

The Three-stage Theory of Exploration

Zoe P. Demery and Jackie Chappell

Centre for Ornithology, School of Biosciences, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

Abstract

Faced with a vast, dynamic environment, some animals and robots often need to acquire and segregate information about objects. The form of their internal representation depends on how the information is utilised. Sometimes it should be compressed and abstracted from the original, often complex, sensory information, so it can be efficiently stored and manipulated, for deriving interpretations, causal relationships, functions or affordances. We discuss how salient features of objects can be used to generate compact representations, later allowing for relatively accurate reconstructions and reasoning. Particular moments in the course of an object-related process can be selected and stored as ‘key frames’. Specifically, we consider the problem of representing and reasoning about a deformable object from the viewpoint of both an artificial and a natural agent.

Keywords: Learning, Animal Cognition, Representations, Deformable Objects

Material from of this paper has formed part of three publications (Demery et al., 2010; Chappell et al., 2012; Arriola-Rios et al., 2012).

1. Introduction

The brain of any animal is finite, so it cannot contain a perfect model of the world around it. Animals receive a variety of information through their sensors, but how useful that information is dependent on several factors including: environmental conditions; the accuracy of information processing and interpretation; the information’s use; and the animal’s behavioural response (Sloman, 2011). An individual must balance the amount of useful, relevant information obtained about its surroundings against the energy expended in gaining it (Chappell and Sloman, 2007). What is unclear is which bits of information are most useful for processing and storing in the brain, or how they can best be represented. Principally, we propose that when an individual gathers information through its senses, it often forms object

representations supported by systematic exploration¹.

Exploration is found throughout the animal kingdom and environmental contexts (e.g. Berlyne, 1960; Glickman and Sroges, 1966). Yet to date, there has been little systematic, quantitative research about it, and its structure in supporting learning mechanisms of different individuals (for more discussion see White, 1959; Rochat, 2001). What research there is, has largely originated from the field of developmental psychology (chapter 1). Here it is widely accepted that exploration is integral to a child’s learning and causal understanding about their world. However, many developmental psychologists advocate probabilistic learning as

¹Cognition does not always rely on internal representations and the degree of detail in any internal representation can vary greatly depending on the situation. Quite complex-looking actions can often be performed by simple mechanisms and small neural architectures (e.g. Webb, 2001; McCrone, 2006). For instance, there can be a lack of detail especially when the environment can largely control an animal’s behaviour, such as in flocking behaviour or in using pheromone trails. Here alternative, but complementary, mechanisms may be more relevant, such as emergency or embodiment (reviewed in Calvo et al., 2008). However, in this chapter we will not consider these cases and are concerned with more complex, cognitively flexible animals.

Email address: zoe@vortices.com (Zoe P. Demery and Jackie Chappell)

the sole mechanism to explain human exploration, contending it cannot be extended to other animals (Schulz and Sommerville, 2006; Schulz and Bonawitz, 2007).

Among the non-human animal researchers, behavioural ecologists view exploration as another instinctive motivation, like foraging or courtship, driven by uncertainty in variable environments (e.g. Miller and Dollard, 1941; Inglis, 1983; chapter 1). The moves towards integrating cognitive processes (e.g. discrepancy reductions or spatial cognitive maps), are discussed in the context of passive animals driven by impulses, rather than conscious individuals, reasoning about the world around them (Toates, 1983; Loewenstein, 1994; Inglis and Langton, 2006). Taking an evolutionary perspective, this field pays slighter attention to the individual's perspective.

Animal cognition researchers have focussed on what the different cognitive capacities of different species are, rather than how they process information to achieve those capacities (Shettleworth, 1993; Thomas, 1996; chapter 1). For example, the 'trap-tube task' is a typical litmus test for causal understanding of gravity, which has revealed much about various species (e.g. Visalberghi and Tomasello, 1998; Penn et al., 2008; Taylor et al., 2009). This test is only a binary measure, however, of whether an individual can complete the task – not what the underlying processes are.

An under-represented approach is the structural-mechanistic model (chapter 1). The aim here is gain answers by studying the pattern or sequence of exploratory behaviour (e.g. Gibson, 1988; Renner, 1990; Sloman, 2008a; Magnani, 2009). The sensorimotor apparatus of an exploring animal is considered, asking how the senses support exploratory learning (e.g. Gibson, 1962; Mesulam, 1998; Shams and Seitz, 2008; Demery et al., 2011).

This perception-orientated approach is a blossoming area particularly in artificial intelligence. Designing such cognitive models requires specific parameters, which can aid animal behaviourists to form concrete, testable hypotheses (for more information see Demery et al., 2010). However, it seems there is not yet a robot/simulation that can form concepts, or generalise information to new situations like some animals. Cognitive roboticists have looked at different learning

mechanisms in isolation with relative success, but few projects have tried combining them into one information-processing system (e.g. Hawes et al., 2010; Markram, 2006; Timmis et al., 2008).

In Part I of this thesis we investigated how the senses can support information-gathering by exploration. In Part II we considered how the environment can influence exploration. Now we will start to consider how such an information-processing system is built. In this chapter, inspired by the designer-based approach used in artificial intelligence, we will take a distinctive top-down perspective. We will first examine some of the general environmental features that make it easier for animals to internalise the surrounding world and make predictions about it, but also some environmental problems that need to be overcome (Section 2). From this, we will consider how these features may be processed by the design of the animal's exploratory system. We integrate elements from each discipline discussed above in a novel way, leading to our Three-stage Theory of Exploration.

We propose that when forming representations, object exploration is not always random, but *structured, selective* and *sensitive* to particular features and salient categorical stimuli in the environment. Further, an exploring individual can follow through three stages of theory formation – the forming (Section 3), the testing (Section 4) and the refining of hypotheses (Section 5). Each hypothesis is probably specific to a particular group ('exploration domain') of processes or affordances, but they may also be generalisable to novel contexts. We will lastly consider how we may test this theory, particularly through a series of comparative behavioural tests on parrots and humans (Section 6) – both being notoriously exploratory and playful throughout their lives (chapters 6 and 7).

2. Predictable environmental features

As a basis for the Three-Stage Theory of Exploration, we first consider what features in the environment are reliably constant enough for an individual to learn about them. Shaw (2008, page 5) suggests,

“The chief end of an intelligent individual is to understand the world around it.”

The word ‘understanding’ implies the animal can make predictions about the world. For this, the animal needs to be able to detect salient features in its environmental niche. It can do this by actively exploring its surrounding environment, often by combining perception and cognitive analysis with action. An exploring animal is driven by its goals. These can be explicit goals, such as foraging or particular problem-solving tasks. Alternatively, they can be implicit goals, such as gathering information by apparently random playful or exploratory behaviour.

Some of the environmental features perceived and learned about may come with environmental problems; these problems must be surmounted first. No two animals face the same environmental problems, as the specific features perceived depend on the life history strategy of each individual, as well as the affordances of each individual’s niche (Greenberg and Mettke-Hofmann, 2001; Mettke-Hofmann et al., 2002). The following subsections present a general list of environmental features and problems, which may be perceived and solved by various exploratory species (see also Chappell et al., 2012).

2.1. Persistence

For an animal to make relatively accurate environmental predictions, a few unchanging features need to persist in the environment for a significant length of time. The information-processing system should assume these properties continue to exist in a range of environmental contexts, which differ in space and time. These features can then be utilised by the animal’s learning mechanisms, to develop a representation of the world.

We expect all individuals to work from the assumption the environment is persistent (enabling them to re-use information), but the strength of this persistency-assumption will vary across species. The parasitoid wasp (*Ammophila dysmica*) displays a sophisticated-looking behaviour sequence (Fabre et al., 1918; Rosenheim, 1987), which is in fact genetically fixed. Any small change to the environment renders the wasp confused and behaving erratically, as it works from the assumption of universal persistency. Rats (*Rattus norvegicus*) in a maze are not confused by changes to layout (e.g. Wilz and Bolton, 1971; Albert and Mah,

1972), as their assumptions of persistency are flexible enough to cope with these changes.

2.2. Regularity

Environmental regularity is the predictable presence of certain features, which can be learned by an animal perceiving a fixed relationship between occurrences². These environmental patterns enable partial representations, particularly when an individual is faced with different causal problems. Causality is a type of regularity, where partaking elements are not always identifiable, but whose manifestation always entails the same consequence.

Animals should have mechanisms capable of detecting and exploiting these patterns. It is likely they categorise the environment, selecting key features linked by predictive relationships³. Among the mechanisms capable of detecting such patterns is ‘associative learning’, which has been studied extensively in several species Shettleworth (for a review of the different mechanisms see 1999). For example, information about food availability will become associated with the act of feeding, so the presence of this information will increase foraging activity. Even very simple animals (e.g. nudibranch, *Hermisenda crassicornis*; Alkon et al., 1982) can detect causal associations and patterns in time and space.

2.2.1. Sequentiality

Sequentiality is a series of features, nearly always perceived in the same order⁴. The first few features can be used to identify the sequence and predict either the following features, or the rules-set needed to process them. Examples include lightning flashes being followed by thunder claps, or listening to the sound of a prey and preparing to chase.

²Regular environmental features can be present as groups of features, in different dimensions, or in a hierarchical structure.

³The information contained within each of these discrete categories can include information about continuous features (e.g. length/height).

⁴These may not be contiguous and can include cause-and-effect learning.

2.2.2. Branching sub-sequences

Further to sequentiality, there can be a succession of sub-sequences. The links here would allow a few options to follow, such as beginnings of other sub-sequences. This forms a branching structure, which becomes layered, modular, and sometimes hierarchical (e.g. Arriola-Rios and Savage, 2007). This branching structure found in the environment, may be reflected in an animal's internal representation (see also Section 4).

The maximum length of a sequence, and the maximum number of branches that can be remembered and manipulated, impose strict limitations upon what the individual can understand, and the types of patterns it is capable of detecting. However, this branching structure would allow more complex individuals to make abstractions. Concepts formed at one stage could be re-used and refined to repeatedly form ever more complex concepts in multiple ways (see Section 5). This allows for progressively specific and parallel processes, such as the models described in Karmiloff-Smith, 1995 (subsection 3.2). For example, the New Caledonian crow (*Corvus moneduloides*) has been shown to be particularly adept at detecting and planning through multiple-steps in an environmental sequence (Wimpenny et al., 2009). Although the end result was a complex display, the constituent steps could be simple mechanisms.

2.3. Consistency

When multiple mechanisms are used to collect and process sensory signals, sometimes they complement each other, by providing different information about the same object of interest. This may provide superfluous information, but it also allows the animal to confirm its own sensory inputs and establish the information is reliable. It is important for the animal that different mechanisms of perceiving or deducing the same thing are consistent with each other. For instance, an individual exploring a potential food item, should receive tactile information in agreement with the position and surface it sees.

2.4. Redundancy

Given the inherent limitations of the brain, animals can only build partial internal representations of

their surroundings. These partial representations may not allow individuals to make perfect predictions for all environmental events⁵. As individuals receive sensory information, they may not succeed at processing relevant, useful⁶ signals. Therefore, we expect there to be errors at different levels of the perceptual and cognitive processes. Thus, information-processing systems should be built in such a way as to tolerate this margin-of-error.

Some animals have more than one mechanism to perform actions, find things, or solve problems. The individual could just use reactive mechanisms, dependent on different layers of sensory-signal filtering, or it could use a combination of different learning mechanisms. While qualitatively different, all of these mechanisms produce similar, valid results, which overcome the problem of environmental redundancy. What is not clear is how each of these mechanisms allows the animal to learn about and retain only the relevant information. How does the animal reconstruct faulty perceptions from new perceptions that convey the same information?

For example, Horner and Whiten (2005) presented a puzzle box to chimpanzees (*Pan troglodytes*) and human children. The puzzle box was designed so that there were different ways of getting into it, although all the ways led to the same result (retrieving the reward out of the box). Children tended to directly imitate, action for action, a demonstrator's way of getting into the box. This included the irrelevant actions, which did not actually help them attain the goal. However, the chimpanzees did not imitate the demonstrator, but *emulated* them. In other words, they only performed their own relevant, functional actions to attain goal. These issues will be discussed below in subsection 2.6.

2.5. Variability

Many environments vary over time; due to changes in climate, geology, competitors' behaviours, co-evolutionary arms races between predators and prey, niche construction (e.g. Sterelny,

⁵An animal (e.g. a nematode, *Caenorhabditis elegans*, or even various forms of bacteria; Qin and Wheeler, 2007; Ben-Jacob, 2009) in a very simple environment can make perfect predictions; but we are not concerned with these cases.

⁶Here, relevant or useful information is that which – if acted on – will influence the animal's evolutionary fitness.

2007) and food availability (e.g. Houston et al., 1980; Kacelnik and Krebs, 1985; Kacelnik and Todd, 1992). Little experience is needed in a relatively static environment, where precocial animals – whose behaviour has been almost completely determined by their genome – just need to survive long enough to reproduce (e.g. wasp described in subsection 2.1). Other individuals are required to adapt to diverse, dynamic environments by learning (for greater discussion see Chappell and Sloman, 2007; also see chapter 5).

The different exploratory mechanisms, for extracting relevant information (subsection 2.4), are likely shaped by experience. The animal should seek out information to reinforce, evolve and, where possible, prove or disprove its current internal representations, particularly if its expectations are violated. Depending on the individual's needs and competences, a specific, relevant subset of experiences allow specific, relevant features of its niche to be captured (e.g. von Bayern et al., 2009). We believe there is continual extension of these 'branches', or 'information blocks', throughout the individual's life. At different developmental stages, a human child likely takes in different aspects of the same overheard conversation, or different aspects of the operation of the same tool, then later adapts accordingly (see discussion in Sloman, 2010).

The kea (*Nestor notabilis*), a New Zealand parrot, has proved very adaptable and cognitively flexible. Huber, Gajdon and colleagues (reviewed in 2006) have documented how kea display quite innovative behaviours in relatively artificial situations, such as lifting restaurant bin lids through various steps. Kea have been shown to learn from previous experiences, even from several months ago. They also seem very exploratory during problem-solving, although there are several individual differences in strategies employed (e.g. Auersperg et al., 2011).

2.6. The cyclical animal-environment interaction

Together these environmental features and problems form a structured universe (Section 2). Parts of this structure can be perceived and understood by animals. The existence of these predictable features reduces the amount of information needed to represent the environment. Once the animal has

processed sufficient features, the remaining important information can be inferred when needed – the question is how.

There is increasing evidence (at least in humans, e.g. Gibson, 1988; Cook et al., 2011) that exploration is not random, but structured, selective and sensitive to particular environmental stimuli. We propose exploration is composed of structured behavioural strategies supported by specific sensorimotor predispositions (e.g. see chapter 2). Thus, we turn our discussion to what internal features an exploratory animal may have to process environmental information. In parallel, we consider how these internal features may manifest themselves on the external world via the animal's exploratory behaviour (Section 6).

Although many of these features have been studied before in isolation, no one has yet integrated them into one information-processing system, or addressed the short-comings of each type of mechanism (chapter 1). Additionally, no one has comprehensively considered how they may follow on from one another over a lifetime. We propose that exploratory animals, from the start to the end of their lives, generally follow three stages of exploration. How they progress from one stage to another may vary between exploration domains (and between individuals) and there may be some overlap between the stages. The actual time-scales for each stage's progression are not clear and are ripe for future research (Section 6). In the first stage, an individual forms their hypotheses about the world (Section 3), then they test these hypotheses in the second stage Section 4. The final stage is the individual extending and refining their hypotheses (Section 5). There are different design features within each stage, discussed below.

3. Stage I: forming hypotheses

We propose in the first stage the animal forms hypotheses about the world by interacting with it. The individual is probably young and just beginning to gather information about the world around it. Consequently, it detects pattern in the environment using simple learning mechanisms, such as probabilistic learning, or trial-and-error (subsection 3.3). These are directed by sensorimotor predispositions (subsection 3.1), which are particularly sensitive to categorical stimuli (subsec-

tion 3.2). Their exploratory behaviour focusses on novel stimuli (subsection 3.5), or any obvious changes that are detected in the environment (subsection 3.4). We will now discuss each of these internal design features of Stage I separately.

3.1. *Sensorimotor predispositions*

A cognitive system consists of a body with a sensory and a motor apparatus that allows for interaction with the environment and a brain to map sensory stimulation onto motor actions (Floreato and Mattiussi, 2008). Merleau-Ponty (2002) was among the first to point out that the perceiver's physical body affects sensation; that perception is not simply a passive recording of environmental stimulation. The senses provide another level of *active* exploration, in addition to what is gained from physical body movement (chapter 3). Objects are invariable in the environment without movement/processes. For instance, bees can only see static snapshots of the world, so they generate their own processes by moving through it – thus perceiving through optic flow (Srinivasan and Gregory, 1992; Loomis and Beall, 1998)

The Gibsons (Gibson, 1977; Gibson, 1988) argued the infant's sensory and perceptual systems have evolved to pick up information that is already available in the environment, so no cognitive reconstruction of information needs to take place. The infant's perception detects invariant information (like that described in subsection 2.2) by differentiation across several experiences. They argued against traditional associative learning mechanisms: the infant is learning a richer series of different levels of signal quality of the same stimuli, rather than binary pairs of associations.

Certainly, it is helpful to look at what the environment offers, then consider how the organism is processing that information (as in Section 2). The environment offers regularities and invariant information, but as these still need to be sorted from the 'noise' (subsection 2.4), we should not discount construction of internal representations altogether. However varied the environments, certain competences will be hard-wired into the genes. It may be more efficient for the individual's brain to store information gathered from its different senses *amodally*, and group it into categories (subsection 3.2), perhaps in some sort of sym-

bolic (rather than probabilistic), hierarchical structure (subsection 2.2.2).

It follows reason that an individual's sensorimotor system is adapted to attend to particularly important information, such as regularities providing information about different objects' affordances. Initial sensorimotor competences should be designed to make use of relevant environmental information, such as salient shapes with their: texture, edges/contours, curvatures and orientations of surface fragments. Object areas such as flat, smooth surfaces are less likely to contain useful information about the object, so more sensorimotor attention should be paid to corners, or areas of high curvature.

Some object properties are difficult to determine without touching them. For example, to determine weight, an individual needs to lift the object (Flanagan and Wing, 1997; Wing and Lederman, 1998). Information gained by haptic exploration does not need to be metrical. The exact properties attended to will depend on what sensorimotor apparatus the animal is equipped with, and on which features of objects or events are relevant to its ecological niche (e.g. psittacine visual fields and bill tip organ described in Part 1).

3.2. *Object categorisation*

To deal with the computational load of processing a dynamic environment (subsection 2.5), with a large amount of potentially redundant information (subsection 2.4), sometimes the animal may pursue a strategy of exploring the environment first, and then switching to exploiting it (e.g. Krebs et al., 1978). However, when other sources of uncertainty are involved, a learning system needs to also have good criteria for selecting environmental features to attend to.

Categorisation enables generalisation in novel conditions from a few experiences, while greatly reducing the computational complexity of perceiving objects and environmental processes. An enormously varied environment can be decomposed collections of object affordances and processes – i.e. exploration domains ('micro-domains' in Karmiloff-Smith, 1995). We suggest decomposition is achieved during exploration by perceptual and motor interactions with the environment (subsection 3.4). By concentrating on environmental

subsets and systematically varying its exploratory behaviour on it, the animal resolves what to attend to and limits the phenomena for which patterns are sought. As a prerequisite, the animal must start with perceptual mechanisms capable of detecting and recording the structures and motions produced by exploratory behaviours (subsection 3.1).

There have been many studies into what different animals can discriminate between (e.g. Kelman et al., 2008; Giret et al., 2009; Wills et al., 2009; Avargues-Weber et al., 2010; Soto and Wasserman, 2012; Wasserman et al., 2012). Pigeons can even tell the difference between art movements, something beyond many humans (Watanabe et al., 1995). However, they are likely rooted in quite simple same/different mechanisms, such as by differentiating between the variability in perceived thresholds of pairs of stimuli (Wright and Katz, 2006; Vermeulen et al., 2009; Wasserman and Young, 2010; Smith et al., 2011).

Alternatively, each object type may be categorised by the different features it possesses (Perone et al., 2008; Hammer et al., 2009). Cows and horses are both ungulates. They have four legs and a tail, but so does a dog; thus more detail is required to sub-categorise. This can form an progressively specific, branching structure, which the animal can build on through learning (2.2.2). This process is aided by a sensorimotor apparatus that especially attends to functional differences between object categories, rather than more aesthetic differences. This same apparatus would also attend more to perceived differences within a category (i.e. the unexpected), rather than between categories.

This kind of structure would allow for parallel processing at all levels, but there would also be interaction between different groups. This kind of idea has been developed through several neural network models (e.g. Quartz, 1999; Carruthers, 2006; Op de Beeck and Baker, 2010; Shanahan, 2012). The problem comes when there is a combinatorial explosion: too much information is stored in this structure that the system actually becomes slower. For instance, there is a huge search space involved in combining different perceptions, motor sequences, and exploratory behaviours to fit specific shapes, sizes, and object relations (Bellman, 1961; Perlovsky, 1998).

3.3. Combination of 'simple' learning mechanisms

When the environment is too variable, evolution cannot discover in advance suitable fixed responses to all needs in all situations (subsection 2.5). Instead, it provides mechanisms of learning and development that use information about the environment. There should be a selection of simple learning mechanisms, each specialised for particular exploration domains, which allow animals to attend to and learn about stimuli from restricted classes of environment (e.g. Karmiloff-Smith, 1995; subsection 3.2). This enables individuals to discover useful new actions, threats and opportunities. However, these discoveries must start from some initial motivations of some form.

The vast majority of the cognitive scientists are transfixed with arguing for each learning mechanism in isolation and do not consider them in tandem. We propose animals start exploring the world around them using a combination of different mechanisms, which includes (but not restricted to) associative learning (e.g. Rescorla, 1968), trial-and-error learning (e.g. Krueger and Dayan, 2009; Muller, 2010), and probabilistic learning, perhaps using some form of Bayesian networks (e.g. Spirites et al., 2000; Tenenbaum and Niyogi, 2003).

Initial exploratory behaviours may appear random, where the animal tries lots of different behaviours on lots of different things it knows little about (trial-and-error). However, when it starts to detect the invariances and regularities (through some crude type of probabilistic learning), it starts to narrow down the different possible causal relationships (of different object categories, subsection 3.2). This will be guided by the sensorimotor selectivities described in subsection 3.1.

Quite a large time investment is needed in these relatively simple learning mechanisms. It is important not to discount the impact of simple learning mechanisms in attaining quite sophisticated competences and behaviours (even for humans; Shanks, 2007). However, there is accumulating evidence about the involvement of top-down processes in human and non-human animals, such as causal reasoning and, more recently, probabilistic-driven causal inferences (subsection 4.2).

3.4. Behaviour structured to maximise information gain

Physical actions on the world are very important for the individual to reveal the processes present and objects' invisible affordances (Ravishankar et al., 2008). It is widely accepted human infants' understanding of object properties derives from both their exploratory behaviour and the information-processing systems generating and modifying their behaviour (Gibson, 1988; Gibson and Pick, 2003; Piaget, 1952; Rochat, 2001). In turn, the representations resulting from such activity alter and direct the actions infants perform on objects (Perone et al., 2008).

The combination of strategies used are likely dependent both on the environmental conditions and the individual's own competences. Children learn much quicker, for instance, when they explore objects themselves, rather than watching others (e.g. Fagard and Lockman, 2010; Heyes, 2011; Jacquet et al., 2012).

There has been much less work on the form and function of the information-gathering aspects of exploration in non-human animals (Kacelnik, 1987; Renner, 1990; Inglis, 1983; Inglis et al., 2001; Power, 2000)⁷. What we do know, suggests animals' sensorimotor behaviour acts to increase the quantity and quality of information gained. Many species show active information gathering. Rats alter the speed and pattern of their whisking behaviour to increase information about shape and texture of objects they contact with their vibrissae (Grant et al., 2009). This has been confirmed by modelling the behaviour in a robot (Pearson et al., 2007). The rats' whisking behaviour is 'designed' to increase the probability of detecting important environmental features.

Some of an animal's exploratory behaviours will involve only its body parts – whether limbs, eyes or the whole body (chapter 3). Others will also involve certain objects, where the individual may repeatedly grab, push, pull or twist the same thing. In more complex manipulations, there can be several objects and object parts involved – for instance in stacking or arranging them (for further

⁷This does not include studies of the current or future fitness benefits of the behaviour usually referred to as 'play' (Pellegrini et al., 2007; Bekoff and Byers, 1998; Held and Spinka, 2011).

examples see Appendix C). Whatever level of exploratory complexity, the behaviours would likely be focussed on any changes detected in the environment, especially in an object's affordances (function).

Different exploration domains are often interleaved, such as in alternating between eating and playing with food. This allows knowledge of different domains to develop roughly in parallel (Bushnell and Boudreau, 1993; subsection 5.3). When switching domains, the individual needs to be able to group bits of information together according to the current domain involved. For example, certain materials such as sticks may have one kind of affordance in the tool-using domain, but others when building a nest. For different species, the objects and their affordances will differ according to individuals' ecological niche, but there are probably some common exploration mechanisms across species (Sloman and Chappell, 2005; Chappell and Sloman, 2007).

It is often difficult to distinguish exploratory behaviour from executive action: is the animal lifting an object to transport it, or to learn its weight? Of course, it may fulfil both goals simultaneously (e.g. Elner and Hughes, 1978), but for researchers to determine when (or whether) an animal is simply collecting perceptual information, they need detailed knowledge about the extent of its sensory realm (Demery et al., 2011). We also need reliable behavioural or physiological 'markers' of exploratory behaviour in non-human animals (Section 6).

3.5. Preference for 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 animals' development (e.g. Heinrich, 1995; Pellegrini et al., 2007). By definition, animals do not have adequate information about novel objects, places and events, so they should prioritise their interaction with them. In particular, animals in dynamic environments might use exploration to experiment with strategies or behaviours in the current environmental context.

Ruff (1986) hypothesised that if the main function of 'examining' behaviour in human infants is to

gather information, it should:

1. decrease in frequency with exposure to a particular object;
2. and occur before other behaviours with new objects.

She found both hypotheses were upheld. Additionally, the latency and duration of examination indicated 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. Perone et al. (2008) presented infants with an image sequence, depicting a hand acting on a colourful toy that produced a sound, followed by either the action-sound pair or the object's appearance changing. They found infants attended more to changes in action than appearance. From an evolutionary perspective, it is not clear why such salience differences exist, but perhaps appearance changes are less likely to have important implications for the object's function.

4. Stage II: testing hypotheses

We propose in the second stage of exploration, the animal tests its internal hypotheses. Its exploratory behaviours are targeted more selectively, according to the specific object and environmental problem presented (subsection 4.3). The animal uses progressively more complex learning mechanisms like causal reasoning (subsection 4.2), so exploratory behaviours become less repetitive (unlike in subsection 3.3). Consequently it can cope with more complex problems, such as an apparatus with multiple object relations (e.g. Miyata et al., 2011). This is aided by a pre-existing biological framework of physical rules, particularly object solidity, continuity and connectedness (subsection 4.1). When its theories are violated, greater exploration ensues (subsection 4.4).

4.1. Sensitivity to physical rules

There are certain aspects of the physical world that can be regarded as constants (subsection 2.1). This includes gravity, the properties of contact, solidity, and biological movement or agency. Exploring animals probably harness these features as de-

faults ('basic physical rules') from birth or hatching. This enables them to have some basic understanding of how the world should work and how objects should interact.

Although somewhat contentious (e.g. Karmiloff-Smith, 1995), the extensive developmental work of Spelke and colleagues (summarised in Spelke, 2000) demonstrates human infants probably have innate systems representing objects, number and space. Developing alternative developmental methods should clarify these ideas further. This has begun with some non-human animal work on juveniles (e.g. Funk, 2002; Zucca et al., 2007; Bird and Emery, 2010). This has been extended to adult animals to show at least some species have such pre-dispositions (e.g. Hauser et al., 1999; Kundey et al., 2010; Cacchione and Call, 2010; Jaakkola et al., 2010; O'Connell and Dunbar, 2005, on solidity). These are fine-tuned and built upon with experience (Section 5).

There is likely, however, a distinction between individuals' implicit (subconscious) and explicit (conscious) understanding of a situation. There is a difference between noting a causal relation is wrong (e.g. a floating block) and predicting how it should be (see also subsection 4.4). Previously, this dichotomy has been differentiated in developmental psychology by presenting infants with impossible events, then measuring looking-times and where they search/explore in the test apparatus.

4.2. Combination of 'complex' learning mechanisms

As animals progress they likely use a selection of successively more complex learning mechanisms (e.g. 'causal reasoning'), coupled with simpler mechanisms (subsection 3.3). An animal may start out gathering information about a particular domain using probabilistic learning. When that no longer yields information, yet there is still a known gap in their knowledge (subsection 5.2), the information-processing system becomes more generative in some way. This allows it to notice deep invariances between examples, which previously were thought to share an object category (see also subsection 3.2 and subsection 4.4). The question is what process drives this internal transition.

Unlike trial-and-error and probabilistic learning, some prediction is possible without haptic exploration. Causal understanding of different struc-

tures' functions is made possible by reasoning through different possibilities (e.g. Blaisdell et al., 2006; Darredeau et al., 2009). An individual may utilise several concepts and mechanisms of causal reasoning, depending on the causal problem they face (e.g. canopy route planning; Tecwyn et al., 2012).

A formal, specific model is yet to be determined. A hierarchical structure may be the answer (subsection 2.2.2). The number of learning strategies used in a particular domain may determine the structure's number of levels. To activate a new level of a set of concepts, lower levels may need to be acquired first, through particular learning mechanisms. For example, relevant objects or relationships could be identified by probabilistic mechanisms that detect correlations (e.g. Chater et al., 2006). Once those correlations have been found, incorporating them as neural network nodes (e.g. Quartz, 1999; Marcus, 2003) allows new pattern learning. This combination of objects and patterns can feed an inference system (e.g. Ackerman and Others, 2004; subsection 5.1) for building new concepts. To varying extents, these mechanisms can be revealed through observing spontaneous exploratory behaviour (Section 6).

4.3. Behaviour more targeted and selective

Assuming the function of exploratory behaviour is to gather information, we would expect its form (and underlying mechanisms) to change with context (see also subsection 3.4). Each type of exploratory behaviour generates perceptual changes best suited to the sensory modality used (subsection 3.1), maximising opportunities to detect relevant features (subsection 2.4).

The type of manipulations human infants employ alter with how various toys' affordances change. When an object's texture changes, the looking and fingering of it increases (Ruff, 1984). In contrast, shape changes lead to increased rotation and transfer of the object between hands. Similarly, while looking at the object, infants are more likely to transfer it between hands, or finger its surface, but while mouthing it, they are more likely to rotate it (Ruff et al., 1992).

Renner and colleagues (Renner and Rosenzweig, 1986; Renner and Seltzer, 1994) showed rats employed different types of exploratory behaviours

when there was potential to gain more information. The relative frequencies, sequences and complexities of different exploratory behaviours varied according to previous experience and particular object characteristics. Heyser and Chemero (2012) showed mouse (*Mus musculus domesticus*) exploration levels and actions displayed depended on functional cues of a novel object, rather than purely visible, aesthetic cues.

4.4. Active 'testing' when expectations are violated

If an organism's current empirical observations do not fit with previously collected information, it should re-initiate exploration to resolve this discrepancy. There is increasing evidence human children use a conditional intervention principle to resolve discrepancies and learn about causes (Gopnik, 1996; Gopnik and Schulz, 2004; Tenenbaum et al., 2006). Their exploration appears to be systematic and sensitive to ambiguous information (see also subsection 5.2). Schulz et al. (2008) presented blocks of a certain category magnetically sticking to a board. When children found properties of new blocks varied within the category, they explored more.

There are several different, inter-linked processes here. The individual needs to detect some aspect of the world is surprising; then commence exploration, focussing on resolving this. The latter may involve re-organisation, or other changes in representations (e.g. see subsection 5.3). There is again much less literature on non-human animals. Some studies have manipulated certain environmental stimuli and measured exploratory behavioural sequences (Bekoff, 1975; Renner, 1990; Kuba et al., 2006). At least in rats and octopi (*Octopus vulgaris*), these revealed behaviour ostensibly similar to human children. As far as we know, there have been no studies into whether non-human animals spontaneously perform their own 'tests' like human children (Povinelli and Dunphy-Lelii, 2001; Buchsbaum et al., 2012). We predict non-human animals, like children, would become less repetitive in their exploratory actions as they develop – they would display greater diversity of exploratory behaviours, rather than repeating a few actions on the same part of the environment. This would allow them to learn what kinds of actions would be most effective in different situations.

The ability to deal with surprises and test unknown elements is limited by the potential of the learning mechanisms utilised and the form of the information representation (e.g. difficult with the simple mechanisms described in subsection 3.3). Some models, such as artificial neural networks, relying on a series of weighted associations, fall short in explaining this behaviour (e.g. Quartz, 1999). These models do not always make clear what associations the simulated individual is acquiring. However, logical, symbolic models allow researchers to produce demonstrations of the simulated individual's exploration pattern (e.g. Arriola-Rios et al., 2012).

5. Stage III: extending and refining hypotheses

In the final stage, we propose animals extend and refine their hypotheses. They analyse their theories and re-use information, combining them in related, but new, environmental situations (subsection 5.3). Individuals can now use causal inference to abstractly fill information gaps (subsection 5.1), extending hypotheses throughout life. This is especially important in a dynamic environment, where the animal should seek to test its current models, particularly when its expectations are violated (like in the previous stage; subsection 4.4). These abilities imply meta-cognitive mechanisms and an endogenous motivation to fill in the identified gaps of information (subsection 5.2).

5.1. Causal inference by abstracting information

Animals explore more efficiently if they are guided by previous knowledge of hidden object features to causally infer potential, new information. Information is usually acquired in a format of restricted use (subsection 3.3). If animals find generalisations by abstracting across different types of information, the information has greater use, particularly for solving causal problems. After subsequent experiences, to use the stored structures for specific functions, animals need to re-organise information into a new generative form. This form has wider scope, so it is both more economical and powerful. It is a deductive system for deriving novel conclusions (i.e. 'representational redescription'; Karmiloff-Smith, 1995).

A prominent example is the transition in human children from using empirically learnt words to generative syntax. An infinite number of sentences can now be generated and understood. Other animals demonstrate a simpler form of this cognitive flexibility to access food resources. They apply elements of existing knowledge about particular behaviours in one environmental context, to an entirely new context (e.g. keas and crows in Auer-sperg et al., 2011).

There is evidence that a number of taxa (e.g. apes, rats and dogs) can make causal inferences using multiple information sources (see also subsection 5.1). Some can harness information from the object itself (e.g. weight; Blaisdell et al., 2006; Call, 2004; Brauer et al., 2006; Hanus and Call, 2008), social cues (e.g. Povinelli et al., 1990), or by exclusion (Aust et al., 2008; Hill et al., 2011; Call, 2006). However, these experiments relied on animals observing the state of the world, or watching others perform actions on objects.

The mechanisms underlying abstraction processes are still debatable. Sidman (e.g. 2000) proposed several concepts (e.g. symmetry or transitivity) are acquired purely as a consequence of reinforcement contingencies. This has been supported by experiments on captive pigeons (*Columba livia*) using successive matching (Urcuioli, 2008). However, these experiments involved a small number of familiar, simple stimuli, presented under tightly constrained learning conditions. Whether they are ecologically valid, or can be extended to the field, is unclear. The natural world can offer a rich array of complex details and contributing factors (discussed in chapter 5), where consequences of actions need to be reasoned out, rather than retrieved from memory. What the reinforcement might be in such cases is not apparent, nor can the mechanisms proposed by Sidman (2000) readily explain a deductive system similar to that described above.

A flexible, abstractive system can re-organise and re-represent information from one domain to generalise others. This later allows for the original domain to become an object of exploratory attention. Individuals can then question issues such as what may occur within that domain (subsection 5.3). Such abstractive and causal inference systems aids information-gathering about invisible affordances or structures. Particularly important information may be several steps along from the initial

exploratory behaviour, or only become apparent once acted upon (e.g. object permanence concept with hidden object; see subsection 4.1).

5.2. *Meta-cognitive ability to recognise gaps in knowledge*

When there is a lack of information in the environment, animals would benefit from being able to track their own level of uncertainty (e.g. Inglis et al., 2001). For instance, when making a decision, an animal should be aware if it lacks sufficient information. There are many forms of meta-cognition (i.e. self-awareness; e.g. Karmiloff-Smith, 1995; Povinelli and Preuss, 1995), but in this chapter we refer specifically to self-directed meta-cognition. It has been difficult to establish evidence of this kind of meta-cognition in animals (for a review see Smith, 2009) and there are few biologically plausible, working models.

There may be some self-organising knowledge stores, reacting automatically to changes and new opportunities. In other cases, the information-processing system may possess a separate sub-system. This could monitor other sub-systems' behaviour and detect opportunities to initiate major re-organisation (e.g. Sussman, 1973). For example, when objects behave contrary to their appearance (as in subsection 4.4), the animal is stimulated to perform more exploration.

In Stage III, animals are capable of representing different scenarios in their past, present and future. After the results of an exploratory behaviour are observed, this capability together with meta-cognition/uncertainty, enables thoughts about what could have been ('counterfactual thinking'). Such abilities have again been described in human children (Beck et al., 2008; Byrne, 2002), but much less in other animals (e.g. Seed et al., 2012). Children explore more and target their behaviours (subsection 4.3) when faced with ambiguous information and their own uncertainty (summarised in Schulz, 2012).

Self-criticism mechanisms can further aid learning by improving problem-solving skills (Sussman, 1973; Sloman, 2008b). These mechanisms may initially be genetically-encoded, but their effects vary according to individuals' experience. Different forms of learning develop throughout life, influenced by the genome *and* the environment. For ex-

ample, humans learning about advanced mathematical concepts need to first develop various forms of representation, allowing understanding of progressively abstract structures and processes. These types of learning differ from earlier forms not just in their content, but in their structure.

5.3. *Extending knowledge by combining domains*

Combining old exploration domains can create new ones. One common simplification by animals is the discovery that two domains, involving different perceptual contents and affordances, can share structures and be unified into a useful, new abstraction. This can be applied to existing exploration domains and form new domains. For instance, in combining information about soil with information about water, to form knowledge about mud. Abstractions about physical rules can be applied to several exploration domains (e.g. cohesiveness; subsection 4.1).

Such combinations are made possible by several domains representing spatial structures and processes – behaviours originally performed at different locations or times, can later be performed together. This can lead to new forms of interaction (e.g. Taylor et al., 2009; Miyata et al., 2011). In some cases, what was previously learnt in separate domains, enable an animal to predict and reason about novel concepts. In other cases, more empirical learning is required, followed by meta-cognition (subsection 5.2) and new forms of theory re-construction (e.g. about mud properties).

It is not clear how the many environmental properties represented in adults as numerical measures (e.g. position, velocity, volume; Rheingold, 1985) can be represented in a young learner. Karmiloff-Smith (1995) emphasised how, in humans, sometimes competences gained at one developmental stage, can be later revised or transformed at another developmental stage.

It is difficult to determine whether this is due to internal re-organisation, or whether a separate meta-cognitive system is required. Cognitive robotics can reveal some alternative models, along with their implications, costs, and demonstrations of what is possible (e.g. Lopes and Oudeyer, 2010). We propose animals extend hypotheses throughout life, as suggested by the high level of neophilia in cognitively flexible and long-lived organisms

(e.g. parrots; Luescher, 2006).

6. Testing exploratory learning

By observing exploratory behaviour, we can gain an insight into how different animals gather information and learn about the world around them. It is challenging to explicitly discriminate between the different underlying learning mechanisms employed at any given moment. Instead cognitive scientists should use exploration to *frame* what an animal is doing and what types of information they may be collecting. Then we can begin to discuss how the exploratory behaviour observed under different conditions might fit into different learning mechanisms.

With this approach in mind, we have presented a general three-stage theory for how different animals can gather different kinds of information throughout their lives. The details likely vary between different species, individuals, exploratory domains and environmental situations. This theory was purposefully designed as a general framework, so it can be applied to a wide range of contexts. The theory is an example of how apparently different forms of thinking from different research fields (discussed in chapter 1) can be integrated to provide a fuller account of the research problem. This provides diverse opportunities for interdisciplinary collaboration and further systematic investigation.

For instance, in Arriola-Rios and Demery (2012), we explored how a parrot may internally represent different forms of objects' compliance, using techniques from artificial intelligence. A working preliminary model of object compliance, using specific physical rules and probabilistic learning, was compared with behavioural results collected from live animals (from experiments discussed in chapter 6)⁸. We showed how a selection of key elements from the environment (a form of categorisation) could be used as a basis for efficiently representing objects and their related processes (e.g. a sponge being squeezed).

⁸This was just a correlational comparison. In this study, we could not confirm if the observed parrot exploration was due to similar underlying mechanisms as those presented in the artificial model.

In the future, we hope to apply this same artificial model to a range of materials and problem-solving tasks (as in e.g. Arriola-Rios and Savage, 2007). A selection of 'built-in properties' that a 'virtual parrot' may have are described in Appendix E. This has been shown to work in a very basic simulated environment using StarLogo TNG version 1.5 (Colella et al., 2001). To check the ecological validity of these artificial models, they should be verified against a wider range of biological examples, under different experimental conditions.

The theory described in this chapter provides a framework for forming testable hypotheses. We chose to study two exemplar exploratory species; kakariki (*Cyanoramphus novaezelandiae*) and humans. There is much less literature on exploration in non-human animals, so in the kakariki tests we focused on the earlier stages of exploration (Section 3 and parts of Section 4). In the human tests, we focused on the later stages of exploration (Section 4 and Section 5), which have been neglected in the human exploration literature. Thus, in chapter 6 we hypothesise kakariki will explore more with:

- the corners and curves of objects over flat surfaces (due to points explained in subsection 3.1);
- increasing object complexity (due to points explained in subsection 3.4);
- novel over familiar objects (subsection 3.5);
- functional, invisible changes in an object over less functional or aesthetic changes, especially if the change occurs within-category (rather than between-category; subsection 3.2);
- unexpected changes in an object (i.e. if the haptic cues contradict the visual cues; subsection 4.4);
- extreme object categories first, but with time will explore the intermediate categories more (subsection 3.2 and subsection 3.4).

Further, in chapter 7 we hypothesise that human children will explore more:

- functional/action changes in an object than aesthetic changes (like the kakariki; due to point explained in subsection 3.2);

- when a change in a physical rule has occurred (subsection 4.1), and the initial exploratory behaviour will be directed at the change's source (subsection 4.2), rather than the (usually now functionless) object area previously habituated to (subsection 3.3);
 - when there is missing or ambiguous information (and causal inferences are needed; subsection 5.1 and subsection 5.2);
 - when older (in terms of behavioural diversity; subsection 4.3) and there will be a developmental shift across different exploratory domains (subsection 5.3).
- Ackerman, B. P., Others, A., Oct. 2004. Building a story representation: the effects of early concept prominence on later causal inferences by children. *Developmental Psychology* 27 (3), 370–80. 4.2
- Albert, D. J., Mah, C. J., 1972. An examination of conditioned reinforcement using a one-trial learning procedure. *Learning and Motivation* 3, 369–388. 2.1
- Alkon, D. L., Lederhendler, I., Shoukimas, J., Feb. 1982. Primary changes of membrane currents during retention of associative learning. *Science* 215 (4533), 693–695. 2.2
- Arriola-Rios, V. E., Demery, Z. P., 2012. Salient features and snapshots in time: an interdisciplinary perspective on object representation. In: *Proceedings of the Symposium on Natural/Unconventional Computing and its Philosophical Significance at the AISB/IACAP World Congress 2012*. pp. 1–6. 6
- Arriola-Rios, V. E., Demery, Z. P., Wyatt, J., Sloman, A., Chappell, J., Jul. 2012. Salient features and snapshots in time: an interdisciplinary perspective on object representation. In: Crnkovic, G. D., Giovagnoli, R. (Eds.), *Studies in Applied Philosophy, Epistemology and Rational Ethics: Computing Nature*. Springer, Berlin, pp. 1–14. (document), 4.4
- Arriola-Rios, V. E., Savage, J., 2007. Knowledge Acquisition and Automatic Generation of Rules for the Inference Machine CLIPS. In: Gelbukh, A. F., Kuri Morales, A. F. (Eds.), *MICA 2007: Advances in Artificial Intelligence*, 6th Mexican International Conference on Artificial Intelligence. Springer, Aguascalientes, Mexico, pp. 725–735. 2.2.2, 6
- Auersperg, A., von Bayern, A. M. P., Gajdon, G. K., Huber, L., Kacelnik, A., 2011. Flexibility in problem solving and tool use of kea and New Caledonian crows in a multi access box paradigm. *PloS One* 6 (6), e20231. 2.5, 5.1
- Aust, U. U., Range, F., Steurer, M., Huber, L., Oct. 2008. Inferential reasoning by exclusion in pigeons, dogs, and humans. *Animal Cognition* 11 (4), 587–597. 5.1
- Avargues-Weber, A., Portelli, G., Benard, J., Dyer, A. G., Giurfa, M., Feb. 2010. Configural processing enables discrimination and categorization of face-like stimuli in honeybees. *Journal of Experimental Biology* 213 (4), 593–601. 3.2
- Beck, S. R., Robinson, E. J., Freeth, M., Apr. 2008. Can children resist making interpretations when uncertain? *Journal of Experimental Child Psychology* 99 (4), 252–270. 5.2
- Bekoff, M., 1975. Animal play and behavioral diversity. *American Naturalist* 109 (969), 601–603. 4.4
- Bekoff, M., Byers, J. A., 1998. *Animal Play: Evolutionary, Comparative, and Ecological Perspectives*. Cambridge University Press, Cambridge, UK. 7
- Bellman, R. E., 1961. *Adaptive control processes: a guided tour*. Princeton University Press, Princeton, NJ. 3.2
- Ben-Jacob, E., Oct. 2009. Learning from bacteria about natural information processing. *Annals of the New York Academy of Sciences* 1178, 78–90. 5
- Berlyne, D. E., 1960. *Conflict, arousal, and curiosity*. McGraw-Hill Book Company, New York, NY, US. 1
- Bird, C. D., Emery, N. J., Jan. 2010. Rooks perceive support relations similar to six-month-old babies. *Proceedings of the Royal Society of London B: Biological Sciences* 277 (1678), 147–151. 4.1
- Blaisdell, A. P., Sawa, K., Leising, K. J., Waldmann, M. R., Jan. 2006. Causal reasoning in rats. *Science* 311 (5763), 1020–1022. 4.2, 5.1
- Brauer, J., Kaminski, J., Riedel, J., Call, J., Tomasello, M., Feb. 2006. Making inferences about the location of hidden food: social dog, causal ape. *Journal of Comparative Psychology* 120 (1), 38–47. 5.1
- Buchsbaum, D. D., Bridgers, S. S., Weisberg, D. D. S., Gopnik, A., Sep. 2012. The power of possibility: causal learning, counterfactual reasoning, and pretend play. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (1599), 2202–2212. 4.4
- Bushnell, E. W., Boudreau, J. P., Aug. 1993. Motor development and the mind: the potential role of motor abilities as a determinant of aspects of perceptual development. *Child Development* 64 (4), 1005–1021. 3.4
- Byrne, R. M. J., Oct. 2002. Mental models and counterfactual thoughts about what might have been. *Trends in Cognitive Sciences* 6 (10), 426–431. 5.2
- Cacchione, T., Call, J., 2010. Intuitions about gravity and solidity in great apes: the tubes task. *Developmental Science* 13 (2), 320–330. 4.1
- Call, J., Jun. 2004. Inferences about the location of food in the great apes (*Pan paniscus*, *Pan troglodytes*, *Gorilla gorilla*, and *Pongo pygmaeus*). *Journal of Comparative Psychology* 118 (2), 232–241. 5.1
- Call, J., Oct. 2006. Inferences by exclusion in the great apes: the effect of age and species. *Animal Cognition* 9 (4), 393–403. 5.1
- Calvo, P., Gomila, T., Gomila, A., 2008. *Handbook of Cognitive Science: An Embodied Approach*. Vol. New York. Elsevier, London. 1
- Carruthers, P., 2006. *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought: massive modularity and the flexibility of thought*. Clarendon Press, London. 3.2
- Chappell, J., Demery, Z. P., Arriola-Rios, V. E., Sloman, A., Feb. 2012. How to build an information gathering and processing system: lessons from naturally and artificially intelligent systems. *Behavioural Processes* 89 (2), 179–186. (document), 2
- Chappell, J., Sloman, A., 2007. Natural and Artificial Meta-Configured Altricial Information-Processing Systems. *International Journal of Unconventional Computing* 2 (3), 211–239. 1, 2.5, 3.4
- Chater, N., Tenenbaum, J. B., Yuille, A., Jul. 2006. Probabilistic models of cognition: conceptual foundations. *Trends in Cognitive Sciences* 10 (7), 287–291. 4.2
- Colella, V., Klopfer, E., Resnick, M., Jun. 2001. *Adventures in Modeling: Exploring Complex, Dynamic Systems with StarLogo*. Teachers College Press, London. 6
- Cook, C., Goodman, N. D., Schulz, L. E., Sep. 2011. Where

- science starts: Spontaneous experiments in preschoolers exploratory play. *Cognition* 120 (3), 341–349. 2.6
- Darredeau, C., Baetu, I., Baker, A. G., Murphy, R. A., Jan. 2009. Competition between multiple causes of a single outcome in causal reasoning. *Journal of Experimental Psychology: Animal Behavior Processes* 35 (1), 1–14. 4.2
- Demery, Z. P., Arriola-Rios, V. E., Sloman, A., Wyatt, J., Chappell, J., 2010. Construct to understand: learning through exploration. In: *Proceedings of the International Symposium on AI-Inspired Biology*. pp. 59–61. (document), 1
- Demery, Z. P., Chappell, J., Martin, G. R., Dec. 2011. Vision, touch and object manipulation in Senegal parrots *Poicephalus senegalus*. *Proceedings of the Royal Society of London B: Biological Sciences* 278 (1725), 3687–3693. 1, 3.4
- Elner, R., Hughes, R. N., 1978. Energy maximization in the diet of the shore crab, *Carcinus maenas*. *Journal of Animal Ecology* 47 (1), 103–116. 3.4
- Fabre, J. H., de Mattos, A. T., Miall, B., 1918. *The wonders of instinct*. The Century Co., Minnesota. 2.1
- Fagard, J., Lockman, J., Jan. 2010. Change in imitation for object manipulation between 10 and 12 months of age. *Developmental Psychobiology* 52 (1), 90–99. 3.4
- Flanagan, J. R., Wing, A. M., Feb. 1997. The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *Journal of Neuroscience* 17 (4), 1519–1528. 3.1
- Floreano, D., Mattiussi, C., 2008. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. MIT Press, Cambridge, MA. 3.1
- Funk, M. S., Sep. 2002. Problem solving skills in young yellow-crowned parakeets (*Cyanoramphus auriceps*). *Animal Cognition* 5 (3), 167–176. 4.1
- Gibson, E. J., Jan. 1988. Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual Reviews of Psychology* 39 (1), 1–42. 1, 2.6, 3.1, 3.4
- Gibson, E. J., Pick, A., 2003. *An Ecological Approach to Perceptual Learning and Development*. Oxford University Press, New York. 3.4
- Gibson, J. J., Nov. 1962. Observations on active touch. *Psychological Review* 69, 477–491. 1
- Gibson, J. J., 1977. The theory of affordances. In: Shaw, R., Bransford, J. (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology*. Erlbaum, Hillsdale, NJ, pp. 67–82. 3.1
- Giret, N., Peron, F., Nagle, L., Kreutzer, M., Bovet, D., Nov. 2009. Spontaneous categorization of vocal imitations in African grey parrots (*Psittacus erithacus*). *Behavioural Processes* 82 (3), 244–248. 3.2
- Glickman, S. E., Sroges, R. W., 1966. Curiosity in zoo animals. *Behaviour* 26 (1), 151–188. 1
- Gopnik, A., Dec. 1996. The Scientist as Child. *Philosophy of Science* 63 (4), 485–514. 4.4
- Gopnik, A., Schulz, L. E., Aug. 2004. Mechanisms of theory formation in young children. *Trends in Cognitive Sciences* 8 (8), 371–377. 4.4
- Grant, R. A., Mitchinson, B., Fox, C. W., Prescott, T. J., Feb. 2009. Active touch sensing in the rat: anticipatory and regulatory control of whisker movements during surface exploration. *Journal of Neurophysiology* 101 (2), 862–874. 3.4
- Greenberg, R., Mettke-Hofmann, C., 2001. Ecological aspects of neophobia and neophilia in birds. In: Nolan, V., Thompson, C. F. (Eds.), *Current Ornithology*. Springer, Boston, MA, pp. 119–178. 2, 3.5
- Hammer, R., Diesendruck, G., Weinshall, D., Hochstein, S., Jun. 2009. The development of category learning strategies: What makes the difference? *Cognition* 112 (1), 105–119. 3.2
- Hanus, D., Call, J., May 2008. Chimpanzees infer the location of a reward on the basis of the effect of its weight. *Current Biology* 18 (9), R370–R372. 5.1
- Hauser, M. D., Kralik, J., Botto-Mahan, C., 1999. Problem solving and functional design features: experiments on cotton-top tamarins, *Saguinus oedipus oedipus*. *Animal Behaviour* 57, 565–582. 4.1
- Hawes, N., Hanheide, M., Sjo, K., Aydemir, A., Jensfelt, P., Gobelbecker, M., Brenner, M., Zender, H., Lison, P., Kruijff-Korabayova, I., Kruijff, G. J. M., Zillich, M., May 2010. Dora the Explorer: a motivated robot. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, pp. 1617–1618. 1
- Heinrich, B., 1995. Neophilia and exploration in juvenile common ravens, *Corvus corax*. *Animal Behaviour* 50 (3), 695–704. 3.5
- Held, S., Spinka, M., 2011. Animal play and animal welfare. *Animal Behaviour* 81, 891–899. 7
- Heyes, C., Sep. 2011. What's social about social learning? *Journal of Comparative Psychology* 126 (2), 193–202. 3.4
- Heyser, C. J., Chemo, A., Mar. 2012. Novel object exploration in mice: not all objects are created equal. *Behavioural Processes* 89 (3), 232–238. 4.3
- Hill, A., Collier-Baker, E., Suddendorf, T., Feb. 2011. Inferential reasoning by exclusion in great apes, lesser apes, and spider monkeys. *Journal of Comparative Psychology* 125 (1), 91–103. 5.1
- Horner, V., Whiten, A., Jun. 2005. Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). *Animal Cognition* 8 (3), 164–181. 2.4
- Houston, A. I., Krebs, J. R., Erichsen, J. T., Jan. 1980. Optimal prey choice and discrimination time in the great tit (*Parus major L.*). *Behavioral Ecology and Sociobiology* 6 (3), 169–175. 2.5
- Huber, L., Gajdon, G. K., Sep. 2006. Technical intelligence in animals: the kea model. *Animal Cognition* 9 (4), 295–305. 2.5
- Inglis, I. R., Nov. 1983. Towards a cognitive theory of exploratory behaviour. In: Archer, J., Birke, L. I. A. (Eds.), *Exploration in Animals and Humans: comparative psychology of animals and humans*. Van Nostrand Reinhold, Workingham, pp. 72–117. 1, 3.4
- Inglis, I. R., Langton, S., 2006. How an animal's behavioural repertoire changes in response to a changing environment: a stochastic model. *Behaviour* 143, 1563–1596. 1
- Inglis, I. R., Langton, S., Forkman, B., Lazarus, J., Sep. 2001. An information primacy model of exploratory and foraging behaviour. *Animal Behaviour* 62 (3), 543–557. 3.4, 5.2
- Jaakkola, K., Guarino, E., Rodriguez, M., Erb, L., Trone, M., Jan. 2010. What do dolphins (*Tursiops truncatus*) understand about hidden objects? *Animal Cognition* 13 (1), 103–120. 4.1
- Jacquet, P. O., Chambon, V., Borghi, A. M., Tessari, A., 2012. Object Affordances Tune Observers' Prior Expectations about Tool-Use Behaviors. *PLoS One* 7 (6), e39629. 3.4
- Kacelnik, A., Jun. 1987. Information primacy or preference for familiar foraging techniques? A critique of Inglis & Ferguson. *Animal Behaviour* 35 (3), 925–926. 3.4
- 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. 2.5

- Kacelnik, A., Todd, I. A., Feb. 1992. Psychological mechanisms and the Marginal Value Theorem: effect of variability in travel time on patch exploitation. *Animal Behaviour* 43 (2), 313–322. 2.5
- Karmiloff-Smith, A., 1995. *Beyond Modularity: A Developmental Perspective on Cognitive Science*, 4th Edition. MIT Press, London. 2.2.2, 3.2, 3.3, 4.1, 5.1, 5.2, 5.3
- Kelman, E. J., Osorio, D., Baddeley, R. J., May 2008. A review of cuttlefish camouflage and object recognition and evidence for depth perception. *Journal of Experimental Biology* 211 (11), 1757–1763. 3.2
- Krebs, J. R., Kacelnik, A., Taylor, P., Sep. 1978. Test of optimal sampling by foraging great tits. *Nature* 275 (5675), 27–31. 3.2
- Krueger, K. A., Dayan, P., Mar. 2009. Flexible shaping: how learning in small steps helps. *Cognition* 110 (3), 380–394. 3.3
- Kuba, M. J., Byrne, R. M. J., Meisel, D. V., Mather, J. A., Jul. 2006. When do octopuses play? Effects of repeated testing, object type, age, and food deprivation on object play in *Octopus vulgaris*. *Journal of Comparative Psychology* 120 (3), 184–190. 4.4
- Kundey, S. M. A., De Los Reyes, A., Taglang, C., Baruch, A., German, R., May 2010. Domesticated dogs' (*Canis familiaris*) use of the solidity principle. *Animal Cognition* 13 (3), 497–505. 4.1
- Loewenstein, G., 1994. The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin* 116 (1), 75–98. 1
- Loomis, J. M., Beall, A. C., Jan. 1998. Visually controlled locomotion: its dependence on optic flow, three-dimensional space perception, and cognition. *Ecological Psychology* 10 (3), 271–285. 3.1
- Lopes, M., Oudeyer, P.-Y., 2010. Guest Editorial: Active learning and intrinsically motivated exploration in robots: advances and challenges. In: *IEEE Transactions on Autonomous Mental Development*. pp. 65–69. 5.3
- Luescher, A. U., 2006. *Manual Of Parrot Behavior*. Blackwell, San Francisco. 5.3
- Magnani, L., 2009. Abduction, affordances, and cognitive niches. In: *Abductive Cognition*. Springer, Berlin, pp. 317–359. 1
- Marcus, G. F., 2003. *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. MIT Press, London. 4.2
- Markram, H., Feb. 2006. The blue brain project. *Nature Reviews Neuroscience* 7 (2), 153–160. 1
- McCrone, J., 2006. Smarter than the average bug. *New Scientist* 190 (2553), 37–39. 1
- Merleau-Ponty, M., 2002. *Phenomenology of Perception*. Routledge, London. 3.1
- Mesulam, M., Jun. 1998. From sensation to cognition. *Brain* 121 (6), 1013–1052. 1
- Mettke-Hofmann, C., Winkler, H., Leisler, B., 2002. The significance of ecological factors for exploration and neophobia in parrots. *Ethology* 108 (3), 249–272. 2, 3.5
- Miller, N. E., Dollard, J., 1941. *Social Learning and Imitation*. International Library of Sociology. Yale University Press, New Haven, CT. 1
- Miyata, H., Gajdon, G. K., Huber, L., Fujita, K., Jan. 2011. How do keas (*Nestor notabilis*) solve artificial-fruit problems with multiple locks? *Animal Cognition* 14 (1), 45–58. 4, 5.3
- Muller, C. A., Mar. 2010. Do anvil-using banded mongooses understand means-end relationships? A field experiment. *Animal Cognition* 13 (2), 325–330. 3.3
- O'Connell, S., Dunbar, R. I. M., Jan. 2005. The perception of causality in chimpanzees (*Pan spp.*). *Animal Cognition* 8 (1), 60–66. 4.1
- Op de Beeck, H. P., Baker, C. L., Jan. 2010. The neural basis of visual object learning. *Trends in Cognitive Sciences* 14 (1), 22–30. 3.2
- Pearson, M. J., Pipe, A. G., Melhuish, C., Mitchinson, B., Prescott, T. J., Sep. 2007. Whiskerbot: a robotic active touch system modeled on the rat whisker sensory system. *Adaptive Behavior* 15 (3), 223–240. 3.4
- Pellegrini, A. D., Dupuis, D., Smith, P. K., Jun. 2007. Play in evolution and development. *Developmental Review* 27 (2), 261–276. 7, 3.5
- Penn, D. C., Holyoak, K. J., Povinelli, D. J., Apr. 2008. Darwin's mistake: explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences* 31 (2), 109–178. 1
- Perlovsky, L. I., 1998. Conundrum of combinatorial complexity. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20 (6), 666–670. 3.2
- Perone, S. S., Madole, K. L., Ross-Sheehy, S., Carey, M., Oakes, L. M., Aug. 2008. The relation between infants' activity with objects and attention to object appearance. *Developmental Psychology* 44 (5), 1242–1248. 3.2, 3.4, 3.5
- Piaget, J., 1952. *The Origins of Intelligence in Children*. International Universities Press, Michigan. 3.4
- Povinelli, D. J., Dunphy-Lelii, S., Jun. 2001. Do chimpanzees seek explanations? Preliminary comparative investigations. *Canadian Journal of Experimental Psychology* 55 (2), 185–193. 4.4
- Povinelli, D. J., Nelson, K., Boysen, S. T., Sep. 1990. Inferences about guessing and knowing by chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology* 104 (3), 203–210. 5.1
- Povinelli, D. J., Preuss, T. M., Aug. 1995. Theory of mind: evolutionary history of a cognitive specialization. *Trends in Neurosciences* 18 (9), 418–424. 5.2
- Power, T. G., 2000. *Play and Exploration in Children and Animals*, 1st Edition. Psychology Press, London. 3.4
- Qin, J., Wheeler, A. R., Feb. 2007. Maze exploration and learning in *C. elegans*. *Lab on a Chip* 7 (2), 186–192. 5
- Quartz, S. R., Feb. 1999. The constructivist brain. *Trends in Cognitive Sciences* 3 (2), 48–57. 3.2, 4.2, 4.4
- Ravishankar, S., Jain, A., Mittal, A., 2008. Multi-stage contour based detection of deformable objects. In: *Computer Vision ECCV 2008*. pp. 483–496. 3.4
- Renner, M. J., 1990. Neglected aspects of exploratory and investigatory behavior. *Psychobiology* 18 (1), 16–22. 1, 3.4, 4.4
- Renner, M. J., Rosenzweig, M. R., 1986. Object interactions in juvenile rats (*Rattus norvegicus*): effects of different experiential histories. *Journal of Comparative Psychology* 100 (3), 229–236. 4.3
- Renner, M. J., Seltzer, C. P., 1994. Sequential structure in behavioral components of object investigation by Long-Evans rats (*Rattus norvegicus*). *Journal of Comparative Psychology* 108 (4), 335. 4.3
- Rescorla, R. A., Aug. 1968. Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology* 66 (1), 1–5. 3.3
- Rheingold, H. L., 1985. Development as the acquisition of familiarity. *Annual Review of Psychology* 36, 1–18. 5.3
- Rochat, P., 2001. *The Infant's World*. Harvard University Press, Cambridge, MA. 1, 3.4
- Rosenheim, J. A., 1987. Nesting behavior and bionomics of a solitary ground-nesting wasp, *Ammophila dysmica* (Hymenoptera: Sphecidae): influence of parasite pressure. In: *Annals of the Entomological Society of America*. Entomological Society of America, pp. 739–749. 2.1
- Ruff, H. A., 1984. Infants' manipulative exploration of objects:

- Effects of age and object characteristics. *Developmental Psychology* 20 (1), 9–20. 4.3
- Ruff, H. A., Feb. 1986. Components of attention during infants' manipulative exploration. *Child Development* 57 (1), 105–114. 3.5
- Ruff, H. A., Saltarelli, L. M., Capozzoli, M., Dubiner, K., 1992. The differentiation of activity in infants' exploration of objects. *Developmental Psychology* 28 (5), 851–861. 4.3
- Schulz, L. E., Jul. 2012. The origins of inquiry: inductive inference and exploration in early childhood. *Trends in Cognitive Sciences* 16 (7), 382–389. 5.2
- Schulz, L. E., Bonawitz, E. B., Jun. 2007. Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology* 43 (4), 1045–1050. 1
- Schulz, L. E., Sommerville, J., 2006. God does not play dice: Causal determinism and preschoolers' causal inferences. *Child Development* 77 (2), 427–442. 1
- Schulz, L. E., Standing, H. R., Bonawitz, E. B., Aug. 2008. Word, thought, and deed: the role of object categories in children's inductive inferences and exploratory play. *Developmental Psychology* 44 (5), 1266–1276. 4.4
- Seed, A. M., Seddon, E. E., Greene, B. B., Call, J., Oct. 2012. Chimpanzee 'folk physics': bringing failures into focus. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (1603), 2743–2752. 5.2
- Shams, L., Seitz, A. R., Nov. 2008. Benefits of multisensory learning. *Trends in Cognitive Sciences* 12 (11), 411–417. 1
- Shanahan, M., Oct. 2012. The brain's connective core and its role in animal cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (1603), 2704–2714. 3.2
- Shanks, D. R. R., Mar. 2007. Associationism and cognition: human contingency learning at 25. *Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology* 60 (3), 291–309. 3.3
- Shaw, J. M., 2008. Unifying Perception and Curiosity. Ph.D. thesis, University of Rochester, Department of Computer Science, The College Arts and Sciences. Rochester, New York. 2
- Shettleworth, S. J., May 1993. Where is the comparison in comparative cognition? *Alternative research programs. Psychological Science* 4 (3), 179–184. 1
- Shettleworth, S. J., 1999. *Cognition, Evolution, and Behavior*. Oxford University Press, New York. 2.2
- Sidman, M., Jul. 2000. Equivalence relations and the reinforcement contingency. *Journal of the Experimental Analysis of Behavior* 74 (1), 127–146. 5.1
- Sloman, A., Dec. 2008a. The well-designed young mathematician. *Artificial Intelligence* 172 (18), 2015–2034. 1
- Sloman, A., Jul. 2008b. Varieties of Meta-cognition in Natural and Artificial Systems Some pressures on design-space from niche-space I apologise For using Linux and L. Tech. rep., University of Birmingham, Cambridge, MA. 5.2
- Sloman, A., 2010. The design-based approach to the study of mind (in humans, other animals, and machines), including the study of behaviour involving mental processes. In: *Proceedings of the International Symposium on AI-Inspired Biology*. pp. 87–94. 2.5
- Sloman, A., Jan. 2011. What's information, for an organism or intelligent machine? How can a machine or organism mean? . In: *Dodig-Crnkovic, G., Burgin, M. G. (Eds.), Information and Computation*. World Scientific Publishing Co. Pte. Ltd., London, pp. 1–32. 1
- Sloman, A., Chappell, J., 2005. Altricial self-organising information-processing systems. *AISB Quarterly* 121, 5–7. 3.4
- Smith, J. D., Sep. 2009. The study of animal metacognition. *Trends in Cognitive Sciences* 13 (9), 389–396. 5.2
- Smith, J. D., Ashby, F. G., Berg, M. E., Murphy, M. S., Spiering, B., Cook, R. G., Grace, R. C., Apr. 2011. Pigeons' categorization may be exclusively nonanalytic. *Psychonomic Bulletin and Review* 18 (2), 414–421. 3.2
- Soto, F. A., Wasserman, E. A., Mar. 2012. Visual object categorization in birds and primates: integrating behavioral, neurobiological, and computational evidence within a "general process" framework. *Cognitive, Affective and Behavioral Neuroscience* 12 (1), 220–240. 3.2
- Spelke, E. S., 2000. Core knowledge. *American Psychologist* 55 (11), 1233–1243. 4.1
- Spirtes, P., Glymour, C. N., Scheines, R., 2000. *Causation, Prediction, and Search*. MIT Press. 3.3
- Srinivasan, M., Gregory, R. L., 1992. How Bees Exploit Optic Flow: Behavioural Experiments and Neural Models [and Discussion]. *Philosophical Transactions of the Royal Society B: Biological Sciences* 337 (1281), 253–259. 3.1
- Sterelny, K., Apr. 2007. Social intelligence, human intelligence and niche construction. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362 (1480), 719–730. 2.5
- Sussman, G. J., 1973. *A Computational Model of Skill Acquisition*. Ph.D. thesis, Elsevier Science Inc., Boston. 5.2
- Taylor, A. H., Hunt, G. R., Medina, F. S., Gray, R. D., Jan. 2009. Do New Caledonian crows solve physical problems through causal reasoning? *Proceedings of the Royal Society of London B: Biological Sciences* 276 (1655), 247–254. 1, 5.3
- Tecwyn, E. C., Thorpe, S. K. S., Chappell, J., Jan. 2012. What cognitive strategies do orangutans (*Pongo pygmaeus*) use to solve a trial-unique puzzle-tube task incorporating multiple obstacles? *Animal Cognition* 15 (1), 121–133. 4.2
- Tenenbaum, J. B., Griffiths, T. L., Kemp, C., Jul. 2006. Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences* 10 (7), 309–318. 4.4
- Tenenbaum, J. B., Niyogi, S., 2003. Learning causal laws. In: *Proceedings of the 25th annual meeting of the cognitive science society*. Citeseer, pp. 1152–1157. 3.3
- Thomas, R. K., Jun. 1996. Investigating cognitive abilities in animals: unrealized potential. *Cognitive Brain Research* 3 (3-4), 157–166. 1
- Timmis, J., Hart, E., Hone, A., Neal, M., Robins, A., Stepney, S., Tyrrell, A., 2008. Immuno-engineering. In: *Hinchey, M., Pagnoni, A., Rammig, F., Schmeck, H. (Eds.), IFIP International Federation for Information Processing*. Springer, Boston, MA, pp. 3–17. 1
- Toates, F. M., 1983. Exploration as a motivational and learning system: a cognitive incentive view. In: *Archer, J., Birke, L. I. A. (Eds.), Exploration in animals and humans*. Van Nostrand Reinhold, Berkshire, England, pp. 55–71. 1
- Urcioli, P. J., Nov. 2008. Associative symmetry, antisymmetry, and a theory of pigeons' equivalence-class formation. *Journal of the Experimental Analysis of Behavior* 90 (3), 257–282. 5.1
- Vermeulen, N., Mermillod, M., Godefroid, J., Corneille, O., Sep. 2009. Unintended embodiment of concepts into percepts: sensory activation boosts attention for same-modality concepts in the attentional blink paradigm. *Cognition* 112 (3), 467–472. 3.2
- Visalberghi, E., Tomasello, M., 1998. Primate causal understanding in the physical and psychological domains. *Behavioural Processes* 42 (2-3), 189–203. 1
- von Bayern, A. M. P., Heathcote, R. J., Rutz, C., Kacelnik, A., Dec. 2009. The role of experience in problem solving and innovative tool use in crows. *Current Biology* 19 (22), 1965–

1968. 2.5
- Wasserman, E. A., Castro, L., Freeman, J. H., 2012. Same-different categorization in rats. *Learning and Memory* 19 (4), 142–145. 3.2
- Wasserman, E. A., Young, M. E., Jan. 2010. Same-different discrimination: the keel and backbone of thought and reasoning. *Journal of Experimental Psychology: Animal Behavior Processes* 36 (1), 3–22. 3.2
- Watanabe, S. S., Sakamoto, J. J., Wakita, M., Mar. 1995. Pigeons' discrimination of paintings by Monet and Picasso. *Journal of the Experimental Analysis of Behavior* 63 (2), 165–174. 3.2
- Webb, B., Dec. 2001. Can robots make good models of biological behaviour? *Behavioral and Brain Sciences* 24 (6), 1033–50–discussion 1050–94. 1
- White, R. W., 1959. Motivation reconsidered: The concept of competence. *Psychological Review* 66 (5), 297–333. 1
- Wills, A. J., Lea, S. E. G., Leaver, L. A., Osthaus, B., Ryan, C. M. E., Suret, M. B., Bryant, C. M. L., Chapman, S., Millar, L., Nov. 2009. A comparative analysis of the categorization of multidimensional stimuli: I. Unidimensional classification does not necessarily imply analytic processing; evidence from pigeons (*Columba livia*), squirrels (*Sciurus carolinensis*), and humans (*Homo sapiens*). *Journal of Comparative Psychology* 123 (4), 391–405. 3.2
- Wilz, K. J., Bolton, R. L., 1971. Exploratory behavior in response to the spatial rearrangement of familiar stimuli. *Psychonomic Science* 24 (3), 117–118. 2.1
- Wimpenny, J. H., Weir, A. S., Clayton, L., Rutz, C., Kacelnik, A., 2009. Cognitive processes associated with sequential tool use in New Caledonian crows. *PloS One* 4 (8), e6471–e6471. 2.2.2
- Wing, A. M., Lederman, S. J., Dec. 1998. Anticipating load torques produced by voluntary movements. *Journal of Experimental Psychology: Human Perception and Performance* 24 (6), 1571–1581. 3.1
- Wright, A., Katz, J. S., Jun. 2006. Mechanisms of same/different concept learning in primates and avians. *Behavioural Processes* 72 (3), 234–254. 3.2
- Zucca, P., Milos, N., Vallortigara, G., Apr. 2007. Piagetian object permanence and its development in Eurasian jays (*Garrulus glandarius*). *Animal Cognition* 10 (2), 243–258. 4.1