EMBODIYING COGNITION:
TOWARDS AN INTEGRATED APPROACH

Palma de Mallorca (Spain), 14-16 december 2006

Despite its diversity, Cognitive Science is still being mainly carried out under a broad cognitivist approach. By cognitivism, we mean the view of cognition as information-processing as conceived by the representational-computational view of the mind: as symbol-crunching. Cognitivism suffers from well known shortcomings (frame-problem, grounding problem, common-code problem, lack of temporal dimension,...) which were identified a long time ago, without much progress towards its overcoming.

Connectionism, when it entered the scene with renewed impetus in the mid-eighties, presented itself as an alternative approach to cognition. However, a number of exchanges between cognitivists and connectionist modellers makes us suspect that, in the end, connectionism belongs to cognitivism, in spite of the obvious architectural differences between symbol systems and connectionist networks. For instance, cognitivist researchers (e.g., Marcus et al., 1999) continue to search for cognitive abilities that, defying statistical explanation, may embarrass connectionists. In turn, connectionist rule-following sceptics (Seidenberg and Elman, 1999) rejoin by finding 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. Put bluntly, debate focuses upon whether cognition boils down to the manipulation of symbolic items according to explicit algebraic rules, as opposed to the manipulation of sub-symbolic 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, in the last 4 years, the speech segmentation debate (Peña et al., 2002; Perruchet et al., in press) have benefited partially from the "classical-connectionist, *within*-paradigm battle to win souls" (Calvo Garzón, 2005). We are not sure that this dialectic can deliver much more significant scientific progress, and we think that it is high time for a real alternative approach to get strong hold.

As a matter of fact, “non-cognitivist” positions in Cognitive Science have existed all along, and accumulated gradually during the last three decades, generally inspired in old functionist or phenomenological views of the mind. In the last ten years, though,“Post-cognitivist” ones have gained in visibility, influence and substantial momentum . Varieties of these non-classical research programmes include ecological perception (Gibson, 1979; Turvey, 1975), situated action (L. Stein,), embodied cognition (Varela, Rosch, Thompson, 1992; Clancey, 1997), distributed cognition (Hutchins, 1995) social mind (Werstch,), perceptual symbol systems (Barsalou, 1999; Schyns et al., 1998, Glenberg, 1997), some forms of connectionism (Grossman, ;Elman et al., 1996; Rolls and Treves, 1999), interactivism (Bickhard and Terveen, 1995), and dynamical systems (Port and van Gelder, 1995; Schöner et al., 1997), and. All these “brand names” stand for different research programmes, each with its corresponding set of basic postulates, followers, journals, academic associations, etc., whose common aim is to develop a, more or less radical, alternative to classical, cognitivist, Cognitive Science.. These approaches may be roughly described as incompatible with at least some central tenet of cognitivism.[1] They are unified in rejecting the metaphor of cognition as a centralized, information-processing mechanism. All of them coincide in viewing cognition as rather an emergent and extended self-organizing phenomenon, explanation of which requires the simultaneous scientific understanding of neural, body and environmental factors as they interact with each other in real time. In short, as we see it, the central idea that underlies post-cognitivism is a denial or radical transformation of the dogma that our minds must be described as *computing* and/or *representing*.[2] However, it’s also true that this diversity of programmes and efforts have gone along generally on their own, without looking for some kind of convergence into a unified, alternative, cognitive paradigm. This dispersion of efforts has made it more difficult for a robust, mainstream alternative to develop, and it has made it easier for cognitivism to apply a “divide and conquer” strategy of rebuttal or assimilation. In need for a better name (and research agenda?), we’ll momentarily refer to these post-cognitivist approaches collectively as

---

[1] See also Varela et al. (1991); and Suchman (1987); Chiel and Beer (1997); Clark (1997); Kelso (1995); Lakoff and Johnson (1999); Thelen and Smith (1994); Turvey and Carello (1995); and Gallagher (2005) to name but a few.

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

Embodied Cognitive Science (Clark, 1997; Núñez & Freeman, 1999; Pfeifer and Scheier, 2001).

As a matter of fact, it has become somewhat inescapable to accept that a final understanding of human intelligence will be embodied and embedded. However, despite this fact, and the aforementioned gained momentum, recent episodes in the Cognitive Sciences suggest that cognitivism might be succeeding in assimilating some post-cognitivist methodologies and insights. Just as an example of how this strategy is carried out, take Vera & Simon (1993a) way to answer Gibson and Brooks' (1991) challenge of "doing without representing" (to use Clark & Toribio's 1995 title): for them, Gibson's affordances and Brooks "navigation without representation" approaches should be seen as an illustration of "orthodox symbol systems"! Commenting on Brooks' line of research, Vera and Simon (1993a) claim 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). Likewise, Vera and Simon (1993a) argue that 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). This very strategy of response, consisting of internalizing the relevant interactive relationships as symbolic states, has become typical and might be recognized in several recent dismissive discussions of Embodied Cognitive Science as a real alternative (Markman & Dietrich, ; Wilson, 2002).

Nevertheless, it should be acknowledged that this "assimilationist" strategy is facilitated by the lack of similar explanatory grip from the challenger. Thus, for instance, embodied cognition lays the stress upon real world, time-pressured, situations as the context in which cognition takes place and makes sense. This ecological dimension is also pivotal to ecological psychology (Gibson, 1979). A Gibsonian approach fits nicely with embodied cognitive science, insofar as affordances allow for a direct reach that avoids the exploitation of inner representational resources. Allegedly, no information-processing is required, simply tuning to environmental invariants.[3] We cannot possibly account for a cognitive agent's behaviour unless we treat it scientifically on a par with

[3] Traditionally, dynamicism has been seen as allied with Gibsonian approaches, while connectionism has been seen as allied with information-processing (Thelen and Bates, 2003).

the environment in which the agent is acting. Unfortunately, it is not crystal clear what we mean when we say that non-cognitivist approaches require a *non-symbolical* interpretation of a cognitive system's ecological interactions (Winograd & Flores, 1986). Does "non-symbolical" mean "sub-symbolical" or "non-representational" *tout court*? How can affordances play a causal role without being "tokened"? Is symbol-crunching the only way these affordances may give rise to an adaptive response? Are embodied patterns of activity "direct" in the sense that Gibsonian affordances are meant to be? (Calvo Garzón, in print). Taking into account that Brooks's subsumption architecture and Gibson's (1979) "direct realism" are standard-bearers against the sort of "indirect realism" that cognitivism instantiates, the question of whether we are talking past each other must be raised, and addressed.[4]

Anyway, we believe that there is a real difference in approach and explanatory scope, and that these different research programmes might resist the assimilationist strategy of cognitivism by developing their synergies and converging into an alternative paradigm. This is not an easy project, and we are well aware that the sure way for that not to happen would be 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 meeting, and its success will depend upon being able to generate the right intellectual climate and a common agenda.

As a matter of fact, the prospects are good, some things are changing, a move in the right direction can already be appreciated insofar as (i) decisions as to what to observe in studying cognition have changed (e.g., ecological studies of vision attend first to the optic array); (ii) the relevance of perception-action and social interaction to cognition

[4] In Vera and Simon's (1993a; 1993b) view, classicism and situated cognitive science need not be antithetical. As a matter of fact, situated action can be interpreted symbolically, or so they claim. It is not clear what is at stake in the dispute between classicist cognitive scientists and their opponents. 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). In view of these comments we cannot conclude much except that the debate is far from settled.

are out of question; (iii) the mathematical and conceptual toolkit of dynamic systems theory appears as a possible common formal medium of explanationc[5]; and (iv) questions are structured differently (e.g., we don't ask about the cognitive agent itself, but about the coupling of dynamic systems, where decentralization precludes the use of strict cognitivist answers). These changes have generated new experimental settings and novel predictions (epitomized by Thelen et al.'s, 2001, well-known dynamic understanding of the A-not-B perseveration error in infants, among many other).

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 efective convergence is the fact that conceptual issues are usually treated by philosophers, and empirical ones by cognitive scientists, separately. We need to put together conceptual analysis of the notions of representation, computation, emergence 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. On the other hand, we would also like to stress that the proof is in the pudding. An embodied cognitive 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 an embodied cognitive science has to offer that we can assess the extent to which we're confronting something truly different from classical cognitive science.

The purpose of the meeting, hence, is to connect fundamental theoretical problems about cognition and representation with recent developments in the empirical science. To achieve this, we need to enhance the intellectual exchange between the philosophy of cognitive science and the cognitive sciences themselves. Unfortunately, the working hypothesis of cognitivism has usually taken a form that may easily be dismissed as a straw man. Caricaturizing the two ends of the cognitive science spectrum has also

---

[5] In a recent special issue of *Developmental Science* (Spencer and Thelen, eds., 2003) researchers that have championed the connectionist and the dynamicist movements in the last decade united their efforts to fight against cognitivism. The risk exists, however, that the efforts at rapprochement between connectionism and dynamicism (Spencer and Thelen, 2003) force dynamicism into a connectionist mould.

prevented genuine progress.[6] It is, however, important to emphasize an ecumenical stance. Different frameworks are useful for different purposes. We would like to think of competing approaches as complementary scientific research programs and understand the criteria and constraints governing our choice of representational (?) framework in the same light.

In these regard, some of the questions we would like to see addressed at the meeting are:

- Is the technical vocabulary of dynamicism adequate enough for the study of cognition, or do we need other tools and theoretical notions?
- What are the sensorimotor grounds of higher cognition?
- What kinds of mental representations, if any, are still necessary as explanatory constructs? What notion of representation is involved?
- What is the nature of mental processes in the extended, social mind? Is causal relationships all there is to explanation?
- How can executive functions operate without a central executive?
- How can mental chronology studies be understood from an embedded, time-constrained, point of view?
- What is the relevance of robotics as simulation?
- Does it make sense, both empirically and conceptually, to integrate some of these approaches in hybrid models?
- Is it possible to go beyond a correlational neuroscience?
- Can new psychological phenomena be properly identified?
- What is it exactly that makes a paradigm or a research agenda fall under the cognitivist umbrella?
- Which agenda would have to be set in order to explote potential synergies?

In addition, there are some other, more specialized questions, that should also be considered. To begin with, in a trivial sense, cognition causally spreads out. However, it is not that clear that distribution per se implies that the cognitive unity object of scientific study must necessarily be the distributed system as such. Traditional analytical tools have been followed in an attempt to define a cognitive agent in terms of necessary and sufficient conditions.[7] We don't think that's a productive strategy, since we're lost in the tacit theoretical commitments being made when deciding upon which conditions to honour. Beyond the inner/outer dichotomy, operational descriptions of sensorimotor coordination are needed. Recent work certainly fits this desideratum (e.g., Thelen et al, 2001). However, if it is to increase its (mechanistic) explanatory value,

---

[6] For an illustration of this sterile dialectic see Vera and Simon (1993b) and Clancey (1993).
[7] See Wilson (2002) for an elaboration of this point.

more constraints from neurobiology must be honoured; constraints that allow for experimental testing, and for the generation of more accurate (and novel?) predictions.

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, forthcoming). 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 neurobiological plausible artificial neural network architectures and algorithms (Rolls and Treves, XXXX) 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, it is also a question of methodological convergence with the fast-growing Neuroscience.

A different, although related issue is whether a mature Embodied Cognitive Science must re-conceptualize the notion of "minimal cognitive behaviour" (Beer, 1996). This relates to the question of what counts as a cognitive system. Other questions relate to the sort of explanation that dynamic systems theory can offer to an embodied cognitive science. Does explanation boils down to the mathematical description of the changes of an extended system over time? And more importantly, how does that relate to the mechanistic components of the coupled system? What notion of explanation are we talking about? How can an embodied cognitive science relate to a (mechanistic) explanation (Bechtel, in press) of the sensorimotor coordination inner mechanisms that give rise to higher-level cognitive activity? It is important to emphasize that the answers to these questions has a direct bearing upon the epistemology of science, with immediate consequences as far as methodologies and the generation of testable predictions is concerned.

Moreover, things get worse when we move to higher-level capacities. Clark (1999), for example, presses upon representation-hungry cases, and questions whether an embodied cognitive science can account for them, beyond simpler forms of adaptive coupling, which involve some sort of online tracking, and that are less problematic. And this relates to question of representation. Intuitively speaking, online forms of covariation do not amount to representations (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 embodied cognitive science? Is the "representation-hungry" problem a byproduct of an information-processing framing of the problem? 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? What do we mean by "subsymbolic"? How can a representation be amodal? (see Markman and Dietrich, 2000, for some insights on these and related issues). Note that this is not "a debate for the sake of appearance", as has been suggested elsewhere (Haselager et al., 2003). Were things to end up being explained in adaptive coupling terms, that may lead to novel predictions, both qualitatively and quantitatively. We want to discuss what those types of novel predictions might possibly look like. In short, we want to know how far a full-blown embodied cognitive science can go (Clark, 1999).

Grush' (2004) concept of emulation, and his dynamical analysis of temporal representation (forthcoming) appears to throw some light upon these issues. Certainly, off-line cognitive activity critically unfolds in time and is embodied, but how does the endogenous exploitation of sensorimotor representations in emulation relate to symbolic offloading? From the perspective of the conceptual tools that dynamicism furnishes embodied cognitive science with, is it not possible to understand mathematically all phenomena, online as well as offline, as adaptive coupling? How do in short online and offline cognitive activity relate to each other?

Traditionally the burden of proof is on the side of non-classical approaches. Cognition is symbol-crunching, unless proved different ("the only game in town"). Although genuine as this dialectic is, we don't want to get lost in philosophical terrain, and are happy *not* to bite the bullet, and assume without further ado that embodied cognitive science is the most promising theory of the mind/brain we have. We want to pull of the existing common thread that non-cognitivist views exploit, namely, the decentralization of cognition. 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" 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 is: Let's unify this research effort into a single framework. Let's look at cognitive and non-cognitive phenomena in the same way, and see how far we can get in terms of generation of predictions and non-classical readings of the data collected. We want to assess the significance of an integrated, embodied cognitive science.

Toni Gomila - toni.gomila@uib.es

Paco Calvo - fjcalvo@um.es

DRAFT 13/10/2006

References

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

Bechtel, W. (in press). Explanation: Mechanism, modularity, and situated cognition. In P. Robbins and M. Aydede (Eds.). Cambridge handbook of situated cognition . Cambridge: Cambridge University Press.

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

Beer, R.D. (1996) Toward the evolution of dynamical neural networks for minimally cognitive behavior. In From Animals to Animats 4: Proc. 4th Int. Conf. on Simulation of Adaptive Behavior (Maes, P. et al., eds), pp. 421–429, MIT Press.

Beer, R.D. (forthcoming) "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, RA (1991) Intelligence without representation, Artificial Intelligence 47, 139–159

Calvo Garzón, F. (2005) "Rules, Similarity, and the Information-Processing Blind Alley", Behavioral and Brain Sciences 28, pp. 17-18..

Calvo Garzón, F. (in press) "Towards a general theory of antirepresentationalism", The British Journal for the

Philosophy of Science.

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.

Clancey, William J. (1993). Situated action: A neuropsychological interpretation (Response to Vera and Simon). Cognitive Science 17(1):87-107.

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

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.

Elman, J.L. and E. Bates, et al. Rethinking innateness. MIT Press, 1996.

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

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

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

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

Grush, R. (forthcoming) "Temporal representation and dynamics", *New Ideas in Psychology*.

Haselager, W.F.G., De Groot, A.D., & Van Rappart, J.F.H. (2003). Representationalism versus anti-representationalism: a debate for the sake of appearance. Philosophical Psychology, 16 (1), 5-24.
Hutchins, E. (1995) Cognition in the Wild, MIT Press
Kelso, S. (1995) Dynamic Patterns, MIT Press
Lakoff, G. and Johnson, M. (1999) Philosophy in the Flesh, MIT Press
Marcus, G., S. Vijayan, S. Bandi Rao, P. M. Vishton, Science 283, 77 (1999) Rule Learning by Seven-Month-Old Infants.
Markman, A.B. and Dietrich, E. (2000) “Extending the classical view of representation”, *Trends in Cognitive Sciences* 4: 470-5.
Nolfi, S. and Floreano, D. (2000). Evolutionary Robotics : The Biology, Intelligence, and Technology of Self-Organizing Machines.MIT Press.Pfeifer, R. & Scheier, C. (2001). Understanding Intelligence. Cambridge, MA: MIT Press.
Peña et al., 2002
Perruchet et al., in press
Pinker and Ullman, 2002
Port and van Gelder, 1995
Ramscar, 2002
Rolls and Treves, 1999
Schöner et al., 1997
Schyns et al., 1998
Seidenberg and Elman, 1999
Smith, Brian Cantwell (1996) On the origin of objects. Cambridge, MA: MIT Press.
Spencer and Thelen, eds., 2003
Suchman, A. (1987) Plans and Situated Actions, Cambridge University Press
Thelen and Bates, 2003
Thelen, E. and Smith, L. (1994) A Dynamic Systems Approach to the Development of Cognition and Action, MIT Press
Thelen et al.’s, 2001
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. et al. (1991) The Embodied Mind, MIT Press
Vera and Simon (1993a)
Vera and Simon (1993b)
Wilson, 2002
Winograd, T. and Flores, C. F. (1986) Understanding Computers and Cognition. A New Foundation for Design. Norwood: Ablex