a strongly visual species, it is easy for us to overlook information which might be unavailable or barely perceptible to our own senses.

We also need to examine the precise behaviours of individual animals in much more detail, paying particular attention to the errors that they make, as these are likely to be revealing of the operation of their information processing systems, much as bugs in software can reveal important features in the structure of the underlying code. If I have learned anything while studying animal cognition, it is that there is enormous variation between individuals. Thus, while the mean performance of subjects is interesting in some respects, in the face of high variation within a population it can be relatively uninformative about the processes involved. The genetic and developmental histories of individuals (Chappell and Sloman, 2007), as well as their experiences throughout life, can have a significant effect on the range and extent of their capabilities (Sloman and Chappell, 2007). It is a complex knot to unpick, but attempting to understand how and why individuals differ in their cognitive capabilities might lead us to important insights about the processes and mechanisms involved.

Finally, it is my opinion that we can gain enormous advantages from working with roboticists and other computer scientists. AI can inspire biology as effectively as the reverse position (Chappell and Thorpe, 2010): having to construct working robots reveals gaps in knowledge and flaws in assumptions quickly, and robots and simulations can allow you to test hypotheses which would be impossible (for practical or ethical reasons) to test in animals (Grant et al., 2009; Webb, 2000). Even if you do not build a working robot or simulation, the exercise of thinking how one might go about it can provide a valuable ‘tool for thinking’.

These are deep and highly complex problems, and the more closely one studies the problem, the more detail and complexity is revealed, like zooming into an image of a fractal, but without the conceptual benefit of knowing that an equation describes the structure. It is certainly true that the more closely researchers have looked at the capabilities of animals, the more advanced capabilities they have found. The bridge between ‘uniquely human’ abilities and those of other animals is being gradually eroded. We cannot solve these problems as isolated disciplines: I hope that we can pool our expertise and continue to be “only interested in everything".

References

Avarguès-Weber, A., Dyer, A. G., Giurfa, M., 2011. Conceptualization of above and below relationships by an insect. Proceedings of the Royal Society B: Biological Sciences 278 (1707), 898–905.

Bellman, R., Mar. 1961. Adaptive control processes: a guided tour. Princeton University Press, Princeton, NJ.

Benjamin, S. P., Zschokke, S., 2004. Homology, behaviour and spider webs: web construction behaviour of *Linyphia hortensis* and *L. triangularis* (Araneae: Linyphiidae) and its evolutionary significance. Journal of Evolutionary Biology 17 (1), 120–130.

Bird, C. D., Emery, N. J., 2009a. Insightful problem solving and creative tool modification by captive nontool-using rooks. Proceedings of the National Academy of Sciences of the United States of America 106 (25), 10370–10375.

Bird, C. D., Emery, N. J., 2009b. Rooks Use Stones to Raise the Water Level to Reach a Floating Worm. Current Biology 19, 1410–1414.

Bluff, L. A., Troscianko, J., Weir, A. A. S., Kacelnik, A., Rutz, C., 2010. Tool use by wild New Caledonian crows *Corvus moneduloides* at natural foraging sites. Proceedings of the Royal Society B: Biological Sciences 277 (1686), 1377–1385.

Cacchione, T., Call, J., 2010. Intuitions about gravity and solidity in great apes: the tubes task. Developmental Science 13 (2), 320–330.

Cacchione, T., Call, J., Zingg, R., 2009. Gravity and solidity in four great ape species (*Gorilla gorilla*, *Pongo pygmaeus*, *Pan troglodytes*, *Pan paniscus*): Vertical and horizontal variations of the table task. Journal of Comparative Psychology 123 (2), 168–180.

Chappell, J., Demery, Z. P., Arriola-Rios, V., Sloman, A., In Review. How to build an information gathering and processing system: lessons from naturally and artificially intelligent systems. Behavioural Processes.

Chappell, J., Kacelnik, A., 2002. Tool selectivity in a non-primate, the New Caledonian crow (*Corvus moneduloides*). Animal Cognition 5 (2), 71–78.

Chappell, J., Kacelnik, A., 2004. Selection of tool diameter by New Caledonian crows *Corvus moneduloides*. Animal Cognition 7 (2), 121–127.

Chappell, J., Sloman, A., 2007. Natural and Artificial Meta-Configured Altricial Information-Processing Systems. International Journal of Unconventional Computing 3 (3), 211–239.

Chappell, J., Thorpe, S., 2010. AI-Inspired Biology: Does AI Have Something to Contribute to Biology? Proceedings of the International Symposium on AI Inspired Biology: A Symposium at the AISB 2010 Convention, Leicester, UK. SSAISB, sp.

Chittka, L., Jensen, K., 2011. Animal cognition: concepts from apes to bees. Current Biology 21 (3), R116–9.

Cook, C., Goodman, N. D., Schulz, L. E., 2011. Where science starts: Spontaneous experiments in preschoolers' exploratory play. Cognition 120 (3), 341–349.

Demery, Z. P., Chappell, J., Martin, G. R., Apr. 2011. Vision, touch and object manipulation in Senegal parrots *Poicephalus senegalus*. Proceedings of the Royal Society B: Biological Sciences.

Desimone, R., Duncan, J., 1995. Neural mechanisms of selective visual attention. Annual Review of Neuroscience 18, 193–222.

Dukas, R., 2002. Behavioural and ecological consequences of limited attention. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences 357 (1427), 1539–1547.

Dukas, R., Kamil, A. C., 2000. The cost of limited attention in blue jays. Behavioral Ecology 11 (5), 502–506.

Dunbar, R. I. M., McAdam, M. R., O'connell, S., 2005. Mental rehearsal in great apes (*Pan troglodytes* and *Pongo pygmaeus*) and children. Behavioural Processes 69 (3), 323–330.

Fragaszy, D., Johnson-Pynn, J., Hirsh, E., Brakke, K., 2003. Strategic navigation of two-dimensional alley mazes: comparing capuchin monkeys and chimpanzees. Animal Cognition 6 (3), 149–160.

Fragaszy, D. M., Kennedy, E., Murnane, A., Menzel, C., Brewer, G., Johnson-Pynn, J., Hopkins, W., 2009. Navigating two-dimensional mazes: chimpanzees (*Pan troglodytes*) and capuchins (*Cebus apella sp.*) profit from experience differently. Animal Cognition 12 (3), 491–504.

Gibson, J. J., 1977. The theory of affordances. In: The ecological approach to visual perception. Lawrence Erlbaum Associates, London, pp. 127–143.

Grant, R. A., Mitchinson, B., Fox, C. W., Prescott, T. J., 2009. Active touch sensing in the rat: anticipatory and regulatory control of whisker movements during surface exploration. Journal of Neurophysiology 101 (2), 862–874.

Held, S. D., Špinka, M., 2011. Animal play and animal welfare. Animal Behaviour 81 (5), 891–899.

Keagy, J., Savard, J.-F., Borgia, G., 2011. Complex relationship between multiple measures of cognitive ability and male mating success in satin bowerbirds, *Ptilonorhynchus violaceus*. Animal Behaviour 81 (5), 1063–1070.

Kundey, S. M. A., de Los Reyes, A., Taglang, C., Baruch, A., German, R., 2009. Domesticated dogs' (*Canis familiaris*) use of the solidity principle. Animal Cognition 13 (3), 497–505.

Martin, G. R., 2007. Visual fields and their functions in birds. Journal of Ornithology 148 (S2), 547–562.

Pearson, M. J., Pipe, A. G., Melhuish, C., Mitchinson, B., Prescott, T. J., 2007. Whiskerbot: A Robotic Active Touch System Modeled on the Rat Whisker Sensory System. Adaptive Behavior 15 (3), 223–240.

Penn, D. C., Povinelli, D. J., 2007. Causal cognition in human and nonhuman animals: a comparative, critical review. Annual Review of Psychology 58, 97–118.

Pepperberg, I. M., 1987. Acquisition of the same different concept by an African gray parrot (*Psittacus erithacus*) - learning with respect to categories of color, shape, and material. Learning & Behavior 15 (4), 423–432.

Perlovsky, L., 1998. Conundrum of Combinatorial Complexity. IEEE Transactions on Pattern Analysis and Machine Intelligence 20 (6), 666–670.

Power, T. G., 2000. Play and Exploration in Children and Animals. Lawrence Erlbaum Associates, London.

Sandoval, C., 1994. Plasticity in Web Design in the Spider *Parawixia bistriata*: A Response to Variable Prey Type. Functional Ecology 8 (6), 701–707.

Seed, A. M., Call, J., Emery, N. J., Clayton, N. S., 2009. Chimpanzees solve the trap problem when the confound of tool-use is removed. Journal of Experimental Psychology: Animal Behavior Processes 35 (1), 23–34.

Sloman, A., 1999. What sort of architecture is required for a human-like agent. In: Wooldridge, M., Rao, A. (Eds.), Foundations of Rational Agency. Kluwer Academic Publishers, Dordrecht, pp. 35–52.

Sloman, A., 2005. The design-based approach to the study of mind (in humans, other animals, and machines), including the study of behaviour involving mental processes. `http://cs.bham.ac.uk/research/projects/cogaff/misc/design-based-approach.html`.

Sloman, A., 2011. What's information, for an organism or intelligent machine? How can a machine or organism mean? In: Dodig-Crnkovic, G., Burgin, M. (Eds.), Information and Computation. World Scientific, New Jersey, pp. 393–438.

Sloman, A., Chappell, J., 2007. Computational cognitive epigenetics. Behavioral and Brain Sciences 30 (4), 375–376.

Spelke, E. S., Kinzler, K. D., 2007. Core knowledge. Developmental Science 10 (1), 89–96.

Tecwyn, E. C., Thorpe, S. K. S., Chappell, J., Jul. 2011. What cognitive strategies do orangutans (*Pongo pygmaeus*) use to solve a trial-unique puzzle-tube task incorporating multiple obstacles? Animal Cognition.

Thorpe, S. K. S., Crompton, R. H., 2005. Locomotor ecology of wild orang-utans (*Pongo pygmaeus abelii*) in the Gunung Leuser ecosystem, Sumatra, Indonesia: A multivariate analysis using log-linear modelling. American Journal of Physical Anthropology 127, 58–78.

Thorpe, S. K. S., Crompton, R. H., 2006. Orangutan Positional Behavior and the Nature of Arboreal Locomotion in Hominoidea. American Journal of Physical Anthropology 131, 384–401.

Thorpe, S. K. S., Crompton, R. H., Alexander, R. M., 2007. Orangutans utilise compliant branches to lower the energetic cost of locomotion. Biology Letters 3, 253–256.

Webb, B., 2000. What does robotics offer animal behaviour? Animal Behaviour 60 (5), 545–558.

Weir, A. A. S., Chappell, J., Kacelnik, A., 2002. Shaping of hooks in New Caledonian crows. Science 297 (5583), 981.