

Slightly modified version of a paper that appeared in *Language Sciences* 26(6): 661-692, 2004. *Distributed Cognition and Integrational Linguistics*, special edition.

Self-directedness, integration and higher cognition

Wayne Christensen

Christensen@ukzn.ac.za

Philosophy, University of KwaZulu-Natal, South Africa

Abstract

In this paper I discuss connections between self-directedness, integration and higher cognition. I present a model of self-directedness as a basis for approaching higher cognition from a situated cognition perspective. According to this model increases in sensorimotor complexity create pressure for integrative higher order control and learning processes for acquiring information about the context in which action occurs. This generates complex articulated abstractive information processing, which forms the major basis for higher cognition. I present evidence that indicates that the same integrative characteristics found in lower cognitive process such as motor adaptation are present in a range of higher cognitive process, including conceptual learning. This account helps explain situated cognition phenomena in humans because the integrative processes by which the brain adapts to control interaction are relatively agnostic concerning the source of the structure participating in the process. Thus, from the perspective of the motor control system using a tool is not fundamentally different to simply controlling an arm.

Keywords: situated cognition, learning, executive cognition, evolution

1 Introduction: self-directedness and higher cognition

The account of self-directedness presented here¹ is a conceptual model of some generic characteristics of the relationship between lower cognition, understood as basic sensory and motor control processes, and high order cognition, understood as including conceptual and strategic reasoning. The model does not attempt to explain specific human abilities such as language or 'theory of mind', but focuses instead on the relationship between executive cognition and learning processes. I argue that the more general cognitive processes characterised by the model are central to higher cognition, and very likely play an important role in the kinds of language phenomena that interest

¹ The account of self-directedness was originally developed in collaboration with Cliff Hooker. See Christensen and Hooker, 2000; 2002. It has also been influenced by the work of Mark Bickhard.

integrationist linguistics. The model has been developed from what can be broadly characterised as a situated cognition perspective, however it follows a general constructivist schema with influences that date as far back as pragmatism (James, 1997) and William Whewell (Butts, 1989), and it also draws heavily on contemporary cognitive neuroscience.

Situated cognition research has emphasised the importance of distributed, locally coordinated action systems and the interactive structuring of behaviour.² As such, it has represented a major break with the traditional, almost universal assumption within cognitive science that action is distinctively (and largely exclusively) explained by internal cognitive systems that can be functionally characterised abstractly in terms of representational computation. However there have long been doubts that the complete rejection of internal representational control advocated by some proponents of the 'New AI'³ would be adequate for higher cognition.⁴ Situated cognition has demonstrated that there is deep continuity in the organising principles for adaptive behaviour across a range running from simple insects to complex 'real world' problem solving in humans. But a balanced, systematic approach to cognitive science must explain both similarities and differences, and situated cognition principles currently offer little insight into the basis of major cognitive differences.

Given that most would agree that the situated cognition paradigm is incomplete, we can distinguish between two main kinds of strategies for approaching higher cognition in the light of situated cognition. One is a conservative strategy, emending the traditional representational approach to cognition to accommodate some of the known features of situated cognition phenomena, but focusing on classical ideas about symbolic cognition to explain higher cognitive abilities, such as a possible need for stable symbols to permit information processing that 'goes beyond the present'.⁵ The other strategy involves extending situated cognition principles to develop a more radical re-theorisation of higher cognition. This can include representation, but remaining continuous with situated cognition requires a systems orientation that is likely to treat representation as one embedded feature of a broad array of processes. This is the approach I am taking, and it arguably offers the best way of capturing the complex mixture of similarities and differences across diverse forms of cognition.

² See e.g. Brooks, 1991; Clark, 1997; Pfeifer and Scheier, 1999. Braitenberg (1984) is a seminal influence on autonomous agent robotics.

³ Especially Beer, 1995; 2000; Brooks, 1991; Hendriks-Jansen, 1996; van Gelder, 1995; 1998

⁴ E.g. Kirsh, 1991; Clark, 1997.

⁵ See e.g. Prinz and Barsalou, 2000; Markman and Dietrich, 2000; Sterelny, 2003.

Symbolic processes are clearly an important feature of higher cognition; however focusing *primarily* on symbolic issues carries with it certain kinds of problems. Basing explanations for higher cognition largely or exclusively on representational models leaves the *similarities* with lower cognition unaddressed, and so the major explanatory challenge posed by situated cognition is largely sidestepped. Higher cognition, no less than lower cognition, involves complex, multidimensional dynamical processes. This is well demonstrated in contemporary real-time strategic management problems such as piloting an aircraft (Hutchins, 1995), and it is also arguably likely to be prominent in the problems characteristic of the evolution of cognition, including foraging, hunting, fighting and social cognition, where multiple factors have to be integrated into real time decision making. Higher cognition did not evolve in libraries, even if the Western intellectual tradition to a significant degree did.

Moreover it is hard to understand how classical symbolic cognition can effectively mesh with distributed dynamical processes. It has inherent, systematic difficulties coping with complex, multidimensional, real-time phenomena. All inputs have to be coded symbolically before they can enter the cognitive stream, and cognitive operations on representations of dynamic multidimensional information face the problem of combinatorial explosion. Contemporary advocates of a 'representational' approach to higher cognition (in my terms, the conservative strategy) often disavow old-style computationalist cognitivism, but this then fails to address the question of how representation participates in cognitive processes, at best offers an incomplete approach to meaning, and is possibly misconceived. The problem is that algorithms are required to guarantee that the representational output of a cognitive process is determined by the representational input, so retaining a classical or quasi-classical conception of representation whilst abandoning algorithms creates major puzzles about meaning and cognitive process. Dietrich and Markman (2000, p.8) claim that representation and algorithms go hand in hand, and the connection is indeed hard to break.

Arguably the only way forward here is to directly tackle the problem of cognitive process. Focusing on representation in the absence of process is not a viable research strategy even on its own terms because there is no way to adequately specify meaning without taking into account the process context in which the representations occur.

Constructivism provides a natural way to approach this issue from a non-formalist perspective, by systematically investigating how higher cognitive processes emerge from and are embedded in lower cognitive processes. Older paradigms such as pragmatism and Piagetianism can provide helpful guidelines for such a project, and there are also potentially fruitful connections to pragmatics-oriented approaches to language.

Recent cognitive neuroscience research has the potential to have a major impact on these issues. It is developing sophisticated integrative 'systems' and cross-disciplinary methodologies⁶ and is producing models of high level cognitive phenomena such as conceptual organisation and strategic cognition. Yet outside of a few high profile examples, such as Damasio's (1994) somatic marker hypothesis and the 'mirror neuron' findings (Rizzolatti and Arbib, 1998), awareness of this research in the wider cognitive science and philosophy community is still comparatively limited. For those interested in situated cognition cognitive neuroscience constitutes fertile ground. Although of necessity it is primarily laboratory-based, the systems oriented picture of cognition it is developing meshes reasonably well at a general level with the distributed approach of situated cognition.

There is thus a promising connection between situated cognition, cognitive neuroscience and evolutionary neurobiology. Animals capable of higher cognition such as apes and humans have evolved, through several intermediary stages, from invertebrates with the kind of simple highly distributed sensorimotor systems that behaviour-based robotics has imitated with some success. So a natural way to extend situated cognition is to take into account the elaborations and extensions of the basic ancestral form of sensorimotor organisation. The self-directedness model I am presenting here attempts to capture broad features of this elaboration, and can be divided into three main components:

- (i) **Basic characteristics of self-directedness.** At the most fundamental level the problem for all adaptive action is to bring needs and opportunities into alignment. Reactive architectures solve this problem by responding to contingencies with 'built in' actions, where the efficacy of these actions is based on higher level evolutionary processes or, in the case of a robot, a design process. *Self-directedness* arises through the addition of integrative processes that provide 'onboard' means to improve the coordination between actions, opportunities and requirements, allowing the agent to act in a more flexible, 'proactive' way. The general pressure driving the evolution of cognition is the need for integrative context sensitivity when modularised reactive rules cease to be effective. There are three fundamental functional capacities involved in self-directedness: *anticipation*, *evaluation* and *action modulation*, which together provide a basis for flexible goal-directed behaviour. The self-directedness model takes an *incremental extension* approach to understanding the basis for these capacities, in contrast with a *hybrid* approach that postulates two distinct kinds of functional system: one distributed, one based on some form of symbolic representation. Specifically, the self-directedness model proposes that increasing flexibility occurs through incremental increases in integration across

⁶ See e.g. Arbib et al., 1998.

subsystems, permitting a given action to be influenced by an increasingly broad range of environmental and internal factors. Simple neuronal mechanisms for habituation, for example, can be elaborated into memory systems that facilitate more complex forms of anticipation. Likewise, integration between anticipative, perceptual and action systems can serve to inhibit responses incompatible with current goals.

- (ii) **Complexity management.** Increasing sensorimotor complexity is required for finely targeted, complex action. But sensorimotor complexity in turn brings with it a cluster of management problems that create further pressure for integration. The problem of *dynamic coherence* concerns the fact that complex action depends on the coordination of activity of multiple systems within the organism and being sensitive to multiple factors in an environmental context. Other problems stem from this: the sensorimotor system must be *self-calibrating* to cope with stochastic variation and adapt to changing functional circumstances. The large number of *degrees of freedom* of a complex sensorimotor system imposes a requirement for a progressive development incorporating extensive experience-dependent plasticity. Since many systems contribute to any given activity, development and learning require active *credit assignment* for localising success and error to specific components of the system together with appropriate modification. Collectively these problems impose a requirement for integrative, high order regulation, which is the platform on which cognition is based.
- (iii) **Context learning.** Integrative regulation gives rise directly to higher order information processing. One of the central pressures driving higher order information processing is *context learning*: learning about the circumstances in which action is performed. Context learning plays a role in solving complexity management problems, and involves integrating information from multiple heterogeneous lower order information processing mechanisms and constructing organised, generalised, flexibly retrievable knowledge. From an information processing perspective no special new representational system is required for this. Rather, increased integration *itself* is substantially responsible for generating higher order information by performing a range of information processing functions, including ambiguity reduction, success and error localization, abstraction, the formation of predictive models and conceptualization. The processes involved in this are arguably central to higher cognition.

It is hardly surprising that the complex information processing involved in cognition *involves* a great deal of integration. However it is less obvious that integration itself is a

fundamental information processing mechanism. To appreciate some of the basic reasons why integration plays a central role in cognition to consider a key difference between biological systems and formal symbolic systems. The theory of formal symbolic systems was developed as an extension of the axiomatic-deductive method, and involved the regimentation of deductive steps in terms of logical rules taking into account only the syntactic form of statements (Sieg, 1999). The goal was to develop a formal representation of mathematical proofs, and these kinds of systems achieve logical precision by building precision in at the bottom level. Reflecting this, the hardware tolerances required for physically implementing an elaborate formal computational system are extremely strict, and modern computers are a triumph of precision engineering. Biological systems, however, are built out of an extraordinarily heterogeneous array of sloppy, unreliable components that largely originated for other purposes. They achieve reliability through redundancy and mutual coordination, using something like a 'velcro' principle to achieve a systemic strength out of many individually unreliable connections. This has some very interesting similarities to William Whewell's 'consilience' model of scientific knowledge, in which the robustness of a theory or hypothesis is based on the way it unites many individually uncertain inductions (Butts, 1989). On the position I am developing here this parallel is no accident: higher cognition faces some of the same thematic problems as biological organisation more generally, and has its roots in bio-regulation.

One of the most basic forms of information processing functions performed by integration is ambiguity reduction (see figure 1).




Figure 1: **C**: context, **R**: receiver, **S**: signal

When the 'meaning' of a signal is ambiguous for a receiver it can use information about context to disambiguate the signal, and the more sources of information about context that the receiver has the greater its potential ability to reduce ambiguity. This principle is extremely simple yet powerful, and appears to be exploited in cognition at a massive level. Interactions between adjacent cells in the primary visual cortex help resolve the orientation of a line stimulus. Integration in a higher visual area such as MT can resolve information about the speed and direction of a stimulus. Integration across sensory

modalities such as vision and hearing can also be involved in determining both the location and identity of a stimulus, which is the basis of the ventriloquist illusion (de Gelder and Bertelson, 2003). In addition to early perception, integration also plays a role in 'top down' cognitive processing involving higher knowledge. For example knowledge about conversational context can influence phoneme perception (Bowers and Davis, 2004; Remez, 2000), and contextual knowledge also plays an important role in memory recall (Neisser, 1997). Integration in the brain reaches its greatest extent in the prefrontal cortex, which is involved in 'executive' control processes, playing a general role in assembling task-specific information and regulating overall brain activity in accordance with current goals and context (Miller, 2000).

The standard interpretation of signal-receiver relationships deriving from information theory is that the receiver is using the signal as a basis for representing the state of the sender. However a much broader range of functional relationships are possible and appear to occur in cognition. A more fundamental biological functional relation is where the receiver is ambiguous about what to do next and uses the signal as a constraint to modify its activity. In this paper I will avoid so far as is possible issues of fundamental semantics, but a more general point can be made. The very simplicity of integration – essentially, the use of multiple constraints to shape action – allows it to play a role in a broad range of functional relations, and permits considerable agnosticism concerning the source of the constraints. This allows cognition to be progressively built up during evolution out of many neural systems initially specialised for disparate functions, and for the ready incorporation of experience-derived information during development and learning.

In what follows I present evidence in support of the self-directedness model, focusing on complexity management and context learning to argue that low order integrative information processing provides a constructive platform for higher cognition. I do not examine the general features of self-directedness any further here; for a broader discussion see Christensen and Hooker (2002).

2 Complexity management

Intelligent, context-sensitive behaviour has obvious adaptive advantages: improved targeting of action gains higher returns and reduces the costs accrued by misses, whilst the ability to perform more complex action opens up new kinds of opportunities. But complex, context-sensitive action requires a complex sensorimotor system, and complexity considerations alone impose substantial constraints on the organisation of cognition. One of the most fundamental constraints is that high functional complexity requires both diversity and coherency: the system should be able to generate functional

states that have a large number of component states, these functional states should show high coherency, and the system should be able to rapidly shift between diverse functional states.⁷ Arguably, a balance of modularity and integration is the only way to achieve this because functional complexity involves high regional heterogeneity *coupled with* high global coherency. The gross architecture of the mammalian nervous system illustrates this principle: action for a mammal typically involves a coordinated suite of activities ranging from the activation and inhibition of specific muscle groups through to the generation of perceptual expectations and modulation of affect and arousal state. Simply performing a movement towards a sensory target presents major coordination requirements, including transforming information from the reference frame of the perceptual modality, such as a head-centred reference frame in the case of audition, through to motor reference frames centred on the body or limb. This creates a requirement for integration between perceptual and body information, and in humans and monkeys the posterior parietal and premotor cortices have been shown to be involved in mediating this (Cohen and Andersen, 2002; Hoshi and Tanji, 2000).

One of the coherence requirements for a complex heterogeneous functional system is that it should be self-calibrating, which in turn requires integration, plasticity and self-regulation. Cross-calibration of spatial mapping in visual, auditory and motor systems in owls provides an example.⁸ As reviewed in Knudsen (2002), owls are capable of very precise auditory localization, achieved in the face of considerable difficulties. Individual owls vary in the shape of their head and ears, and as an owl develops it may experience hearing loss and changes in its nervous system. As a result the auditory system must be calibrated based on experience using cross-correlation with visual and motor information. Experimental evidence shows that the spatial mappings have high levels of plasticity. An owl that has one of its ears plugged will initially mislocalize sound sources towards the direction of the open ear, but recovers accurate localization after some weeks of experience. Owls also regain accurate spatial localization ability after the imposition of prismatic spectacles that displace the visual field. Owls learn to orient so that they see the sound source through the spectacles, and also adjust their flight and strike behaviour.

Experience-based functional modification can be more extensive than just calibration, however. Complex sensorimotor systems face what is referred to as the *degrees of freedom problem*: with many kinematic, dynamic and other degrees of freedom they

⁷ This formulation has some similarities to *logical depth* in algorithmic complexity theory (Bennett, 1985). Collier and Hooker (1999) develop an analysis of complexity in open and biological systems that draws on logical depth. Tononi and Edelman (1998) and Tononi et al. (1998) apply a similar principle to understanding consciousness.

⁸ For a general review of cross-modal plasticity see Bavelier and Neville (2002).

possess astronomical state possibilities, yet effective control rests on finding a small subset of these possibilities. Since full search of the space is impossible, particularly given the constraints of real-time interaction, the most plausible explanation for how effective control is achieved is via 'bootstrapping' processes in which an initial set of 'seed skills' provide basic competencies at birth that are elaborated into locomotor and manipulation skills based on experience (see e.g. Giszter et al., 2001). The animal begins its life with simplified motor possibilities and interaction problems, and the control architecture functionally adapts to achieve simple interaction skills. As these are mastered new degrees of freedom open up and functional control organisation is refined progressively. Human adults show the same progressive development when they acquire novel complex skills: they initially simplify the problem by holding many of the control parameters constant while they attempt to perform the task. Later refinement occurs by exploring variation in the parameters initially held constant (Ivanchenko and Jacobs, 2003).⁹

The degrees of freedom problem suggests that cognitive systems require high levels of plasticity and powerful mechanisms for self-regulation. There is a great deal of evidence that such regulated plasticity exists, here I will cite just a few examples. Hsieh et al. (2001) show that the brain areas engaged during speech perception differ between English and Chinese speakers, suggesting that the functional neural circuits involved in speech perception are significantly shaped by linguistic experience. There is also evidence for anatomical and functional differences in the brains of professional musicians as compared with non-musicians (Münste et al., 2002). In a review of the neuropsychological literature on brain damage, Stiles (2000) argues that the traditional view—that plasticity is a transient reactive phenomenon permitting the recovery of a maturational program after insult—should be replaced by a view of normal cognition as stabilised plasticity. Plasticity extends through life and is the basis for neural development as an active, constructive process.¹⁰

Understanding the nature of the regulation that shapes plastic processes into useful functional organisation is therefore a crucial issue. Again, sheer complexity imposes substantial constraints and a requirement for extensive integration. Adaptive development and learning requires that the parts of the system that contribute to success and error should be modified appropriately, but for a complex sensorimotor

⁹ For related developmentally oriented perspectives see also Bickhard, 1992; Karmiloff-Smith, 1992; 1998; Quartz and Sejnowski, 1997; Smith and Thelen, 1993. Vygotsky and Piaget are of course the major intellectual ancestors for developmental approaches to cognition.

¹⁰ For a sample of the range of current research on brain plasticity see the following reviews: Berardi et al., 2003; Ewing-Cobbs et al., 2003; Krubitzer and Kahn, 2003; Rosenzweig, 2003; Ungerleider et al., 2002.

system engaged in complex interaction *localizing* the sources of success and error is far from trivial (referred to in learning theory as the *credit assignment problem*). The strategy described above of simplifying the initial problem and adding complexity through progressive development is only one part of the solution, since it reduces but does not eliminate the credit assignment problem.

Determining success and error requires integrating a number of broad sources of information, including physiological state, body configuration, motor state and perceptual information about the environment. The information processing is therefore highly complex and involves many neural systems; here I will briefly outline some of the evidence on reward, emphasising integrative and higher order reward processing.¹¹ The orbitofrontal cortex and amygdala are involved determining the type and magnitude of reward. For instance the orbitofrontal cortex assembles visual, taste, olfactory and touch information about reward value (Rolls, 2000). Neural systems that are involved in processing higher order reward information include the dopaminergic system, which appears to produce a reward prediction error rather than specify reward per se (Schultz, 2000).¹² The medial orbitofrontal and ventromedial prefrontal cortices also process generalised reward and risk information (Bechara et al., 1994, Bechara et al., 1997), as does the anterior cingulate cortex (Allman et al., 2001). For instance some neurons in the orbitofrontal cortex respond to relative value: it has been shown in monkeys that in the presence of any two food types they fire in response to the more preferred food (Schultz, 2000). The anterior cingulate cortex has been shown to be involved in monitoring conflict between responses (Bothvinick et al., 1999) and response selection (Turken and Swick, 1999), and, along with the medial prefrontal cortices, to respond to winning and losing streaks, thus suggesting that it is involved in processing extended reward expectations (Akitsuki et al., 2003). Such complexity is to be expected on the basis of the credit assignment problem: information processing that is both wide-ranging and fine-grained is required for isolating specific components of reward value and behaviour and relating them to one another.

The predictive properties of the dopaminergic system point to a need for active information processing for credit assignment in a complex system (in this case active processing of temporal relations). Other evidence for this comes from motor control research. The generation of error information for reaching movement requires the integration of goal, motor command and proprioceptive information. In order to compensate for errors during self-generated movement it has been argued that an internal model is required that predicts the action consequences based on an efference

¹¹ For a general review of reward systems see Schultz, 2000.

¹² See also Montague and Sejnowski, 1994; Schultz et al., 1997; Suri et al., 2001.

copy of the motor command (e.g. Wolpert et al., 1995). However Blakemore et al. (1998) show that the efference motor command copy and a single internal model are not enough; effective prediction also requires a model of the context in which the action occurs. This indicates an important general principle of credit assignment: in sufficiently complex situations passive rules are inadequate for effective credit assignment, active information processing such as context estimation may be required to disambiguate the significance of incoming information.

The final aspect of complexity management I will discuss concerns the role of hierarchical control organisation. Classical control hierarchies, which emphasise top-down regulation, face an informational bottleneck because all command information must flow from the top. Avoiding such bottlenecks was one of the major reasons for shifting to parallel distributed architectures within robotics (Brooks, 1986; 1991). However the integration requirements of a complex sensorimotor system generate a need for hierarchy (see Brooks, 1997). An animal engaged in rapid complex interaction with the environment needs to be capable of organism-wide parametric setting of functional systems to current and expected interaction circumstances, which includes active focus on the major interaction parameters and suppressing or otherwise ignoring currently irrelevant information. This requires integrative higher order control. To adapt an example from Giszter et al. (2001), a cat that is about to pounce on a mouse should focus on the overall balance and preparation required for the leap, not become distracted by an obstacle that brushes its leg.

The gross organisation of the vertebrate motor system exhibits a hierarchy of layered competencies rather than the strict top down hierarchy of traditional AI systems. Central pattern generators in the spinal cord produce much of the coordination required for basic activities such as walking (Dietz, 2003), whilst higher neural structures contribute refinement and goal-directedness (Gazzaniga et al., 1998; Giszter et al., 2001). This helps avoid informational bottlenecks because it reduces the amount of information that must flow up and down the hierarchy. However it remains a substantial problem to understand just how this works. The continuous multidimensional kinematic and dynamic properties of the motor system must somehow be organised so as to permit higher level processes that organise action into discrete units that can be arranged into ensembles and sequences for complex flexible action such as object manipulation. Loeb et al. (2000) argue that one mechanism for forming units is for higher controllers to learn muscle co-activation synergies. Johansson (1998) outlines another mechanism, in which the formation and retrieval of internal models mediates between planning processes and motor implementation.

To summarise this section, the requirements of complexity management impose substantial constraints on the organisation of cognition. A balance of modularity and integration is required in order to produce behaviour that is both diverse and coherent. Stochastic variation requires that a complex heterogeneous functional system should be self-calibrating, which in turn creates a requirement for significant plasticity and self-regulation. The degrees of freedom problem means that highly complex functional organisation can only be constructed via a progressive development process that incorporates information from experience. This requires very high levels of plasticity and powerful forms of self-regulation, which in turn places emphasis on the credit assignment problem as a central issue. Many neural systems are involved in this, and evidence from motor control suggests that in some circumstances credit assignment requires active information processing in the form of predictive models of action and context. The demands of larger scale coherence in action require high order control that manages tactical and strategic coordination problems. However such higher control processes must be able to interface effectively with lower systems. Again, motor control evidence suggests that this is in part an active process, involving the grouping or ‘unitization’¹³ of lower dynamics and the formation of predictive models that mediate between planning and motor systems.

To summarise the summary, complexity considerations suggest that cognition is integrative, active, and involves the construction of high order information.

3 Context learning

Context can have a range of effects on learning. If an animal learns a response association in a particular context and it is presented the stimulus in a very different context, response to the stimulus is normally significantly reduced (Gluck and Meyers, 2001: ch.7). Similarly an animal that is conditioned to a stimulus in context X, then has the conditioned response extinguished in a second context Y, will regain the response if the stimulus is presented again in context X (Bouton and King, 1983; Pearce and Bouton, 2001). The context serves as an *occasion setter* providing information to the animal about how it should respond to the stimulus. In mammals this kind of sensitivity to context appears to depend on the hippocampus: rats with a lesioned hippocampus will respond at the same level when a conditioned stimulus is presented in a new context (Gluck and Meyers, 2001: ch.7).

¹³ I borrow the term from perceptual learning (Czerwinski et al., 1992; Schyns et al., 1998). The Schyns et al. research is discussed below.

Defining context learning is not a simple matter since the phenomenon has to a substantial degree been characterised indirectly via its effects on associative learning. Indeed, perhaps the most straightforward way to interpret context learning is to view the 'context' as a bundle of stimuli that come to be used as conditional predictors of reward in the same way as the orthodox associative relationship – i.e., as 'just another CS'. However there is evidence that this is not correct as a general picture, and that information about context has a modulatory effect on associative learning but is not itself always governed by reward association (Pearce and Bouton, 2001). The definition of context learning I will adopt here is correspondingly broad, namely as any kind of learning about the setting within which action occurs.

However hard it may be to define, context learning opens up a broad spectrum of information processing possibilities. It is a fundamental mechanism for promoting response flexibility because it permits different responses to a single type of stimulus in different circumstances. For example, in an experiment by Collett and Kelber (1988, discussed in Gallistel, 1994) honeybees were trained to visit two huts in succession. Inside the huts were identical, containing two blue and two yellow cylinders in a square formation. In one hut a sucrose reward was placed between the yellow cylinders, whilst in the other hut it was placed between the blue cylinders. Despite the fact that there were no perceptual cues inside the huts to differentiate them the bees would search in the appropriate place in each hut, suggesting that they used broader information about the context of each hut to solve the problem.

Context learning is also useful for resolving other kinds of uncertainty, such as when the stimulus is perceptually ambiguous. Collett et al. (1997) examined the ability of orchid bees to locate individual flowers in their range. Bee vision has poor resolution, a plant may look different when approached from different angles, and different individual plants may look very similar. The study found that bees are able to use local landmarks as contextual clues to guide their behaviour, helping to explain how they can resolve the ambiguity in recognising a specific plant. Collett et al. hypothesise that the bees recognise a plant-in-context through a dynamic process in which contextual landmarks prime the memory of the plant as the bee approaches the location. They also point out that the ability to use context information to influence the performance of arbitrary tasks has implications for the organisation of the mediating neural systems: context information must both be broadcast widely across the bee brain *and* be able to be focused on a restricted set of neural sites (ibid., p.350).

The capacity of context learning to facilitate uncertainty reduction points to a wider role for it in making the credit assignment problem tractable. Identifying what went wrong (or right) may depend on tracing relationships beyond what is immediately perceptually

obvious. As Schultz et al. (1997) observe, a rat that reaches a dead end in a maze may have made the incorrect choice several turns previously. Unless it can associate the outcome with the actual turn that leads to the dead end it cannot learn the maze. In general, the more extensive and differentiated the contextual information that a learner can bring to bear on a learning problem, the greater its chances are of being able to localize sources of success and error in the problem. Amnesiac patients are able to learn simple associative tasks but they are impaired at trial and error learning, which may be explained by the fact that they are unable to acquire enough context information to disambiguate successful from unsuccessful trials (Gluck and Meyers, 2001: ch.7).

Rich contextual information also allows an animal to respond flexibly to environmental or motivational changes. The classic evidence for this comes from the 'cognitive map' experiments of the 1930s and 40s. In the latent learning paradigm (Blodgett, 1929, described in Tolman, 1948) rats are first exposed to a maze for a number of days (in the range of 6 to 10), but not rewarded. During this period the behaviour of the rats appears random, however in the next phase of the experiment a food reward is placed at the end of the maze leading to a marked change of behaviour. The rats begin to run the maze to reach the food with low error rates. This indicates that during the non-rewarded period the rats have been acquiring contextual information about the organisation of the maze (a 'cognitive map'). The rats can then use this information to guide their behaviour when they are motivated to do so. Indeed, in some versions of the experiment such rats show lower error rates than rats that have been rewarded from the beginning of their exposure to the maze, indicating that unrewarded exploration can be a more efficient process for acquiring maze information than goal-directed learning (Tolman and Honzik, 1930). This is possibly because they acquire contextual information more widely and experience less reward-interference as they learn the relational features of the context. Their 'cognitive map' would then be more detailed, integrated and stable.

Another experimental paradigm demonstrates more directly flexible response to changing motivational circumstances (the experiment is by Spence and Lippitt, described in Tolman, 1948). Rats are exposed for several days to a simple Y-shaped maze in which one arm of the Y contains water and the other arm contains food. In the initial phase the rats are satiated for both food and water prior to exposure to the maze. In the second phase of the test the rats are divided into two subgroups, one of which is made hungry and the other thirsty. When the rats are then reintroduced to the maze the hungry rats go directly to the arm that contains food, whilst the thirsty rats go to the arm that contains water.

At this point it is worth mentioning some (comparatively) recent criticism of the claim that animals have cognitive maps. In an influential paper Bennett (1996) argues that no

experiments to that date have in fact demonstrated the presence of a cognitive map in animals. Drawing on the animal navigation literature Bennett canvasses a number of definitions of cognitive maps that all interpret the concept rather literally in terms of spatial representations used for navigation. Following O'Keefe and Nadel (1978), Bennett distinguishes between *route following* and map-based navigation, where the former involves memory of a particular path to a goal. He argues that demonstrating map-based navigation requires both showing that the animal has taken a shortcut to the goal across an unfamiliar area—where it cannot have acquired a route memory—and that it could not have performed the task by dead reckoning. Bennett claims that no cognitive map experiments have conclusively eliminated both route memory and dead reckoning hypotheses for apparent short-cut taking.

However Tolman (1948) clearly has a more general concept of a cognitive map than a literal representation of space. He introduces the notion by contrasting it with stimulus-response models of behaviour, and one of the experiments he uses to illustrate cognitive maps is not a navigation experiment at all – it is a pattern matching experiment. Tolman refers to the Spence and Lippitt experiment described above as “probably the best experiment demonstrating latent learning” (p.195), but as a navigational problem it is relatively trivial. And in his concluding remarks he makes wide-ranging comments on potential implications of cognitive maps for social psychology. Tolman's conception of a cognitive map thus appears to be any kind of model that captures larger scale relational information.

This makes a substantial difference to Bennett's argument. A lot depends on what a 'route memory' really is: if it is interpreted as a strict stimulus response procedure then it is not necessary to show that the animal has never crossed a spatial area before to demonstrate the presence of a cognitive map. It is only necessary to show that the animal can use acquired information flexibly under changed circumstances, and this is just what the original cognitive map experiments do show. A 'short-cut' can include choosing the best path known to reach the current goal, rather than the path that has been most rewarded in the past. If route memory is conceived of more broadly to encompass the possibility that it might be used flexibly then it corresponds to what Tolman calls a 'strip map' (p.193).

Tolman's notion of cognitive maps may be very broad compared to the idea that a cognitive map is a representation of space but I suggest that it is also the more basic and cognitively significant conception. It relates to behavioural flexibility in general, not just spatial navigation, and it points to relational information processing abilities that may be fundamental to cognition more widely. Several recent computational models provide insight into mechanisms that may be involved in the acquisition of organised information.

Schultz et al. (1997) use a predictive reward model based on the dopaminergic system to show how a 'value map' of an animal's environment could be acquired. As the animal experiences its environment the predictive reward system acquires information about reward value of the objects in the environment. In the model the 'animal' is a simple artificial creature that is able to look around its environment. As it does so its reward system generates predictions of the rewards it would obtain if it approached particular objects. This information can be used to make choices between particular approach options. Such a predictive reward system could, in a more complex model or a real animal, be coupled with 'imagination' to generate extended predictions. Gluck and Meyers (2001, p.195) describe a simple neural network that can be used to solve navigation problems. The network acts as a forward model: given a current location and a movement it learns to predict what the next location will be. If this network were coupled with a predictive reward network it could be used to evaluate possible moves and iteratively find a path to a goal.

Eichenbaum (2001) presents a 'memory space' model of declarative memory that draws on James (1890), Tolman (1932) and Bartlett (1932), and points to a simple yet somewhat more powerful process for acquiring organised cognitive map-like knowledge.¹⁴ Declarative memory is composed of semantic and episodic memory¹⁵, and in Eichenbaum's model episodic memory is the 'gateway' for the formation of semantic memory. Specifically, semantic memory is acquired by linking multiple experiences that share some of the same information. The informational links across episodes constitute the animal's 'memory space': a larger organisation of information that can be used to generate solutions to new problems. The hippocampus is proposed to mediate the acquisition of sequential context-specific information involved in the formation of episodic memories, and to play a role in identifying common features across episodes as part of the formation of generalised semantic memory.

The model has been tested on rats using transitive inference as the measure for the acquisition of organised information (Dusek and Eichenbaum, 1997). Rats are presented with pairs of odours and taught to discriminate between them, where the rewarded discriminations have the following hierarchical structure: $A > B$, $B > C$, $C > D$, $D > E$. The test for transitive inference is the animal's choice when presented with the pair BD , which is not encountered during the training phase. Normal rats rapidly learned the discrimination pairs and, when presented with the transitive inference probe, made the correct $B > D$ choice as accurately as they did on the training pairs. In fact 88% of the rats made the correct choice on the first presentation of the transitive inference pair. In contrast, whilst

¹⁴ See also Eichenbaum, 2000; Greene et al., 2001; Wallenstein et al., 1998.

¹⁵ Roughly, knowledge about properties of the world and autobiographical knowledge respectively.

rats with a disconnected hippocampus were able to learn the rewarded discrimination pairs at the same rate as the intact rats, they were unable to solve the transitive inference discrimination, performing at chance level.

Internal predictive models and the kind of integrative information processing characterised in the ‘memory network’ model of the hippocampus provide an indication of some of the relatively simple mechanisms that can be coupled to generate increasingly complex higher order information processing. Conventional associative learning theory treats stimulus, action and reward as part of a single function: the probability of producing an action in response to a stimulus is determined by the frequency with which that stimulus-response association has been rewarded in the past. However a cognitive architecture that couples a complex reward system with multiple internal models and integrative memory has the potential to partially disentangle perceptual, value and action information. A predictive reward system capable of parametric adjustment to different reward regimes can generate multiple value maps based on contextual factors that signal a change in reward regime. A forward model of context can be run under a variety of motivational settings to generate varying predictions of the reward outcome of particular actions. A ‘memory network’ can link salient information across interaction episodes to extract generalised relational information beyond specifically rewarded stimulus-stimulus or stimulus-action conjunctions. In combination such mechanisms can greatly expand the capacity for context-sensitive, goal directed behaviour.

I conclude this section by commenting on the relation between this broader approach to cognitive maps and the way in which representational explanation is conceived more generally in cognitive science.¹⁶ In one respect Tolman’s ‘cognitive map’ label is misleading inasmuch as it encourages the idea of a static ‘objective’ map. However forward models are dynamic and, since they are trained through interactive experience, based substantially on egocentric information. But once they are trained they do not have to be driven by current perceptual information, they can be run in simulation. They are also capable of interpolation: they will generate a prediction for input values that have never been experienced before. Eichenbaum’s ‘memory network’ model indicates another way in which broader relational information can be built up out of egocentric information by integrating across multiple experiential episodes.

It is interesting that Eichenbaum’s (2001) presentation of his ‘memory network’ theory of declarative memory should begin with William James, since the theory itself is rather

¹⁶ I am not however attempting to present any theory of representation per se – cf. Grush (1997), who does use internal models as the basis for a theory of representation.

pragmatist in nature. A central pragmatist idea was to show how higher knowledge is constructed out of basic information derived from interaction with the world, and it contrasted itself with what it regarded as ‘intellectualist’ philosophy that took ‘objective’ representational knowledge as fundamental.¹⁷ I shall term this general explanatory model ‘constructivist’ (not to be confused with antirealist forms of constructivism).¹⁸ Piaget is another constructivist exemplar; systematically applying the same basic strategy to developmental psychology in an attempt to show how generalised conceptual knowledge is progressively built up out of lower order egocentric, context-laden information.

With this in mind we can characterise a heuristic problem with typical approaches to representational explanation in cognitive science and philosophy. The tendency is to assume that, if a problem is being solved cognitively at all, then it is being solved by representational processes that directly reflect the principal features of the problem being solved. Questions about the potential cognitive basis of a particular behaviour are transformed into questions about representational format, where the format is read directly off the problem. For example the problem of disentangling value, perception and action information is connected with Sterelny’s (2003) concept of decoupled representation. The models described above might, in fact, be considered mechanisms for generating decoupled representation. However Sterelny appears to treat the issue of ‘decoupled representation’ as a question of representational format, whereas these models suggest that the decoupling is a systemic property and the representations are actually dynamically stabilized constructs. Another example of systemic decoupling occurs in the role of the prefrontal cortex in inhibiting prepotent actions. This is one of the most important forms of cognitive decoupling, and whilst it certainly is dependent on information processing that occurs in the prefrontal cortex and elsewhere, the decoupling is a complex product of the major regulatory connections between the prefrontal cortex and other brain regions. It is not a property inherent in representations in the prefrontal cortex, which are in fact highly dynamic – see Duncan (2001).

¹⁷ See e.g. James, 1997.

¹⁸ Pragmatism is an important but ambiguous example of constructivism. The ambiguity lies in the difference between claiming that there is nothing more to truth than successful action—an idea that causes pain to many philosophers—and the idea that ultimately all knowledge is based on differences in interaction with the world. The latter principle is surely correct, but the former suggests an eliminativist reduction of truth. Eliminativism can be avoided via a clarified constructivist position in which a substantial concept of truth is captured as a higher order construct. Bickhard (1993) develops a theory of representation along these lines. It is arguable that the constructivist interpretation best captures the spirit of early pragmatism, which showed great interest in the actual dynamics of knowledge acquisition in preference to sterile abstracted epistemic and metaphysical debates.

Questions about representational format are undoubtedly an important issue in cognition, but treating cognitive problems in general as problems of representational format obscures the role of indirect and systemic factors. Moreover this produces an excessively sharp split between representational and non- or 'locally' representational processes, with the latter often serving as a catch-all null hypothesis. For instance the contrast between route memory and cognitive maps drawn by O'Keefe and Nadel (1978) reflects a distinction in traditional cognitive science between procedural and declarative knowledge.¹⁹ Yet from a constructivist perspective this distinction is not strict, and the three 'cognitive map' models that I have just discussed illustrate specific types of information processing mechanisms that can play a role in constructing generalised information from context and action specific information. A critique of cognitive map hypotheses by Wang and Spelke (2002), much in the spirit of Bennett, further illustrates the problem. They treat evidence *for* use of perspectival information in navigation as evidence *against* cognitive maps, but this doesn't follow for the reasons discussed above. Their characterisation of the contrast is revealing:

...animals, including humans, navigate primarily by representations that *are momentary rather than enduring, egocentric rather than geocentric, and limited in the environmental information that they capture*. Uniquely human forms of navigation build on these representations.

(Wang and Spelke 2002, p.376, italics added)

4 Self-directed anticipative learning

Rats are likely to share many of their cognitive traits with the early mammals, which, along with birds, underwent a major adaptive shift associated with the evolution of thermoregulation. Mammals and birds tend to have larger brains than comparably sized reptiles. However they also require a great deal more food because the requirements of temperature homeostasis are such that resting mammals and birds expend between 5 and 10 times as much energy as do comparably sized reptiles (Allman, 1999). Mammals and birds are locked into a reciprocal set of opportunities and demands. They are capable of higher activity levels, and the increase in brain size allows this action to be more sophisticated, thereby improving their ability to acquire food. But they also need these resources to maintain the underlying physiological and neural infrastructure. The cost of error therefore increases.

Mammals and birds thus face the complexity management problems described above in sharp form.²⁰ They need to acquire complex interaction skills, and they need

¹⁹ Cf. the discussion in Sterelny, 2003: pp.40-45.

²⁰ This leaves open the status of the bee evidence I outlined. Many people find it jarring that evidence concerning honeybees should be introduced in a discussion of cognition. The answer to

sophisticated mechanisms to determine success and error in the process of building these skills and acquiring information about the opportunities and threats in the environment. There was probably therefore strong evolutionary pressure for enhanced context learning ability. I will now describe a progressive learning process termed *self-directed anticipative learning* that may play an important role in context learning. The central idea is quite simple, but it extends the systemic approach to cognition and learning developed thus far.

Self-directed anticipative learning (SDAL) involves a reciprocal interplay between interaction and anticipative model construction. In the initial phase the learner doesn't understand the task structure and doesn't know what is relevant. The learner is alert and attempts to gather information widely, looking for possibly relevant clues to the structure of the situation. As the learner engages interactively with the situation it begins to gather more perceptual information about the relational features of the situation together with some success and error information. But with still-limited relational information about the situation it is difficult to assign success and error specifically. The learner is acquiring improved relevance information, however, which guides more specific attention and interaction. More focused interaction should, in turn, generate better relational information specific to the task, which will then improve the capacity to assign success and error. The learner is in effect progressively constructing a relational model of the task situation, incorporating a mixture of relevancy, perceptual and motor information. The model is abstractive because task-irrelevant information 'fades' from attention and the model. If the learner acquires a good situation model it can quickly attend to important features of the context, and respond flexibly to changes in task parameters.

From an information processing perspective the key features of SDAL are:

Integration of multiple overlapping sources of information: In order to achieve successful completion of the task the learner must integrate information about performance from multiple overlapping sources (including perception, affect and motor information, and repeated interaction with the task environment). The significance of having multiple sources of information is that it provides orientation and gradient information that helps distinguish better from worse situations and focus action. The more convergent sources of information the learner has about

this I think is that some of the complexity pressures that give rise to cognition arise much earlier than is usually supposed. Thus, the foraging problems that bees face present some of the problems of ambiguity and uncertainty reduction that drive the evolution of context learning mechanisms. They do not however face the same kind of sensorimotor management problems that mammals and birds do. Whatever the underlying reasons may be, though, honeybee research is showing surprisingly high levels of information processing complexity in a simple invertebrate architecture – see Menzel and Giurfa, 2001.

performance the more effective its actions can be. At the initial stage of the learning process the sources of information are impoverished, whilst at the concluding stage they have been sufficiently enriched to converge on a solution.

Improved anticipation and error detection: As the learner interacts it generates information that allows it to construct anticipative models of the interaction process. This allows the learner to: a) improve its recognition of relevant information, b) perform more focused activity, and c) to evaluate its performance more precisely. Indeed, error itself is a rich source of context-sensitive information that can be used to further refine anticipations.

Rochat et al. (1999) present evidence from human infant development that illustrates at least part of such a process. Infants at 2, 4 and 6 months interacted with an adult female stranger in a 1-minute peekaboo game that was either organised or disorganised. 2-month infants gazed and smiled at the adult equally regardless of how organised the peekaboo game was. However 4- and 6-month old infants smiled more and gazed less in the organised peekaboo condition as compared with the disorganised condition. Rochat et al. interpret this as a progression from diffuse sensitivity to a social partner to an emerging sensitivity to the narrative envelope of social exchange. Below I discuss a perceptual classification task involving feature learning that involves extended interaction between top-down and bottom up processes. Thornton (1999) investigates a learning process very similar to what I am calling SDAL in child conceptual development. She argues that,

[t]he dynamic interaction between the strategies children bring to the task and the detailed structure of the task in hand acts to redirect the child's attention to new aspects of a problem, opening up the possibility of new discoveries that were not prefigured in the child's original approach.

(Ibid., p.588)

Two cognitive neuroscience models provide further indication of some of the kinds of information processing components that could, in cooperation, generate SDAL-type processes. Dehaene et al. (1998) present a 'global workspace' hypothesis for effortful cognitive tasks. According to this hypothesis difficult tasks provoke the activation of a class of neurons with heavily interconnected long-range axons that are distributed across the brain but are especially common in the prefrontal cortex. These neurons assemble information from all brain regions, allowing it to be processed in an integrative, flexible way. Such a process could serve as the basis for the assembly of contextual information for anticipative learning and for executive guidance in the learning process. Gluck and Meyers (2001) present a model of cortico-hippocampal interaction that views the hippocampus as functioning to compress redundant information and differentiate

predictive information. This processed information is then passed on to the neocortex and cerebellum where it is stored in long term memory.²¹ This kind of information processing could serve as the basis for the formation of contextual models, whereby information that is irrelevant or adds no predictive value for the learner comes to be suppressed or ignored, whilst information that is predictive within a context or differentiates between contexts is accentuated. Over time the information that a learner retrieves from long term memory during task performance comes to be increasingly well organised in relation to the important features of the task.

The SDAL model appeals to relatively generic cognitive abilities but is able to explain rapid learning in complex domains. For instance Regier (2003) argues that a similar concept of 'accelerated learning' could explain features of word acquisition that have been thought to require innate domain-specific mechanisms. The characteristics of repeated exposure to the situation coupled with changing attention biases and expectations mean that SDAL may be an important process for extracting situational and cross-situational invariances, and hence for the formation of generalised, 'portable' knowledge.

I have presented evidence that suggests that SDAL occurs in child development, but does it occur in non-human animals? I am not aware of direct evidence for progressive learning with SDAL-characteristics in non-human animals, however it could be playing a role in 'memory network' forms of context learning. An important point to note is that the power of this type of learning process is highly sensitive to quantitative factors such as the ability to assemble context information. Cortex size is likely to be a major limiting factor on this, as indicated by recent accounts of cortical function such as the global workspace model. Other limiting factors include attention, working memory, gross informational capacity of long term memory, and the capacity for organised storage and retrieval of memory. SDAL might therefore form part of the explanation for the common observation that intelligent, flexible behaviour in birds and mammals scales in proportion to brain size.²²

I will now review the key features of self-directedness and bring together the main points from the discussion so far. According to the self-directedness model increasing sensorimotor and behavioural complexity generates integration pressure. This drives the evolution of higher order control processes that serve to coordinate the activity of

²¹ The model is not incompatible with Eichenbaum's 'memory space' account, and indeed it might explain part of the underlying basis by which the integration across episodes occurs. More generally, it is likely that the hippocampus performs multiple related information processing roles – see e.g. Gluck et al., 2003.

²² I discuss experimental evidence supporting this in the next section.

multiple systems within the organism and to manage larger scale needs – actions – environmental relations. Two main effects arise:

- Animals become more self-directed, in the sense that they are capable of more complex goal-directed action that integrates opportunities and needs
- The information processing conditions for higher cognition are generated

The second point is probably the most controversial. The predominant view is that there is a qualitative difference between lower and higher cognition, and the most common interpretation of this difference is that higher cognition exhibits linguistic or quasi-linguistic representational structure, whereas lower cognition exhibits no structure or at most simple associative structure. Much of the evidence I have presented so far, however, demonstrates that lower cognition exhibits more complex organisation than simple associative learning. The degrees of freedom problem forces a developmental structure to cognition that depends on sophisticated integrative regulation. The credit assignment problem creates a need for internal models and context learning. The ‘cognitive map’ evidence discussed indicates that mammals, specifically rats, are able to extract stable relational information about context that permits flexible response to changing and novel circumstances. SDAL-type learning processes involving repeated exposure to a problem situation with cumulatively formed expectations and skills may play an important role in the formation of abstracted, ‘portable’ knowledge.

Three further features are especially noteworthy. Firstly, in both motor control and context learning there is a need for *unitization* – the creation of informational groupings. In motor control the reason is that the continuous, multidimensional dynamics of the motor system must be converted into a form amenable to the performance of articulated action. In context learning there is a need to group together or suppress redundant information, and to differentiate key features within a context and that mark the differences between contexts. Secondly, whereas traditional symbolic approaches to information processing assume the prior existence of fundamental information processing units, in these cases the units are dynamic constructs embedded in integrative processes. Thirdly, whilst unitization helps to generate articulation, it is integration that generates the major information processing characteristics that shape cognition, including success and error localization and abstraction.

5 General supportive evidence

At the broadest level the self-directedness model makes several core claims:

- (i) Degree of integrative regulation is the major parameter governing the transition from reactive to intelligent proactive agency

- (ii) Higher cognition emerges out of the principles of integrative neural sensorimotor control

At this level of generality there are currently three major models for the basis of higher cognition. These are the classical general computational theory of mind, the massive modularity hypothesis, and the hypothesis that language is responsible for higher cognition.²³ All three are variants of symbolic computationalism, with the massive modularity hypothesis fragmenting the computational system into multiple domain-specific modules, and the language hypothesis proposing that natural language is the vehicle for symbolic computation. I will now outline several sources of evidence that support the self-directedness model.

The first is the correlation between brain size and problem solving flexibility.²⁴ This has been measured by the ‘transfer index’: an experimental measure of behavioural flexibility devised by Rumbaugh (1970) and applied to diverse primate species. The transfer index is a reversal learning problem (Bitterman, 1960), and is rather like a simple version of the Wisconsin Card Sorting Task used in developmental neuropsychology as a measure of executive function. The animal is first presented with two stimuli and rewarded for choosing one of them. After being trained to one of two levels of competency (67% or 84%) the reward is switched to the other stimulus. ‘Transfer’ refers to the extent to which learning on the first task affects performance on the second task. ‘Negative transfer’ occurs where increased training on the first task *decreases* performance on the second task, and is interpreted as difficulty on the part of the animal in inhibiting the previously learned response. ‘Positive transfer’ is where increased training on the first task *improves* performance on the second task. Rumbaugh interprets this as resulting from the animal conceptualizing the tasks (the problem in each case is choose the currently rewarded response). However it might be put more neutrally as indicating that the animal has extracted contextual information from the first task that helps it perform the second task. This can be as simple as marking the two contexts as different, helping to disambiguate the second task. But it can also involve more complex relational information that picks out similarities between the two tasks. Performance on the transfer

²³ I do not include social intelligence hypotheses here because they have no direct implications for a theory of cognitive architecture. Insofar as some versions do make architectural claims for a ‘metarepresentational’ module they are inconsistent with the evidence below.

²⁴ The self-directedness model is similar in many important respects to the ‘differentiation and construction’ account of the evolution of intelligence presented by Gibson (2002), who also points towards the correlation between brain size and problem solving flexibility. These two accounts may be largely complementary.

index across primate species was found to correlate with absolute brain size (Gibson et al., 2001).²⁵

A second source of support comes from evidence that 'general intelligence' in humans is a real phenomenon (sometimes referred to as 'factor g' or 'gF'). The performance of an individual across a range of intelligence tasks tends to be strongly correlated: an individual who performs well on one kind of task is likely to perform well on other kinds of tasks. In a review Deary (2001) reports a current consensus for a three-stratum hierarchy model of intelligence ability differences: A general cognitive factor accounts for around 50% of the variance in performance on a broad range of mental tests given to a large sample of the population. There is also a subsidiary set of variance attributed to 'group factors' such as verbal comprehension, spatial cognition, memory and processing speed, and finally some variance is accounted for by very specific mental abilities. This strong correlation across mental abilities is difficult to explain using a 'massive modularity' account of cognitive architecture, and it is not easily accounted for by hypotheses that postulate language processing as the factor responsible for cognitive integration, given that many intelligence tests are specifically designed not to tap into language ability.

On its own factor g might be taken as evidence for the classical general computational theory of mind. However in conjunction with the transfer index evidence for the correlation between intelligence and brain size in primates it more strongly supports the self-directedness model. The general computational theory explains cognitive flexibility in terms of a general symbolic computational architecture, and it has to postulate a major functional discontinuity corresponding to the evolutionary advent of this architecture – Newell's 'Great Move' (Newell, 1990). The self-directedness model on the other hand explains cognitive flexibility as a graded function of neural integration capacity. The quantitative correlation between brain size and flexible intelligence supports the latter hypothesis.

Imaging evidence concerning the neural underpinnings of general intelligence provides more direct support for the connection between neural integration and intelligence. Using positron emission tomography Duncan et al. (2000) show that tasks that are highly correlated with general intelligence produce specific activation in the lateral frontal cortices. At first glance this might seem to conflict with the primate evidence that intelligence correlates with *overall* brain size, however the frontal cortex has been shown to play a general regulative role in cognition, performing executive functions that involve

²⁵ Performance did not correlate with the Encephalization Quotient measuring the relationship between brain and body size (Jerison, 1973).

the selective inhibition and activation of other brain regions based on current task demands (e.g. Cohen and Servan-Schreiber, 1992; Goldman-Rakic, 1995; Miller, 2000).

Taken together this evidence supports the thesis of the self-directedness model that behavioural flexibility and higher cognition are directly connected to neural integration capacity. The self-directedness model explains the flexibility of intelligence, which the massive modularity hypothesis cannot, and it explains the quantitative connection between brain size and intelligence, which is problematic for the general computational theory and for the language hypothesis.

An important qualification to this should be noted, however. Research by Premack (1983) and Oden, Thompson, & Premack (1988; 1990) suggests that, whilst language trained chimpanzees are able to spontaneously solve higher relational problems, such as matching half an apple with half a glass of milk, chimpanzees without language training do not. In one respect this supports a constructivist picture of cognition, and in discussing the research Thompson (2000, p.382) notes that his own analysis is similar to Clark's (1998) idea that words can act as scaffolds to make particular features of the world more concrete. But these results also suggest that there may be more to higher relational cognition than generic integrative learning. The picture is complicated by the fact that infant chimpanzees are perceptually *sensitive* to such relations, indicating that they have an underlying cognitive ability to discriminate these relationships, but that something further is required before higher relational properties can be used to make judgements. Still, it is very plausible that language has a scaffolding interrelationship with higher cognition in humans, and language provides a powerful external medium for the unitization phenomenon I discussed earlier. It would be a mistake, though, to conclude that higher cognition is *based* on language, and I will now present evidence that the generic processes I have been characterising also play a central role in higher cognition, and that language is far from being the only medium for unitization.

6 Specific evidence for integration in higher cognition

In this section I will briefly present two kinds of evidence showing that the features of low order cognition highlighted by the self-directedness model are also important characteristics of higher cognition. This research questions the traditional emphasis on categorical representations and points to the involvement of much more fluid, integrative processes.

Traditional theories of concept learning assume that learning occurs within a fixed feature space. However research by Schyns et al. (1998) and Goldstone et al. (2000) indicates that although current perceptual feature recognition constrains the concepts

that can be acquired, the concepts acquired also influence the learning of features. This suggests that there may be no fixed set of fundamental features. The experiments designed to demonstrate this involve presenting subjects with irregular figures for which there are no obvious component parts. The subjects are taught to classify the figures according to experimenter-determined categories through trial and error: the subjects are not explicitly told the category rules. In the early stages of the learning process the subjects acquire some ability to classify the figures correctly but are not able to consciously identify any basis for classification. They have, however, become sensitized to specific stimulus dimensions of the figures. In later stages classification becomes nearly errorless, though the subjects are still not able to verbally identify the basis for classification. At this point they have (unconsciously) localized the key diagnostic features for classification. These may be particular shape characteristics, or it can involve the unitization²⁶ of an array of features. In the final stage of the process the subjects are able to verbally identify the diagnostic features, and they 'see' the figures in terms of 'objects' determined by these features. One of the more interesting aspects of the experiments, however, is that different groups of subjects are presented with the same stimulus sets but required to classify them differently. Subjects learn to identify the particular features that are diagnostic for the task they are given and do not 'see' the features that subjects trained on other tasks see.

In a review of concept research Solomon et al. (1999) argue that researchers have greatly overemphasized categorisation and that this has produced a distorted theoretical picture. They argue that concepts serve multiple functions that interact to affect conceptual structure and functioning. They cite a study by Medin et al. (1997) that indicated that landscapers use different conceptual structures when presented with categorisation and reasoning problems. For categorisation the landscapers used goal-derived categories based on landscaping, such as shade trees, weed trees, and so on. When asked to make inferences about trees, however, they used a taxonomic scheme based on biological relatedness. In a more recent experiment Shafto and Coley (2003) required undergraduate students and commercial fishermen to sort marine animals into groups. Novices sorted the animals according to similarity of appearance, whilst experts sorted according to a variety of factors, including commercial, ecological and behavioural characteristics. They conclude that, "[e]xpertise appears to involve knowledge of multiple relations among entities and context-sensitive application of those relations" (p. 641).

In an intuitive sense these results are obvious. Of course experts can make finer distinctions and reason more flexibly than novices: they know more. Still, there may be important theoretical implications to be had from them. On a traditional cognitivist

²⁶ As noted earlier, this perceptual learning research is in fact the source for the term.

account 'knowing more' means more or less just having a larger database. In theory it is supposed to involve higher level knowledge organisation, including input chunking, scripts, frames and so on, but there have never been convincing mechanisms for flexible higher order learning in the cognitivist framework. The account I am developing here, based largely on cognitive neuroscience evidence, points to a significantly different picture of what 'knowing more' involves. Specifically, with respect to neural systems learning involves systemic interactions across multiple brain regions leading to system-wide activity dependent modifications that, in addition to storing information, to some degree refine and tune the cognitive architecture itself. With respect to behaviour there is an unfolding interaction process in which acquired information generates changed behaviour patterns that in turn modifies what is learned, in substantial part by changing the ability to detect relevant information. Thus, perception, memory and executive control exhibit a dynamic interplay between bottom-up and top-down processes. Feature perception, conceptual organisation and executive frontal processes are all reciprocally tuned during context learning to facilitate the flexible assemblage of higher conceptual and strategic information according to situational demands. Experts don't just know more, they see, think and interact differently.

The account I am presenting is clearly not anti-representational and nor does it claim that all the structure that participates in cognition is developmentally constructed. The electrical depolarization properties of neurons and molecular pathways for synaptic and other forms of neural plasticity provide basic mechanisms for information processing. The cell types and major connectivity patterns for neural regions such as the cerebellum or hippocampus are genetically specified. Multiple coordinate systems are required for sensorimotor control and learning. There are many specialised information processing abilities that appear early in human infant development, such as perceptual expectations about the numerical and type identity of objects (Xu and Carey, 1996). This is to be expected on the basis of the arguments I introduced in section 2, because the same complexity management problems that force an extended plastic development also ensure that there must be considerable 'built in' organisation: too much plasticity is unmanageable.

However the account does differ in fundamental ways from standard representational approaches to cognition. It is systemically oriented, treating cognition as structured by multiple neural systems and interaction with the environment. For instance one prominent feature of higher cognition is *information portability*: the ability to store information, transfer it across contexts, and use it in a variety of cognitive processes. In a classical approach this is explained in terms of the properties of symbols, but in the present account a large part of the explanation lies with integrative neural architecture and interactive, abstractive learning processes. This has important advantages. The very

plurality of neural systems, and the idiosyncrasy of the information they process, suggests that the articulation of higher cognition is not based on a general coding scheme in the way that a navigation chart might be based on the Mercator projection system. Moreover strict formal symbolic computation is an inflexible form of information processing, highly prone to combinatorial explosion. Indeed, the problem of combinatorial explosion is one of the arguments used to justify the massive modularity hypothesis (e.g. Carruthers 2002, p.705). It can just as easily be used, however, to argue that cognition is not based on symbolic computation at all in the classical sense.

The integrative, constructive approach has the benefit of being able to explain how higher cognition is able to function effectively in complex dynamical domains. As Schyns et al. (1998) observe, traditional feature space theories of concept acquisition face a difficult design choice with respect to the nature of the basic features represented in the space. Choosing low level features increases power in the sense that many more kinds of objects can be represented by the system. However it also increases the search space for object identification and learning, which can result in computational intractability. Choosing high level features can greatly improve speed and tractability, but at the price of representational flexibility. To be functional such systems must necessarily be highly domain specific. Models that incorporate feature learning do not face this trade-off, or at least have a powerful strategy for making it tractable. They can couple initial flexibility with later-acquired specificity that accelerates the learning process. This is another version of the degrees of freedom problem with the same kind of solution.

Dynamical coherence and complexity management are just as much features of higher cognition as they are of lower cognition. Higher cognition often involves the ability to keep track of many factors simultaneously: think of holding a conversation with a good friend who is offering to sell you a car that you may or may not be interested in buying. The same kinds of complexity management strategies involving integrative success and error localization and predictive learning that occur in lower cognition are likely to also play a key role in higher cognition. This has a number of important implications. Higher cognition is not screened off from lower cognition and dynamical interaction; it is deeply entwined with it. It is therefore crucial to understand how cognition is embedded in deeper processes, and also vital that a multidisciplinary approach is pursued. Cowley (this volume) and Ross (this volume) provide very different perspectives on how such a constructivist embedding might be carried out, and both may be providing an important part of the picture.

As a non-specialist it would be foolish to try to make specific proposals about language research. However the general thrust of this account supports an emphasis on the importance of pragmatics in language. It does not necessarily support the hypothesis

that syntax can be fully explained as the developmental outcome of usage constraints, which some formulations of integrationist linguistics appear to suggest. But it does support the view that research on interactive dynamics in language development and use provides a very productive window onto language. It also questions the cognitive significance that is often loaded onto syntax, such as the idea that it is the vehicle for higher relational cognition.²⁷ Many of the most central features of higher relational cognition cannot be explained by syntax – for example context sensitive lexical retrieval.²⁸ Palma (this volume) comments that the contextual dynamics of language may be scientifically intractable, and this kind of argument against holistic approaches has a long history.²⁹ As a general claim this is a risky position, to say the least. One irony is that behaviourists argued that any kind of explanation for behaviour that went beyond stimulus-response relationships was scientifically intractable, and Chomsky (1959) played a decisive role in undermining this idea. Another irony is that cognitive neuroscience is now developing sophisticated methodologies for investigating ‘systems-level’ phenomena, and this is providing a great deal of concrete evidence for older ideas concerning the importance of holistic processes in cognition. On the other hand the difficulty of such ‘integrationist’ research should not be underestimated. Developing workable, productive research paradigms is a formidable problem, not least because the phenomena of interest are not readily captured in highly controlled situations. Interdisciplinary connections can help ameliorate this, however, because convergence between several partial perspectives provides greater confidence and direction than a single perspective alone.

7 Conclusion

I will conclude by saying more about the connection between this account and situated cognition, since I have said very little directly about the kinds of scaffolding phenomena that situated cognition usually focuses on. To begin with, it is an account systematically based on embodiment, tracing the relationship between increasing sensorimotor complexity and the formation of cognitive properties. Moreover the degrees of freedom problem and the credit assignment problem provide an underlying theoretical rationale for the importance of scaffolding in high as well as low order cognition. It may be that one of the most critical limiting factors in the evolution of cognition is the capacity to regulate plasticity. Basic neurophysiology—the neural toolkit present in all animals—exhibits a rich array of activity-dependent plasticity mechanisms. It is not lack of plasticity that is the problem. Rather, the problem may be guidance mechanisms for

²⁷ See e.g. Spelke, 2003.

²⁸ See Thompson-Schill et al. (1997) for evidence on the role of the prefrontal cortex in lexical retrieval.

²⁹ See e.g. Wertheimer, 1938.

extended functional development. An organised environment provides an external scaffold for cognitive development, and it seems likely that increasingly rich parental and social scaffolding was an important factor in making it possible to extend development and open up cognitive plasticity in human evolution. Furthermore the high level of regulated plasticity in the brain helps explain how it is so easy to incorporate environmental organisation into our cognitive practices. For example, the brain makes a distinction between 'far space' (the space beyond reaching distance) and 'near space' (the space within reaching distance). There are different neural systems involved in processing each type of spatial information, and dissociations are possible in which, for example, awareness of 'near space' is lost whilst awareness of 'far space' remains intact. Berti and Frassinetti (2000) show that simply holding a stick causes a remapping of far space to near space. In effect the brain, at least for some purposes, treats the stick as though it were part of the body.

Acknowledgements

During the period when this paper was written I was a postdoctoral research fellow with the Konrad Lorenz Institute For Evolution and Cognition. A number of people have provided me with helpful comments and I am especially grateful to Luca Tommasi for many suggestions and requests for clarification. I would also like to thank the organisers of *Integrational Linguistics and Distributed Cognition*, Durban 2003, for a very enjoyable conference.

References

- Akitsuki, Y., Sugiura, M., Watanabe, J., Yamashita, K., Y., S., Awata, S., Matsuoka, H., Maeda, Y., Matsue, Y., Fukuda, H. & Kawashima, R., 2003. Context-Dependent Cortical Activation in Response to Financial Reward and Penalty: An Event-Related Fmri Study. *Neuroimage* 19, 1674-85.
- Allman, J.M., 1999. Evolving Brains. W. H. Freeman/Scientific American.
- Allman, J.M., Hakeem, A., Erwin, J.M., Nimchinsky, E. & Hof, P., 2001. The Anterior Cingulate Cortex: The Evolution of an Interface between Emotion and Cognition. *Annals of the New York Academy of Sciences* 935, 107-117.
- Arbib, M., Érdi, P. & Szentagothai, J., 1998. *Neural Organization: Structure, Function and Dynamics*. MIT Press.
- Bartlett, F.C., 1932. *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press.
- Bavelier, D. & Neville, H.J., 2002. Cross-Modal Plasticity: Where and How? *Nature Reviews Neuroscience* 3, 443-452.
- Bechara, A., Damasio, A.R., Damasio, H. & Anderson, S.W., 1994. Insensitivity to Future Consequences Following Damage to Human Prefrontal Cortex. *Cognition* 50, 7-15.
- Bechara, A., Damasio, H., Tranel, D. & Damasio, A.R., 1997. Deciding Advantageously before Knowing the Advantageous Strategy. *Science* 275, 1293-1294.
- Beer, R.D., 1995. Computational and Dynamical Languages for Autonomous Agents. In: Port, R.F. & Van Gelder, T., *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press, Cambridge, MA.
- Beer, R.D., 2000. Dynamical Approaches to Cognitive Science. *Trends in Cognitive Sciences* 4, 91-99.

- Bennett, A.T.D., 1996. Do Animals Have Cognitive Maps? *The Journal of Experimental Biology* 199, 219-224.
- Bennett, C.H., 1985. Dissipation, Information, Computational Complexity and the Definition of Organization. In: Pines, D., *Emerging Syntheses in Science. Proceedings of the Founding Workshops of the Santa Fe Institute.* Addison West Publishing Company, Redwood City, Calif.
- Berardi, N., Pizzorusso, T., Ratto, G.M. & Maffei, L., 2003. Molecular Basis of Plasticity in the Visual Cortex. *Trends in Neurosciences* 26, 369-378.
- Berti, A. & Frassinetti, F., 2000. When Far Becomes Near: Remapping of Space by Tool Use. *Journal of Cognitive Neuroscience* 12, 415-420.
- Bickhard, M.H., 1992. Scaffolding and Self Scaffolding: Central Aspects of Development. In: Winegar, L.T. & Valsiner, J., *Children's Development within Social Contexts: Research and Methodology.* Erlbaum, Hillsdale, N.J.
- Bickhard, M.H., 1993. Representational Content in Humans and Machines. *Journal of Experimental and Theoretical Artificial Intelligence* 5, 285-333.
- Bitterman, M.E., 1960. Toward a Comparative Psychology of Learning. *American Psychologist* 15, 704-712.
- Blakemore, S.J., Goodbody, S.J. & Wolpert, D.M., 1998. Predicting the Consequences of Our Own Actions: The Role of Sensorimotor Context Estimation. *Journal of Neuroscience* 18, 7511-7518.
- Blodgett, H.C., 1929. The Effect of the Introduction of Reward Upon the Maze Performance of Rats. *University of California Publications in Psychology* 4.
- Bothvinick, M., Nystrom, L.E., Fissell, K., Carter, C.S. & Cohen, J.D., 1999. Conflict Monitoring Versus Selection-for-Action in Anterior Cingulate Cortex. *Nature* 402, 179-181.
- Bouton, M. & King, D., 1983. Contextual Control of the Extinction of Conditioned Fear: Tests for the Associative Value of the Context. *Journal of Experimental Psychology: Animal Behavior Processes* 9, 248-265.
- Bowers, J.S. & Davis, C.J., 2004. Is Speech Perception Modular or Interactive? *Trends in Cognitive Sciences* 8, 3-5.
- Braitenberg, V., 1984. *Vehicles: Experiments in Synthetic Psychology.* MIT Press, Cambridge, MA.
- Brooks, R.A., 1986. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation* 1, 14-23.
- Brooks, R.A., 1991. Intelligence without Representation. *Artificial Intelligence* 47, 139-159.
- Brooks, R.A., 1997. From Earwigs to Humans. *Robotics and Autonomous Systems* 20, 291-304.
- Butts, R.E. ed. 1989. *William Whewell: Theory of Scientific Method,* Hackett Publishing Company, Indianapolis/Cambridge.
- Carruthers, P., 2002. The Cognitive Functions of Language. *Behavioral and Brain Sciences* 25, 657-726.
- Chomsky, N., 1959. A Review of Skinner's "Verbal Behavior". *Language* 35, 26-58.
- Christensen, W.D. & Hooker, C.A., 2000. An Interactivist-Constructivist Approach to Intelligence: Self-Directed Anticipative Learning. *Philosophical Psychology* 13, 5-45.
- Christensen, W.D. & Hooker, C.A., 2002. Self-Directed Agents. In: McIntosh, J.S., *Naturalism, Evolution & Intentionality.* University of Calgary Press, Calgary.
- Clark, A., 1997. *Being There: Putting Brain, Body and World Together Again.* MIT Press, Cambridge, MA.
- Clark, A., 1998. Magic Words: How Language Augments Human Computation. In: Carruthers, P. & Boucher, J., *Language and Thought: Interdisciplinary Themes.* Cambridge University Press, New York.
- Cohen, J.D. & Servan-Schreiber, D., 1992. Context, Cortex, and Dopamine: A Connectionist Approach to Behavior and Biology in Schizophrenia. *Psychological Review* 99, 45-77.
- Cohen, Y.E. & Andersen, R.A., 2002. A Common Reference Frame for Movement Plans in the Posterior Parietal Cortex. *Nature Reviews Neuroscience* 3, 553-562.
- Collett, T.S., Fauria, K., Dale, K. & Baron, J., 1997. Places and Patterns - a Study of Context Learning in Honeybees. *Journal of Comparative Physiology A* 181, 343-353.
- Collett, T.S. & Kelber, A., 1988. The Retrieval of Visuo-Spatial Memories by Honeybees. *Journal of Comparative Physiology A* 163, 145-150.
- Collier, J.D. & Hooker, C.A., 1999. Complexly Organised Dynamical Systems. *Open Systems and Information Dynamics* 6, 241-302.

- Czerwinski, M., Lightfoot, N. & Shiffrin, R.M., 1992. Automatization and Training in Visual Search. *American Journal of Psychology* 105, 271-315.
- Damasio, A.R., 1994. *Descartes' Error: Emotion, Reason and the Human Brain*. Hayrer Collins, New York.
- De Gelder, B. & Bertelson, P., 2003. Multisensory Integration, Perception and Ecological Validity. *Trends in Cognitive Sciences* 7, 460-467.
- Deary, I.J., 2001. Human Intelligence Differences: Towards a Combined Experimental–Differential Approach. *Trends in Cognitive Sciences* 5, 164-170.
- Dehaene, S., Kerszberg, M. & Changeux, J.-P., 1998. A Neuronal Model of a Global Workspace in Effortful Cognitive Tasks. *Proc. Natl. Acad. Sci. U. S. A.* 95, 14529–14534.
- Dietrich, E. & Markman, A.B., 2000. Cognitive Dynamics: Computation and Representation Regained. In: Dietrich, E. & Markman, A.B., *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*. Lawrence Erlbaum, Mahwah, NJ.
- Dietz, V., 2003. Spinal Cord Pattern Generators for Locomotion. *Clinical Neurophysiology* 114, 1379-1389.
- Duncan, J., 2001. An Adaptive Coding Model of Neural Function in Prefrontal Cortex. *Nature Reviews Neuroscience* 2, 820-829.
- Duncan, J., Seitz, R.J., Kolodny, J., Bor, D., Herzog, H., Ahmed, A., Newell, F.N. & Emslie, H., 2000. A Neural Basis for General Intelligence. *Science* 289, 457-460.
- Dusek, J.A. & Eichenbaum, H., 1997. The Hippocampus and Memory for Orderly Stimulus Relations. *Proc. Natl. Acad. Sci. USA* 94, 7109–7114.
- Eichenbaum, H., 2000. A Cortical–Hippocampal System for Declarative Memory. *Nature Reviews Neuroscience* 1, 41-50.
- Eichenbaum, H., 2001. The Hippocampus and Declarative Memory: Cognitive Mechanisms and Neural Codes. *Behavioural Brain Research* 127, 199-207.
- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M. & Tanila, H., 1999. The Hippocampus, Memory, and Place Cells: Is It Spatial Memory or a Memory Space? *Neuron* 23, 209-226.
- Ewing-Cobbs, L., Barnes, M.A. & Fletcher, J.M., 2003. Early Brain Injury in Children: Development and Reorganization of Cognitive Function. *Developmental Neuropsychology* 24, 669-704.
- Gallistel, C.R., 1994. Space and Time. In: Mackintosh, N.J., *Animal Learning and Cognition*. Academic Press, New York.
- Gazzaniga, M.S., Ivry, R.B. & Mangun, G.R., 1998. *Cognitive Neuroscience: The Biology of the Mind*. W. W. Norton & Company, New York.
- Gibson, K.R., 2002. Evolution of Human Intelligence: The Roles of Brain Size and Mental Construction. *Brain, Behavior, and Evolution* 59, 10-20.
- Gibson, K.R., Rumbaugh, D. & Beran, M., 2001. Bigger Is Better: Primate Brain Size in Relation to Cognition. In: Falk, D. & Gibson, K.R., *Evolutionary Anatomy of the Primate Cerebral Cortex*. Cambridge University Press, Cambridge.
- Giszter, S.F., Moxon, K.A., Rybak, I.A. & Chapin, J.K., 2001. Neurobiological and Neurorobotic Approaches to Control Architectures for a Humanoid Motor System. *Robotics and Autonomous Systems* 37, 219-235.
- Gluck, M.A., Meeter, M. & Myers, C.E., 2003. Computational Models of the Hippocampal Region: Linking Incremental Learning and Episodic Memory. *Trends in Cognitive Sciences* 7, 269-276.
- Gluck, M.A. & Myers, C.E., 2001. *Gateway to Memory: An Introduction to Neural Network Models of the Hippocampus and Learning*. MIT Press, Cambridge, MA.
- Goldman-Rakic, P.S., 1995. Architecture of the Prefrontal Cortex and the Central Executive. In: Grafman, J., Holyoak, K.J. & Boller, F., *Structure and Functions of the Human Prefrontal Cortex*. The New York Academy of Sciences, New York.
- Goldstone, R.L., Steyvers, M., Smith, J.S. & Kersten, A., 2000. Interactions between Perceptual and Conceptual Learning. In: Dietrich, E. & Markman, A.B., *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*. Lawrence Erlbaum, Mahwah, NJ.
- Greene, A.J., Spellman, B.A., Dusek, J.A., Eichenbaum, H.B. & Levy, W.B., 2001. Relational Learning with and without Awareness: Transitive Inference Using Nonverbal Stimuli in Humans. *Memory & Cognition* 29, 893-902.
- Grush, R., 1997. The Architecture of Representation. *Philosophical Psychology* 10, 5-24.

- Hendriks-Jansen, H., 1996. *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution and Human Thought*. MIT Press, Cambridge, MA.
- Hoshi, E. & Tanji, J., 2000. Integration of Target and Body-Part Information in the Premotor Cortex When Planning Action. *Nature* 408, 466-470.
- Hsieh, L., Gandour, J., Wong, D. & Hutchins, G.D., 2001. Functional Heterogeneity of Inferior Frontal Gyrus Is Shaped by Linguistic Experience. *Brain and Language* 76, 227-252.
- Hutchins, E., 1995. How a Cockpit Remembers Its Speeds. *Cognitive Science* 19, 265-288.
- Ivanchenko, V. & Jacobs, R.A., 2003. A Developmental Approach Aids Motor Learning. *Neural Computation* 15, 2051-2065.
- James, W., 1890. *The Principles of Psychology*. Holt, New York.
- James, W., 1997. What Pragmatism Means. In: Menand, L., *Pragmatism - a Reader*. Random House, New York.
- Jerison, H.J., 1973. *Evolution of the Brain and Intelligence*. Academic Press, New York.
- Johansson, R.S., 1998. Sensory Input and Control of Grip. *Novartis Found Symp* 218, 45-59; discussion 59-63.
- Karmiloff-Smith, A., 1992. *Beyond Modularity: A Developmental Perspective on Cognitive Science*. MIT Press.
- Karmiloff-Smith, A., 1998. Development Itself Is the Key to Understanding Developmental Disorders. *Trends in Cognitive Sciences* 2, 389-98.
- Kirsh, D., 1991. Today the Earwig, Tomorrow Man. *Artificial Intelligence* 47, 161-184.
- Knudsen, E.I., 2002. Instructed Learning in the Auditory Localization Pathway of the Barn Owl. *Nature* 417, 322-328.
- Krubitzer, L. & Kahn, D.M., 2003. Nature Versus Nurture Revisited: An Old Idea with a New Twist. *Progress in Neurobiology* 70, 33-52.
- Loeb, E.P., Giszter, S.F., Saltiel, P., Mussa-Ivaldi, F.A. & Bizzi, E., 2000. Output Units of Motor Behavior: An Experimental and Modeling Study. *Journal of Cognitive Neuroscience* 12, 78-97.
- Markman, A.B. & Dietrich, E., 2000. Extending the Classical View of Representation. *Trends in Cognitive Sciences* 4, 470-475.
- Medin, D.L., Lynch, E.B., Coley, J.D. & Atran, S., 1997. Categorization and Reasoning among Tree Experts: Do All Roads Lead to Rome? *Cognitive Psychology* 32, 49-96.
- Menzel, R. & Giurfa, M., 2001. Cognitive Architecture of a Minibrain: The Honeybee. *Trends in Cognitive Sciences* 5, 62-71.
- Miller, E.K., 2000. The Prefrontal Cortex and Cognitive Control. *Nature Reviews Neuroscience* 1, 59-65.
- Montague, P.R. & Sejnowski, T.J., 1994. The Predictive Brain: Temporal Coincidence and Temporal Order in Synaptic Learning Mechanisms. *Learning & Memory* 1, 1-33.
- Münste, Thomas F., Altenmüller, E. & Jäncke, L., 2002. The Musician's Brain as a Model of Neuroplasticity. *Nature Reviews Neuroscience* 3, 473-478.
- Neisser, U., 1997. The Ecological Study of Memory. *Philosophical Transactions of the The Royal Society of London B: Biological Sciences* 352, 1697-1701.
- Newell, A., 1990. *Unified Theories of Cognition*. Harvard University Press.
- O'keefe, J. & Nadel, L., 1978. *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford.
- Oden, D.L., Thompson, R.K.R. & Premack, D., 1988. Spontaneous Transfer of Matching by Infant Chimpanzees (*Pan Troglodytes*). *Journal of Experimental Psychology: Animal Behavior Processes* 14, 140-145.
- Oden, D.L., Thompson, R.K.R. & Premack, D., 1990. Infant Chimpanzees (*Pan Troglodytes*) Spontaneously Perceive Both Concrete and Abstract Same/Different Relations. *Child Development* 61, 621-631.
- Pearce, J.M. & Bouton, M.E., 2001. Theories of Associative Learning in Animals. *Annual Review of Psychology* 52, 111-139.
- Pfeifer, R. & Scheier, C., 1999. *Understanding Intelligence*. MIT Press, Cambridge, MA.
- Premack, D., 1983. The Codes of Man and Beast. *The Behavioral and Brain Sciences* 6, 125-137.

- Prinz, J. & Barsalou, L.W., 2000. Steering a Course for Embodied Representation. In: Dietrich, E. & Markman, A.B., *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*. Lawrence Erlbaum, Mahwah, NJ.
- Quartz, S.R. & Sejnowski, T.J., 1997. The Neural Basis of Cognitive Development: A Constructivist Manifesto. *Behavioural and Brain Sciences* 20, 537-596.
- Regier, T., 2003. Emergent Constraints on Word-Learning: A Computational Perspective. *Trends in Cognitive Sciences* 7, 263-268.
- Rizzolatti, G. & Arbib, M., 1998. Language within Our Grasp. *Trends in Neuroscience* 21, 188-94.
- Rochat, P., Querido, J.G. & Striano, T., 1999. Emerging Sensitivity to the Timing and Structure of Protoconversation in Early Infancy. *Developmental Psychology* 35, 950-7.
- Rolls, E.T., 2000. The Orbitofrontal Cortex and Reward. *Cerebral Cortex* 10, 284-294.
- Rosenzweig, M.R., 2003. Effects of Differential Experience on the Brain and Behavior. *Developmental Neuropsychology* 24, 523-40.
- Rumbaugh, D.M., 1970. Learning Skills of Anthropoids. In: Rosenblum, L., *Primate Behavior: Developments in Field and Laboratory Research*. Academic Press, New York.
- Schultz, W., 2000. Multiple Reward Signals in the Brain. *Nature Reviews Neuroscience* 1, 199-207.
- Schultz, W., Dayan, P. & Montague, P.R., 1997. A Neural Substrate of Prediction and Reward. *Science* 275, 1593-9.
- Schyns, P.G., Goldstone, R.L. & Thibaut, J.-P., 1998. The Development of Features in Object Concepts. *Behavioral and Brain Sciences* 21, 1-54.
- Shafto, P. & Coley, J.D., 2003. Development of Categorization and Reasoning in the Natural World: Novices to Experts, Naive Similarity to Ecological Knowledge. *J Exp Psychol Learn Mem Cogn* 29, 641-9.
- Sieg, W., 1999. Formal Systems, Properties Of. In: Wilson, R.A. & Keil, F., *The Mit Encyclopedia of the Cognitive Sciences*. Bradford Books.
- Smith, L.V. & Thelen, E. (eds., 1993. *A Dynamics Systems Approach to the Development of Cognition and Action*, Bradford Books, Cambridge, MA.
- Solomon, K.O., Medin, D.L. & Lynch, E., 1999. Concepts Do More Than Categorize. *Trends in Cognitive Sciences* 3, 99-105.
- Spelke, E., 2003. What Makes Us Smart? Core Knowledge and Natural Language. In: Gentner, D. & Goldin-Meadow, S., *Language in Mind: Advances in the Study of Language and Thought*. Bradford Books/MIT Press, Cambridge, MA.
- Sterelny, K., 2003. *Thought in a Hostile World*. Blackwell Publishing.
- Stiles, J., 2000. Neural Plasticity and Cognitive Development. *Developmental Neuropsychology* 18, 237-272.
- Suri, R.E., Bargas, J. & Arbib, M.A., 2001. Modeling Functions of Striatal Dopamine Modulation in Learning and Planning. *Neuroscience* 103, 65-85.
- Thompson, R.K.R. & Oden, D.L., 2000. Categorical Perception and Conceptual Judgments by Nonhuman Primates: The Paleological Monkey and the Analogical Ape. *Cognitive Science* 24, 363-396.
- Thompson-Schill, S.L., D'Esposito, M., Aguirre, G.K. & Farah, M.J., 1997. Role of Left Inferior Prefrontal Cortex in Retrieval of Semantic Knowledge: A Reevaluation. *PNAS* 94, 14792-14797.
- Thornton, S., 1999. Creating the Conditions for Cognitive Change: The Interaction between Task Structures and Specific Strategies. *Child Development* 70, 588-603.
- Tolman, E.C., 1932/1951. *Purposive Behavior in Animals and Men*. University of California Press, Berkeley, CA.
- Tolman, E.C., 1948. Cognitive Maps in Rats and Men. *The Psychological Review* 55, 189-208.
- Tolman, E.C. & Honzik, C., 1930. Introduction and Removal of Reward, and Maze Performance in Rats. *University of California Publications in Psychology* 4, 257-275.
- Tononi, G. & Edelman, G.M., 1998. Consciousness and Complexity. *Science* 282, 1846-1851.
- Tononi, G., Edelman, G.M. & Sporns, O., 1998. Complexity and Coherency: Integrating Information in the Brain. *Trends in Cognitive Sciences* 2, 474-484.
- Turken, A.U. & Swick, D., 1999. Response Selection in the Human Anterior Cingulate Cortex. *Nature Neuroscience* 2, 920-924.

- Ungerleider, L.G., Doyon, J. & Karni, A., 2002. Imaging Brain Plasticity During Motor Skill Learning. *Neurobiology of Learning and Memory* 78, 553-564.
- Van Gelder, T., 1995. What Might Cognition Be, If Not Computation? *The Journal of Philosophy* 92.
- Van Gelder, T., 1998. The Dynamical Hypothesis in Cognitive Science. *Behavioral and Brain Sciences* 21, 615-627.
- Wallenstein, G.V., Hasselmo, M.E. & Eichenbaum, H., 1998. The Hippocampus as an Associator of Discontiguous Events. *Trends in Neuroscience* 21, 317-323.
- Wang, R.F. & Spelke, E.S., 2002. Human Spatial Representation: Insights from Animals. *Trends in Cognitive Sciences* 6, 376-382.
- Wertheimer, M., 1938. Gestalt Theory. In: Ellis, W.D., *Source Book of Gestalt Psychology*. Harcourt, Brace and Co, New York.
- Wolpert, D.M., Ghahramani, Z. & Jordan, M.I., 1995. An Internal Model for Sensorimotor Integration. *Science* 269, 1880-1882.
- Xu, F. & Carey, S., 1996. Infant Metaphysics: The Case of Numerical Identity. *Cognitive Psychology* 30, 111-153.