WAYNE CHRISTENSEN

SELF-DIRECTEDNESS: A PROCESS APPROACH TO COGNITION

ABSTRACT. Standard approaches to cognition emphasise structures (representations and rules) much more than processes, in part because this appears to be necessary to capture the normative features of cognition. However the resultant models are inflexible and face the problem of computational intractability. I argue that the ability of real world cognition to cope with complexity results from deep and subtle coupling between cognitive and non-cognitive processes. In order to capture this, theories of cognition must shift from a structural rule-defined conception of cognition to a thoroughgoing embedded process approach.

1. INTRODUCTION

Despite the fact that processes predominate in the natural world, thoroughgoing process theories are rare in science and philosophy; for many reasons structures tend to take the focus of attention. In this paper I will argue that the emphasis on structures has led to deep problems in cognitive science, and that contemporary evidence favours a systematically process-oriented approach. Developing such an approach poses enormous challenges because it requires us to rethink the central issues for understanding cognition, but the payoff is a more integrated perspective that helps reframe traditional problems and opens up a more productive relationship between high-level theory and the various empirical sciences of cognition.

2. COGNITION SANDWICHED BETWEEN STRUCTURES

Standard approaches to cognition emphasize structures far more than processes. In particular, the following structural assumptions are typical in cognitive science:[1]

- Representations are structures (symbols) that have the function of corresponding to features of the world.
- Cognitive processes are operations on representations.

Axiomathes **14:** 157–175, 2004.

- Cognitive architectures are structures that specify possible cognitive processes.
- The world has a structural ontology consisting of objects, properties and categories.

It has been very difficult to think about cognition in any other way. Cognitive processes are intuitively thought of as transformations between perception, ideas and action. As such, it is clear that they can vary in how well they work, with accurate perception and reliable judgement of consequences being clearly important for cognitive success. These factors in turn seem to essentially require that cognitive processes satisfy certain kinds of relations: truth and the principles of valid inference being primary candidates. These are then taken to be the normative standards for cognition.

In explaining normative standards it is natural to think that there must be an encompassing cognitive system that defines possible meaningful relationships and rules for cognitive processes (e.g., the principles of rationality, the computational architecture of classical artificial intelligence, or Chomsky's universal grammar), and basic structures that serve as *meaning bearers*. In a complementary way, the world is thought of as consisting of objects with defining properties that interact in a law-like manner and are organised into hierarchical categories, such as natural kinds.[2] Cognitive processes must be rationally structured so that they can represent these relationships.

So, once we have bought into the compelling idea that cognition is systematic we are then forced from cognitive processes upwards to general structures that define the systematicity of the processes, and downwards to the specific structural types that participate in the processes. Cognition therefore seems to be necessarily sandwiched between structures below and above. And when it works *well* – that is, it is rational – the structural articulation of cognition isomorphically mirrors the structured relationships in the world.

3. COGNITIVISM AND THE PROBLEM OF FLEXIBILITY

Cognitivism is the modern elaboration of these intuitions, and normativity enforces its structuralist orientation with an iron grip: cognition can only be rational if it is rule-governed. However there is a fundamental tension between the idea that cognition is rule-governed and another important property of intelligence: *flexibility*. Complex cognition is highly context sensitive in at least two interrelated ways. Firstly, the interpretation of meaning is shaped by background knowledge and knowledge of the context. When dining in a restaurant, for example, small cues like placing the knife and fork together on a half finished plate can indicate both that you want the plate removed and that you were not happy with the meal, so reducing the waiter's expectations concerning the tip. Secondly, the control of action can be made sensitive to many simultaneous factors in a situation. Whether you order the salmon can depend on a combination of vital factors, including the way it is prepared, the sauce, the selection of white wines, which desert you would prefer, the evening meal you had the previous day, whether the restaurant has a reputation for sea food, and the counter-temptation of the lamb cutlet. I shall generically refer to this ability to take many factors into account when deciding action as *situational awareness*, and argue that it is a crucial factor in most complex cognitive activities because it facilitates flexible, context-sensitive action.

At first glance situational awareness seems to be just a matter of accurately representing the world in the manner envisioned by the traditional normative account of cognition described above. Yet, curiously, cognitivist artificial intelligence (AI) – which has carried the rule-governed approach to its logical conclusion by modelling cognitive processes as formal symbolic processes such as algorithms and heuristic searches of problem spaces – has faced apparently intractable difficulties in trying to understand situational awareness. One key difficulty concerns the fact that something like situational awareness is required for cognition at all. Early AI approaches attempted to model reasoning abilities that only employed general principles in order to solve problems, such as Newell and Simon's General Problem Solver program (Newell and Simon 1972). However, it became apparent that most human problem solving involves 'knowledge-based' rather than 'general' reasoning. The AI approach to modelling so-called knowledge-based reasoning has been to program AI agents with large databases encoding 'common sense' and 'domain specific' information. Such a database is assumed to constitute a representation of a problem domain, and the problem solving process proceeds by searching the space in some way. However this strategy has led directly to further problems, in particular computational intractability due to combinatorial explosion.

Combinatorial explosion occurs because the computational task of updating perception and calculating the effects of changes in the world grows exponentially as the complexity of the agent's representation of the world is increased. The well known 'frame problem' is related to this and concerns the issue of representing a dynamic environment. To keep computation manageable it has been assumed that the agent should restrict its calculation of states of the world to only those properties that are changed by an event. Thus, to use a standard example, when a robot calculates the effects that result when it moves into a room, it should *not* have to compute the fact that the colour of the room will remain the same. Without such an assumption it would seem that for *each* event that occurs the agent would need to search its complete knowledge base and recalculate the entire state of the world. The trouble with implementing the idea lies in finding a way of pre-specifying what will change and what won't for any given event. One strategy has been to build in a default assumption that properties persist, unless there is an explicit rule to the contrary. However this leads either to implausibly simplistic assumptions about the world or an inordinate number of rules: most properties change in at least some circumstances.

Indeed, whilst the motivation for such a strategy seems straightforward enough within the computational framework in which it arose, in the light of more recent research emphasising cognition in realistic situations it looks rather dubious. To assume prevailing stasis as a way of achieving computational tractability is effectively an attempt to represent a dynamic world by assuming that it is not very dynamic. Clearly not everything changes all the time, or cognition would be impossible. But on the other hand the real-world problem contexts that demand intelligence are *characteristically* highly dynamic. Predominantly static environments are not ones that require or would give rise to intelligence. The crux of the issue can be seen more clearly if we contrast the frame problem with the human ability to be sensitive to relevance. A human performing a skilled activity at a high standard is able to focus quickly on *just* the relevant factors in a situation, and to quickly to shift attention as changes in the situation make new factors relevant. Attempting to restrict calculation to actual change in the environment is, per se, at best a very crude approximation to the ability to detect relevance: situational factors can change and yet not be relevant to the agent given the agent's goals, whilst a change in goals can make new factors relevant to the agent that have not in themselves changed. So the frame problem as understood in AI is really just the tip of the iceberg, so far as the problem of effective cognitive functioning in a complex environment is concerned.

Two principle strategies for solving the frame problem have been pursued in AI: (1) finding ontologies for knowledge representation that aptly express relationships in the world, and (2) finding effective search heuristics. An important thing to note is that these are general computational strategies. As such, there is something essentially misconceived about them: in environments complex enough to require intelligence the information that is relevant to action varies in a highly context-dependent way.

Knowing what is relevant is therefore inherently something that cannot be achieved via a general computational strategy. Situation-specific learning must be involved. This is supported by the empirical fact that in highly novel situations people have great difficulty in effectively differentiating between relevant and irrelevant information.

A possible response to this might be to propose a layered solution that involves first solving the representational ontology and heuristic computational issues, and then implementing learning processes that refine the agent's knowledge about particular situations. The difficulty is that, no matter how the representational ontology is constructed, realistic cognitive situations have enough significant variables to present astronomical numbers of possible states. And since AI learning techniques rely on searching the represented possibility space, learning itself becomes intractable. In other words there is simply no way around the problem of combinatorial explosion. It might be protested that this simply can't be true, since we have an existence proof that cognition is possible in situations of at least the complexity that humans face; namely, human cognition itself. But this fails to consider the possibility that human cognition might operate according to very different principles to those assumed by cognitivist AI.

To give some flesh to this prospect we can consider briefly some of the key tenets of alternative paradigms, in particular pragmatism and situated cognition research. Perhaps the most basic principle of situated cognition is that an adaptive agent doesn't have to represent the world internally in order to successfully interact with it. Brooks (1991) argued that a critical problem with the cognitivist AI approach was the fact that all information used to control behaviour had to first pass through a central representational module. Robots built according to this principle became paralysed as they attempted to update perception and calculate what action to take. He advocated an alternative design paradigm known as behaviour-based robotics, which entirely dispensed with internal representation of the world and instead relied on well-tuned actions triggered by environmental stimuli.

However, although Brooks originally proposed that behaviour-based robotics could serve as a model for intelligence generally, it soon became apparent that it doesn't scale up very well. One of the basic reasons is very simple: some kinds of tasks require complex sequences of actions, and it is difficult or impossible to find external cues that will evoke the actions in the right order.[3] Consequently, in these circumstances some form of internal control of action sequencing is required. This brings us back to the issue of representation, and one stream of research in robotics has sought to develop hybrid systems that combine elements of the behaviour- based and classical approaches in an effort to gain the benefits of each. However this would seem to have only limited prospects, since the original problems with AI have not been solved.

The pragmatist tradition suggests the basis for an alternative route. AI robots became bogged down because to do anything they require extremely detailed, and hence computationally expensive, representational models of the world. Behaviour-based robots gained real-time effectiveness by dispensing with representations, but at the cost of being unable to perform tasks that require internal ordering of action. In contrast, pragmatism does not eschew representation, but it shares with behaviour-based robotics the idea that the world is not represented in cognition as a collection of objects with intrinsic properties and law-like behaviours. Rather, representation emerges out of the control of action. The agent learns the properties of objects by interacting with them, and is particularly focussed on learning those factors that affect action success.[4] This is very different to the cognitivist approach to representation, and the differences have important implications for understanding the way in which real cognitive agents are able to manage complex relationships. For cognitivism representation is *prior* to action, whereas for pragmatism it emerges *out of* action and is learned in an action- relevant way. Thus, the interaction process acts as a filter helping to determine the relevance of the information that goes into the *formation* of representations.

The nature of the representations on this interpretation is also inherently action-relevant. The base representational format for cognitivism is explicit description, whereas for pragmatist approaches it is some form of control function. J. J. Gibson's theory of affordances, for example, claims that we perceive the world in terms of opportunities to interact with it (Gibson 1977). Empirical evidence provides support for this: both infants and adults perceive objects and spatial relations in relation to their body size and orientation, and respond to stimuli prospectively in terms of the actions they may need to perform (e.g., Bertenthal and Clifton 1997). Control-based representations can be much more efficient than explicit description because they are able to rely on implicit regularities, for example through implicit predication (Bickhard, this volume), and can exploit distributed sensorimotor and affective information.

Clearly humans do *acquire* knowledge in a form that makes possible explicit description, referred to as declarative knowledge. But it is learned in a rich sensorimotor context and there is evidence that deeper non-declarative forms of knowledge are required in order for declarative knowledge to appropriately influence action (e.g., Bechara et al. 1997). Research in cognitive linguistics also suggests that conceptualisation is deeply shaped by embodiment (e.g., Lakoff 1987). Moreover a further advantage of control-based representation is that it is easy to understand how the burden of control can be distributed between the agent and the environment, and redistributed through learning. Research on human cognition in real world contexts has shown that humans can simplify the complexity of the cognitive problems they must solve during a complex task by using environmental organisation – such as the layout of the navigation deck on a warship or an aeroplane cockpit – as part of the problem solving process (Hutchins 1995a, b).

Situated cognition research thus suggests that there are at least two complementary ways in which real cognitive agents avoid the computational complexity problems that afflict AI, or at least transform them into a tractable form.[5] Firstly, in many cases they avoid the need for representing the environment by allowing the environment itself to organise their behaviour. Secondly, where representation is required, it is constructed from interaction in a way that is inherently related to action, and is guided by learning processes that help build in relevance. At the most general level, then, we can see that situated cognition gains an advantage by treating cognitive processes as deeply coupled to non-cognitive processes. In particular, it treats the organisation of cognition as being in part the result of interaction, which means that understanding how cognition can focus on the relevant aspects of interaction is much less problematic.

On the other hand, from a traditional perspective the price is high. Formal computation provides a framework that appears to satisfy the key normative requirements that a theory of cognition must satisfy, by showing how cognitive processes can be both representational and rule-governed. On the pragmatist account it appears that the formation of representations must be a non-cognitive process, since *ex hypothesi* the process is not itself mediated by representations and rule-governed cognitive processes. Moreover, since the properties of the pragmatist representations lie beyond the scope of cognitive norms, it would appear that the behaviour of these representations in cognitive processes must also be cognitively ill-defined. It is therefore not at all clear how normative cognitive standards could be applied to pragmatist cognition. I will examine this problem in more detail in the next section, but first I will consider one final strategy for rescuing cognitivism because it highlights just how important the issue of flexibility is.

Another approach to the problem of combinatorial explosion that remains within the classical framework is to try to bypass it by compartmentalising cognition into modules, with each module being responsible only for some specialised cognitive domain. Computational complexity is reduced because the cognitive module is equipped with just a limited set of domain specific knowledge. It is also possible to use representational formats for each domain that facilitate the type of computations performed in the domain (e.g., visual or auditory perception), further speeding cognition. And it has been further argued that evolution should have furnished us with a suite of cognitive modules apt for solving the problems we characteristically face.

However there are a number of problems with the so- called 'massive modularity' principle as a solution for cognitive science, one of the more striking being that it abandons, or at least greatly reduces, any attempt to capture the generality or flexibility of cognition. Instead, cognitive problems are claimed to fall within well-defined domains for which we are equipped with pre-specified solutions. It is certainly true that human intelligence shows considerable domain specific organisation: a chef may be able to make highly intelligent choices in a kitchen, but this will not translate in any direct way to an ability to be an effective platoon commander in combat. But it is also true that in many normal situations we effortlessly blend information from a variety of types of sources, and are able to call up various types of information as it becomes relevant. A simple shopping trip to the supermarket, for instance, calls on numerical knowledge, cooking knowledge, experience with the products available, ability to deal with parking lots, queues and checkout personnel, and so on.

One possibility which might seem plausible is that, although successfully completing the overall shopping trip requires the input of knowledge from a variety of sources, it can be broken down into a series of sub-problems, such as parking, finding the right entrance, selecting the products, paying at the cash-register, each of which could be managed by a domain-specific computational module (spatial, numerical, social, etc.). But whilst there are some situations that might be treated in a strictly modular way, modularity cannot be a general solution to effective behaviour in realistic situations. Many of the phases of the shopping process involve some varying *mixture* of kinds of knowledge, with a high degree of interdependency between background factors across domains.

For example, having invited a friend to dinner whom I want to impress, I plan to cook my most effective dish. Unfortunately the supermarket doesn't have the required ingredients, and nor does it have the ingredients for my second most effective dish. The butcher next door has some high quality steak, but I'm not entirely confident about making an appropriate sauce, and including a decent bottle of red wine will push the total cost up to the point where I will be without money for the last two days before payday. A Thai green curry with beer or a Riesling will be much cheaper, but might not create the intimate social atmosphere I'm aiming for. A simple gorgonzola gnocchi accompanied by an aged Semillon could set a more appropriate mood. But the steak will have more class. Can I make the sauce? Can I borrow money from a friend 'till payday? I'm strongly tempted to gamble on the sauce, but I also know that almost all of my friends are even more likely to have spent all their money than I am, and the only one who isn't is out of town. Reluctantly, I toss the packet of gnocchi into my shopping basket and head for the cheese aisle.

Fodor (2000) argues that abduction (inference to the best explanation) is an important feature of human reasoning, and consequently the massive modularity thesis cannot be right. Nor, indeed, can any variant of computational cognitivism. The thrust of his argument is that abduction requires calling on information from many different domains for a given inference. The domains therefore *can't* be strongly modular in the way that the massive modularity thesis supposes, and cognitive science has no theory of how knowledge can be organised in the global ways required for abduction. My supermarket scenario is designed to suggest that this argument is essentially correct. But it is also too limited: the problems concern flexibility in intelligence generally, not just abduction, especially if abduction is interpreted to be a special, epistemically demanding cognitive process, perhaps distinctive to science as Carruthers (forthcoming) proposes. To the contrary, the ability to call on many different forms of knowledge fluidly and appropriately as the situation unfolds is a feature of most human practical reasoning.

Insofar as cognition does show domain organisation, extensive evidence in developmental and cognitive neuroscience suggests that these domains are permeable: their inner workings are accessible to each other and to higher executive processes that are not domain specific. A process as important to cognition as learning the properties of an object, for instance, requires the coordination of information from different modalities, such as visual shape and tactile shape. Learning about objects thus requires integrating information across modalities and through time. Moreover it also appears to be the case that the principles governing the specialisation of cognitive organisation in the higher brain areas are not primarily geared towards separation into complete problem types, which is what is required for domain specific modularity in the sense assumed by the massive modularity thesis. Rather, they seem to concern *aspects* of complex behaviour – numerical, social, culinary, and so on – that are *combined* to produce situationally appropriate action. For example, individuals with brain lesions may lose the ability to recognise tools but not animals, or vice versa, suggesting that specialised regions process these forms of knowledge.

Nevertheless there is no barrier in normal cognition to thinking about both animals and tools at the same time or combining the information in problem solving, and there is evidence that retrieval of different types of semantic information is managed by executive systems in the prefrontal cortex (Thompson-Schill 1997). Moreover, a study by Tranel et al. (1997) has shown that the category related specialisation can be explained in terms of perceptual and use-related factors, suggesting that the specialisation is the result of learning rather than innate.

4. COGNITION ENMESHED IN PROCESSES

Cognitivist AI has struggled to explain how agents could act intelligently in complex environments because on the one hand this requires extensive situational awareness, and on the other hand such awareness seems to immediately produce computational intractability. In the previous section I argued that pragmatist approaches might provide a way out of the problem because they treat cognition as being intimately coupled with other processes, thus helping to explain how cognition could be tuned to the relevant aspects of interaction. Cognitivism hasn't looked for such coupling, and it is inherently difficult for it to explain.

The reason for this lies in the fact that formal systems are essentially closed. If cognition is formal computation there are only two ways in which a non- cognitive process can be related to a cognitive process: as an implementation, or as an input that is first represented and then enters the computational stream. Any other kind of influence can only interfere with cognition, since (a) an influence that results in the production of an undefined symbol (or symbols) cannot be treated as meaningful by the computational rules that transform symbolic statements into other statements, (b) an influence that manifests as a change in the rules that define the computations that the system can perform will result in false, invalid or meaningless output, and (c) any causal effect that does not change the symbol properties or computational rules (such as painting the computer green) will have no computational influence at all. The whole point of the development of formal systems theory has been to eliminate the need for subjective (i.e., not formally defined) interpretation in proofs by regimenting the inferential steps as simple mechanical rules (Seig 1999). But this makes it apparent why any formalist theory of cognition will face the problem of computational intractability: all coupling between the cognitive system and the world must be mediated by representations, and realistic cognitive situations have an extremely large number of potentially relevant factors. To be able to represent the inputs and solutions to the problems that may arise the system has to have an extremely large representational space, but then it faces the problem that all the cognitive processes it must perform involve searching that space.

4.1. *Reciprocal coupling between low and high order cognition*

At this point it is worth examining the issue of the formation of higher order cognitive organisation, including concepts, models etc. The ability to be sensitive to relevance is connected to the ability to form higher order cognitive organisation because it involves, at least in part, grouping into like and unlike, important and unimportant. The basic principle by which cognitivism explains higher cognitive organisation is composition: simple representational constituents are combined to form complex expressions. However there is evidence that suggests some important forms of cognition do not follow principles of compositionality. Traditional theories of speech perception have favoured building block models in which auditory perception is processed for phonetic categorization and phonemes are assembled into words and higher grammatical structures. But recent evidence has shown that phoneme perception is strongly sensitive to higher cognitive influence. That is, the perception of a sound pattern as a phoneme is greatly influenced by the linguistic context and the expectations of the listener (Remez 2000).

On this basis, Remez argues that the reification of linguistic entities as ‘things’ that are assembled is mistaken, and that a ‘system of differences’ approach better explains the nature of information processing in speech perception. Though Remez doesn’t quite describe it like this, the ‘system of differences’ idea treats the information processing ability of the system as an emergent and variable property affected by learning. The problem of reducing the ambiguity in a perceptual signal is achieved by the imposition of constraints restricting the possible interpretations, and these constraints can be assembled in higher cognition and used to shape early perception. For example, without contextual information, to interpret a group of sounds as a word an English speaking listener has 100 000 possible words to choose from. With sentence and pragmatic information about speaker intentions the possibilities can be much reduced, sometimes to the point where no sound signal at all is required in order to guess the word – e.g., “A stitch in time saves…” (ibid., p. 108). Thus, increases in the higher knowledge that an individual can bring to bear in a situation help to improve the sharpness of the perceptual differentiations the individual can make. Neisser (1997) points to the same principle as a crucial factor in the accuracy of memory recall.

The principle of convergent constraints provides a way for high order cognition to improve low order cognitive processes such as auditory perception and word recall. Goldstone et al. (2000) demonstrate an even stronger top-down effect in visual perception (one that is implicit in Remez' discussion) whereby early perceptual feature recognition is learned, and the learning is guided by high order conceptual knowledge. They cite neurological evidence showing that cortical areas involved with early perceptual processes can be flexibly and context sensitively tuned by training, and behavioural evidence that in areas as diverse as radiology and beer tasting experts make different perceptual differentiations compared to novices (pp. 193–194). Experimentally they show two related effects in perceptual learning: dimension sensitization and unitization. Dimension sensitization involves paying increasing attention to stimulus dimensions (e.g., colour, shape, etc.) that are important for classifications. These attention biases can occur in early perception, where they are the result of perceptual learning, or late perception, where they are based on strategic judgment concerning the situation. Indeed, the prior learning of the individual can affect the way in which a whole is decomposed into parts. Unitization involves grouping together a complex of features into a single functional whole. For instance the letter *A* may initially be perceived as a group of oriented lines, but after extended learning come to be perceived as a functional unit (ibid., p. 215).

Collectively this evidence suggests that, rather than being combinatorial, cognition – at least in these cases – involves mutual influence between sub-featural, featural and higher order cognition. Rather than being stuck with a fixed set of representational primitives, cognition is able to construct new representational abilities based on higher order task demands and conceptual learning. Furthermore, to a significant extent the robustness and precision of cognition depends, not on stable representational atoms and algorithms that are built into the fundamental architecture, but on flexible cognitive processes that construct higher order contextual information which then serves to refine lower order processes in a reciprocal and open ended cycle of adjustment.

The reciprocal modulation between low and high order cognition indicates that there is a significant amount of plasticity in cognition, and that this plasticity is regulated by functional activity. Evidence from neuroscience concerning activity-dependent plasticity in neurons and neural systems provides some explanation for how this is possible. Hebbian learning is a simple example: the synaptic connection strength between two neurons increases if their activity is correlated, and decreases if it is uncorrelated. But there are also many other kinds of activity-dependent plasticity, ranging from changes in dendritic morphology based on activity in the immediate vicinity (Quartz and Sejnowski 1997), regulation of plasticity in a cortical region by reward systems (Schultz 2000), through to global synaptic scaling mechanisms (Desai et al. 2002). The sheer extent of activity-dependent plasticity in the brain suggests that a great deal of cognitive functional organisation is shaped by the activities performed by the individual (e.g., Stiles 2000). Direct evidence for this comes from neuroimaging studies showing that musicians have functional and anatomical brain differences related to music processing (Münte et al. 2002).

Such activity dependent plasticity helps explain how the reciprocal connection between low and high order cognition is possible, but it is very problematic for cognitivism for several reasons. It violates the ‘virtual machine’ hypothesis, which is a straightforward corollary of the assumption that cognition is formal computation, and from a methodological point of view is used as a reason to ignore neural phenomena as being merely a matter of implementation. But activity-dependent plasticity criss-crosses the supposed boundary between cognition and implementation, changing lower level processes based on higher functional processes, and in turn changing the nature of the higher processes. Activity-dependent plasticity also violates the cognitivist structure-content distinction, which ensures that cognitive processes are only sensitive to syntax and are hence algorithmically implementable. More subtly, the structure-content distinction reflects the normative idea that cognition does – or should – operate according to general rules that employ placeholders for representations of particulars. However the separation between rules and content also makes the system inflexible: precisely because of this, the system can’t modify the rules it uses. In contrast, the *content* of what a person sees can be used to change the information processing used to make visual discriminations.

Thus, although the presence of an interconnection between content and structure violates some deep assumptions about computation, it has important advantages for learning: it allows the functional organisation of the system to be tuned by the nature of its activities. There are a large variety of neural processes that show activity-dependent plasticity, and this provides an extremely large number of degrees of freedom for the functional organisation of cognition. These degrees of freedom are selectively locked down or released by regulative signals operating with a range of spatial, temporal and pathway specific characteristics. Stepping back from the details, the overall regulative loop from activity to architectural change can build properties of the organism and environment into the functional organisation of the neural system. The selective regulation ensures that the various types of plasticity are only expressed in restricted circumstances. This helps to minimise the potential complexity the cognitive system must deal with at any given time by pre-tuning it to its context, whilst preserving a background capacity to adjust to different functional contexts.

4.2. *The process context of intelligent action*

Just as the structuralist rule-based conception of cognition is flawed, so too is the cognitivist world ontology, especially as it relates to intelligent action. Cognitivism treats the world as an agent-independent, hierarchically structured collection of objects. But from the point of view of intelligent action the world is an agent-related complex of interdependent processes. Of course, objects are an important part of an agent's environment, and object tracking, categorisation and modelling is a significant part of cognition. But objects are of derivative, not primary importance, and tracking process relationships presents different cognitive demands to that of recognising objects per se.

The situations that form the context for intelligent action are much *more* than just assemblages of objects because intelligent agents are biological organisms, and as such fundamentally constituted in and shaped by processes.[6] The most basic biological process is the reproduction cycle, but interconnected processes define virtually all aspects of life. For humans these include the sleep-wake cycle, the hunger-satiation cycle, the working day, the budgetary cycle between paydays, the dynamics of social relations, and so on. These processes are multidimensional, often continuously valued, and interdependent. And it is the process interdependencies that are the source of the constraints that differentiate better and worse action. Metabolism is the fundamental process that makes action possible for a biological organism, and also determines one of the primary goals for organismic behaviour: acquire food. For many humans activity in the workplace that is useful, or apparently useful, to the larger processes of the employer organisation is the means for acquiring food and other resources. But there is more to human life than sustaining metabolism. Human goals are such that quality of life requires taking part in other processes, including recreational, social and intellectual, and these can compete with work and basic metabolism for the individual's resources. Thus, any given situation within which cognitive problems occur – the workplace, the supermarket, dinner with a friend, etc. – exists within a complex network of process relationships. A single action in a particular situation can have ramifications for many of these relationships, so the challenge for the individual is to find the best way of balancing the interdependencies.

Cognition is thus enmeshed in processes, and it is this that makes flexibility such a central feature of intelligence. As factors in a situation shift, action needs to be reorganised. After spending five minutes in the queue at the checkout I remember that my flatmate is going to give me money for the telephone bill tonight, and I'm fairly sure that the final demand from the telephone company won't arrive before payday. This means that I can put off actually paying the bill until then, and since I can live off the telephone money, I can afford the steak. I leave the queue, return the gnocchi and gorgonzola to the shelves, and go looking for a good bottle of red.

5. SELF-DIRECTEDNESS, NORMS AND FLEXIBILITY

A great deal of recent cognitive science has abandoned the strict formalist framework of cognitivist AI, but has failed to supply any alternative normative conception of cognition. For a philosopher steeped in the traditional rationality project this trend amounts to a descent into the maelstrom.[7] In Section 3 I raised the objection against pragmatism that it does not appear to be normatively evaluable by traditional standards. Unless an alternative theory of cognitive norms can be supplied this is a critical problem because there is no account of the distinction between better and worse cognition, and any theory that cannot explain this distinction is at the very least seriously incomplete.

Pragmatist approaches do appeal to a normative standard: the ability to achieve effective interaction. The trouble is that per se this appears to be a merely functional rather than cognitive norm. Traditionalists are right to want distinctively cognitive norms, because it is only with such norms that cognition can be effectively distinguished from other kinds of functional processes. But they are wrong to look for an entirely sui generis theory of such norms. Cognitive processes have evolved as specialisations of deeper biological processes, and even in their highest order form they are intimately connected to other kinds of processes. It follows that whilst cognition may have some distinctive normative attributes, its norms are entwined with functional biological norms.

Elsewhere in collaboration I have developed a theory of self-directed agents that characterises the evolutionary emergence of cognition as a capacity deeply integrated with other biological processes. The central idea is very similar to Romanes' claim that the evolution of intelligence is marked by a transition from reflexive behaviour to more flexible action under higher 'mental' control (Romanes 1904). The account of self-directedness focuses on the ability to integrate multiple factors in action production, permitting action to be sensitive to a broader range of conditions and hence to be able to succeed in complex variable situations where reflex action cannot be relied on (Christensen and Hooker 2000, 2001). There is not the space to present the details of this theory; for the current purposes I shall focus on a few key aspects of it.

The first point to make is that although animal cognition research has traditionally focused on stimulus-response learning, there is evidence that more complex integrative abilities are phylogenetically surprisingly ancient. Invertebrates such as locusts and honeybees possess integrative neural systems and are capable of multidimensional tradeoffs and abstract learning (Raubenheimer 1999; Menzel and Giurfa 2001). Self-directedness arises from integration involving a number of key interrelated organism processes: interaction in a complex environment, affect, motor control, and development. Complex context-sensitive action requires coordination between these processes. Foraging problems can require the capacity for time and space awareness, and choice involving multidimensional tradeoffs. Affective processes must detect conditions relevant to the animal's viability, including states such as hunger, tiredness and the presence of things in the environment that should be approached or avoided. Complex motor control requires coordination between multiple sensorimotor maps, dynamic control and predictive modelling. Development involves building and maintaining systems of sensorimotor coordination and the learning of skills. The elaboration and integration of all these processes has been achieved in brain evolution through the addition of new structures that refine and coordinate the function of pre-existing structures (Arbib et al. 1998).

With regard to learning, the capacity for memory and anticipation plays an important role in facilitating higher cognitive processes such as goal and concept construction. However a distinctive claim of the theory of self-directedness is that this learning is highly relational, and primarily builds knowledge of situational interdependencies rather than context-independent representations. Representation is involved in this learning, but higher cognition is not representationally *driven*. Rather, cognition involves multiple processes, including semantic and non-semantic, conscious and non-conscious, which work in coordination. Learning helps refine the coordination, balancing low order implicit processes with higher order control to achieve effective situational activity. Situational awareness develops through experience as the agent learns to represent and attend to the key factors that must be controlled, but just as importantly the agent learns *not* to attend to factors that are irrelevant to the agent's goals.

This account turns the cognitivist structural sandwich inside out. Complex interacting micro, meso and macro processes allow high order cognitive organization to form and dissolve. Rules in cognition only appear as embedded constructs: as low order behavioural stimulus-response relations or as codified high order strategic and policy-level reasoning. Moreover it is very plausible that all cognitive rules should be thought of as systematically defeasible. There are no convincing general rules for epistemology or ethics, for instance, and the reason for this is, arguably, directly related to the extremely complex, multidimensional conditions of human existence. In seeking generality and rigour philosophers have fetishized mathematico-deductive argument and largely ignored practical knowledge. But reasoning processes modelled on mathematico- deductive argument are inherently unsuited for real time, real world cognition.

6. CONCLUSION

The position of cognitivism as the theoretical centre of cognitive science has been eroded by a growing stream of counter evidence, new non-symbolic and non-computational modelling techniques, and some spectacular efforts to launch rival paradigms. However the in-principle cognitivist arguments to the effect that only cognitivism can capture core cognitive properties, and hence must ultimately be correct, have not received a great deal of direct challenge. This paper has attempted to mount a counter-argument that cognitivism misframes the core cognitive issues by focusing on representations and rules rather than the adaptive embeddedness and flexibility of cognition, and that it is unable to capture these properties. In order to understand these aspects of cognition we must develop process-based theories that directly model the complex coupling between cognition and deeper biological processes.

ACKNOWLEDGEMENTS

During the period of writing this paper I have been supported as a Postdoctoral Research Fellow at the Konrad Lorenz Institute for Evolution and Cognition. The theory of self-directedness was originally developed in collaboration with Cliff Hooker, and the ideas presented here have also benefited greatly from many discussions with Mark Bickhard.

NOTES

[1] Or at least have been until recently – see below.

[2] For an extended discussion and critique of this assumption see Lakoff 1987.

[3] See Bryson 2000, Brooks 1997.

[4] For a contemporary theory of representation in the pragmatist tradition see Bickhard 1993; this volume.

[5] I will suggest some more in the next section.

[6] When artificial intelligent agents become possible it is likely that they will need to replicate many of the features of biological agents and will be just as process-bound.

[7] It might be objected that the traditional normative project is prescriptive rather than descriptive (in contrast to empirical cognitive science), but if cognition isn't – even in its ideal form – a well-defined representational process then the traditional norms simply don't apply, even prescriptively. If the problem were that real cognition is a sloppy representational-computational process then a prescriptive program to eliminate the slop is viable. But the problem being raised here is that *rigorous*representational computation theory is inherently unable to explain the properties of cognition. Ipso facto, the norms associated with representational computation cannot be directly applied to real cognition. At the very least they must be supplemented or revised in some way.

REFERENCES

Arbib, M., P. Érdi, and J. Szentagothai: 1998, *Neural Organization: Structure, Function and Dynamics*, Cambridge, MA: MIT Press.

Bechara, A., H. Damasio, D. Tranel, and A. Damasio: 1997, 'Deciding Advantageously before Knowing the Advantageous Strategy', *Science* **275**, 1293–1294.

Bertenthal, B. I. and R. Clifton: 1997, 'Perception and Action', in D. Kuhn and R. Siegler (eds.), *Handbook of Child Psychology*, New York: Wiley.

Bickhard, M. H.: 1993, 'Representational Content in Humans and Machines', *Journal of Experimental and Theoretical Artificial Intelligence* **5**, 285–333.

Brooks, R. A.: 1991, 'Intelligence Without Representation', *Artificial Intelligence* **47**, 139–159.

Bryson, J. J.: 2000, 'Cross-Paradigm Analysis of Autonomous Agent Architecture', *Journal of Experimental and Theoretical Artificial Intelligence* **12**, 165–190.

Carruthers, P.: forthcoming, 'Distinctively Human Thinking: Modular Precursors and Components', in P. Carruthers, S. Laurence, and S. Stich (eds.), *The Structure of the Innate Mind.*

Christensen, W. D. and C. A. Hooker: 2000, 'An Interactivist-Constructivist Approach to Intelligence: Self-directed Anticipative Learning', *Philosophical Psychology* **13**, 5–45.

Christensen, W. D. and C. A. Hooker: 2001, 'Self-directed Agents', in J. S. McIntosh (ed.), *Naturalism, Evolution & Intentionality*, Calgary: University of Calgary Press, pp. 19–52.

Christensen, W. D. and M. H. Bickhard: 2002, 'The Process Dynamics of Normative Function', *Monist* **85**, 3–28.

Desai, N. S., R. H. Cudmore, S. B. Nelson, and G. G. Turrigiano: 2002, 'Critical Periods for Experience-dependent Synaptic Scaling in Visual Cortex', *Nature Neuroscience* **5**, 783–789.

Fodor, J.: 2000, *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*, Cambridge, MA: Bradford Books/MIT Press.

Gibson, J. J.: 1977, 'The Theory of Affordances', in R. Shaw and J. Bransford (eds.), *Perceiving, Acting and Knowing*, New York: Wiley, pp. 67–82.

Goldstone, R. L., M. Steyvers, J. S. Smith, and Kersten, A.: 2000, 'Interactions between Perceptual and Conceptual Learning', in E. Dietrich and A. B. Markman (eds.), *Cog-*

nitive Dynamics: Conceptual and Representational Change in Humans and Machines, Mahweh, NJ: Lawrence Erlbaum, pp. 191–228.

Holyoak, K. J. and J. E. Hummel: 2000, 'The Proper Treatment of Symbols in a Connectionist Architecture', in E. Dietrich and A. B. Markman (eds.), *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*, Mahwah, NJ: Lawrence Erlbaum, pp. 229–264.

Hutchins, E.: 1995a, *Cognition in the Wild*, Cambridge, MA: MIT Press.

Hutchins, E.: 1995b, 'How a Cockpit Remembers its Speeds', *Cognitive Science* **19**, 265–288.

Lakoff, G.: 1987, *Women, Fire and Dangerous Things*, Chicago: University of Chicago Press.

Menzel, R. and M. Giurfa: 2001, 'Cognitive Architecture of a Minibrain: The Honeybee', *Trends in Cognitive Sciences* **5**, 62–71.

Münte, Thomas F., F. Altenmüller, and L. Jäncke: 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 Royal Society of London B: Biological Sciences* **352**, 1697–1701.

Newell, A. and H. A. Simon: 1972, *Human Problem Solving*, Englewood Cliffs, NJ: Prentice-Hall.

Payton, D.: 1990, 'Internalized Plans: A Representation for Action Resources', in P. Maes (ed.), *Robotics and Autonomous Systems, Special Issue on Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pp. 89–104.

Quartz, S. R. and T. J. Sejnowski: 1997, 'The Neural Basis of Cognitive Development: A Constructivist Manifesto', *Behavioural and Brain Sciences* **20**, 537–596.

Raubenheimer, D. and S. Simpson: 1999, 'Integrating Nutrition: A Geometrical Approach', *Entomologia Experimentalis et Applicata* **99**, 67–82.

Remez, R. E.: 2000, 'Speech Spoken and Represented', in E. Dietrich and A. B. Markman (eds.), *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*, Mahwah, NJ: Lawrence Erlbaum, pp. 93–115.

Romanes, G.: 1904, *Animal Intelligence*, London: Kegan Paul, Trench, Trubner & Co.

Schultz, W.: 2000, 'Multiple Reward Signals in the Brain', *Nature Reviews Neuroscience* **1**, 199–207.

Sieg, W.: 1999, 'Formal Systems, Properties of', in R. A. Wilson and F. Keil (eds.), *The MIT Encyclopedia of the Cognitive Sciences*, Cambridge, MA: Bradford Books.

Stiles, J.: 2000, 'Neural Plasticity and Cognitive Development', *Developmental Neuropsychology* **18**, 237–272.

Thompson-Schill, S. L., M. D'Esposito, G. K. Aguirre, and M. J. Farah: 1997, 'Role of Left Inferior Prefrontal Cortex in Retrieval of Semantic Knowledge: A Reevaluation', *PNAS* **94**, 14792–14797.

Tranel, D., C. G. Logan, R. J. Frank, and A. R. Damasio: 1997, 'Explaining Category-related Effects in the Retrieval of Conceptual and Lexical Knowledge for Concrete Entities: Operationalization and Analysis of Factors', *Neuropsychologia* **35**, 1329–1339.

Wilson, R. A. and F. Keil: 1999, *The MIT Encyclopedia of the Cognitive Sciences*, Cambridge, MA: Bradford Books.