DIRECTIONS FOR AN EMBODIED COGNITIVE SCIENCE: TOWARDS AN INTEGRATED APPROACH

Toni Gomila - toni.gomila@uib.cat

Paco Calvo - fjcalvo@um.es

INDEX

1. Cognitivism in a blind alley

Is cognitive activity more similar to a game of chess than to a game of pool? In order to answer this question we need to know first what are the relevant differences, cognitively speaking, between a game of pool and a game of chess. Questions of this sort (Kirsh, 1991; Haugeland, 1998) highlight the contraposition between those aspects of cognition where a rule-governed, formal approach appears to apply, and those other aspects of cognition that are deeply rooted in the physical nuts-and-bolts of the interacting agent. Whereas chess is a *formal* game that can be played regardless of details of physical implementation (think of Kasparov's legal defeat to Deep Blue), and where *rule-governed* manipulations of symbolic states suffice for the purposes of conforming to the rules of chess, in the case of pool, the actual striking of the cue ball with a nice solid hit cannot be dispensed with. Simply, pool is not a formal game. Rather than the rule-governed manipulation of inner states, real-time physical interactions need to be honoured if a game that conforms to the rules of pool is to be played. On the other hand, insofar as *digital* systems can be described by abstracting away from details of implementation, another way to emphasize the distinction (Haugeland, 1998) is by noting that whereas chess is digital, pool is not. In short, we may say that formal features identified at an algorithmic level of description are critical to a game of chess (even Kasparov's reliable offline mental rehearsal of the lost game would lack any relevant twists). However, an analogous formal level description would not work in the case of pool[1].

The assumption that cognition is like a game of chess has been the driving force of research programs in cognitive science since the inception in the mid 1950s of the cognitivist paradigm, writ large (Chomsky, 1959; cf. the canonical history of the cognitive revolution: Gardner, 1985). Classical cognitivism takes as its starting point the concepts of representation and computation. Very roughly, models of the mind are likened to a von Neumann architecture in such a way that cognitive processing boils down to the computational manipulation of representational inner states. That the brain is a piece of biological hardware, and the mind is the software running on top, means that it can be modelled as a Turing machine. Cognition consists in the rule-governed manipulation of symbols that Newell and Simon's (1972) *Physical Symbol System* epitomized.

---

[1] Of course, it goes without saying that you can play a virtual *game of pool* on a computer, but this shares with real pool nothing other than the commercial name. In fact, there cannot be a *Deep Pool* counterpart of Deep Blue. It would have to be a robot, with a vision system, and a cue mechanism, etc., and this would be different from Deep Blue. Put bluntly, the robot would actually have to *do* something.

The computer metaphor with its hardware/software divide is held quite literally by proponents of classical cognitivism. As Fodor and Pylyshyn (1988) claim, “the symbol structures in a Classical model are assumed to correspond to real physical structures in the brain and the combinatorial structure of a representation is supposed to have a counterpart in structural relations among physical properties of the brain. For example, the relation ‘part of’, which holds between a relatively simple symbol and a more complex one, is assumed to correspond to some physical relation among brain states. This is why Newell (1980) speaks of computational systems such as brains and Classical computers as ‘physical symbol systems’ ” (p. 13). And we would add, this is why cognitivism is better at spelling out the algorithms and heuristics that Deep Blue deployed in its victory over Kasparov than at trying to cash out the physically constrained scenario of a game of pool. Information-processing of discrete abstract symbols delivers the goods in the former case quite straightforwardly, or so the story goes (cf. Dreyfus, 1992), but not in the latter.

Connectionist theory, on the other hand, has traditionally been not as uneasy with less abstract, more context-dependent tasks. Unfortunately, the way cognitivism has been learned and taught in the 1980s and 1990s has sometimes assumed a temporal sequence, with classical theories (Newell and Simon, 1972) as the pinnacle of cognitivism, followed in the 1980s by neural network theory as the alternative to cognitivism, courtesy of the *bible* of connectionism (Rumelhart, McClelland et al.,, 1986; McClelland, Rumelhart et al., 1986), with hybrid models (Anderson, 1993) pulling together both ends, in an attempt to exploit the best of both worlds. However, from a historical point of view, all working hypotheses developed more or less at the same time. In fact, the basics of neurocomputation date back to the 1940s with McCulloch and Pitts (1943), and Hebb (1949). Unfortunately, although Rosenblatt’s (1959) perceptron was a breath of fresh air, Minsky and Papert’s

(1969) devastating critique of two-layer neural networks and the delta rule pushed the field out of the spotlight. It was only in the 80s with the deployment of the backpropagation learning algorithm (the generalized delta rule) in multilayered networks that connectionism re-emerged as a viable candidate to explain cognition.

Such a time line has fostered the illusion that, properly speaking, cognitivism *is* Newell and Simon (1972) and their physical symbol system hypothesis, Chomsky (1980) and his theory of rules and representations, Marr (1982) and his theory of vision paired with the well-known threefold distinction between computation, algorithm and implementation, and Fodor and Pylyshyn (1988) with their seminal critique of connectionism. Connectionism was, as a matter of fact, believed for almost two decades to be *the* alternative to cognitivism. Certainly, Chomsky, Newell and Simon, Marr, and Fodor and Pylyshyn (among others, it goes without saying) represent orthodoxy in cognitive science,[2] but the discipline is on the move, and nowadays connectionist theory just represents the alternative to *classical* cognitivism, not to cognitivism. In fact, in a broader sense, as we shall show next, connectionism belongs to cognitivism.

We strongly believe that if we are to exploit scientifically the similarities between cognitive activity and a game of chess or a game of pool, cognitivism with a twist of neural networks is helpless. A more drastic change of focus is needed. Rather than between classicist and connectionist contenders, the critical contrast to be currently drawn is between cognitivism (both classicist and connectionist) and post-cognitivism: between a view of cognition as abstract computation *versus* a view of cognition as interactive,

[2] See Garham (2008) for a historical overview. It's ironic that neither Chomsky nor Fodor really espouse cognitivism as "the" theory of cognition; they keep it restricted to some narrow areas of the human mind. See Chomsky (1996) and Fodor (2000).

embodied and embedded. Hence, this *Handbook of Embodied Cognitive Science*, where we aim to provide a panoramic view of the richness, variety, and potential of current research programmes working within this broad, post-cognitivist approach, and assess their potential convergence. We proceed in this introduction as follows: In the remainder of this section, we map out the blind alley into which, we believe, cognitivism has been driven. This allows us in section 2 to introduce the different alternative research programmes that feed post-cognitivism and recognize the problems that such an alternative must meet as a challenger candidate to explain cognition and behaviourr (another, superior, "game in town"). Section 3 discusses the commonalities among these different programmes that t coincide in calling cognitivism into question, and highlights the milestones required in order to make progress towards an integrated approach. Finally, we address the question of how to scale-up from bodily-based cognitive interaction to higher cognition, within this alternative approach. As this set of issues unfolds, a short review of the contributions that shape the volume will be provided along the way, as the best way to show how they all cohere.

So, first of all, what do we make of the standard way of framing the debate of the post-behaviourism era? According to classical cognitivism, symbols are stored in memory and retrieved and transformed by means of algorithms that specify how to compose them syntactically and how to transform them. The supposed manifest systematicity and inferential coherence of human thought, among other things, calls for this working hypothesis as an inference to the best explanation (Fodor and Pylyshyn, 1988). Underneath this conception is the idea that thought can be understood as some form of logic-like inferential processing -- bluntly, what Deep Blue does. But the result of embracing classicism is the detachment of central cognitive processes from the perceptual and motor systems. The latter reduce to input and output modules that feed the system and output the system's response (say, 1…e5 in response to white's opening 1.e4), respectively. The propositions that cognitive psychology posits, or the Fodorian picture that results in the philosophy of psychology, have the same result, namely, the endorsement of the view that cognition is information-processing as conceived by the representational-computational view of the mind, or, as we may say, some form of symbol-crunching according to algebraic rules.

The "hundred step" constraint (Feldman & Ballard, 1982), graceful degradation, neurobiological plausibility, pattern recognition and content-based retrieval of information were all interpreted as reasons to turn to connectionism (Rumelhart, McClelland et al., 1986). From the connectionist perspective, cognition was seen as the emergent outcome of the interconnectivity of numerous basic processing units connected in parallel within an allegedly biologically plausible neural network. Under this lens, the retrieval of information has as a consequence a flow of inhibitions and excitations throughout an entire network of weighted connections, which are shaped as the weights are modified in response to the statistical regularities that the network is fed.

Unfortunately, things do not look that different from the connectionist perspective than they do from the classical position, despite the different technical jargon. If cognition amounted to some form of symbol crunching according to algebraic rules under classicism, connectionism now attempts to explain cognition in terms of the computational manipulation of subsymbols, according to statistical rules. Connectionism is in fact a form of cognitivism, in spite of the obvious architectural differences between symbol systems and connectionist networks (serial v. parallel, discrete v. distributed, etc.) Whereas orthodox classical cognitivism assigns symbolic content to the sort of physical entities that get stored in von Neumann architectures, connectionist cognitivism assigns subsymbolic content to the sort of physical entities that are fully distributed and superposed on the network's weight matrix.

Hurley (1998) makes a similar point when she warns against the 'cognitive sandwich' metaphor implicitly endorsed by many cognitive scientists. Classicism is committed to a *sandwich* architecture insofar as it understands cognition 'proper' as the filling in between a perception-action *bun*. But cognitive sandwiches need not be Fodorian. A feedforward connectionist network conforms equally to the sandwich metaphor. The input layer is identified with a perception module, the output layer with an action one, and hidden space serves to identify metrically, in terms of the distance relations among patterns of activation, the structural relations that obtain among concepts. The hidden layer this time contains the meat of the connectionist sandwich. In this way, we may say that, in the worst case, connectionism amounts to a hypothesis as to the implementation on top of which classical algebraic rules operate. In the best case, it amounts to an algebraic variation of a classical algorithm, insofar as symbols are incorporated into the neural network, either in the teaching pattern of the learning algorithm itself, in the case of supervised learning, or in the input encoding, as the network is fed with patterns of activation as training proceeds (Marcus, 2001). Of course, there is a lot to be said in response to Marcus' line of argument in terms of *implementational* versus *eliminative* connectionism. Marcus' criticism echoes Fodor and Pylyshyn's (1988) attack (connectionism can only account for the systematicity, productivity, etc., of thought insofar as it implements a classical model in doing so; otherwise, it becomes eliminativist, and fails to explain the character of thought). In doing so, it is subject to the same sort of criticisms that Fodor and Pylyshyn have encountered (e.g., Chalmers, 1990; Elman, 1998). However, although self-supervised learning algorithms of the sort employed by Elman (1990; 1998) may be less suspected of furnishing the network with symbols, insofar as the teaching pattern is this time the next input pattern in the training pool, and although fully distributed input encodings may reflect the statistical regularities of the environment in more subtle ways than localist encodings, we shall not pursue this line any further, since connectionist computations are also conceived as abstract, just as the inputs and outputs are codified quite apart of the real details of perception and action.

Granting this setting then, why is cognitivism in a blind alley? The recent exchanges in the empirical and modelling literature between classicist and connectionist cognitivists demonstrate the reason. Sympathisers of classicism (e.g., Marcus et al., 1999; Marcus, 2001) continue to search for cognitive abilities that, defying a statistical explanation under the Chomskian poverty of the stimulus lens, may embarrass their connectionist foes. In turn, connectionist rule-following sceptics (Seidenberg and Elman, 1999) rejoin by showing the informational richness of ecological data, that can be exploited statistically and allow connectionist networks to remain computationally adequate. In this way, the architecture of cognition is in dispute, but assumptions about its computational underpinnings remain unchallenged. The debate focuses upon whether cognition boils down to the manipulation of symbolic items according to explicit algebraic rules, as opposed to the manipulation of subsymbolic items according to implicit statistical rules. The dispute, however, is entirely internecine warfare among proponents of a generalized, representationalist information-processing paradigm. The past-tense debate (Pinker and Ullman, 2002; Ramscar, 2002), the systematicity debate (Fodor and Pylyshyn, 1988; Elman, 1998), the algebra-versus-statistics debate (Marcus et al., 1999; Seidenberg and

Elman, 1999), and, more recently, the speech segmentation debate (Peña et al., 2002; Perruchet et al., 2004; Laakso and Calvo, 2008) have all re-enacted the "classical-connectionist, *within*-paradigm battle to win souls" (Calvo, 2005).

As the reader familiar with these debates knows, the connectionist's overall strategy is to show that stimuli are not that poor after all! Although things are never black and white, the debate has moved along these lines since the re-emergence of connectionism in the mid 1980s. The debate has been fruitful insofar as contributions have filled in empirical gaps at algorithmic levels of description. This is, however, a "cognitive decathlon" (Anderson and Lebiere, 2003) where we might never be able to declare a winner! Perhaps we are stuck in a never-ending dialectic of positing challenges to connectionism and then trying to account for them statistically, forever and ever (Calvo, 2003). We are not sure that this dialectic can deliver much more significant scientific progress. Worse still, focussing attention on this project becomes a way to ignore the deep roadblocks that cognitivism has stumbled upon (the frame problem, the grounding problem, the common code problem, etc.), and which only arise from a cognitivist standpoint. Progress on these problems, then, seems conditional on jettisoning cognitivist assumptions in the first place (Gomila, 2008).

The truth is, if one considers some of the conditions that a successful theory of cognition must satisfy according to Newell's (1980) classical paper (flexible behaviour, real-time performance, adaptive behaviour, vast knowledge base, dynamic behaviour, knowledge integration, natural language, learning, development, evolution, and brain realization), it is easy to realize that little progress has been made within cognitivism, either classical or connectionist, on most of these problems (Anderson and Lebiere, 2003). Hence, the blind alley! We think it is high time to consider ways to make real progress in all these critical challenges, ways to get out of the blind alley; and to put those ways that already show the direction of progress in the foreground. Development, real-time performance, flexible, adaptive and dynamic behaviours, evolution and brain realization, to name but a few, are dimensions that post-cognitivist theories of cognition aim at accounting for, and where their successes, even at this early stage of development, clearly outperform their cognitivist competitors. The present volume will provide both evidence for this claim and reflection on how to make further progress.

The time is ripe indeed for a real alternative approach to cognitivism to establish itself on a firm basis. As the song goes, "The times they are a-changin'."

2. Alternative approaches to cognitivism

The philosophical interest in the notion of embodiment (Clark, 1997; Hurley, 1998; Haugeland, 1998; Noë, 2004; Shapiro, 2004; Gallagher, 2005; Wheeler, 2005; Rowlands, 2006; Chemero; forthcoming) goes back to the move from a Cartesian framework into phenomenology, especially with Heidegger (1962) and Merleau-Ponty (1962), and with Wittgenstein (1953). More generally, as a more or less radical scientific alternative to cognitivism, post-cognitivism (understood as the vindication of embodiment for the understanding of cognition), has existed for quite some time, but accelerated gradually during the last three decades, gaining visibility, influence and substantial momentum, since the 1990s.[3] A number of research programmes clearly fall under the umbrella of post-cognitivism. These include ecological psychology (Gibson, 1966, 1979; Turvey and Carello; 1995), behaviour-based AI (Brooks 1986; 1991; 1999; Pfeifer & Scheier, 1999;

[3] Anderson (2003) provides a good entry point to the reader unfamiliar with the literature.

Beer, 1990), embodied cognition (Varela, Rosch, Thompson, 1991; Clancey, 1997; Ballard, 1991), distributed cognition (Hutchins, 1995), perceptual symbol systems (Barsalou, 1999; Glenberg, 1997), some forms of connectionism (Freeman, 1999; Rolls and Treves, 1999), interactivism (Bickhard and Terveen, 1995), and dynamical systems theory (Kelso, 1995; Port and van Gelder, 1995; Erlangen & Schöner, 2002), to name but a few. In this Handbook, we have tried to represent all these research programmes, and to show the additional diversity in the area. We have also considered whether the various different programmes are converging into a unified approach, and which conceptual and foundation issues would be required to facilitate this development.

At a minimum, all these approaches conceive of cognition and behaviour in terms of the dynamical interaction (coupling) of an embodied system that is embedded into the surrounding environment. As a result of their embodied-embedded nature, cognition and behaviour cannot be accounted for without taking into account the perceptual and motor apparatus that facilitates the agent's dealing with the external world in the first place, and to try to do so amounts to taking this external world also into account. This tells directly against the aforementioned cognitive sandwich and other forms of "methodological solipsism" (Fodor, 1980). Cognition is not a matter of crunching symbolically or subsymbolically, but of interacting, of coupling. To understand a cognitive system we need to take as the unit of analysis the 'system' embedded into its surrounding environment; a kind of interaction in analogy with the biological notions of species and habitat. In a sense to be further spelt out below, understanding cognition involves understanding the coupled system as such, and not the mind/brain in itself.

Areas of research where post-cognitivist principles have been applied successfully include (cognitive) neuroscience (Skarda & Freeman, 1987; Jeannerod, 1997; Damasio, 1994; Chiel and Beer, 1997), AI and evolutionary robotics (Arkin, 1998; Murphy, 2000; Nolfi & Floreano, 2000), cognitive anthropology (Hutchins, 1995; Suchman, 1987), cognitive linguistics (Lakoff and Johnson, 1980; 1999; Langacker, 1987, 1991; Regier, 1996; Tomasello, 1998), motor control and learning (Thelen & Smith, 1994), enactivism in the philosophy of perception (Thompson, 1995; O'Regan & Noë, 2001; Noë, 2004), neurophenomenology (Hanna and Thompson, 2003), education (Resnick, 1994; Greeno, 1996), and even social psychology, a field traditionally less closed to post-cognitivist methodologies that is attaining increasing attention (Semin and Smith, 2002). It is no exaggeration to say that virtually all of the connections in Gardner's (1985) well-known *cognitive hexagon* are at present being thoroughly explored via the same computational, linguistic and behavioural methodologies that helped shape the cognitive hexagon in the first place. Neuroimaging data, moreover, are helping pile up a whole new set of evidence, although if a post-cognitivist framework is to be developed in all its consequences, a thorough revision in the cognitive neurosciences and the accompanying neuroimaging methodologies must be accomplished **(Anderson; Haselager et al., this volume)**.

Illustrations of insightful applications within post-cognitivism include using pen and paper (Norman, 1993; McClelland et al., 1986), counting with one's fingers, and drawing Venn diagrams. These actions permit us to offload cognitive cargo into the world. Other examples include gesturing while speaking (Iverson & Goldin-Meadow, 1998), ballistic interception (Smeets & Brenner, 1995), the time course of motor response (Erlangen & Schöner, 2002), epistemic actions (Kirsh and Maglio, 1994) such as helping yourself in a game of cards by laying the hand out in a particular order, body–based metaphors in thinking and reasoning (Lakoff and Johnson, 1980, 1999), and many others.

Accordingly, post-cognitivist approaches have applied and developed new formal instruments, such as the theory of dynamical systems, imported from the physical sciences. Limb movement is a classic example in the literature, as in the 'HKB model' (Kelso et al., 1998) of finger coordination. Kelso (1995) studied the wagging of index fingers and a number of properties were successfully described and predicted dynamically. The phenomenon could be explained as a property of a non-linear dynamical system that achieves self-organization around certain points of instability. Post-cognitivist approaches lay the stress upon real world, time-pressured, situations as the context in which cognition and behaviour take place and make sense. This ecological dimension is pivotal, for instance, in ecological psychology (Gibson, 1979). A Gibsonian approach fits nicely with post-cognitivism, insofar as affordances allow for a direct reach that avoids the exploitation of inner representational resources. Allegedly, no information-processing, no abstract symbol crunching, is required, but simply tuning to environmental invariants through context-sensitive cue extraction and physical adjustment that do not involve a centralized process of control (**Smeets and Brenner, this volume**). We cannot possibly account for a cognitive agent's behaviour unless we treat it scientifically on a par with the environment in which the agent is acting.

Similarly, the interactivist developmental psychology of Piaget (1928, 1955) has also been a driving influence, with its processes of accommodation and assimilation that drive the reorganization of the system and the emergence of new cognitive abilities. This interactivist approach has been renewed by Thelen and Smith's dynamic systems approach to development, with its emphasis on decentralized motor development (discharged of the rationality, teleology and systematicity of Piaget's approach). Remarkable contributions from this approach include the induction of steps in infancy courtesy of a motorized treadmill (Thelen and Smith, 1994) and the well-known, although still highly controversial, explanation of the "A-not-B error" in infancy (Thelen et al., 2001). In the first case, a spring-like biomechanical response underlies stepping. In the second one, Thelen et al. (2002) go exhaustively over the literature on the A-not-B error, and offer yet another non-cognitivist explanation of motor control and development in that context. In their view, the A-not-B error can be perfectly explained in terms of the dynamical evolution of the coupling of perception, movement and memory, with no need to invoke information-processing concepts or operations (cf. Luo, 2007).

Several chapters in the Handbook provide extensive reviews of remarkable research from an embodied view of cognition: **Beer** offers a state of the art status report on the field of robotics; **Droll and Hayhoe** document the interest in studying visually guided motor control and the results obtained; **Metteyard and Vigliocco** extensively review experimental and neurophysiological research on the sensorimotor and propioceptive involvement in semantic understanding; **Bergen and Feldman, Glenberg, and Núñez**, offer a panorama of the ways in which abstract thought can be grounded in basic sensorimotor abilities.

It must be noted, though, that no single claim carries the full weight of the post-cognitivist research programme. In fact, taxonomies and dimensions of embodiment and embeddedness abound in the literature. We need to be aware of the existing diversity, as witnessed by the plurality of ways of understanding the very notions of embodiment and embeddedness. In this regard, Wilson (2002) still provides a useful starting point by making the following distinctions:

(i) Cognition being situated (Chiel & Beer, 1997; Clark, 1997; Pfeifer & Scheier, 1999; Steels & Brooks, 1995; Beer, 2000; Port & van Gelder, 1995; Thelen & Smith,1994). The emphasis is upon the maintenance of a competency as inputs and outputs relevant to the cognitive process keep impinging on the agent, as opposed to counterfactual thinking, the execution of an offline plan, etc.

But, (ii) cognition is also time-pressured (Brooks, 1991; Pfeifer & Scheier, 1999; Port & van Gelder, 1995). As Kirsh and Maglio (1994)'s research on the game of Tetris (and Scrabble, see Maglio et al., 1999) nicely illustrates, players help themselves to the manipulated external environment ('epistemic actions', such as rotating a Tetris piece as it falls on the screen) in order to ease perceptual processing. In time-pressured tasks, the efficiency of rotating the piece, including over-rotations and corrections in real time, by contrast with *imagining* the potential fit of the piece in each specific context, is manifest.

As in the case of epistemic actions, cognitive agents also (iii) off-load cognitive work onto the environment. An example is Ballard et al.'s (1997) "minimal memory strategy" (**Droll and Hayhoe, this volume**), where subjects are asked to reproduce patterns of coloured blocks, and where visual fixation and re-fixation serves to embody and approximate the experimental task with minimal demands in terms of storage. Glenberg and Robertson's (1999) compass-and-map task is another case in point.

On the other hand, laying the stress somewhere else, it has been argued that (iv) the environment itself is part and parcel of cognition (Beer, 1995; Greeno & Moore, 1993; Thelen & Smith, 1994; see Clark, 1997, for discussion). In Clark and Chalmers' (1998) view, cognition spreads out into the world in a non-trivial way. As they argue, any worldly dimension that contributes to the achievement of a cognitive task and which would count as cognitive, had that contribution come from endogenous processes, should count as a cognitive input. Put bluntly, skin and skull are irrelevant to the identification of a cognitive process. Or as **Richardson et al. (this volume)** point out, organism-environment systems, and not organisms as such, are the proper units of analysis.

Wilson (2002) distinguishes two more views: (v) Cognition as action, with the fields of vision (Churchland, Ramachandran, and Sejnowski, 1994; Ballard, 1996) and memory (Glenberg, 1997) being actively explored. A representative illustration of 'cognition as for action' is the identification of the dorsal and ventral visual pathways with "what"/"how" neural routes, instead of "what"/"where" ones (Goodale & Milner, 1992). However, as **Smeets and Brenner (this volume)** convincingly argue, such a dichotomy is still a tributary of the cognitivist idea of perception as the construction of a visual scene representation (attributed to the "what" route). Instead, what the evidence suggests is that all vision is for action, while different cues may be useful for different tasks, even if they turn out to be inconsistent. One way or the other, the idea of cognition as for action explicitly drives a number of chapters in the Handbook.

Finally, (vi) offline cognition is body-based, as in mental imagery, and in general, sensorimotor functions are exploited for approximating offline competencies (Dennett, 1995; Glenberg, 1997; Barsalou, 1999; Grush, 2004). Several chapters in this volume also deal with such an idea mainly in relation with meaning and the "grounding problem". **Bergen and Feldman, Metteyard and Vigliocco, Glenberg, and Núñez** all offer complementary approaches to explaining how abstract meaning gets its hold on sensorimotor interaction with the environment. **Sanz et al. (this volume)** also tackles these issues through his discussion of the hierarchy of control and the "internal model" kind of control, with its use of an "efferent copy", which allows the system to have expectancies on its interaction in a fast way.

These six notions do not stand or fall together. In fact, Wilson (2002) takes (i), (ii), (iii) and (v) to be true, finds (iv) somewhat problematic, and considers (vi) to be the most interesting, although the less explored, conception. We very much agree that special care needs to be taken with (vi), if post-cognitivism is to scale up at all (see section 4, below). But this six-fold taxonomy is not the only one available, and the emphasis can equally be laid upon another axis. For example, **Berkeley (this volume)** adopts Ziemke's (2003) distinction between structural coupling, historical embodiment, physical embodiment, "organismoid" embodiment, and organismic embodiment. **Sharkey and Sharkey (this volume)** also relate this classification of notions of embodiment to a "weak" and "strong embodiment" dichotomy, reminiscent of the "weak/strong AI" distinction, as well as a distinction between mechanistic and phenomenal embodiment. The latter is the stronger version, characteristic of the phenomenological tradition, where the focus is on how the body is felt from within **(Gallagher, this volume; Sheets-Johnstone, this volume**). **Clark (this volume)** defends instead a functionalist understanding of embodiment, where the same cognitive organization may be multiply realized across different triplets of bodies-brains-environments.

Richardson et al. (this volume) offer yet another six principles central to the ecological perspective: (vii) Organism-environment systems as the proper units of analysis; (viii) a call for the definition of environmental realities at the ecological scale; (ix) behaviour as emergent and self-organized; (x) perception and action as continuous and cyclic; (xi) information as specificational, and finally, (xii) perception as of affordances.

In any case, regardless of the minutiae, all these different notions of embodiment may roughly be described as incompatible with at least some central tenet of cognitivism, and *prima facie* reciprocally compatible. They are unified in rejecting the metaphor of cognition as a centralized, information-processing mechanism, but still in the business of interactive control. However, as **Anderson (this volume)** reminds us, it cannot be taken for granted that they are somehow convergent, and different predictions of each have to be developed to find out how they fare.

Summing up, whereas an information-processing agent counts as a computational system insofar as its state-transitions can be accounted for in terms of manipulations on abstract, amodal, representations, with the related problems of framing, grounding, binding and the like, the central idea that underlies post-cognitivist programmes is a denial or radical transformation of the dogma that our minds must be described as *computing* and/or *representing*, understood as symbol/subsymbol-crunching. However it remains to be determined whether they are just a heterogeneous cluster of approaches that happen to coincide in the rejection of cognitivism, or whether these different notions and approaches can converge into a unified view of cognition. In other words, whether they only allow for a negative characterization of cognition (what cognition is not), or whether they share a common positive ground. To this question we now turn.

3. Post-cognitivism in the making: Common ground and conceptual issues

A post-cognitivist interactive and extended architecture is an empirical working hypothesis that needs to be made explicit in operational terms. It is only by looking at the details of what post-cognitivism has to offer that we can assess the extent to which we are confronting something truly different from classical cognitive science. We do think that there are reasons to answer this question in the affirmative, despite the obvious differences of emphasis, point of view, notions of embodiment, and areas of research among the different trends. In this section, we, first, show the common threads that shape the fabric of post-cognitivist programmes, as illustrated by the chapters in the volume, and second, underline the conceptual problems that need to be addressed in order to develop this common ground into a well-founded research programme.

Although not every author in this volume would agree with every item, this list, in our opinion, captures the central tendencies of the post-cognitivist approach:

- Rather than the topical emphasis on embodiment, even though the interaction is made possible by the body **(Bickhard, this volume)**, it seems to us that interactivism and dynamicism are the central postulates of post-cognitivism. What really unites post-cognitivist approaches is an interactivist and dynamic view of cognition, such that to understand the cognitive system attention has also to be paid to the context or environment in which it moves, evolves, develops, and the time-course of the interaction, at the different time scales at which it unfolds. This sort of robust, but flexible, interaction is what the term "coupling" refers to.

- This dynamic interaction depends upon the body, in a way that has still to be made more precise and committed – not just physical interaction but social interaction as well **(Gallagher; Sharkey and Sharkey, this volume)**. Thus, interaction for coupling happens at all levels of physical aggregation.

- This emphasis on interaction brings sensorimotor aspects to the center of the study of cognition. As **Pfeifer et al. (this volume)** and **Richardson et al. (this volume)** forcefully argue, the informational structure a system can exploit depends on its bodily constitution in terms of sensors and effectors, materials, morphology, etc.

- Higher cognition is to be understood as constructed from this basic set of restrictions and allowances (plus maybe some new form of control: see next section).

- This standpoint breaks apart the "cognitive sandwich", and makes it clear that perception is active ("enactive"), and action is perceptually guided. Several chapters of this volume develop this theme: **Ballard** presents his latest work on perceptually guided action in a simulated environment; **Smeets and Brenner** show how in taking into account that perception is for action various puzzles dissolve, especially those derived from the idea that perception consists in building a visual representation; **Droll and Hayhoe** review research on visually guided action; and **Beer, and Pfeifer et al.** review the progress made in robotics in this regard.

- There is also a growing consensus that proper explanation requires the simultaneous scientific understanding of neural, bodily and environmental factors as they interact with each other in real time. The time course of the interaction and the activation turns out to be crucial to the explanation. New formal methods are called for to deal with this requirement, and it is in this regard that the application of dynamical systems theory proves especially appropriate. A model of mental activity must respect the same principles of non-linearity, time-dependence and continuity that are generally invoked in explanations of bodily interactions and neural activity (Freeman, 2000).

- All of the strands also coincide in viewing cognition as an emergent, self-organizing phenomenon, arising out of the local activity of distributed units; no global plan is required, and there is no single location in the system in control where everything comes together. The notion of criticality of dynamic systems theory helps in this respect by making the notion of emergence non-mysterious. As a matter of fact, it is a rather general natural phenomenon.

- The interactivism of the approach also involves an extended, situated, view of cognition; a clear way to make this point concrete is by saying that the unit of analysis is the system-cum-normal environment, and not the system in isolation.

- This set of basic assumptions naturally implies a transformation of the research questions: what should be studied shifts according to a naturalistic, ecological, setting. Instead of artificial tasks in laboratory settings, post-cognitivism studies how the cognitive system deals with its contextual demands: grasping, reaching, interception, navigation, problem-solving in the real world, etc. But also attention must be paid to consciousness. In so doing, post-cognitivism introduces more complex tasks and settings. In so doing, though, it also avoids artificial complexity, such as the requirement to fixate on a point in visual laboratory experiments, which requires inhibitory control of spontaneous saccades **(Droll and Hayhoe, this volume)**.

- The formal toolkit of modeling and simulation has been deeply renovated. Logicist approaches recede and are replaced by formal and mathematical approaches more appropriate to dealing with the interactivity and time dependence of the processes. It is no surprise that mathematical dynamic systems theory, and evolutionary algorithms, have been resorted to (Schneegans and **Schöner, this volume**), as well as connectionist modelling that allows for nonlinearity and neurobiological plausibility (**Stemme and Deco, this volume**). An appeal to kinetics has also been useful in accounting for the nature of the forces and movements the body has to exert (**Smeets and Brenner, this volume; Richardson et al., this volume**). Explanation thus becomes understanding a behaviour as a trajectory in a state space, rather than identifying its ballistic cause.

- Also, neurobiogical plausibility is a must (**Stemme and Deco, this volume**), as much as interest in bodily detail is, in the configuration of forces and torques (**Pfeifer et al., this volume**). Moreover, phenomenology, how the processes are experienced, constitutes part of the explananda, something to account for.

This basic set of common postulates, though, is not without problems. For instance, an approach such as the one outlined has to offer a viable notion of scientific explanation. Can explanation boil down to the mathematical description of the range of changes an extended system can experience over time? How does this state-space characterization of the coupled system relate to its mechanistic components; components that interact causally as well as informationally? How can an embodied cognitive science relate to an explanation of the inner mechanisms of sensorimotor coordination that give rise to higher-level cognitive activity? It is important to emphasize that the answers to these questions have a direct bearing upon the epistemology of science, with consequences as far as methodologies and the generation of testable predictions are concerned.

Of course, there is disagreement over many other aspects as well. We have already taken stock of the plurality and ambiguity of the notion of embodiment, even though its different aspects need not be seen necessarily as contradictory. However, in some cases they may be in conflict, especially as regards the level of intrinsic dependence between cognition and body. **Sheets-Johnstone (this volume)** opts for a more strict dependency of the mind on the specifics of the body, along the phenomenological tradition, while **Clark (this volume)** defends a form of extended functionalism that departs from more radical readings of embodiment such as Noë's (2004) or Shapiro's (2004).

Another area where further clarification is required concerns the hot issue of whether the notion of representation still has a role to play in cognitive explanation. Someone might take away the wrong idea that, by eschewing the information-processing notions of computation and representation, we are throwing out the representational baby with the bathwater (Hayes, Ford and Agnew, 1994). Nevertheless, as **Bickhard (this volume)** points out,[4] embodiment is necessary for representation, and therefore, for cognition. His interactivist project derives representation from action and interaction, which only makes sense if the system is embodied so as to interact with its medium. In like vein, the reader can see how it bears upon the aforementioned problems. This takes us back to the concept of circular causal flow. Notice once again the contrast with the cognitive sandwich. As the output of a connectionist system exerts no influence upon the input patters of activation, the system is engaged in no interaction whatsoever **(Bickhard, this volume)**.

Moreover, doubts may arise as to the convergence of different methodologies. Thus, evolutionary considerations (Sheets-Johnstone, 1990; Nolfi and Floreano, 2000) are

[4] For an insightful analysis of the problem of representation see Bickhard and Terveen (1995).

congenial with the post-cognitivist take on the aforementioned problems. The modelling of toy embodied, embedded systems by means of evolutionary algorithms constitutes a promising approach (Nolfi & Floreano, 2000) insofar as fitness is measured globally, and no *a priori* decisions as to what belongs to the (cognitive) system need to be made in advance (Beer, in press). However, we must ask to what extent these modelling strategies add to neurobiologically plausible models. Although, for obvious reasons, these models will be of little use in the generation of quantitative predictions (see Beer, forthcoming), the reasons for concern run deeper. Insofar as computational neuroethology (Beer, 1990) honours critical biomechanical and ecological aspects, it is certainly a move in the right direction, but we need to know whether neurobiologically plausible artificial neural network architectures and algorithms (Rolls and Treves, 1998) will converge with the statistical analyses of these models. That is, it is *not* simply a question of being able to generate quantitative as opposed to qualitative testable predictions. Rather, it is also a question of methodological convergence with the fast-growing neurosciences.

In closing, it must be emphasized that post-cognitivism aims not merely at cashing out the posits of the information-processing paradigm in trendy mathematical terms, but rather at articulating a brand new way to understand cognition. This is not an easy project, and we are well aware that it is certain to be jeopardized if one of these research programmes tried to become hegemonic; in other words, the only way for that to happen is real synergistic convergence. Contributing to this goal is the aim of this Handbook, and its success will depend upon being able to generate the right intellectual climate and a common agenda. This seems to be required because, although at first sight the articulation of ecological, dynamic, interactive, situated and embodied approaches within one single framework may look pretty straightforward, one reason for the lack of progress on effective convergence seems to reside in the fact that conceptual issues are usually treated by philosophers, and empirical ones by the rest of the cognitive science community, separately. We need to put together conceptual analysis of the notions of representation, computation, emergence, embodiment, and the like, with empirical work that allows us to bring together ecological, dynamic, interactive, situated and embodied approaches to the scientific study of cognition. The effort will be comprehensive insofar as it succeeds in unifying a conceptual/empirical framework for the cognitive sciences that allows for conceptual constraints upon the experimental paradigms and contrasting hypotheses, on the one hand, and whose empirical results inform further theoretical developments, on the other.

Thus, for this unified approach to consolidate, a systematic and forward looking approach is also needed, beyond the temptation to just identify post-cognitivism as the alternative to cognitivism. In a controversial paper that appeared in *Science* in the 1960s, Platt (1964) asked what it is that allows some disciplines to make substantial progress in very little time (think of molecular biology, for instance), whereas other areas of research (think this time maybe (?) of cognitive psychology) advance at a slower pace. In Platt's view, it is not a matter of the intrinsic difficulty of the subject (theoretical physics) or of the money injected in the area (high-energy physics). It is instead an intellectual matter that makes up the divide. Whereas all scientific disciplines in their application of the scientific method accumulate inductive inferences in support of the working hypotheses, only in some fields this is done *systematically* ("formally and explicitly and regularly"). In our view, if we want cognitive science to be problem-oriented, rather than method-oriented, we must be willing to call into question the grounds of post-cognitivism itself in order to make progress. In fact, we need to be systematic, not just in the search for crucial sets of experiments to help us decide between cognitivist and post-cognitivist hypotheses, but rather in the empirical comparison of different post-cognitivist hypotheses (**Anderson, this volume**). This is a point that too often goes unnoticed, and unless it is focused upon more thoroughly, the question as to whether post-cognitivism is moving towards an integrated approach cannot be definitely answered in the positive.

It is in this context that the shift of paradigm that post-cognitivism represents should be submitted to critical scrutiny. In the next section, we thus review efforts at discounting post-cognitivism as just a form of cognitivism. In doing so, we also address the problem of scaling up -- the ultimate challenge, if post-cognitivism is to present itself as a comprehensive and viable alternative.

4. Scaling up: Higher-level cognitive processes

Despite the aforementioned momentum, and the degree of convergence in a unified approach, recent episodes in cognitive science suggest that cognitivism might still resist by assimilating some post-cognitivist methodologies and insights as complementary to its main thrust. We envisage two distinct strategies such a reaction may take. On the one hand, it could be claimed that post-cognitivism is not so much an alternative view of cognition, but the right approach to deal with low-level cognitive processes, those involved in sensorimotor coordination, while cognitivism is still the right way to approach high-level, symbolic, cognition – even allowing for some kind of grounding of the latter on the former. On the other hand, it could be claimed that the sort of interactivism put forward by post-cognitivism still requires some sort of internal stand-in to account for the causal powers of the mind/brain states involved, so that, in the end, even at the basic sensorimotor level, the commitment to internal representations and computations over them is inescapable. Of course, these strategies are not equally challenging for post-cognitivism. The latter, in fact, tries to absorb the post-cognitivist principles into a basic cognitivist architecture, while the former amounts to an acceptance of some form of hybridism. In what follows we address these two issues in reverse order.

Just as an illustration of how this latter strategy may be carried out, consider the response by Vera and Simon (1993a; 1993b) to the Gibsonian and Brooks's (1991) challenges of "doing without representing" (Clark and Toribio, 1994). In Vera and Simon's view, cognitivism and post-cognitivism need not be antithetical. As a matter of fact, as Vera and Simon argue, Gibson's affordances and Brooks "navigation without representation" approaches should be seen as an illustration of "orthodox symbol systems". As Vera and Simon claim, the information-processing paradigm does not ignore the medium in which cognitive activity takes place. In their view, "the thing that corresponds to an affordance is a symbol stored in central memory denoting the encoding in functional terms of a complex visual display, the latter produced, in turn, by the actual physical scene that is being viewed" (p. 20). Gibsonian affordances, "far from removing the need for internal representations, are carefully and simply encoded internal representations of complex configurations of external objects, the encodings capturing the functional significance of the objects" (p. 41). Commenting on Brooks' line of research, Vera and Simon (1993a) assert that "[sensory] information is converted to symbols which are then processed and evaluated in order to determine the appropriate motor symbols that lead to behavior" (p. 34).

This very strategy of response, consisting of internalizing the relevant interactive relationships as symbolic states, has proliferated and might be recognized in several recent dismissive discussions of post-cognitivism as a real alternative (for instance, Markman and Dietrich, 2000). Although from these comments we cannot conclude much except that the debate is far from settled, it should be acknowledged that this “assimilationist” strategy is facilitated by the lack of a similar explanatory grip on the part of the post-cognitivist challenger. Unfortunately, it is not crystal clear what we mean when we say that post-cognitivist approaches require a *non-symbolic* interpretation of a cognitive system’s ecological interactions (Winograd & Flores, 1986). Does for example “non-symbolic” mean “sub-symbolic” or “non-representational” *tout court*?

This relates to the question of Representation with a capital ‘R’. Intuitively speaking, online forms of co-variation do not amount to representation (think of the classical example of the sunflower’s solar tracking behaviour, which is interpreted in purely reactive, non-cognitive, terms – Cantwell Smith, 1996). The explanation would be that there is some exogenous feature that the sunflower manages to keep track of adaptively. But does the distinction between adaptive coupling and “other things” make sense in a full-blown post-cognitivist science? In case representations and mediating states don’t vanish altogether, what properties do they have? Must they endure? Does it make sense to talk of enduring states at all? How can a representation be amodal? We are far from reaching a consensus here (Beer, 2000; Brooks, 1991; Clark 1997; Keijzer, 2002; Markman and Dietrich, 2000; van Gelder, 1995, Vera & Simon, 1993). We take it, anyway, that what is needed is rather a different notion of representation, relative to the sort of processes it sustains –so, non-syntactically individuated, not subsymbolically constituted – not as an internal reflection of an external feature, but as what allows for the coupling or interaction. It thus has normativity **(Bickhard, this volume)**, although it has none of the features of cognitivist representation.

This revision of the notion of representation may also be instrumental in foreclosing another interpretation of post-cognitivism as a sort of neo-behaviourism. Thus, someone may argue that the move from cognitivism to post-cognitivism involves a shift, indeed a U-turn, in the status of cognitive science itself (Ramsey, 2007). This claim presupposes both that (i) the emergence of cognitivism, as a reaction to behaviourism, capitalized on the concept of representation, and that (ii) the materialization of post-cognitivism involves a return to some form of pre-cognitivist behaviourism. As Ramsey (2007) puts it, after the cognitive revolution, a "revolution in reverse" (p. 223) is now taking place.[5] We believe that a premature endorsement of this reading may risk misinterpreting the new explanatory principles and models of post-cognitivism. The real issue is whether the traditional Western ways to think of mental representation up until now are the only way to conceive of such internal mediating states. For instance, if one takes **Schneegans and Schöner's (this volume)** notion of "dynamic fields," the representational jargon adds little to its functioning.

At this point, guidance could be found in the way the parallel problem of how the genome codes the genetic information and controls its expression has been dealt with by development systems theory (Gray, Griffiths and Oyama, 2001). In this area, as well, the standard view of representation drives the idea of a genetic coding of phenotypic characters, while when one realizes that the right unit of analysis is the genome-cum-environment (since the ontogenetic process depends upon the stability of the environment in supplying the chemicals needed), the idea that particular segments code particular

[5] For an argument that contests Ramsey's analysis of where cognitive science is heading, to the effect that both (i) and (ii) may be called into question, see Calvo and García (submitted).

characters by themselves loses its grip, because the information lies in the interaction between genetic sequences and robust enabling environment (which may involve the cell as well). To put in a more cognitive way, that the genetic information is context-dependent. Of course, this does not diminish at all the reactive causal powers of the DNA sequences; it just underlines the fact that their informational interpretation depends upon the normal environment in which the epigenetic process takes place.

Anyway, the question of the representational nature of the mediating internal states is definitely key to understanding higher level cognition, dependent upon the offline activation of such internal states. It is in the context of higher-level capacities that ‘representation-hungry’ cases have to be addressed (Clark, 1999), that representations appear to be explanatorily unavoidable. In the same vein, Wilson (2002) concluded that what’s distinctive in human cognition is the possibility of offline, symbolic, processing, decoupled from current spatio-temporal context. **Ballard (this volume), and Berkeley (this volume)** also seem to assume that cognitivism may still be useful as regards symbolic, higher-level processes, as far as postulating abstract, amodal, representations may keep its explanatory grip. The pressing question, then, is whether, and how, post-cognitivism can account for such a higher-level cognition, whether it is able to go beyond simpler forms of adaptive coupling, which only involve some sort of online tracking, and for which has already proved valid, to account as well for offline, higher-level, cognition.

As **Wilson (this volume)** points out, understanding our higher level abilities has to make evolutionary sense, and in this regard post-cognitivism is better placed than cognitivism. While cognitivism establishes a deep divide between animal cognition and human cognition (given the lesser degree of systematicity, productivity, and flexibility of the former), post-cognitivism tries to overcome all the big traditional dichotomies and thus stresses the elements of continuity with animal cognition, given the shared basic sensorimotor abilities. An evolutionary account is needed, then, to provide an account of how abstract, decoupled, symbolic, thought can emerge out of these basic abilities, and what special conditions of humans restrict the emergent level of cognition to our species. **Haselager et al. (this volume)** further argue that effective control may be achieved by control systems that co-evolve in relation to constraints in terms of embodiment and embeddedness. Evolutionary considerations again allow for the "cognitive fit" of the extended system.

Wilson (this volume) singles out motor control, analogy and imitation as the key aspects from which such an account can be worked out. To take them in reverse order: it is clear that our ultrasociality has had something to do with our cognitive make-up, and it is beyond doubt that imitation may be a special kind of social learning in our species, allowing for the social scaffolding that introduces the new members of the species to the social ways of thinking of the group (Tomasello, 1999). The successful notion of a "mirror system" as the brain structure that supports such a competence may have something to do with how we are able to go along with others. But as **Gallagher (this volume)** forcefully argues, the dominant interpretation of the workings of such a system, as simply the basis of imitation, is deeply contentious (suffice it to remember that macaques, where mirror neurons were first identified, using a single neuron paradigm, do not imitate, what seems to shortcut the link between such neurons and imitation). He proposes what we could term a post-cognitivist view of social interaction as a kind of bodily grounded intentional understanding, while reinterpreting the role of the mirror system in allowing it.

As regards analogy, **Wilson** refers to the pioneering work of Lakoff on conceptual metaphors, which has been successfully applied to areas as symbolic and abstract as mathematics (**Núñez, this volume; Glenberg, this volume**). Other contributors also work along this approach (**Bergen and Feldman, this volume**). In a way, this approach could be seen as compatible with the kind of symbolic, cognitivist processes, in that it offers a solution to the grounding problem for abstract, amodal representation, which, once constituted could then be worked according to syntactic processes (Barsalou, 1999; for a recent monograph on such an approach, de Vega, Glenberg & Graesser, 2008). However, such an approach is ambiguous as regards its commitment to cognitivism. We think it ill-advised to try to have it both ways, as an account of the grounding of cognitivist representations which are then submitted to computational processes. As we noted in the first section, such a project runs the risk of getting trapped in the blind alley. In addition, such a cognitivist interpretation fares poorly with the importance of imagination in such processes as depicted by these contributions (present even in the most abstract of problem solving; see Arp, 2008). In fact, what they suggest naturally is an analogical view of internal states, not as perceptual images, but as internal, dynamical, maps (Gomila, 2008).

A fashionable alternative reading would consist in placing such imaginative conceptual abilities in the context of simulation theories, as imagination is currently accounted for as simulation: visual imagination (Kosslyn, 1994), motor imagination (Jeannerod, 1997), empathic imagination (Goldman, 2006), etc. As a matter of fact, this is one of the ways abstract content is supposed to be grounded (in this volume, **Bergen and Feldman, Glenberg, Metteyard and Vigliocco, and Núñez** refer to this possibility). But as **Anderson (this volume)** rightly points out, sensorimotor activations associated with higher level cognition cannot be viewed in simulationist terms without further ado. His "massive redeployment hypothesis" contends that this is an instance of a general phenomenon of re-use of structures for new functions. Maybe more important in this regard is to realize that such abilities have to be seen in the context of the third aspect Wilson mentions: a change in the nature of cognitive control, given that imagination involves precisely the sort of voluntary control that she views as needed for detached, decoupled, abstract, thinking. Following Grush (2004), we think that such abilities are to be better conceived from the point of view of the internal model control architecture (**Sanz et al., this volume**; Gomila, 2007), which postulates that the controller sends an efferent copy to an internal emulator, an internal model of the interaction between systems and environment; such an internal model is what would support offline cognition, being accessed top-down, that is, non-stimulus driven. Of course, this is not the only possible control architecture, and **Sanz et al. (this volume)** offers a wholesale view of control architectures in terms of levels of control, but the "internal model" architecture seems to be a good place to start an account of higher cognition. The question is how to conceive of such internal models, but we think that neither classicism nor connectionism offer plausible suggestions, given that such internal models have to work in real time and along the same dimensions as the bodily interactions (sensory feedback, anticipation of propioceptive cues, etc.). On the contrary, proposals such as the dynamic field (**Schneegans and Schöner, this volume**) or neurobiologically plausible Hebbian networks (**Stemme and Deco, this volume**) provide an illustration of how such internal models can be conceived such that they can be viewed as representations, or thought to compute, in a non-contentious, mathematical way. It is clear anyway that complex cognition requires a complex integrated system, and that such a system requires forms of control that are not purely distributed and reactive (without being committed, for such reason, to postulating a "central executive" or anything homunculus-like).

Thus, instead of opting for a hybrid view of human cognition, the possibility exists to reinterpret higher level, decoupled, cognition, in terms of post-cognitivist principles. In our view, cognitivism, both in its classical version and in its connectionist form, is unable to deal fully with the dichotomy regarding cognition that the games of chess and pool example served to illustrate at the outset of this chapter. It is only when post-cognitivist models of cognition enter the picture that an answer to our opening question can begin to be given. Roughly speaking, the working hypothesis of post-cognitivism is that higher level cognitive activity never goes completely formal. Instead, it remains a 'game of pool', we may say, in which non-formal perception-action activation patterns are ubiquitous.

As Kuhn taught us, though, it is only when a new paradigm is ready that the old one will begin to be overcome. It is tempting to ask whether there is anything here really deserving to be called a "new paradigm" yet. We would actually like to try to avoid too much talk in terms of paradigms and paradigm shifts. What really matters is the theoretical significance of the new framework. Our point here is: let's unify this research effort into a single framework. We want to assess the significance of an integrated, embodied cognitive science, and invite others to explore the path. Enjoy the ride.

Acknowledgements

We are grateful to Mike Anderson, Aarre Laakso, Peter Slezak and Meg Wilson for helpful comments and suggestions on a previous version of this chapter. Preparation of the manuscript was supported by DGICYT Project HUM2006-11603-C02-01 (Spanish Ministry of Science and Education and European Union FEDER Funds).

References

Anderson, J. (1993). Rules of the Mind. Hillsdale, NJ: Erlbaum.

Anderson, J. & Lebiere, C. L. (2003). The Newell test for a theory of cognition. Behavioral and Brain Science 26, 587-637

Anderson, Michael L. (2003) “Embodied Cognition: A field guide”. Artif. Intell. 149(1): 91-130.

Arkin, Ronald C. (1998) Behavior-Based Robotics, The MIT Press, Cambridge, MA

Arp, R. (2008) Scenario Visualization: an evolutionary account of creative problem solving. MIT Press.

Ballard, D.H., "Animate vision," Artificial Intelligence Journal 48, 57-86, 1991

Ballard, D. H. (1996). On the function of visual representation. In K. A. Akins (Ed.), Perception (pp. 111-131). Oxford: Oxford University Press.

Ballard, D., Hayhoe, M., Pook, P., & Rao, R. (1997). Deictic codes for the embodiment of cognition. Behavioral and Brain Sciences, 20, 723-767.

Barsalou, L.W. (1999). Perceptual symbol systems. Behavioral and Brain Sciences, 22, 577-609.

Beer, R.D. (1990) Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology, Academic Press.

Beer, R.D. (1995). A dynamical systems perspective on agent-environment interaction. Artificial Intelligence 72:173-215.

Beer, R.D. (2000). Dynamical approaches to cognitive science. Trends in Cognitive Sciences 4(3):91-99.

Beer, R.D. (in press) “Beyond Control: The Dynamics of Brain-Body-Environment Interaction in Motor Systems”, in D. Sternad (Ed.), Progress in Motor Control V: A Multidisciplinary Perspective. Springer.

Bickhard, M. H. and Terveen, L. (1995) Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution. (Elsevier Scientific, Amsterdam).

Brooks, R. A. (1986) "A Robust Layered Control System for a Mobile Robot", IEEE Journal of Robotics and Automation, Vol. 2, No. 1, pp. 14–23

Brooks, RA (1991) Intelligence without representation, Artificial Intelligence 47, 139–159

Brooks, R.A. Cambrian Intelligence, MIT Press, Cambridge, MA, 1999.

Calvo, P. (2003) “Non-classical connectionism should enter the decathlon”, Behavioral and Brain Sciences 26, 603-604.

Calvo, P. (2005) “Rules, Similarity, and the Information-Processing Blind Alley”, Behavioral and Brain Sciences 28, pp. 17-18.

Calvo, P. and García, A. (submitted) “Where is cognitive science heading?”

Chalmers, D. J. (1990). Why Fodor and Pylyshyn were wrong: The simplest refutation. Proceedings of the twelfth annual conference of the Cognitive Science Society (pp. 340-347).

Chemero, A. (forthcoming) Radical Embodied Cognitive Science, MIT Press.

Chiel, H. and Beer, R. (1997) The brain has a body: adaptive behaviour emerges from interactions of nervous system, body and environment Trends Neurosci. 20, 553–557.

Chomsky, N. (1959) Review of Skinner’s Verbal Behavior, Language 35, no. 1, 26-57.

Chomsky, N. (1980) Rules and Representations. New York: Columbia University Press and Oxford: Basil Blackwell Publisher.

Chomsky, N. (1996) Powers and Prospects. Reflections on human nature and the social order. South End Press.

Churchland,P. S.;Ramachandran,VS;Sejnowski,T. J. (1994). “A critique of pure vision” pp. 23-60. In Large-scale neuronal theories of the brain , edited by Koch and Davis, Bradford Book/MIT Press, Cambridge MA

Clancey, W. J. (1997). Situated Cognition, Cambridge, Cambridge University Press.

Clark, A. (1997) Being There: Putting Brain, Body and World Together Again, MIT Press.

Clark, A. (1999) "An Embodied Cognitive Science?" Trends In Cognitive Sciences 3:9:1999 p. 345-351

Clark, A. and Chalmers, D. (1998) “The Extended Mind”, Analysis 58: 1, p.7-19.

Clark, A. and Toribio, J. (1994) “Doing Without Representing?" Synthese 101: 401-431

Damasio, A. (1994) Descartes' Error: Emotion, Reason, and the Human Brain, Putnam Publishing.

de Vega, M., Glenberg, A., & Graesser, A. (eds.) (2008) Symbols and Embodiment. Debates on meaning and cognition. Oxford University Press.

Dennett, D. (1995) Darwin's Dangerous Idea: Evolution and the Meanings of Life, New York: Simon & Schuster, London: Penguin.

Dreyfus, H. (1992). What computers still can’t do. Massachusetts: MIT. ...

Elman, J.L. (1990). Finding structure in time. Cognitive Science, 14, 179-211.

Elman, J. L. (1998) Generalization, simple recurrent networks, and the emergence of structure. In M.A. Gernsbacher and S.J. Derry, editors, Proceedings of the Twentieth Annual Conference of the Cognitive Science Society. Mahwah, NJ: Lawrence Erlbaum Associates.

Erlhagen, W. and Schöner, G. (2002) “Dynamic field theory of movement preparation”, Psychological Review, 109:545-572, 200

Feldman, J.A. y Ballard, D.H. (1982). Connectionist models and their properties. Cognitive Science, 6, 205-254.

Fodor, J. (2000) The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology, MIT Press.

Fodor, J. (1980) “Methodological solipsism considered as a research strategy in cognitive psychology”, Behavioral and Brain Sciences, 3: 63-110.

Fodor, J. and Pylyshyn, Z. (1988). Connectionism and Cognitive Architecture. Cognition 28: 3-71.

Freeman, W.J. 1999 How Brains Make Up Their Minds, Weidenfeld and Nicolson, London.

Freeman, W.J. (2000) Neurodynamics; An Exploration of Mesoscopic Brain Dynamics. London UK: Springer-Verlag.

Gallagher, S. (2005) How the Body Shapes the Mind, Oxford University Press.

Gardner, H. (1985) The Mind’s New Science: A History of the Cognitive Revolution, Basic Books.

Garnham, A. (2008) “Cognitivism”, en Symons, J. y Calvo Garzón, F. (eds.) The Routledge Companion to Philosophy of Psychology, Routledge Philosophy Companions.

Gibson, J. J. (1966) The Senses Considered as Perceptual Systems. Boston: Houghton Mifflin.

Gibson, J.J. (1979) The Ecological Approach to Visual Perception. Lawrence Erlbaum.

Glenberg, A. M. (1997) “What memory is for”, Behavioral & Brain Sciences, Vol 20(1) pp. 1-55.

Glenberg, A. M.; Robertson, D. A (1999) “Indexical understanding of instructions”, Discourse Processes, Vol 28(1), pp. 1-26.

Goldman, A. (2006) Simulating minds: the philosophy, psychology and neuroscience of mindreading. Oxford University Press.

Gomila, A. (2007) “Executive awareness: de-constructing the homunculus”, Estudios de Psicología, 28: 211-230.

Gomila, A. (2008) “Mending or abandoning cognitivism?”. In de Vega, Glenberg and Graesser (2008), pp. 789-834.

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. Trends in Neurosciences, 15, 20-25.

Gray, R., Griffiths, P & Oyama, S. (eds.) (2001) Cycles of contingency. MIT Press.

Greeno, J. (1996) On claims that answer the wrong questions (Response to Anderson et al.) Educational Researcher 25, 5–11

Greeno, J. C, and Moore, J.L. (1993) “Situativity and Symbols: Response to Vera and Simon”, Cognitive Science 17(1): 49-59.

Grush, Rick (2004). The emulation theory of representation: motor control, imagery, and perception. Behavioral and Brain Sciences 27:377-442.

Hanna, Robert and Evan Thompson. 2003. "Neurophenomenology and the Spontaneity of Consciousness." In The Problem of Consciousness: New Essays in Phenomenological Philosophy of Mind, ed. Evan Thompson. Calgary, Alberta, Canada: University of Calgary Press, pp. 133-161.

Haugeland, J. (1998) Having Thought: Essays in the Metaphysics of Mind. Harvard University Press

Ford, K. M., Hayes, P. J. and Agnew, N. (1994) On babies and bathwater: A cautionary tale. AI Magazine, 15(4):14-26.

Hebb, D.O. (1949) The Organization of Behavior, John Wiley.

Heidegger, M. (1962) Being and Time, New York: Harper & Row.

Hurley, S. L. (1998), Consciousness in Action, Cambridge: Harvard University Press.

Hutchins, E. (1995) Cognition in the Wild, MIT Press

Iverson, J. M. & Goldin-Meadow, S. (1998) Why people gesture as they speak. Nature, 396, 228.

Jeannerod, M. (1997) The cognitive neuroscience of action. London: Blackwell.

Keijzer, F. (2002) "Representation in dynamical and embodied cognition", Cognitive Systems Research, 3, 275-288.

Kelso, S. (1995) Dynamic Patterns, MIT Press

Kelso, S., Fuchs, A., Lancaster, R., Holroyd, T, Cheyne, D., & Weinberg, H. (1998) Dynamic cortical activity in the human brian reveals motor equivalence, Nature, 392: 814-818.

Kirsh, D. (1991) Foundations of AI: The Big Issues. Artificial Intelligence.

Kirsh, D. and Maglio, P. (1995) On Distinguishing Epistemic from Pragmatic Actions. Cognitive Science.

Maglio, P., Matlock, T., Raphaely, D., Chernicky, B., & Kirsh D. (1999) Interactive skill in Scrabble. In Proceedings of Twenty-first Annual Conference of the Cognitive Science Society. Mahwah, NJ: Lawrence Erlbaum.

Kosslyn, S. (1994) Image and Brain: the resolution of the imagery debate. MIT Press.

Laakso, A. and Calvo, P. (2008) A connectionist simulation of structural rule learning in language acquisition. Proceedings of the 30th Annual Conference of the Cognitive Science Society.

Lakoff, G., & Johnson, M. (1980). Metaphors we live by. Chicago, IL: University of Chicago Press.

Lakoff, G. and Johnson, M. (1999) Philosophy in the Flesh, MIT Press

Langacker, R. W. (1987) Foundations of cognitive grammar: Theoretical Prerequisites. Stanford, CA: Stanford University Press.

Langacker, Ronald W. (1991): Foundations of Cognitive Grammar. Vol II: Descriptive Application. Stanford: Stanford University Press.

Luo, J. (2007) “Object Permanence as Relational Stability or How to Get Representation from the Dynamics of Embodiment”, Proceedings of the 29th Annual Conference of the Cognitive Science Society.

Marcus, G.F. (2001) The algebraic mind: Integrating connectionism and cognitive science, MIT Press.

Marcus, G.F., Vijayan, S., Bandi Rao, S. y Vishton, P. M. (1999) “Rule learning in seven-month-old infants”, Science 283: 77-80.

Markman, A.B. and Dietrich, E. (2000) “Extending the classical view of representation”, *Trends in Cognitive Sciences* 4: 470-5.

Marr, D. (1982) Vision, W.H. Freeman.

McClelland, J.L., Rumelhart, D.E. and the PDP Research Group (1986) Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 2, MIT Press.

McCulloch, W. y Pitts, W. (1943) “A logical calculus of the ideas immanent in nervous activity”, Bulletin of Mathematical Biophysics 7: 115-133.

Merleau-Ponty, M. (1962). Phenomenology of Perception. C. Smith (translator). Routledge & Kegan Paul.

Minsky, M. and Papert, S. (1969). Perceptrons. MIT Press, Cambridge.

Murphy, R.R. 2000. Introduction to AI robotics. Cambridge, MA: The MIT Press.

Newell, A.(1980). Physical symbol systems. Cognitive Science, 4, 135-183.

Newell, Allen, and Herbert A. Simon. 1972. Human Problem Solving. Englewood Cliffs, NJ: Prentice Hall.

Noë, A. (2004) Action in Perception, The MIT Press.

Nolfi, S. and Floreano, D. (2000). Evolutionary Robotics : The Biology, Intelligence, and Technology of Self-Organizing Machines.MIT Press.

Norman, D.A. (1993) Things That Make Us Smart, Addison–Wesley.

O'Regan, K. & Noë, A. (2001) "A sensorimotor account of vision and visual consciousness." Behavioral and Brain Sciences 24 (5): 883-917.

Peña, M., Bonatti, L., Nespor, M., y Mehler, J. (2002) "Signal-Driven computations in Speech Processing", Science 298: 604-607.

Perruchet, P., Tyler, M. D., Galland, N., & Peereman, R. (2004). Learning nonadjacent dependencies: No need for algebraic-like computations. Journal of Experimental Psychology: General, 133, 573–583.

Pfeifer, R., and Scheier, C. (1999). Understanding intelligence. Cambridge, Mass.: MIT Press.

Piaget, J. (1928). The Child's Conception of the World. London: Routledge and Kegan Paul.

Piaget, J. (1955). The Child's Construction of Reality. London: Routledge and Kegan Paul.

Pinker, S. y Ullman, M. (2002) “The past and future of the past tense”, Trends in Cognitive Sciences 6: 456–463.

Platt, J. (1964) “Strong InferenceCertain systematic methods of scientific thinking may produce much more rapid progress than others”, Science 146, Number 3642.

Port, R. y Van Gelder, T. (1995) Mind as Motion, MIT Press.

Ramscar, M. (2002) “The role of meaning in inflection: why the past tense doesn’t require a rule”, Cognitive Psychology, 45: 45–94.

Ramsey, W.M. (2007) Representation reconsidered, New York: Cambridge University Press.

Regier, T. (1996). The human semantic potential: Spatial language and constrained connectionism. Cambridge, MA: MIT Press.

Resnick, M. (1994) Turtles, Termites and Traffic Jams, MIT Press.

Rolls, E.T. & Treves, A. (1998) Neural Networks and Brain Function, Oxford: Oxford University Press.

Rosenblatt, F. (1959) Principles of Neurodynamics, Spartan Books.

Rowlands, M. (2006) Body Language: Representation in Action, MIT Press.

Rumelhart, D.E., McClelland, J.L. and the PDP Research Group (1986) Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1, MIT Press.

Seidenberg, M. y Elman, J.L. (1999) “Networks are not hidden rules”, Trends in Cognitive Sciences 3: 288–289.

Semin, G. R. & Smith, E. R. (2002). Interfaces of social psychology with situated and embodied cognition. Cognitive Systems Research, 3, 385-396.

Shapiro, L (2004) The Mind Incarnate (MIT Press).

Sheets-Johnstone, Maxine. 1990. The Roots of Thinking. Philadelphia: Temple University Press.

Skarda, C.A. and Freeman, W.J. (1987) ‘How brains make chaos to make sense of the world, Behavioral and Brain Sciences 10, pp. 161-95.

Smeets, J.B.J. & Brenner, E. (1995) The visual guidance of ballistic arm movements. In: Multisensori control of posture. T. Mergner, F. Hlavacka, Eds. (pp: 191-197). Plenum Press, New York.

Smith, Brian Cantwell (1996) On the origin of objects. Cambridge, MA: MIT Press.

Steels, L. and Brooks, R.A. (eds.) (1995) The Artificial Life Route to Artificial Intelligence: Building Embodied Situated Agents, Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.

Suchman, A. (1987) Plans and Situated Actions, Cambridge University Press

Thelen, E. and Smith, L. (1994) A Dynamic Systems Approach to the Development of Cognition and Action, MIT Press

Thelen, E. Schöner, G., Scheier, C. and Smith, L. (2001) The dynamics of embodiment: A field theory of infant perseverative reaching. Brain and Behavioral Sciences, 24:1-33.

Thompson, E. (1995) Colour Vision: a study in Cognitive Science and the Philosophy of Perception. Routledge.

Tomasello, M. (ed.) (1998) The New Psychology of Language: Cognitive and Functional Approaches to Language Structure. L. Erlbaum.

Tomasello, M. (1999) The cultural origins of human cognition. Harvard University Press.

Turvey, M. and Carello, C. (1995) Some dynamical themes in perception and action, in Mind as Motion (Port, R. and Van Gelder, T., eds), pp. 373–401, MIT Press

Varela, F., Rosch, E., & Thompson, E. (1991) The Embodied Mind, MIT Press

Vera, A.H. y Simon, H.A. (1993a) “Situated Action: A Symbolic Interpretation”, Cognitive Science 17: 7-48.

Vera, A.H. y Simon, H.A. (1993b) “Situated Action: Reply to William Clancey”, Cognitive Science 17: 117-33.

Wheeler, M. (2005) Reconstructing the Cognitive World: the Next Step, MIT Press.

Wilson, M. (2002). Six views of embodied cognition. Psychonomic Bulletin and Review, 9, 625-636.

Winograd, T. and Flores, C. F. (1986) Understanding Computers and Cognition. A New Foundation for Design. Norwood: Ablex

Wittgenstein, L. 1953. Philosophical investigations. (Anscombe, G.E.M., trans.). Oxford: Basil Blackwell.

Ziemke, T. (2003) What's that thing called embodiment? Proceedings of the 25th Annual Meeting of the Cognitive Science Society. Philadelphia, PA: Lawrence Erlbaum.