Computationalism, Connectionism, Dynamicism and Beyond: Looking for an Integrated Approach to Cognitive Science

Víctor M. Verdejo

Abstract Cognitive scientists are nowadays apparently required to choose between at least three different competing schools or general approaches: the computational, the connectionist and the dynamicist. More than three decades of unresolved paradigm fight encourage an alternative view: that each of these general approaches offer, not different explanations, but explanations of different aspects of cognitive phenomena. In this paper, I articulate this view by showing that each general approach can be taken to promote research primarily within a particular level of explanation. Failure to appreciate this fact has frequently led to largely incomplete accounts within each school. I argue that, if the articulation offered is sound, it supports the statement of an integrated programme for cognitive science where all the aforementioned general approaches have their place. Finally, I illustrate this analysis via a central theme for a clash of rival explanations in cognitive research, namely, systematicity.

Cognitive science is a discipline in continuous evolution where different and conflicting research strategies are permanently brought to the fore. As a consequence of discussion in the last 30 years or so, cognitive scientists are now apparently required to choose between at least three different schools or overall approaches: the computational, the connectionist, and the (embodied) dynamicist. By quite general assent, these approaches are understood as being rival views on cognition. It has seemed therefore fair to describe this situation in roughly Khunian terms (e.g. Schneider 1987; Chemero 2009): cognitive science cannot consistently be taken to include all these different approaches. On the contrary, each school stands for a

V.M. Verdejo (✉)
Universidade de Santiago de Compostela, Praza Mazarelos s/n,
15782 Santiago de Compostela, Spain
e-mail: vmverdejo@gmail.com

different paradigm and the competing explanations they provide involve a typical scientific antagonism, one that results in a sort of radical paradigm fight. The conceptual and empirical discrepancies are so deep that, after this fight is resolved, at most one of these general approaches could turn out to be correct. The resulting view is somehow puzzling and disappointing. After many years of hard work in the cognitive sciences we might yet lack a clear answer to the question: what are the central theses that define cognitive science as a genuine discipline?

In this paper, I wish to make plausible the idea that this way of seeing cognitive research and the different approaches appeared in its wake is not only wrong but also itself highly pernicious for the correct assessment of its merits and achievements. Even though computationalism, connectionism and dynamicism are frequently seen as exclusive alternatives, much clarification can be gained by studying to which extent they articulate different aspects of one and the same scientific programme.[1]

This idea is certainly not new. Several authors have proposed analyses of cognitive science that, on the one hand, agree that computationalism, connectionism and dynamicism should be seen as three schools that may be united within a single framework (e.g. Cordeschi and Frixione 2007; Dawson 2013). On the other hand, other authors have actually developed specific aspects of this general idea in different fields, such as implementational connectionism (e.g. Marcus 2001), the theory of vision and visualizing (e.g. Pylyshyn 2003) or the representational approach to dynamicism (e.g. Grush 2003).

In this paper, however, I wish to explore the possibility of an integration of computationalism, connectionism and dynamicism in two steps: First, I will argue that in rough outline, the computational, the connectionist and the dynamicist models are plausibly considered as promoting lines of research focused and perhaps frequently 'misfocused' on one particular level of explanation. More precisely, I shall defend that these models correspond roughly to investigation developed primarily within one of Marr's three levels distinction (Marr 1982). Secondly, I will show that, if this is true, there is little reason to think that the aforementioned approaches to cognitive science are incompatible with each other. To the contrary, they are better seen as different lines of research belonging to a unified and integrated general inquiry into the nature of the mind.

Here is the route of the discussion to follow. In the first section, I outline the fundamental traits of Marrian explanation and present the view that computationalism, connectionism and dynamicism focus, or focus primarily, on research at one of these levels. In the second section, I show how concentration on one of Marr's levels has easily led in many cases to the problem of incomplete accounts in cognitive research. Section 3 is devoted to present what I, following Marr's own understanding of his

[1]Theoretical precision might require broadening the analysis to other cognitive schools as well. Extended, enacted, embedded, situated and distributed models might call for specific attention. However, simplicity and also the presumption that these other models might be reduced to or reasonably included into the ones here considered justify restricting our discussion to computationalism, connectionism and dynamicism.

tripartite distinction, present as integrated approaches in cognitive research, that is, approaches that take into consideration the complexity of cognitive phenomena at all levels of explanation. I finally illustrate, in the fourth section, the possibility of such an analysis with one central case in the history of cognitive studies, namely, the explanation of systematicity.

1 Different Approaches or Approaches at Different Levels?

Theorizing in cognitive science has probably so many perspectives as scholars working in the field. Nonetheless, it seems to be unobjectionable that, currently, there are at least three different and allegedly incompatible general strategies that may be held once engaged in the project of the cognitive study of the mind. On the one hand, there is the classical computational approach. This kind of approach capitalizes on the hypothesis that the mind is like a digital computer with discrete compositional symbols as mental representations (Newell and Simon 1972; Newell 1980). An alternative general approach emerges from consideration of connectionist networks. Connectionism challenges the classicist computational approach via the postulation of context-dependent and distributed mental representations in biologically plausible networks (McClelland and Rumelhart 1986; Rumelhart and McClelland 1986). A third and more recent alternative is found in so-called dynamicism or dynamical systems theory. This general approach takes the mathematical models of dynamics as the paradigm of cognitive research and thus provides explanations in terms of nonlinear differential equations (Horgan and Tienson 1994; van Gelder 1998).

I think it is plausible to consider the foregoing general approaches as focusing on and promoting research within one of Marr's levels (Marr 1982, chap. 1). For many years now, computational research has been used to the idea that there are at least three different levels of cognitive description or explanation. This is the seminal idea introduced by David Marr's influential computational account of vision. The three levels in question are, first, the one that specifies what the system does and why (Marr's level 1). Secondly, there is the level that states the algorithm and the representation (Marr's level 2). Finally, the lowest level takes into consideration the realization of the system in hardware structures (Marr's level 3). Remarkably, whereas much has been questioned in cognitive research, the explanatory importance of Marr's level-distinctions is, ultimately, pretty much untouched.[2]

[2]Marr was not alone in distinguishing levels of explanation in cognitive research. For instance, Newell (1986) or Pylyshyn (1984) provided a classification of cognitive explanations very similar to the one introduced by David Marr. These days, levels of explanation are, although largely unquestioned, also largely unattended to in the literature. This is arguably an unfortunate feature of contemporary cognitive theorising (see Verdejo and Quesada 2011 for an illustration of this point).

In this paper, I present an analysis of the aforementioned overall approaches in terms of Marrian levels so as to make plausible that, under certain standard readings, these approaches (1) centre research predominantly within one of these levels, taking the other levels to be irrelevant or else secondary; (2) as a consequence of (1), they often lead to incomplete and flawed accounts of cognition. Point (2) shall be the issue of the next section. In the remainder of this section, I will offer some considerations in favour of the plausibility of (1). Computational, connectionist and dynamical approaches can be analysed in terms of many different characteristic features. In what follows, I will try to show that at least some of these commonly attributed essential or distinctive features are explained because or appealing in the light of the kind of research to be expected at one of Marr's levels.

The introduction of (Marrian) levels of explanation for the analysis of computationalism, connectionism and dynamicism has been exploited by other authors as well (see e.g. Horgan and Tienson 1994; Cordeschi and Frixione 2007 and Dawson 2013). These analyses might differ, even greatly, with respect to the one presented here. It is important to note however that the existence of alternative analyses is no threat to the here presented line of argument. The argument only requires that, in the light of some essential or distinctive traits of each school, it is plausible to locate their contribution as belonging primarily to one of Marr's levels. It is only to be expected however that each of these schools (a) also exhibit commitments and substantial lines of research concerning other levels of explanation, and (b) may be judged as belonging to other levels in the light of different traits.

To begin with, then, let us discuss computationalism. Several ideas might be taken to be central to the models that characterize this school: serial processing, modularity, internal representations or an input-output analysis. Many, if not all, of these theoretical ingredients inevitably lead to the postulation of discrete symbolic structures on which different operations are defined. On the computational approach, cognitive states and processes are essentially defined in terms of the manipulation of these symbolic structures.

But note that discrete symbolic structures are precisely what one finds at high level characterizations of cognitive functions corresponding to Marr's level 1. This is especially clear if one considers that language and logical calculus (as opposed to e.g. perception and action) were the central concerns for the development of early cognitive research inspired by Turing-von Neumann architectures. Discrete symbolic structures, such as the ones found in the formulation of linguistic and logical problems (Newell and Simon 1972), are the hallmark of cognitive theorising at Marr's level 1. If we take cognitive science to consist of, or consist primarily of research at this level, it is only to be expected that the representational and the implementational level are taken to involve discrete representational and physical structures of the kind that result in a physical symbol system (Newell 1980). Whereas it is of course part of the computational view that there is an algorithmic and physical implementation of level-1 categories, it is nonetheless assumed that the structures identified at level 1 are entirely analogous to or provide the heuristics for the ones found at lower (algorithmic and physiological) levels.

Let us now move to connectionism. As is known, connectionist approaches appeared as a reaction to the computational dictum by emphasizing such things as the importance of fast, brain-like parallel processing, graceful degradation in the context of noisy inputs, the appeal to statistic and environmental data or the merits of associationistic mechanisms. The central tenet of connectionist models, however, has to do with the recourse to artificial neural networks or parallel distributed processing networks. Indeed, all of the aforementioned features are better seen as a consequence of the endorsement of this kind of networks as the core aspect of cognitive theorizing.

The adoption of connectionist networks is, patently, a fundamental assumption about the appropriate kind of representation, namely, distributed, multilayered and sub-symbolic representation (Smolensky 1988). But then, this is an alternative to the classicist computational model from the point of view of Marr's level 2. Connectionist models provide flexible context-dependent representations, which may be structurally very different from the categories identified at level 1 and which are arguably better suited for specific cognitive tasks such as fast pattern-recognition or unsupervised learning. On the other hand, if we put our bets on the algorithmic level of explanation as the genuine source for the understanding of cognitive phenomena then we might expect that the adoption of connectionism is or must be *definitive* regarding aspects certainly belonging to other levels, such as the endorsement of unstructured and purely context-sensitive functions (a claim at Marr's level 1) or biological plausibility (an implementational requirement at level 3).

The newest trend in cognitive science is dynamicism or dynamic systems theory. As with connectionism and computationalism, several aspects of this school might be emphasized: stability or self-organization of cognitive systems, agent-environment coupling or the quantification and nonlinearity of the variables relevant in the analysis. The fundamental trait of this kind of approach, however, is the statement of mathematical explanations of behaviour in terms of sets of differential equations that describe the evolution of a cognitive system in a state space or dynamic field.

It is natural to find this kind of approach appealing on the assumption that cognitive research has to be pursued primarily at Marr's implementation level 3. For one thing, according to such an assumption, cognitive systems are understood as being physical systems that continuously evolve over time. The assumption is precisely that cognitive systems are not different in nature (although they are certainly different in complexity) from a purely physical system –such as Watt centrifugal governor (van Gelder 1995). It is because a cognitive system is assumed to be a fundamentally physical system that the mathematics and dynamical language of physics is taken to be its appropriate explanatory tool.

For another thing, it is only in the light of such an assumption that we may find very natural indeed to endorse central ideas that accompany standard statements of dynamicist approaches (e.g. Calvo Garzón 2008; Chemero 2009). Two examples of these central ideas are, on the one hand, embodiment – i.e. the thesis that physical nature (as opposed to abstract functional analysis) is constitutive of cognitive systems–, and on the other hand, anti-representationalism – i.e. the view that representations at the algorithmic level have a residual explanatory role in cognitive research. If explanations at the implementational level are *the* genuine kind of explanations, then it follows that specifications of functions (at level 1) and of representations (at level 2) are dispensable or secondary. From this general point of view, dynamic models are naturally located within Marr's level 3.[3]

2 The Problem of Incomplete Accounts

Alas, and as I will be arguing in what follows, the fact that computationalism, connectionism and dynamicism focus on a particular Marrian level turns, more often than not, on a 'misfocus' as regards *bona fide* cognitive research. More precisely, it is easy to see many aspects of the dialectics confronting these schools as arising from the incompleteness of their accounts regarding the rest of levels of explanation.

Thus, for instance, computational models crystallized in Jerry Fodor's (1975, 2008) seminal Language of Thought Hypothesis (LOT henceforth). Frequent interpretations of this approach however, overemphasize the characterization of (specifically linguistic) cognitive processes from a high-level point of view, indeed, from a level at which common sense psychology could be resolutely vindicated (Fodor 1987). The idea was simple as it was controversial: if a computer parses sentences via a programming language, such as List Processing programming language (LISP), then a mind must also comprehend sentences via some still to be specified programming mental language. Standard developments of this idea easily lead to the impression that consideration of research at the algorithmic and implementational levels is subsidiary. The rule of thumb appears to be to extrapolate psychologically real representational categories and physical symbols from rough characterizations of grammatical parse-dependent functions. From this perspective, and although a LOT strategy arguably amounts to a general cognitive model applicable to a variety of domains (Verdejo 2012), it may seem that the classicist computational approach to cognition is confined to abstract linguistic or logical operations, and is therefore empirically unwarranted and insensitive to the pragmatic, contextual and physical concomitants of cognition (Dreyfus 1992; Hurley 2001; Chemero 2009). In short, much of the dialectical opposition to computational approaches can be seen as

[3]Admittedly, some authors have argued that dynamicism can be seen as contributions to Marr's level 1 (Cordeschi and Frixione 2007) or the algorithmic level 2 (Horgan and Tienson 1994). These approaches take as a defining feature of dynamical systems the statement of mathematical models from an abstract point of view. However, this interpretation of dynamicism arguably leaves unexplained the fundamentally embodied and anti-representational character of recent developments. This is not the place to develop this claim further. For present purposes, it suffices that the location of this school at Marr's level 3 is plausible in the light of the dynamicist traits just mentioned in the main text.

stemming from the mistake of considering, however tacitly, that serious research at the algorithmic and physical level is secondary or dispensable.

Now, as noted, it was a central part of connectionism to resist classicist computational models by focusing on Marr's level 2 and offering non-purely-linguistic or sub-symbolic alternatives (Smolensky 1988). However, the 'misfocus' of connectionism on the algorithmic level generally prevented the approach to properly capture the importance of level-1 characterizations. This was used by computationalists to mount a long-lived challenge to connectionism regarding systematicity (Fodor and Pylyshyn 1988): either connectionism does not explain the systematicity functions or else it explains them only as an implementation of computational models. The source of this and similar challenges is the absence of a clear statement of the function, if any, connectionism was showing how to compute.

On the other hand, a paradigmatic criticism of connectionism stems from its inability to account satisfactorily for the relation between connectionist and neuronal, psychologically and physiologically real networks (e.g. McLaughlin and Warfield 1994). In this respect, predominant connectionist advertising has failed to see that, qua contribution to the algorithmic level of explanation, the neuron-like, implementational virtues were often *not* a reasonable requirement for connectionist networks.

Finally, dynamic mathematical models, originally presented as belonging to the same family as connectionist models (van Gelder 1995), have engendered original explanations of traditional cognitive problems in perception, developmental psychology, sensorimotor tasks and much else. Unfortunately, focus or 'misfocus' on issues at Marr's level 3 has also brought dynamicism into a picture where, at the algorithmic level, anti-representationalism per se is an (empirically doubtful) scientific desideratum (Grush 2003). As a consequence of this, and as in some ramifications of the connectionist school, the absence of clear level-1 characterizations inevitably leads to neglect many high-level categories and to encourage some form of controversial behaviourism or eliminativism (e.g. Ramsey 2007).

3 An Integrated Approach to Cognitive Science

It seems to me that if the foregoing characterization is roughly correct, it automatically suggests an alternative conception of cognitive science where (a) the problem of incomplete accounts of cognitive phenomena does not arise in the first place and (b) computationalism, connectionism and dynamicism are prima facie compatible with each other after all. Partial views in cognitive research that focus or 'misfocus' on a particular level of explanation can be plausibly amended by views in which all levels of explanation are integrated. In fact, the possibility of an integrated account of cognition was arguably what Marr had in mind when he distinguished between levels of explanation in the first place.

Note that, according to Marr, his distinction responds to the real complexity of cognitive phenomena. Thus, the levels must be carefully attended to because “explication of each level involves issues that are rather independent of the other two” (Marr 1982, p. 25). More importantly in the context of this paper, all three levels are levels “at which an information processing device must be understood before one can be said to have understood it completely” (Marr 1982, p. 24). Marr’s fundamental contention was therefore not that one particular level is methodologically or explanatorily prior to the others but that full explanation of cognitive phenomena must involve correct explanations at all three levels. Unsurprisingly then, Marr thought that even the highest level 1, the level of what he called the computational theory, is a level at which we “should look out for the physical constraints [at level 3] that allow the process to do what it does” (Marr 1982, p. 103).

What these considerations suggest is that the best insights of computationalism (regarding level 1), connectionism (at level 2) and dynamicism (at level 3) may be integrated and thus make justice to the extraordinary complexity of cognitive phenomena via an extraordinarily rich repertoire of explanatory models. This is, to be sure, easier to say than to carry out. For this reason, I will introduce in the final section of this paper an illustration of how this guiding idea might go as regards a long-standing dialectic in cognitive research, that is, the systematicity debate.

4 The Case of Systematicity

Let us now consider, as an illustration of the view here presented, the case of an integrated account of systematicity. The illustration will be divided into two stages. On the one hand, I will introduce the fundamental explanatory schema of systematicity phenomena generally. At a second stage, I will present a series of considerations to the effect that an integration of cognitive models is clearly possible and certainly desirable in such a case.[4]

To a first approximation, systematicity is that property in virtue of which if a subject S has a given capacity C, then as a matter of fact, S has another, structurally related, capacity C’. The legendary example is that if S has the capacity to produce/understand the sentence “Mary loves John” then S, as a matter of fact, has also the capacity to produce/understand the sentence “John loves Mary”. Standard doctrine takes it that the correct explanation of systematicity involves compositionality (Fodor 1987, 2008; Fodor and Pylyshyn 1988). In other words, in order to provide a correct explanation of systematic states and processes one has to state, satisfactorily, a constituency relation between a corpus of primitive elements and the corresponding complexes made out from these primitives. Now, a first step in showing the possibility of an integrated approach in this case is

[4]Some authors will be willing to cast doubts on systematicity itself. In what follows, however, I will analyse the explanations that each school provides of the phenomenon and therefore I will assume that systematicity is an (empirically confirmed) explanatory target for each of the schools under analysis.

to note that constituency relations, of the sort that explain systematicity, can be consistently stated among different kinds of elements, depending on the level at which our explanation is operating.

Apparently, however, the explanations of systematicity phenomena that computationalism, connectionism and dynamicism have favoured exhibit radically opposing views. Classical computationalism has taken systematicity as perhaps the strongest evidence for the existence of a LOT, a linguistic representational system of discrete symbolic elements apt for combination and recombination (Fodor 1987; Fodor and Pylyshyn 1988). On the other hand, connectionist theorizing found that alternative, non-linguistic kinds of representational schemes –such as tensor-product networks (Smolensky 1990; Smolensky et al. 1992) or Gödel numbering (van Gelder 1990)– could provide all the sensitivity to structure that systematicity demands. Embodied dynamicism, finally, went on to postulate categories within dynamic theory in order to articulate a suitable interpretation of systematicity. The key idea is that systematicity is a mathematical relation between points or basins of attraction in state space (Horgan and Tienson 1994; Calvo Garzón 2004).

Considered in isolation, the resulting accounts within each school constitute a vivid example of the problems of incomplete accounts. First, misfocus on the linguistic structures that allowed the explanation of systematicity at Marr's level 1, made computational approaches end up with the implausible commitment that algorithmic and neurophysiological counterparts of those structures should be actually 'linguistic' in a very unpalatable sense. This was expressed sometimes as the requirement that systematicity-explaining constituency relations involved a sort of concatenative or spatio-temporal relation among constituents (e.g. Fodor and McLaughlin 1990), as if LOT was in an implausible literal sense necessarily a written or spoken language at the algorithmic and physiological levels.

To great avail, connectionism introduced the alternative sub-symbolic versions of representations for the explanations of systematicity. It was proved mathematically that classical explanations were equivalent to subsymbolic kinds of representations such as tensor-product networks (Smolensky et al. 1992). However, proponents of connectionist networks usually fell prey to a wrong assessment of their real findings: they were not offering alternative general models of cognition at all levels but, quite differently, alternative models at the algorithmic Marrian level 2 for already specified systematicity functions.

Finally, the anti-representational and embodied versions of dynamicism led to the idea that systematicity should be treated as a physical process accounted for by appeal to the mathematical tools and language of differential equations (Calvo Garzón 2004). The dynamicist proposal however introduced considerable confusion as to what kind of systematicity these models are actually explanations of –what the systematicity function they are modelling is – and left mysterious the real and obvious representational demands at level 2 of systematicity phenomena at large.

Careful scrutiny however can show that all these competing explanations of systematicity are only superficially incompatible. Take the 'Mary-loves-John' case. At the level of the function that the system computes (Marr's level 1), the problem and the explanatory categories are obviously linguistic in nature. What we identify at this level is a pair of systematically related capacities, namely, the capacity to understand/produce “Mary loves John” and the capacity to produce/understand “John loves Mary”. At this level, constituency relations of the sort that explain systematicity are stated between words or concepts and sentences or propositional thoughts.

At the level of the algorithm, however, we do not need to appeal to the very same categories that explained the phenomenon at the level of the function. It suffices that the representations postulated at this level are actually representations relevant for the explanation of how the system performs the function identified at level 1. A number of algorithmic possibilities would be available. These include programming languages and vector binding. It is not necessary that the representations involved mimic the very same linguistic structure appealed to in the explanation of the phenomenon at the highest level.

Finally, the continuous physical evolution of the system at the lowest implementational level may leave room for a dynamical insight in terms of nonlinear differential equations that define basins of attraction or a series of points in state space. But the consideration of this possibility does not exclude the existence of linguistic explanatory categories at level 1 or representational categories at level 2. Indeed, these accounts can be seen as showing central dynamical aspects of the way in which level-1 and level-2 categories are physically implemented.

Now we are in a position to see clearly what is gained by adopting an integrated approach in the case of systematicity. In an integrated approach, the constituency relations relevant for the explanation of systematicity are highlighted as involving different sorts of facts at different levels. The full complexity of cognitively relevant constituency relations is thus uncovered and carefully analysed. On an integrated account, constituency relations may be consistently stated between words or lexical categories (at level 1), activation vectors (at level 2) and points or basins of attraction in a dynamical space (at level 3). In the terms I am proposing, this means that an integration of the basic explanatory models under study –the computational, the connectionist and the dynamicist– is clearly possible.

In addition, one such integration is arguably highly desirable. The combination of rigorous accounts at different levels may serve to rectify, contrast and constrain the validity of assumptions at every level. For instance, computational models are presumably right to insist that it is compositional symbols, that is, semantically-cum-syntactically characterized representations that explain systematicity. What advocates of computationalism usually neglect, however, is that since what we are concerned with is cognitive and, to that extend, psychologically real compositionality, a correct statement of the corresponding representations is certainly more involved than, say, high-school grammatical analysis. Algorithmic and physical levels have to be carefully attended to.

The available algorithms, in turn, must take into consideration the full range of phenomena that count as systematicity phenomena (Cummins et al. 2001), and hence include, for instance, the systematicity found in vision and visual abilities. This means that the systematicity algorithm and representation is very likely of a connectionist kind (as opposed to a programming language kind) and will never be 'linguistic' in a very literal sense (Verdejo 2012).

Finally, the physical constraints of the organism, plausibly modelled in dynamic terms, have a substantial role, on the one hand, in the specification of physically plausible mechanisms. Functions and algorithms that the organism computes will be translated into dynamical systems with their own restrictions regarding relevant physical variables and inter-dependence among those variables. On the other hand, the anti-representational leanings of dynamical systems (Calvo Garzon 2008) may prevent the postulation of unnecessary representational categories and hence contribute to a more demanding specification of their explanatory role.

In sum, when systematicity is considered as an 'integrated' cognitive phenomenon, explanations in terms of compositionality and constituency relations result in a substantial, and substantially complex, inter-level account. This is precisely what we get if we buy in for an approach to systematicity of the sort here defended.

5 Conclusion

In this paper, I have undertaken the task of showing, in broad and brief outline, that there is a promising alternative to the opposing terms in which different approaches to cognitive science are usually presented. The alternative suggests that many central aspects of the unresolved confrontation between computationalists, connectionists and dynamicists are better seen as a failure to appreciate the integrative relations between researches at different levels of explanation. From these considerations, the sketch of an integrated approach emerges, one in which empirically grounded explanation at different Marrian levels results in a powerful and rich framework for cognitive science.

Acknowledgement I am grateful to the audience of the EPSA11 conference in Athens and to an anonymous referee for their valuable comments and suggestions on earlier drafts of this paper. This research work has received financial support from the Ministerio de Economía y Competitividad (Spanish government), via the research project FFI2009-08828/FISO, and from the Generalitat de Catalunya (Catalan government), through the consolidated research group GRECC (SGR2009-1528).

References

Calvo Garzón, F. (2004). Context-free versus context-dependent constituency relations: A false dichotomy. In S. Levy & R. Gayler (Eds.), *Proceedings of the American Association for Artificial Intelligence* (pp. 12–16). Menlo Park: AAAI Press.

Calvo Garzón, F. (2008). Towards a general theory of antirepresentationalism. *British Journal for the Philosophy of Science, 59*, 259–292.

Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
Cordeschi, R., & Frixione, M. (2007). Computationalism under attack. In M. Marraffa, M. De Caro, & F. Ferretti (Eds.), *Cartographies of the mind* (pp. 37–49). Dordrecht: Springer.
Cummins, R., Blackmon, J., Byrd, D., Poirier, P., Roth, M., & Schwarz, G. (2001). Systematicity and the cognition of structured domains. *Journal of Philosophy, 98*, 167–185.
Dawson, M. R. W. (2013). *Mind, body, world: Foundations of cognitive science*. Edmonton: Athabasca University Press.
Dreyfus, H. L. (1992). *What computers still can't do*. Cambridge, MA: MIT Press.
Fodor, J. A. (1975). *The language of thought*. New York: Thomas Y. Crowell.
Fodor, J. A. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
Fodor, J. A. (2008). *LOT2: The language of thought revisited*. Oxford: Oxford University Press.
Fodor, J. A., & McLaughlin, B. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition, 35*, 183–205.
Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28*, 3–71.
Grush, R. (2003). In defense of some 'Cartesian' assumptions concerning the brain and its operation. *Biology and Phylosophy, 18*, 53–93.
Horgan, T., & Tienson, J. (1994). A nonclassical framework for cognitive science. *Synthese, 101*, 305–345.
Hurley, S. (2001). Perception and action: Alternative views. *Synthese, 129*, 3–40.
Marcus, G. F. (2001). *The algebraic mind*. Cambridge, MA: MIT Press.
Marr, D. (1982). *Vision*. San Francisco: Freeman.
McClelland, J. L., Rumelhart, D. E., & the PDP Research Group. (1986). *Parallel distributed processing* (Vol. 2). Cambridge, MA: MIT Press.
McLaughlin, B., & Warfield, T. (1994). The allure of connectionism re-examined. *Synthese, 101*, 365–400.
Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4*, 135–183.
Newell, A. (1986). The symbol level and the knowledge level. In Z. Pylyshyn & W. Demopoulos (Eds.), *Meaning and cognitive structure* (pp. 31–39). Norwood: Ablex.
Newell, A., & Simon, H. (1972). *Human problem solving*. Englewood Cliffs: Prentice Hall.
Pylyshyn, Z. W. (1984). *Computation and cognition*. Cambridge, MA: MIT Press.
Pylyshyn, Z. W. (2003). *Seeing and visualizing*. Cambridge, MA: MIT Press.
Ramsey, W. (2007). *Representation reconsidered*. New York: Cambridge University Press.
Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (1986). *Parallel distributed processing* (Vol. 1). Cambridge, MA: MIT Press.
Schneider, W. (1987). Connectionism: Is it a paradigm shift for psychology? *Behaviour Research Methods, Instruments, and Computers, 19*, 73–83.
Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioural and Brain Sciences, 11*, 1–74.
Smolensky, P. (1990). Tensor product variable binding and the representation of structure in connectionist systems. *Artificial Intelligence, 46*, 159–216.
Smolensky, P., Legendre, G., & Miyata, Y. (1992). Principles for an integrated connectionist/symbolic theory of higher cognition. Institute of Cognitive Science, University of Colorado, Technical Report 92-08.
Van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science, 14*, 355–384.
Van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy, 92*, 345–381.
Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioural and Brain Sciences, 21*, 615–628.
Verdejo, V. M. (2012). Meeting the systematicity challenge challenge: A nonlinguistic argument for a language of thought. *Journal of Philosophical Research, 37*, 155–183.
Verdejo, V. M., & Quesada, D. (2011). Levels of explanation vindicated. *Review of Philosophy and Psychology, 2*, 77–88.