

How to build an information gathering and processing system: Lessons from naturally and artificially intelligent systems

Chappell, Jackie; Demery, ZP; Arriola-Rios, V; Sloman, Aaron

DOI:

[10.1016/j.beproc.2011.10.001](https://doi.org/10.1016/j.beproc.2011.10.001)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Chappell, J, Demery, ZP, Arriola-Rios, V & Sloman, A 2012, 'How to build an information gathering and processing system: Lessons from naturally and artificially intelligent systems', *Behavioural Processes*, vol. 89, no. 2, pp. 179-186. <https://doi.org/10.1016/j.beproc.2011.10.001>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

NOTICE: this is the author's version of a work that was accepted for publication in *Behavioural Processes*. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in *Behavioural Processes*, 89(20): 179---186. DOI: <http://dx.doi.org/10.1016/j.beproc.2011.10.001>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

How to build an information gathering and processing system: lessons from naturally and artificially intelligent systems

Jackie Chappell^{a*}, Zoe P. Demery^a, Veronica Arriola-Rios^b, Aaron Sloman^b 2012.
Behavioural Processes, **89**(20): 179-186.

DOI: <http://dx.doi.org/10.1016/j.beproc.2011.10.001>

* To whom correspondence should be addressed

Email: j.m.chappell@bham.ac.uk

^a School of Biosciences, University of Birmingham, Edgbaston, Birmingham, UK. B15 2TT j.m.chappell@bham.ac.uk, zxd878@bham.ac.uk

^b School of Computer Science, University of Birmingham, Edgbaston, Birmingham, UK. A.Sloman@cs.bham.ac.uk, V.Arriola-Rios@cs.bham.ac.uk

Abstract

Imagine a situation in which you had to design a physical agent that could collect information from its environment, then store and process that information to help it respond appropriately to novel situations. What kinds of information should it attend to? How should the information be represented so as to allow efficient use and re-use? What kinds of constraints and trade-offs would there be? There are no unique answers. In this paper, we discuss some of the ways in which the need to be able to address problems of varying kinds and complexity can be met by different information processing systems. We also discuss different ways in which relevant information can be obtained, and how different kinds of information can be processed and used, by both biological organisms and artificial agents. We analyse several constraints and design features, and show how they relate both to biological organisms, and to lessons that can be learned from building artificial systems. Our standpoint overlaps with Karmiloff-Smith (1992) in that we assume that a collection of mechanisms geared to learning and developing in biological environments are available in forms that constrain, but do not determine, what can or will be learnt by individuals.

Keywords: exploration; play; information processing; learning;

1 Introduction

For what purposes do animals need to acquire and use information? At one extreme organisms are merely *acted on* by the environment, which provides them with nutrients and toxins, and subjects them to various forces and abiotic conditions (such as temperature, pressure, and humidity). Such organisms can benefit or suffer as a result, but their ability to alter the effect of these conditions is severely limited. At another extreme, organisms can act on the environment, by planning and performing actions of varying kinds and degrees of complexity to avoid future harm or achieve future gain. In these circumstances, selection of behaviour requires the use of information. There are enormously diverse types of information and types of information processing capabilities.

In this paper, we discuss some examples in the middle of this range of organisms, pose some new questions and suggest some new examples of information processing capabilities. We also illustrate how methods, problems, concepts and theories from Artificial Intelligence (AI) can help biologists and psychologists make progress, advancing work already done by other authors (e.g. Gibson, 1988; Gibson and Pick, 2000; Karmiloff-Smith, 1992). This is just a tiny region in a huge field of research into the possible requirements and designs of different biological information processing systems. It intersects with existing work on motivation, play, learning, perception and development, and should help us understand how such systems have evolved and how they develop within an individual's lifetime.

1.1 What do we mean by “information”?

The word “information” has two main uses:

- a) Shannon's use (Shannon, 1948) refers to a syntactic measure of communicable signals that ignores what the signals refer to, and
- b) The everyday use (Sloman, 2011a) refers to semantic content that is *about* something that actually exists or could exist¹.

We use “information” (and related words, e.g. “concept”, “symbol”, “meaning”, “content”) in an informal way in sense b), as do most biologists, psychologists and engineers when discussing how organisms or machines (“agents”) can acquire factual information, construct theories, make predictions, draw conclusions, or adopt goals referring to something in the environment or within themselves. Information in sense b) matters to an organism or machine if the use of information can provide some benefit to them.

In each case there must be an “information bearer” (representation) in the individual (e.g. chemical or neural signals), or in the environment (e.g. pheromone trails), or straddling the organism and its environment. New information bearers are often constructed in the process of using old information, such as when reasoning, planning, or forming goals. Information bearers need to be acted upon or used, in order to use the information content.

1.2 Plan for this paper

We present some of the requirements for intelligent agents, natural or artificial, that acquire and use information, showing how modes of thinking from different research fields can inform each other in designing working systems (e.g. robots) and in the search for explanations of learning and development in organisms. In this context it is useful to consider biological evolution and human engineers as playing similar roles, namely producing designs for an immature agent that can learn to cope with new environments, including situations not encountered by evolution (Dennett, 1981; McCarthy, 2008).

2 Theoretical requirements

Imagine that you are given the task of designing a physical agent that will explore the surface of another planet, collecting and sending information as it travels (e.g. Mars Rovers²). It will be exposed to a wide variety of unpredictable conditions, so it cannot be pre-programmed for every possible contingency (see Inglis et al. 2001;

¹ There are special cases of things that do not exist in space-time (numbers, theories etc.), but we will ignore those here.

² <http://marsrover.nasa.gov/mission/>

McNamara and Houston, 2009 for further discussion). It will have to operate autonomously, so it must 'decide' for itself where to go and what to do, as well as self-diagnosing and fixing its own hardware or software problems. What are the general requirements of the task of exploration and learning? We sketch some answers from the standpoint of the designer of such a system who is inspired by biological evidence. These design considerations can, in turn, inspire developments in biological theorising, including suggesting new research questions (see for example Hawes, 2011). Thus, in this section and Section 3, we will illustrate our points with examples from a hypothetical artificial agent (the Mars Rover-type robot mentioned earlier, hereafter 'Rover') and a biological agent, the New Caledonian crow (*Corvus moneduloides*, hereafter 'NCC'). Since in both cases we still lack information about the precise mechanisms or computational processes underlying behaviour, these examples will be purely illustrative examples, in order to clarify a broader theoretical point. We suggest linking the study of animal cognition with robotics research not because there are existing robots whose information processing mechanisms might be offered as explanations of animal competences, but because thinking like a robot designer about mechanisms required in a robot can suggest research questions regarding such mechanisms in animals.

We proceed to specify some types of information that an animal or machine can acquire, some ways in which the information can be used, and some transformations of stored information required to extend the generality and power of the agent's competences, in order to illustrate some of the variety of ways in which biological theory and AI overlap.

2.1 Persistence of information

Some uses of information are transient: the individual acquires and immediately uses the information (e.g. while guiding grasping) without retaining any record of it. For other purposes, persistence of information is required. In this paper, we will mainly discuss types of information acquisition and exploration that result in persisting information content.

If there is some change in the animal or in the external environment as a result of acquiring information (in other words, an "information bearer" is formed or modified as discussed in Section 1.1), and it persists for even a short time, it can enable further uses of the information content. Without such persistence, detection of change is impossible. Persistence is also needed in order to re-use information about enduring objects, properties, locations or processes, such as approaching predators, escaping prey, and so on. This persistent information can be re-used through well-known processes of learning such as classical or operant conditioning, or other probabilistic forms of learning (see Shettleworth, 2010 for an excellent discussion). For example, information related to the availability of food will become associated with the act of feeding such that the presence of this information will increase frequency of food-seeking and feeding behaviours.

Rover would probably need persistent information of many kinds in order to increase its efficiency. For example, it would need to store information about the geographical locations of interesting samples, in order to return to that location later to collect the samples. Likewise, a NCC fishing for Cerambycid larvae in rotten logs might store information about the presence or absence of 'frass' from the larvae on the surface of the log, and re-use that information to identify logs which contain active larvae.

Persistent records can also enable discovery of useful abstractions. This is often referred to as "concept formation" (Martin, 2007; Zentall, 2008). Possession of such a concept allows new questions to be asked, and new goals or generalisations about some portion of the world to be formulated. More complex examples include discovery of useful relations (between, inside, touching: e.g. Bird and Emery, 2010), functions (owner, initiator, container or location of), properties of matter (rigid, compliant, elastic, viscous: Bushnell and Boudreau, 1993) and conditional properties (e.g. a compliant branch supports an orangutan's weight only if it is above a minimal stiffness). In some of these examples, classical models of learning will suffice to explain the observed behaviours, but in others, they are not sufficient (see Section 2.2.2 for further discussion).

2.2 Coping with variability of the environment

2.2.1 Pre-configured competences

However varied the environments, certain competences will have to be hard-wired initially. Evolution provides many fully pre-configured responses for different situations, such as feeding behaviours (e.g. suckling, begging, pecking). When the environment is too variable, evolution (like a human designer) cannot discover in advance suitable fixed responses to all needs in all situations. Instead, it provides mechanisms of learning and development that use information about the environment. This enables individuals to discover useful new actions, threats or opportunities (see Section 2.2.2). However, these discoveries must start from some form of pre-configured competences.

Initial competences should be designed to make use of relevant environmental information, for instance detecting shapes of salient objects, their edges or contours, the curvature and orientation of surface fragments, and their texture. The exact properties will depend partly on what sensory apparatus the agent is equipped with, and on which features of objects or events are relevant to its ecological niche. Some objects' properties are difficult to determine without touching them. For example, if the agent needed to determine the hardness of a material, it could apply pressure to the surface of the object, while estimating weight would require it to lift the object (Flanagan and Wing, 1997; Wing and Lederman, 1998). Information gained in this way does not need to be metrical. For example, Rover might only need

two categories of weight (“light enough to carry” or “too heavy”) to enable decisions to be made about transporting samples. In other cases, orderings may suffice.

Rover might need to be pre-programmed with some basic competences (such as the ability to grasp a rock or to avoid exceeding its safe loading capacity), or it might damage itself irreparably, or fail to start collecting samples. However, it could extend those competences by learning more through exploration (e.g. learning how to identify, classify and select the most relevant rock samples). We know that while it can take several years for NCC to achieve adult-level competency at tool use, juveniles perform stereotyped object manipulation patterns which appear to be precursors of adult tool behaviour (Kenward et al., 2006).

Explaining how such discoveries can be made is very difficult. The methods of robot designers (identifying requirements, then generating, testing and debugging designs to meet those requirements) can usefully complement empirical research into observable behaviours and brain mechanisms (Webb, 2001). For instance, a learning task can be broken down into many sub-tasks. One of the tasks could be acquiring and storing information about the layout of the terrain, and the opportunities and dangers it affords. Various mechanisms for doing this both in animals and in robots have been studied (e.g. using SLAM algorithms for Simultaneous Localisation and Mapping; Bailey and Durrant-Whyte, 2006), but will not be discussed further here.

2.2.2 Reducing complexity by partitioning what has to be learnt

Well-designed learners need to be able to cope with many potential sources of variability and dynamic change in the environment (McNamara and Houston, 2009; Shettleworth, 2010). This includes spatio-temporal changes caused by seasonal or climatic changes, geological changes, new behaviours of intra- or inter-specific competitors, variations in food availability (e.g. Houston et al., 1980; Kacelnik and Krebs, 1985; Kacelnik and Todd, 1992), co-evolutionary arms races between predators and prey, and even niche construction (Sterelny, 2007).

Such learning processes are potentially combinatorially explosive (Bellman, 1961; Perlovsky, 1998) because of the huge search spaces involved in combining, for instance, different perceptual tests, motor sequences, and modulations of generic actions to fit specific shapes, sizes, and relations of objects and the processes they generate. In some circumstances it is possible for the animal to reduce computational load by pursuing a strategy of exploring the environment first and then switching to exploiting it. For example, Krebs et al. (1978) showed that great tits learning to choose between exploiting two foraging posts, each with a fixed probability of reward, stopped exploring and remained at one feeding post after a number of trials close to that predicted by the optimal solution. However, in this case, the environment was relatively simple and did not vary over the course of the experiment: the situation is much more complex in non-stationary environments

and other sources of uncertainty (Cohen et al., 2007). In addition, a learning system needs to have good criteria for selecting things to attend to. Evolution seems to have provided many species with initial motivations and learning mechanisms specialised for restricted classes of environment (e.g. Karmiloff-Smith, 1992). Furthermore, in mechanisms of learning such as classical conditioning, animals attend to and learn about stimuli that convey information (e.g. Rescorla, 1968). The selection of relevant information³ to learn may be reduced further by decomposing an enormously varied environment into a collection of object “affordances” (action possibilities, or what the environment provides for the organism) and processes or “exploration domains” (“microdomains” in Karmiloff-Smith, 1992). We suggest that decomposition is achieved by perceptual and motor interactions with the environment during exploration (see Section 4 and Power, 2000 for an extended discussion).

By selecting a physical subset of the environment, and systematically varying actions performed on it, the agent solves (temporarily) the problem of what to attend to and limits the range of phenomena within which patterns are sought. Some of the actions will involve only the agent’s body-parts, for example moving limbs or digits, controlling eye movements or moving the whole body. Others will also involve selected objects, or types of object, such as repeatedly grabbing and pulling, pushing or twisting the same thing, or rearranging a group of objects. All this will only work if the agent starts off with perceptual mechanisms capable of detecting and recording the structures and motions produced by the exploratory actions (see Section 2.2.1). For example, Rover might have to explore the effects of its own actions when it pushes piles of loose gravel-like material around, separating movement of the material from self-generated movement of its own body, thus discovering how to pile the material in a stable heap.

Our anecdotal observations of animals and children suggest that exploration domains are often inter-leaved, for example alternating between eating and playing with food. This allows knowledge of different domains to be developed roughly in parallel (Bushnell and Boudreau, 1993). However, when switching domains, the individual needs to be able to group items of information together according to the current domain involved. For example, for a tool using species like NCC, materials such as twigs and grass stems have one kind of affordance in the tool using domain (inserting into holes to retrieve food items) and other kinds of affordance when building a nest. For different species, the objects in the environment and their affordances will differ according to their ecological niche, but there seem to be some common exploration mechanisms across species (Chappell and Sloman, 2007; Sloman and Chappell, 2005; see Section 4).

³ Here, relevant information is that which – if acted on – will influence the animal’s evolutionary fitness.

2.3 Abstracting information

Usually information is acquired in a format that is only of restricted use. Finding generalisations across cases, using abstraction, extends the use, though usually only within a range of contexts closely related to the learning situations. However, after a succession of changes, making stored structures usable but very particular and shallow, learners (unconsciously) reorganise the information into a new generative form. This is both more economical and more powerful because it is wider in scope – a sort of deductive system in which novel conclusions can be derived. This illustrates what Karmiloff-Smith (1992) calls “Representational Redescription”. The transition in human children from using empirically learnt words and phrases to using generative syntax, allowing a potentially infinite class of sentences to be understood or generated, is a well-known example. As a simpler example, this new kind of generative system potentially allows animals to apply elements of existing knowledge about how to perform a particular action in one context to an entirely new context, in order to access a new food resource (e.g. Gajdon et al., 2006). The mechanisms underlying the processes of abstraction are still subject to debate. Sidman (see Sidman, 1990; 2000 for example) has proposed that several kinds of concepts (such as equivalence relationships like symmetry and transitivity) might be acquired purely as a consequence of reinforcement contingencies. Experiments using successive matching in pigeons (Urcuioli, 2008) support this theory and further suggest that pigeons might be forming stimulus-temporal location compounds. However, these experiments involved a small number of familiar, simple stimuli, presented in tightly constrained learning environments. It is hard to see how that sort of concept learning could enable an animal to classify situations encountered in the wild where the configuration of details can be novel and consequences of possible actions need to be *worked out* rather than retrieved from memory. It is not clear what the reinforcement might be in such cases, nor can the mechanisms proposed by Sidman (1990; 2000) readily explain Karmiloff-Smith’s (1992) representational redescription.

While Rover would initially have to learn anew how to handle each new kind of material it handled, it might eventually represent them in more abstract categories. This would allow it to identify quickly whether a particular material is in the category that requires scooping or grasping, for example. Similarly, NCC might learn something about the abstract affordances of hooks, enabling them to manufacture a hook from a novel material (Weir et al., 2002).

For each such exploration domain, Karmiloff-Smith (1992) suggests that a) the learning depends on innate mechanisms, although *what* is learnt depends on the environment, and b) that exploration goes through characteristic phases, but different domains are explored at different ages. Finally, she suggests that c) learning within a domain initially produces behavioural competence, followed by a succession of revisions of what is learnt. This reorganises and re-represents the information to generalise the competence, later allowing the competence itself to

become an object of attention, often manifested in abilities to answer questions about the domain, such as what can and cannot occur (see Section 4).

2.4 Extending knowledge by combining domains

New domains can be composed by combining old domains, such as combining play with sand and play with water. Such combinations are possible because so many of the domains involve spatial structures and processes: actions originally done at different locations or times can be done together. This can lead to new forms of interaction (e.g. Miyata et al. 2010). In some cases, what was previously learnt in separate domains provides a basis for predicting the results of composition. In other cases, more empirical learning is required, followed by new forms of conceptual revision, theory construction and meta-cognition. For example, learning about mud after learning about sand and about water.

One common simplification is the discovery that two domains, involving different perceptual contents and affordances, can nevertheless share structures and be unified in a useful new abstraction. This can create new domains, or be applied to existing domains. Abstractions such as order, containment, contiguity, motion and causation can be applied to several exploration domains, but showing how a robot can discover and use them has yet to be determined. It is not clear how the many environmental properties represented by adults as numerical measures (e.g. position, velocity, volume) can come to be represented in a young learner.

2.5 Development of meta-cognition

There are many different forms of meta-cognition (e.g. Karmiloff-Smith, 1992 Chapter 5; Povinelli and Preuss 1995), but in this paper we refer specifically to self-directed meta-cognition. It would be beneficial for agents to be able to detect and monitor their own level of uncertainty, for example, by detecting that they do not have sufficient information to make a decision. However, as yet there are no plausible working models for either artificial agents or biological organisms, and the evidence for meta-cognition of this kind in animals has been difficult to establish (see Smith, 2009 for a review). There may be some self-organising knowledge stores, able to react automatically to changes and new opportunities. In other cases, the information processing architecture may include a separate sub-system, with meta-cognitive competences monitoring other sub-systems' behaviour and detecting opportunities to initiate major reorganisation (e.g. Sussman, 1975). For example, objects which do not behave as the agent expects - rocks that appear to be solid but crumble as soon as Rover grasps them, or materials which bend when force is applied but do not return to their original shape when released by a NCC - might stimulate the agent to learn more.

A special sub-set of meta-cognitive competences includes sophisticated self-criticism mechanisms that can drive learning, such as improving problem-solving skills (Sussman, 1975; Sloman 2011b). The mechanisms for generating new forms of learning may initially be genome-based products of evolution, but their effects can vary according to an individual's experience. Forms of learning can develop throughout life, partly influenced by the genome and partly by what has been learnt about learning itself at earlier stages. In humans, for example, learning about more advanced mathematical concepts and techniques requires developing the ability to represent more and more complex structures and processes using new forms of representation, such as propositions with more than one quantifier, and increasingly complex truth-preserving transformations. These types of learning differ from earlier forms not just in their content, but also in their structure.

The work of Karmiloff-Smith (1992) on "representational re-description" suggests that, in some cases at least, the competences gained empirically during a period of learning and development can later be revised or transformed. However, it may be difficult to determine empirically whether that occurs by a process of internal re-organisation within a complex form of representation, or whether it requires a separate meta-cognitive system. Experiments in AI could at least reveal some alternatives, and their implications and costs, along with demonstrations of what is possible.

3 Constraints on exploration and information processing systems

In Section 2 we outlined some theoretical requirements of an exploration and information processing system. We need to consider what constraints evolution imposes upon such systems. There is no single answer, as the constraints depend on the life history strategy of the species, as well as the affordances of the individual's niche (Greenberg and Mettke-Hofmann, 2001; Mettke-Hofmann et al., 2002).

3.1 Constraints on energy expenditure

Exploration requires energy expenditure, on which there is strong selection pressure, so more efficient exploration strategies would provide a selective advantage. One way in which this might be achieved is by combining exploration with other activities that are vital to survival, such as foraging (e.g. Krebs et al., 1978; Cohen et al., 2007; Houston et al., this issue). Since Rover will have to balance time (and thus energy) spent exploring its environment with that spent completing its tasks, it would be helpful to have some mechanism for collecting information while pursuing its assigned tasks (see discussion in Hawes et al., 2010). Another way of scheduling the costs of exploration is to concentrate this activity into periods of the animal's life when other time pressures are low. For instance, many altricial species have an intense period of exploratory or play activity as infants or juveniles, when parents provide food and protection (Power, 2000; Held and Špinková, 2011).

3.2 Constraints on timing of exploration

A framework comparing patterns of development and learning (Chappell and Sloman, 2007; Sloman and Chappell, 2005), distinguishes pre-configured competences (present at birth) and meta-configured competences (based on forms of development that are the products of previous learning). These differ in their costs and benefits. Pre-configured competences emerge quickly, but tend to be less variable. Meta-configured competences are more variable, and can be better tailored to new environments, but require more time to develop. This means that they are potentially risky for the animal if the competence is required early in life (e.g. predator escape behaviour).

As mentioned earlier, altricial NCC seem to spend at least some of their juvenile period exploring tool-related objects and actions (Kenward et al., 2006) and learning how to use and control tools. In contrast, Rover is effectively a precocial agent with no parents nearby to shelter and protect it. Thus, it might spend an initial period acquiring skills close to a known, safe location before carrying out real tasks further way.

While the brain is undergoing development during a protracted period of parental dependency, neural connectivity can be shaped by the current environment, as happens in humans (Supekar et al., 2009) and some non-human animals (e.g. song learning in birds: De Groof et al., 2010). However, this may raise additional problems if the re-organisation requires some systems to have developed or matured before others (Bushnell and Boudreau, 1993).

The relationship between life-history strategy and exploration costs and timing is complex. Longer-lived species have more time to acquire sophisticated competences through interaction with the environment, but they may also have a greater *need* for such competences, because they may experience more environmental variability during their life-span. In turn, the complexity of their competences generates another source of variability (Sterelny, 2007). These types of species extend and fine-tune their knowledge throughout life and neophilia is often present into late adulthood (e.g. parrots; Luescher 2006).

3.3 Constraints on level of abstraction

Some complex environments contain self-repeating elements, leading to redundancy of information. More abstract concepts are applicable in a wider range of contexts. However, they may lead to over-generalisation errors (Marcus et al., 1992), such as wrongly identifying a camel as a kind of horse (Tenenbaum et al., 2006). The impact of those errors will depend on the relative cost of false positives and false negatives. If the cost of false positives is high, abstraction may be constrained. For example, Rover might use an abstraction formed from a constellation of topological features

to identify safe surfaces to move on, avoiding having to probe the surface continuously. However, if it wrongly identifies a darker region as a type of rock rather than a hole, the consequences could be serious. Conversely, if the costs of false positives are low and the benefits of not having to respond differently to each item in a broad category are high, then processes favouring abstraction will be under strong selection pressure.

However, not all exploration domains are equally amenable to abstraction. Simple domains with very low variation in relevant information allow a small number of abstractions to suffice, with a correspondingly small number of associations to be learnt. At the other extreme in domains with high variation, there will be many potential invariances to learn, particularly when the variation arises from complexity. Indeed, it is in these situations, in which complex structures contain simpler components in various relationships, that learners have the most to gain by going beyond mere abstraction and developing a generative theory (e.g. a theory of syntax in human language). One important use of exploration is discovering the type of structural variability in a domain as a precursor to discovering the generative theory that unifies the varied possibilities.

4 How do animals fulfil the requirements of exploration?

Following our discussion of some of the theoretical requirements of an information processing system and the evolutionary constraints it faces, and the increasing evidence that – at least in humans - exploration is not random (Gibson, 1988; Cook et al., 2011), we propose that exploration is composed of structured behavioural strategies supported by specific sensory and motor predispositions. Here we discuss some of the features we could investigate in testing this proposal in animals.

4.1 Pre-dispositions: 'safe' defaults

There are certain aspects of the physical world that can be regarded as constants, such as the effect of gravity, the properties of contact, solidity and connectedness and of biological movement or agency. We can expect many non-human animals to possess mechanisms that cope with or use these features as defaults from birth or hatching. There is evidence that at least some organisms have such pre-dispositions, which are fine-tuned and built upon with experience (e.g. Cacchione and Call, 2010; Cacchione et al., 2009; Kundey et al., 2009 on solidity). The work of Elizabeth Spelke and others (Spelke and Kinzler, 2007) has argued that human infants have systems representing actions, objects, number, space and possibly social partners, although the developmental standpoint of Karmiloff-Smith challenges some of the conclusions. Developing alternative working designs should help us clarify alternative hypotheses. This has begun with some non-human animal studies (e.g. Bird and Emery, 2010; Funk 2002; O'Connell and Dunbar, 2005).

4.2 Behaviour structured to increase information gained

Agents can gain valuable information from perceiving and acting on objects around them. It has been suggested that human infants' understanding of actions and object properties derives from the combination of their exploratory behaviour and the information processing systems generating and modifying their behaviour (Gibson, 1988; Gibson and Pick, 2000; Piaget, 1952; Rochat, 2004). In turn, the representations that result from such exploratory activity alter and direct the kinds of actions infants perform on objects (Perone et al., 2008).

If the function of such exploratory behaviour is to gather information, we would expect the form of exploratory behaviour (and the information processing systems underlying it) to change with context, maximising opportunities for gaining relevant information. In human infants, the types of manipulation used alter depending on how the affordances of a series of toys change: looking at and fingering of objects increases when texture changes, but actions such as rotation and transferring the object between the hands increases when shape changes (Ruff, 1984). Similarly, infants are more likely to transfer an object between hands or finger the surface while looking at it, but are more likely to rotate the object while mouthing it (Ruff et al., 1992). Each kind of action generates perceptual changes that are best suited to the sensory modality used, and may maximise the individuals' opportunities for detecting relevant features.

There has been much less work on the form and function of the information gathering aspects of exploration in non-human animals (Inglis, 1983; Inglis et al., 2001; Power, 2000; see Kacelnik, 1987 and Renner, 1990 for critiques), as opposed to the current or future fitness benefits of behaviour usually referred to as "play" (Pellegrini et al., 2007; Bekoff and Byers, 1998; Held and Špinka, 2011). What we do know, suggests that their sensorimotor behaviour acts to increase the quantity and quality of information gained. Many species show active information gathering (Karmiloff-Smith, 1992). Rats, for instance, alter the speed and pattern of their whisking behaviour to increase information about shape and texture of objects that they contact with their vibrissae (Grant et al., 2009), which has been confirmed by modelling the behaviour in a robot (Pearson et al., 2007). So, the parameters of the rats' whisking behaviour are "designed" in such a way to increase the probability of detecting important environmental features.

It is often difficult to distinguish exploratory behaviour from executive action: is the animal lifting an object because it is trying to transport it, or is it gaining information about the object's weight? Of course, it may be able to fulfil both goals at the same time (e.g. Elner and Hughes, 1978), but it can be difficult to determine when (or whether) an animal is collecting perceptual information, without having more detailed information about the extent of its sensory realm (Demery et al., in press). We need more detailed information on animals' sensory systems, as well as more

reliable behavioural or physiological ‘markers’ of exploratory behaviour in non-human animals.

4.3 Exploration directed towards novelty

Neophilia has been shown to be an important aspect of exploration in non-human animals (Greenberg and Mettke-Hofmann, 2001; Mettke-Hofmann et al., 2002), and is often associated with the juvenile phase of an animal’s development (Heinrich, 1995). Novel objects, places and events are – by definition – items about which the animal does not have adequate information, and so animals should prioritise interaction with these items. In particular, animals in dynamic environments might use exploration to experiment with strategies or behaviours that are effective in the current environmental context (Pellegrini et al., 2007).

For instance, Ruff (1986) hypothesised that if the main function of “examining” behaviour in human infants is to gather information, it should a) decrease in frequency with exposure to a particular object, and b) occur before other behaviours when an object is new. She found that both hypotheses were upheld and that the latency and duration of examining seemed to indicate different features of the process, with latency reflecting the time it takes to activate the information gathering system.

Not all aspects of novelty may be equally salient. For example, Perone and colleagues (2008) presented infants with a sequence of images depicting a hand acting on a colourful toy, which produced a sound (e.g. a purple sphere that squeaks when squeezed), followed by an image in which either the action-sound pair or the appearance of the object changed. They found that infants attended more to changes in action than to changes in appearance. In evolutionary terms, it is not clear why such differences in salience exist, but it may be that changes in appearance are less likely to have important implications for the object or event function, than changes in action.

4.4 Active 'testing' when expectations are violated

If an organism’s current empirical observations do not fit with the information it has collected, it should re-initiate exploration in order to resolve the discrepancy: do we observe this process in biological organisms? There is increasing evidence that human children can use a conditional intervention principle to learn about causes and are sensitive to ambiguous information (Gopnik, 1996; Gopnik and Schulz, 2004; Schulz et al., 2007; Tenenbaum et al., 2006). Furthermore, children’s exploration appears to be systematic in many respects, and is thus capable of both detecting discrepancies between observed data and stored information, and of resolving those discrepancies. For example, when children are shown that blocks of a certain category (defined by a linguistic label or by appearance) stick magnetically to a board, they test the properties of new blocks more extensively when they find

that properties vary *within* the category than *between* categories (Schulz et al., 2008). There are several different, inter-linked processes at work here. The individual needs to detect that some aspect of the world is surprising and then begin exploration, focussing on resolving the discrepancy, which may involve re-organisation, or other changes in representations.

Do any non-human animals show similar ‘testing’ behaviours? There seems to be evidence that a number of taxa (e.g. apes, rats and dogs) can use information from a number of sources to make causal inferences. Some can use evidence from the object itself (auditory information, displacement of other objects, weight etc.: Blaisdell et al., 2006; Bräuer et al., 2006; Call, 2004; Hanus and Call, 2008), through social cues (e.g. Povinelli et al., 1990), or by exclusion (Aust et al., 2008; Call, 2006; Hill et al., 2011). However, these experiments relied on animals observing the state of the world, or watching others perform actions on objects. As far as we know, there have been no studies on whether non-human animals spontaneously perform their own ‘tests’, as shown for human children (Schulz et al., 2008). However, some studies have manipulated certain environmental stimuli and measured exploratory behavioural sequences (Kuba et al., 2006; Renner 1990; Bekoff 1975), which revealed behaviour, at least in rats and octopuses, that appears to be similar to human children. We would predict that non-human animals, like children, would become less repetitive in their exploratory actions as they develop. In other words, they would show a greater diversity of exploratory behaviours rather than simply repeating a small number of actions on the same part of the environment, because they learn what *kinds* of actions will be most effective in each situation.

5 Conclusions

We have discussed information gathering in biological and artificial organisms, and ways in which taking a ‘designer stance’ (Dennett, 1981; McCarthy, 2008) can help to both clarify theoretical issues and suggest lines of enquiry for behavioural experiments on humans and non-human animals (Chappell and Thorpe, 2010). Studying exploration and information acquisition is difficult partly because it intersects with – and has important consequences for – many other aspects of an animal’s behaviour. It is part of learning, as well as physical maturation, and it can help to provide the content and structure for cognition. We need also to take note of both the successes and mistakes of an individual to help illuminate the underlying processes. Understanding the animal’s sensory world is also of vital importance, since it acts both as a conduit and a filter for the information that is gathered, but many details remain to be studied (e.g. Demery et al., in press). There are such substantial differences not only across environments, but also between and within species, so there are probably many different kinds of mechanisms operating in combination in different exploration domains. Future research should investigate

how all these elements *interact* quantitatively and qualitatively. Demonstrable working models should suggest new research questions. However, some questions may only be answered when the brain mechanisms are better understood (e.g. Reynolds and O'Reilly, 2009).

We can make progress on the theoretical issues by using the analytical tools and techniques from AI and robotics (for example Hawes et al. 2010; Markram, 2006; Pardowitz and Dillman, 2007; Saegusa et al. 2008;), particularly to analyse the environmental requirements. Thinking about the ways in which one might implement such features in a robot (or actually building a robot), forces one to consider details and interactions that may otherwise have been overlooked, and can also generate specific, testable hypotheses for behavioural tests with animals (Chappell and Thorpe, 2010; Pearson et al., 2007; Prescott et al., 2009; Webb, 2001).

6 Acknowledgements

We thank our colleagues in the Cognitive Adaptations Research Group and the Intelligent Robotics Laboratory for helpful discussions on these topics, and of course Alex Kacelnik, with whom two of us have had many stimulating and productive conversations.

7 References

- Aust, U., Range, F., Steurer, M., Huber, L. 2008. Inferential reasoning by exclusion in pigeons, dogs, and humans. *Anim. Cogn.* 11, 587-597.
- Bailey, T., Durrant-Whyte, H. 2006. Simultaneous localization and mapping (SLAM): Part II. *Robotics & Automation Magazine, IEEE.* 13, 108-117.
- Bekoff, M. 1975. Animal play and behavioral diversity. *Am. Nat.* 109, 601-603.
- Bekoff, M., Byers, J.A. 1998. *Animal play: Evolutionary, comparative, and ecological perspectives.* Cambridge University Press, UK.
- Bellman, R. 1961. *Adaptive Control Processes: A Guided Tour.* Princeton University Press, Princeton, NJ.
- Bird, C.D., Emery, N.J. 2010. Rooks perceive support relations similar to six month-old babies. *Proc. R. Soc. B.* 277, 147-151.
- Blaisdell, A.P., Sawa, K., Leising, K.J., Waldmann, M.R. 2006. Causal reasoning in rats. *Science.* 311, 1020-1022.
- Bräuer, J., Kaminski, J., Riedel, J., Call, J., Tomasello, M. 2006. Making inferences about the location of hidden food: Social dog, causal ape. *J. Comp. Psychol.* 120, 38-47.
- Bushnell, E.W., Boudreau, J.P. 1993. Motor development and the mind: The potential role of motor abilities as a determinant of aspects of perceptual development. *Child Dev.* 64, 1005-1021.
- 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. *J. Comp. Psychol.* 123, 168-180.

- Cacchione, T., Call, J. 2010. Do gorillas (*Gorilla gorilla*) and orangutans (*Pongo pygmaeus*) fail to represent objects in the context of cohesion violations? *Cognition*. 116, 193-203.
- Call, J. 2004. Inferences about the location of food in the great apes (*Pan paniscus*, *Pan troglodytes*, *Gorilla gorilla*, and *Pongo pygmaeus*). *J. Comp. Psychol.* 118, 232-241.
- Call, J. 2006. Inferences by exclusion in the great apes: the effect of age and species. *Anim. Cogn.* 9, 393-403.
- Chappell, J., Sloman, A. 2007. Natural and artificial meta-configured altricial information processing systems. *IJUC*. 3, 211-239.
- Chappell, J., Thorpe, S. 2010. AI-inspired biology: does AI have something to contribute to biology? in: Chappell, J., Thorpe, S., Hawes, N., Sloman, A. (Eds.), *Proceedings of the International Symposium on AI Inspired Biology: A Symposium at the AISB 2010 Convention, SSAISB*.
- Cohen, J. D., McClure, S. M., Yu, A. J. 2007. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Phil. Trans. R. Soc. B*, 362, 933-942.
- Cook, C., Goodman, N. D., Schulz, L. E. 2011. Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120, 341-349.
- De Groof, G., Verhoye, M., Poirier, C., Leemans, A., Eens, M., Darras, V.M., Van der Linden, A. 2010. Structural changes between seasons in the songbird auditory forebrain. *J. Neurosci.* 29, 13557-13565.
- Demery, Z. P., Chappell, J., Martin, G. R. In press. Vision, touch and object manipulation in Senegal parrots *Poicephalus senegalus*. *Proc. R. Soc. Lond. B*.
- Dennett, D.C., 1981. *Brainstorms: Philosophical Essays on Mind and Psychology*. MIT Press.
- Elnor, R. W., Hughes, R. N. 1978. Energy Maximization in Diet of Shore Crab, *Carcinus maenas*. *J. Anim. Ecol.* 47, 103-116.
- Flanagan, J.R., Wing, A.M. 1997. The role of internal models in motor learning and control: evidence from grip force adjustments during movements of hand-held loads. *J. Neurosci.* 17, 1519-1528.
- Funk, M. 2002. Problem solving skills in young yellow-crowned parakeets (*Cyanoramphus auriceps*). *Anim. Cogn.* 5, 167-176.
- Gajdon, G.K., Fijn, N., Huber, L. 2006. Limited spread of innovation in a wild parrot, the kea (*Nestor notabilis*). *Anim. Cogn.* 9, 173-181.
- Gibson, E. J. 1988. Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Ann. Rev. Psychol.* 39, 1-42.
- Gibson, E.J, Pick, A.D. 2000. *An Ecological Approach to Perceptual Learning and Development*. Oxford University Press, New York.
- Gopnik, A. 1996. The scientist as child. *Phil. Sci.* 63, 485-514.
- Gopnik, A., Schultz, L. 2004. Mechanisms of theory formation in young children. *TICS*. 8, 371-377.
- 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. *J. Neurophys.* 101, 862-874.

- Greenberg, R., Mettke-Hofmann, C. 2001. Ecological aspects of neophobia and neophilia in birds. *Curr. Ornith.* 16, 119-178.
- Hanus, D., Call, J. 2008. Chimpanzees infer the location of a reward on the basis of the effect of its weight. *Curr. Biol.* 18, R370-R372.
- Hawes, N., Hanheide, M., Sjöo, K., Aydemir, A., Jensfelt, P., Göbelbecker, M., Brenner, M., Zender, H., Lison, P., Kruijff-Korbayov, I., Kruijff, G.-J., Zillich, M. 2010. Dora the explorer: A motivated robot. in: *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*.
- Hawes, N. 2011. A survey of motivation frameworks for intelligent systems. *AI.* 175, 1020-1036.
- Heinrich, B. 1995. Neophilia and exploration in juvenile common ravens, *Corvus corax*. *Anim. Behav.* 50, 695-704.
- Held, S.D.E., Špinka, M. 2011. Animal play and animal welfare. *Anim. Behav.* 81, 891-899.
- Hill, A., Collier-Baker, E., Suddendorf, T. 2011. Inferential reasoning by exclusion in great apes, lesser apes, and spider monkeys. *J. Comp. Psychol.* 125, 91-103.
- Houston et al. THIS ISSUE
- Houston, A., Krebs, J., & Erichsen, J. 1980. Optimal prey choice and discrimination time in the great tit (*Parus major* L). *Behav. Ecol. Sociobiol.* 6, 169-175.
- Inglis, I.R., 1983. Towards a cognitive theory of exploratory behavior, in: Archer, J., Birke, L.I.A. (Eds.), *Exploration in animals and humans*. Van Nostrand Reinhold, Workingham, pp. 72-117.
- Inglis, I.R., Langton, S., Forkman, B., Lazarus, J. 2001. An information primacy model of exploratory and foraging behaviour. *Anim. Behav.* 62, 543-557.
- Kacelnik, A. 1987. Information primacy or preference for familiar foraging techniques? A critique of Inglis and Ferguson. *Anim. Behav.* 35, 925-926.
- Kacelnik, A., Krebs, J.R. 1985. Learning to exploit patchily distributed food, in: Sibly, R.M., Smith, R.H. (Eds.), *Behavioural Ecology: Ecological Consequences of Adaptive Behaviour*. Blackwell Scientific Publications, Oxford, pp. 189-205.
- Kacelnik, A., Todd, I. 1992. Psychological mechanisms and the marginal value theorem - effect of variability in travel time on patch exploitation. *Anim. Behav.* 43, 313-322.
- Karmiloff-Smith, A. 1992. *Beyond Modularity: A Developmental Perspective on Cognitive Science*. MIT Press, Cambridge, MA.
- Kenward, B., Rutz, C., Weir, A. A. S., Kacelnik, A. 2006. Development of tool use in New Caledonian crows: inherited action patterns and social influences. *Anim. Behav.* 72, 1329-1343.
- Krebs, J.R., Kacelnik, A., Taylor, P. 1978. Tests of optimal sampling by foraging great tits. *Nature.* 275, 27-31.
- Kuba, M.J., Byrne, R.A., Meisel, D.V., Mather, J.A. 2006. When do octopuses play? Effects of repeated testing, object type, age, and food deprivation on object play in *Octopus vulgaris*. *J. Comp. Psychol.* 120, 184-190.
- Kundey, S.M.A., de Los Reyes, A., Taglang, C., Baruch, A., German, R. 2009. Domesticated dogs' (*Canis familiaris*) use of the solidity principle. *Anim. Cogn.* 13, 497-505.

Luescher, A. U. 2006. Manual of parrot behaviour. Blackwell Publishing, Oxford, UK.

Marcus, G.F., Pinker, S., Ullman, M., Hollander, M., Rosen, T.J., Xu, F. 1992. Overregularization in language acquisition. *Monographs of the Society for Research in Child Development*. 57, 1-182.

Markram, H. 2006. The blue brain project. *Nat. Rev. Neurosci.* 7, 153-160.

Martin, A. 2007. The representation of object concepts in the brain. *Ann. Rev. Psychol.* 58, 25-45.

McCarthy, J. 2008. The Well-Designed Child. *AI*. 172, 2003-2014.

McNamara, J., Houston, A. 2009. Integrating function and mechanism. *TREE*, 24, 670-675.

Mettke-Hofmann, C., Winkler, H., Leisler, B. 2002. The significance of ecological factors for exploration and neophobia in parrots. *Ethology*. 108, 249-272.

Miyata, H., Gajdon, G.K., Huber, L., Fujita, K. 2010. How do keas (*Nestor notabilis*) solve artificial-fruit problems with multiple locks? *Anim. Cogn.* 14, 45-58.

O'Connell, S., Dunbar, R.I.M. 2005. The perception of causality in chimpanzees (*Pan spp.*). *Anim. Cogn.* 8, 60-66.

Pardowitz, M., Dillmann, R. 2007. Towards life-long learning in household robots: the Piagetian approach. in: *Proc. 6th IEEE International Conference on Development and Learning*, London, UK.

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. *Adapt. Behav.* 15, 223.

Pellegrini, A., Dupuis, D., Smith, P. 2007. Play in evolution and development. *Dev. Rev.* 27, 261-276.

Perlovsky, L. 1998. Conundrum of combinatorial complexity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 20, 666-670.

Perone, S., Madole, K.L., Ross-Sheehy, S., Carey, M., Oakes, L.M. 2008. The relation between infants' activity with objects and attention to object appearance. *Dev. Psychol.* 44, 1242-1248.

Piaget, J. 1952. *The Origins of Intelligence in Children*. International Universities Press, New York, NY.

Povinelli, D.J., Nelson, K.E., Boysen, S.T. 1990. Inferences about guessing and knowing by chimpanzees (*Pan troglodytes*). *J. Comp. Psychol.* 104, 203-210.

Povinelli, D.J., Preuss, T.M. 1995. Theory of mind: evolutionary history of a cognitive specialization. *TINS*. 18, 418-424.

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

Prescott, T.J., Pearson, M., Mitchinson, B., Sullivan, J.C.W., Pipe, A.G. 2009. Whisking with robots: From rat vibrissae to biomimetic technology for active touch. *IEEE Robotics and Automation Magazine*. 16, 42-50.

Rescorla, R.A. 1968. Probability of shock in the presence and absence of CS in fear conditioning. *J. comp. Physiol. Psychol.* 66, 1-5.

Renner, M.J. 1990. Neglected aspects of exploratory and investigatory behavior. *Psychobiol.* 18, 16-22.

- Reynolds, J.R., O'Reilly, R. C. 2009. Developing PFC representations using reinforcement learning. *Cognition*. 113, 281-292.
- Rochat, P. 2004. *The Infants World*. Harvard University Press.
- Ruff, H.A. 1984. Infants' manipulative exploration of objects: Effects of age and object characteristics. *Dev. Psychol.* 20, 9-20.
- Ruff, H.A. 1986. Components of attention during infants' manipulative exploration. *Child Dev.* 57, 105-114.
- Ruff, H.A., Saltarelli, L., Capozzoli, M. 1992. The differentiation of activity in infants' exploration of objects. *Dev. Psychol.* 28, 851-861.
- Saegusa, R., Sakka, S., Metta, G. Sandini, G. 2008. Sensory prediction learning - how to model the self and environment. in: *Proceedings of the 12th IMEKO TC1-TC7 joint Symposium on 'Man Science and Measurement' (IMEKO2008)*, 3rd-5th September, pp. 269-275.
- Schulz, L.E., Gopnik, A., Glymour, C. 2007. Preschool children learn about causal structure from conditional interventions. *Dev. Sci.* 10, 322-332.
- Schulz, L.E, Standing, H.R., Bonawitz, E.B. 2008. Word, thought, and deed: The role of object categories in children's inductive inferences and exploratory play. *Dev. Psychol.* 44, 1266-1276.
- Shannon, C. 1948. A mathematical theory of communication. *Bell System Technical Journal*. 27, 379-423 and 623-656.
- Shettleworth, S. J. 2010. *Cognition, Evolution and Behaviour*. New York: Oxford University Press.
- Sidman, M. 1990. Equivalence relations: Where do they come from? In: Blackman, D. E., Lejeune, H. (Eds.), *Behavior analysis in theory and practice: Contributions and controversies*. Lawrence Erlbaum Associates, New Jersey, pp. 93-114.
- Sidman, M. 2000. Equivalence relations and the reinforcement contingency. *J. Exp. Anal. Behav.* 74, 127-146.
- Slovan, A. 2011a. 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 Publishing Co.
- Slovan, A. 2011b. Varieties of meta-cognition in natural and artificial systems, in: Cox, M.T., Raja, A. (Eds.), *Metareasoning: Thinking about Thinking*. MIT Press, Cambridge, MA.
- Slovan, A., Chappell, J. 2005. The Altricial-Precocial Spectrum for Robots. in: *Proceedings IJCAI'05, Edinburgh* pp. 1187-1192.
- Smith, J. D. 2009. The study of animal metacognition. *TICS*, 13, 389-396.
- Spelke, E.S., Kinzler, K.D. 2007. Core knowledge. *Dev. Sci.* 10, 89-96.
- Sterelny, K. 2007. Social intelligence, human intelligence and niche construction. *Phil. Trans. R. Soc. Lond. B.* 362, 719-730.
- Supekar, K., Musen, M., Menon, V. 2009. Development of large-scale functional brain networks in children. *PLoS Biol.* 7, e1000157.
- Sussman, G.J. 1975. *HACKER: A Computer Model of Skill Acquisition*. Elsevier Science Inc.
- Tenenbaum, J.B., Griffiths, T.L., Kemp, C. 2006. Theory-based Bayesian models of inductive learning and reasoning. *TICS*. 10, 309-318.

- Urcuioli, P. J. 2008. Associative symmetry, "anti-symmetry", and a theory of pigeons' equivalence-class formation. *J. Exp. Anal. Behav.* 90, 257-282.
- Webb, B. 2001. Can robots make good models of biological behaviour? *Behav. Brain Sci.* 24, 1033-1050.
- Weir, A. A. S., Chappell, J., Kacelnik, A. 2002. Shaping of hooks in New Caledonian crows. *Science*, 297, 981.
- Wing, A.M., Lederman, S.J. 1998. Anticipating load torques produced by voluntary movements. *J. Exp. Psychol.: Human Percep. Perf.* 24, 1571-1581.
- Zentall, T.R., Wasserman, E.A., Lazareva, O.F., Thompson, R.K.R., Rattermann, M.J. 2008. Concept learning in animals. *Comp. Cogn. Behav. Rev.* 3, 13-45.