

The systematicity challenge to anti-representational dynamicism

Víctor M. Verdejo

Received: 24 February 2014 / Accepted: 20 October 2014 / Published online: 17 December 2014
© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract After more than twenty years of representational debate in the cognitive sciences, anti-representational dynamicism may be seen as offering a rival and radically new kind of explanation of systematicity phenomena. In this paper, I argue that, on the contrary, anti-representational dynamicism must face a version of the old systematicity challenge: either it does not explain systematicity, or else, it is just an implementation of representational theories. To show this, I present a purely behavioral and representation-free account of systematicity. I then consider a case of insect sensorimotor systematic behavior: communicating behavior in honey bees. I conclude that anti-representational dynamicism fails to capture the fundamental trait of systematic behaviors *qua* systematic, i.e., their involving exercises of the same behavioral capacities. I suggest, finally, a collaborative strategy in pursuit of a rich and powerful account of this central phenomenon of high cognition at all levels of explanation, including the representational level.

Keywords Dynamicism · Behavioral systematicity · Levels of explanation · Representation · Systematicity challenge

1 Introduction

New trends in cognitive research capitalize on anti-representational dynamicism, which these days has taken much of the role once played by connectionism as the hallmark of heterodoxy. Connectionism is apparently not at the centre of the controversy anymore. This seems to be so in spite of the fact that connectionism arguably

V. M. Verdejo (✉)
Department of Philosophy, University College London, Gower Street,
London WC1E 6BT, England
e-mail: v.verdejo@ucl.ac.uk

provided the research context that gave rise to and supported dynamical theory in the first place.¹ Prominent quarters in cognitive science agree that cognition is essentially accountable in terms of the assumptions and mathematical language of dynamics. As many authors have pointed out, dynamicism is not necessarily anti-representational. However, dynamicism, unlike its connectionist and classicist predecessors, “forms a powerful framework for developing models of cognition that sidestep representation altogether” (Van Gelder 1998, p. 622). More importantly, the hopes associated with dynamical approaches to cognition crystallize vividly in a growing consensus that assumes anti-representationalism as a working hypothesis (e.g., Varela et al. 1991; Van Gelder 1995; Port and van Gelder 1995; Thompson 2007; Calvo Garzón 2008; Chemero 2009).

Remarkably, the anti-representational dynamicist turn has led to a number of influential accounts in several fields such as robotics (Brooks 1991), motor control and coordination (Haken et al. 1985; Kelso 1995; Beer 1995), developmental psychology (Thelen and Smith 1994; Thelen et al. 2001), perception categorization (Beer 2003), imagined action (Van Rooij et al. 2002), to name a few. In this paper, I concentrate on a particular explanatory target, namely, systematicity. At present, we lack a fully developed or explicitly articulated dynamicist account of systematicity. However, proponents of non-representational dynamicism naturally attempt to extend their models to phenomena of high cognition and hence try to deal with what apparently are “representation hungry problems” (Clark and Toribio 1994; Clark 1997, pp. 166–170). In this context, systematicity is indeed one decisive and central touchstone for the consideration of a promising and overarching cognitive research program of high cognition. Unsurprisingly, systematicity has been explicitly acknowledged as a genuine explananda in dynamicist developments (Horgan and Tienson 1994, pp. 328–333; Petitot 1995; Van Gelder 1998, §6.9; Calvo Garzón 2004). The clear presumption, therefore, is that dynamicism may offer a rival non-representational explanation of systematicity and contribute as a third party to the debate carried out by classicist and connectionist contenders for the last 30 years or so.

As is known, the ‘systematicity debate’ has consisted of a relentless classicist/connectionist debate. Since Fodor and Pylyshyn’s seminal paper (1988), the dialectics surrounding systematicity may make it seem that we are, as it were, stuck in a closed loop. On the one hand, supporters of classical schemes of mental representation have emphasized again and again the need of a compositional system of symbols in the account of systematicity phenomena (e.g., Fodor and McLaughlin 1990; McLaughlin 1993, 2009; García-Carpintero 1995; Aydede 1997; Fodor 1997; Hadley 2004). On the other hand, proponents of connectionist sub-symbolic schemes of mental representation have tried to respond to the different challenges posed by classicists in a variety of ways (e.g., Van Gelder 1990; Smolensky 1990; Smolensky et al. 1992; Matthews 1994; Cummins 1996; Cummins et al. 2001) (see Aizawa 2003 for detailed discussion).

¹ Van Gelder for instance takes connectionism to be “a particular subcategory of the wider class of dynamical systems” (Van Gelder 1995, p. 370), and Horgan and Tienson assume that the dynamical framework provides “the natural mathematical characterization of connectionist networks” and take this fact as a way of “further articulating and supporting the [dynamical] framework” (Horgan and Tienson 1994, p. 305). See also Dorffner (1999) for a general articulation of the relation between connectionism and dynamicism.

Whether or not connectionist models support a form of analogue computation, both classicist and connectionist competitors in the debate are versions of representational computationalism, that is to say, the view that cognition consists in the manipulation of representations according to rules. We are indeed familiar with the view that (radical or extreme) dynamicism parts company with representational computationalism (e.g. Van Gelder and Port 1995; Eliasmith 1996; Chemero 2009; Fresco 2012). In this context, the desiderata associated with anti-representational dynamicism, considered as a general theory of cognition, clearly includes the aim of offering a rival and new kind of explanation of systematicity. Some authors have even shown sympathy to the view that dynamicism may not only offer a rival account, but entirely reformulate the debate out of the representational-computational paradigm so as to resolve or somehow stop the dialectical loop between classicist and connectionist parties. This seems to be Calvo Garzón's standpoint when he writes:

It is noteworthy, however, that both hypotheses, the classical and the connectionist, fall neatly within the information-processing paradigm. [...] 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. [...] In view of this scenario, I contend, we may need to consider turning to questions concerning the role that potential contenders, such as Dynamic Systems Theory (DST) [...] may play in the future. (Calvo Garzón 2004, p. 14)

Calvo Garzón's clear suggestion in this passage is that anti-representational DST is called to provide, in the context of discussions about systematicity, a way out of the "information-processing blind alley", as he calls it, represented by classicist and connectionist versions of computationalism.²

It is my view that anti-representational dynamicism will probably not offer a fully satisfactory account of systematicity of the sort suggested in Calvo Garzón's quotation. In this paper, I will argue, more precisely, that any explanations of systematicity offered on the anti-representational dynamicist model should be ready to meet a renewed version of the old systematicity challenge. Fodor and Pylyshyn's (1988) celebrated challenge can be stated thus: either connectionist accounts do not explain systematicity or, if they do, they are (mere) implementations of (and hence no real alternative to) accounts in terms of symbolic representation. Without endorsing any particular position regarding the outcome of this challenge, in this paper I show that it is revisited in the anti-representational version encouraged by dynamicism: either

² As several authors have pointed out, computation must not be confused with information processing (Piccinini and Scarantino 2011; Fresco 2013). In addition, it is not clear that anti-representational dynamicism should stand against a generic computational paradigm, one that includes some form of non-digital (analogue) or non-representational computation. In the context of the present discussion, anti-representational dynamicism is considered in opposition to the family of theories that embrace a form of representational computationalism—conceived broadly as including classicist or symbolic and connectionist or sub-symbolic variants—which involves (semantic) information processing of the sort required by the recourse to representations. More on the target notion of representation below.

non-representational dynamicism does not explain systematicity, or else, if it does, it is just an implementation of representational accounts.³

I do not mean to suggest that the presented line of reasoning is, if sound, conclusive against the merits of an anti-representationalist stance generally. Anti-representational approaches may shed light on a variety of aspects of cognition and may also have a regulative role for the explanatory useful postulation of representations. Furthermore, the discussion will involve certain obvious limitations: this critical assessment will be based upon an analysis of just one particular simple case and in relation to the specific problem of systematicity. Sympathizers of anti-representational dynamicism may still argue for the irrelevance of the case, the problem, or both. Nonetheless, the suggestion is that the analysis to follow spells out (1) the nature of the difficulties anti-representational dynamicism must face in the account of systematicity and (2) the traits of the debate we are forced to face once anti-representational dynamicism is under serious consideration as a general alternative model of systematic cognition.

The paper is structured in the following way. The next section presents the issue under scrutiny via a neutral characterization of dynamicism, representation and systematicity. This clarification task is especially demanding in the case of the notion of systematicity because an explicit *non*-representational characterization of this notion is lacking in the literature. In Sect. 3, I will analyze a particular case of systematic sensorimotor behavior not suspicious of involving tententious computational or representational assumptions: systematic behavior in the honey bee. I will also carefully examine in which sense dynamicist approaches may provide an explanation of such behavior and conclude that anti-representational dynamicist accounts fail to explain the fundamental trait of systematic behaviors *qua* systematic, i.e., their involving the exercise of the *same* behavioral capacities. As a conclusion (Sect. 4), I will suggest a distinctive way out of the pose-challenge/respond-to-challenge dialectics: to look for a unified, though rich and complex, cognitive science where different levels of explanation—including notably the representational level—result in powerful accounts of cognitive phenomena, such as systematicity.

³ Though we lack a ready-made statement of the implementation relation, in this paper I follow standard practice in the literature (e.g., Fodor and Pylyshyn 1988; Fodor and McLaughlin 1990; Marcus 2001; McLaughlin 2009) in considering implementation as (a) a relation between (categories or processes at) higher levels and (categories or processes at) lower levels of explanation (b) which is asymmetrical (if A implements B, then B does not implement A), one-many (A_1, A_2, \dots, A_n may be different candidate implementations for B), and transitive (Fodor and Pylyshyn 1988, p. 68) (if A implements B and B implements C, then A implements C). In the specific context of this paper, the target high level in question is the representational level, broadly conceived. Anti-representational dynamicism would offer implementations of (classicist or connectionist) representational accounts in dynamical systems of differential equations located at a level or levels below the representational level. That classicist and connectionist representational theories may be implemented in non-representational systems is something explicitly acknowledged already in Fodor and Pylyshyn's seminal paper (1988, p. 10, emphasis theirs):

[...] just as it is possible to implement a *Connectionist* cognitive architecture in a network of causally interacting nonrepresentational [e.g., neurological, or molecular, or quantum mechanical] elements, so too it is perfectly possible to implement a *Classical* cognitive architecture in such a network.

2 Dynamicism, representation and systematicity. The issue neutrally described

This is the question that concerns us here: can dynamicism account for systematicity phenomena without recourse to representations? A proper statement of the question requires an explanation of the meaning of the terms involved.

2.1 Dynamicism

It is hard to downplay the rising importance of dynamicism or dynamical systems theory or the dynamical hypothesis in cognitive science over the last years. Since scholars are familiar with the fundamental traits of this school, a brief characterization will suffice for present purposes.

Dynamicism is the general view that cognitive behavior can be accounted for in terms of the mathematical models of dynamics. More precisely, the dynamicist thesis claims that cognitive behavior is to be explained in terms of sets of (nonlinear) differential equations that provide the values of a number n of variables as changing over time and which define an n -dimensional state space or dynamic field.⁴ In a sense, this thesis is a truism: since, by general assent, cognitive behavior is physically implemented, it must be accountable in terms of the dynamical language of physics. The truism disappears when it is claimed that cognition can and should be modeled in such terms across the board. Thus, dynamicism turns on a commitment to explanations of cognitive phenomena essentially involving (numerical) quantification of variables, a metric of time, analysis of the interdependence between variables, and focus on differential equations of change of those variables over time. In the literature, there are several features that can be emphasized or added to the picture: stability or self-organization (of behavior under certain conditions) (e.g., [Kelso 1995](#)), real time modeling (as opposed to ‘ersatz’ time modeling) (e.g., [Van Gelder and Port 1995](#)), continuity in state-space evolution (e.g., [Calvo Garzón 2008](#)), agent-environment coupling (e.g., [Beer 1995](#); [Chiel and Beer 1997](#)) or quantitative character (of the variables and behavior in the system) (e.g., [Van Gelder 1998](#)). Careful delineation of all these related aspects goes clearly beyond the purposes of our discussion.

Crucially, some theorists take it that dynamicism amounts to or else strongly encourages anti-representationalism, that is, the thesis that cognition is not representational computation. The question of whether dynamicism really entails or otherwise naturally demands anti-representationalism is contentious. As is known, even paradigmatic cases of dynamical models can be taken to involve representations in fundamental ways ([Bechtel 1998](#); [Grush 2003](#)). Assessment of such deep questions concerning dynamicism is beyond the reasonable scope of this paper. Here, I will explore the nature of dynamicist explanations of systematicity only insofar as they are within the scope of anti-representational dynamicism (see also Sect. 1 and fn. 2).

⁴ [Van Gelder \(1995\)](#) famously illustrated dynamical systems via the Watt centrifugal governor for controlling the speed of steam engines. In general, the best examples of dynamical modeling are what in control theory are known as closed-loop control systems (see [Grush 2003](#) for detailed discussion).

2.2 Representation

Although more ambitious statements are available, by representation I understand a state, set of states or process (hence, something out there in the world, plausibly in somebody's head) which stands in for or carries information about other states or events. That which the state or process stands in for or carries information about counts as the content of the representation. Representations, so understood, involve (a) a possible range of contents associated with them; (b) these contents being associated in a determined way (in accordance with a certain representational scheme); and (c) with conditions on proper functioning or manipulation and correct representation. Thus, I follow many authors (e.g., [Clark and Toribio 1994](#); [Clark 1997](#); [Van Gelder 1995](#); [Bechtel 1998](#); [Grush 2003](#); [Chemero 2009](#)) in taking a version of Haugeland's tripartite characterization (cf. [Haugeland 1991](#), p. 62) as a baseline notion of representation. Although such a notion of representation involves a commitment to (information-processing) computationalism, this capsule-form definition has the merit of being quite neutral as regards the many different ways in which the notion of representation might be ultimately understood. For instance, it remains neutral as to whether representational contents are internal or external to the computing system, whether representation is primarily language-like or analogue, or whether representation should be characterized in terms of strong decouplability (of the potentially absent target represented and the representing state) or else in terms of a weak (feedback-dependant) decouplability (as in emulation theory). For present purposes, we can postpone a statement of the precise nature of representation by appealing to the propounded baseline and all-embracing conception of representation.

Generally, and in the particular case of systematicity, I assume that the legitimate postulation of representations requires the specification of a real explanatory task not achievable by alternative theoretical means. Even though one can find a priori arguments in favor of the classical representations in cognitive science (e.g., [Davies 1991](#)), I also assume that whether we are right in considering a state, set of states or process of an organism as fulfilling such a task is, ultimately, an empirical question. There seems to be a large consensus on these points among classicists and dynamicists alike (e.g., [Burge 2010](#); [Van Gelder 1995](#), p. 352; [Bechtel 1998](#); [Beer 2014](#)).

2.3 Systematicity

Our discussion must rely on a neutral notion of systematicity. More precisely, the issue at stake requires that we describe systematicity without appealing, either explicitly or implicitly, to the notion of representation. This turns out to be not a very easy task. The reason is that representation-free characterizations of systematicity are, if not utterly absent, quite *unusual* in the literature: the paradigm cases of systematicity phenomena are linguistic or conceptual cases, that is to say, cases that very obviously fit the representational scheme (see e.g., [Fodor 1987](#); [Fodor and Pylyshyn 1988](#); [McLaughlin 1993, 2009](#)). Some authors would even suggest that systematicity should be stated in terms of “abilities to have mental representations with propositional contents” ([McLaughlin](#)

2009, p. 254). However, a purely behavioral characterization of systematicity is very much needed for a neutral assessment of the systematicity issue generally (see e.g., García-Carpintero 1995; Verdejo 2012). A purely behavioral characterization prevents prejudice in the assessment of rival explanations at the representational level and is patently required when engaged in the project of assessing anti-representational theories. What I propose for present purposes is, therefore, to characterize systematicity in terms of *causally and nomologically related behaviors*.

By nomologically and causally related behaviors (or nomologically related behaviors for short) I understand cognitive behaviors for which conditionals of a certain form are true in virtue of/ justified by appeal to causal laws, or laws involving causal relations. By ‘law’ in this context I refer simply to *ceteris paribus* empirical generalizations or regularities which are counterfactual supporting (cf. Aizawa 2003, pp. 28–29; McLaughlin 2009, pp. 252–253). For present purposes, we do not need to specify the nature of the causal relation that the law involves. Suffice it to say that the causal relation in question would prevent the law to consist of a merely accidental regularity or generalization. The form of the conditionals is as follows: if behavior (of type) A occurs, then behavior (of type) B (at least potentially) also occurs. The conditionals do not express a purely causal relation between behaviors A and B. They express a law or generalization, or an instance of a law or generalization, that connects A and B and which involves a causal relation. Thus, the causal relation need not, and typically should not be read off the conditional itself. The conditionals do not say, and in the cases to be considered it is generally not the case, that A behaviors cause or bring about B behaviors.

Thus, the propounded basic characterization is this: given two nomologically and causally related cognitive behaviors, A and B, they are systematically related behaviors insofar as they are exercises of the same cognitive behavioral capacity or set of behavioral capacities. The relevant truism of this characterization is that any pair of behaviors whatsoever, A and B, might be nomologically related behaviors without their being systematically related. Let us state this more precisely. The characterization assumes that every pair of cognitive behaviors, A and B, are, respectively, the behavioral result of a set of cognitive behavioral capacities $A(c_1^a, c_2^a, \dots, c_n^a)$ and $B(c_1^b, c_2^b, \dots, c_n^b)$. Provided that there is a causal law that connects A and B, to render A and B systematic requires, in addition, that at least one of the capacities involved in the production of A is the same as one of the capacities involved in the production of B—i.e., $\exists c_i^a c_u^b (c_i^a = c_u^b)$.⁵

Some examples will help to elucidate this notion of systematicity. Thus, a subject S’s riding a bicycle (behavior A), causally entails S’s (potential) perceptual identification of bicycles (behavior B). However, the capacities responsible for A (having to do with limb movement, motor-coordination or equilibrium) are in this case completely distinct from the capacities responsible for B (such as memory retrieval or recognition

⁵ The approach must therefore make room for a notion of complex behavior (i.e., behavior that arises out of more than one behavioral capacity) but note that this is not to enter into a discussion about compositionality. In the systematicity debate, compositionality is a property of (structured) representations. For present concerns, granted that we do not wish to beg the question against the anti-representationalist, and even if a connection can be pursued regarding the proposed approach, compositionality issues must be put aside.

of bicycle-defining perceptual traits). Similarly, S's eating a burger (behavior A) is nomologically and causally related to, as it might be, S's digesting it (behavior B). These behaviors would however fail to be systematic because none of the behavioral capacities for A (such as capacities for biting, chewing or swallowing) are the same as the capacities intervening in the production of B (such as the ones involved in nutrients decomposition, chemical alteration and absorption).

Now, a paradigmatic case of systematic behavior is the following: to utter the sentence 'John loves Mary' and to utter the sentence 'Mary loves John'. These behaviors are nomologically related ones: in actual language-users, if one exhibits one of them, then one must (at least potentially) exhibit the other. Following our proposed characterization, this pair of behaviors is, in addition to its being nomologically related, also a systematic pair: the behaviors at stake correspond to the same behavioral capacities, namely, the capacities to utter the words 'Mary', 'John', and 'loves'.

The distinction between merely causally and nomologically related behaviors and systematically related ones is crucial for a correct assessment of the systematicity debate. The distinction reflects the 'intrinsic-connection' requirement for systematicity appealed to by Fodor and allies (e.g., Fodor 1987, p. 149; Fodor and Pylyshyn 1988, p. 37) where, to put it in Aizawa's terms, "the claim regarding intrinsic connections concerns cognitive capacities or competences" (2003, p. 92). In a similar vein, McLaughlin analyzes systematicity as involving a "constitutive basis" for the possession of capacities (McLaughlin 1993, §2). In the context of our discussion, we can gloss the use of these expressions as expressing the fundamental point that in order for a pair of behaviors to be systematic it is not enough that their presence or occurrence is regularly and causally connected in nature. In addition, there must be a common explanation of this fact. In the terms I am proposing, the explanation concerns the existence of common behavioral capacities involved in the production of the target behaviors. We may state the foregoing points in terms of Behavioral Systematicity (BS):

(BS)

A given pair of behaviors A and B is systematic to the extent that:

- (a) A and B are causally and nomologically related behaviors (they comply with causal laws of the conditional form: if A, then (potentially) B).
- (b) A and B are exercises of at least one common behavioral capacity.
- (c) The fact that (b) is (part of) an explanation of the fact that (a).⁶

BS invokes only behaviors and behavioral capacities and is, therefore, a representation-free notion of systematicity of the sort we need in order to be neutral with respect to anti-representationalist views. Clearly, what is crucial for BS is the correct, empirically contrasted, identification of *behavioral capacities as manifested in overt or observable (types of) behavior*. Several counterfactuals may be used to establish the existence of such capacities for any pair of systematic behaviors A and B. The counterfactuals would show that one does not find organisms that exhibit A without exhibiting B, or that if an organism ceases to exhibit A, then it ceases to exhibit B. These would be

⁶ The term 'explanation' should be read in ontological terms as meaning 'true explanation'. It goes without saying that one such explanation must offer a way of properly discriminating behavioral capacities. More on this below.

empirically contrasted counterfactuals to the effect that A and B constitute kinds of cognitive behaviors that, as Fodor and Pylyshyn put it, come “in structurally related clusters” (Fodor and Pylyshyn 1988, p. 49). One may of course deny, or otherwise exhibit extreme skepticism towards the existence or the widespread existence of systematicity phenomena (e.g., Johnson 2004; Gomila et al. 2012). The ongoing dialectics, however, requires that one accepts some neutral characterization of systematicity to advance the present discussion.

Several caveats may prevent much misunderstanding in the discussion to follow. In the first place, the presented notion of systematicity is not to be considered as involving a definition with necessary and sufficient conditions. I side here with McLaughlin in thinking that “no statement of such [noncircular necessary and sufficient] conditions, no definition [of systematicity], is to be had” (McLaughlin 2009, p. 252). In the context of our discussion, it is enough that we have a sufficiently clear, coherent and handy notion of behavioral systematicity applicable to central cases.

Admittedly however, the target notion may be put to the test by a number of limiting cases in which *the individuation of behaviors and behavioral capacities* is not deemed appropriate. For instance, my catching this baseball at 75 km/h (behavior A), is certainly an exercise of the same behavioral capacity as catching this baseball at 75.5 km/h (behavior B). I doubt however that one is really tempted to judge that these behaviors are, in any cognitively relevant sense, systematically related (instead of considering them as the same kind of behavior *simpliciter*). We may refer to these cases as *trivially systematic*: BS is fulfilled for A and B but only because $A=B$. Since we do not want the number of systematic behaviors under consideration to increase absurdly, a minimally interesting reading of the notion requires excluding cases of trivial systematicity. Similarly, note that condition (c) rules out cases of what we can call *idle systematicity*, that is, cases in which (a) and (b) in BS are fulfilled for A and B but only because the capacity considered to be involved is a catch-all, explanatorily irrelevant capacity. For example, eating, climbing, flying or running might all be viewed as exercises of some general capacity of survival or as involving the capacity of breathing. But these behaviors are systematically related only in an idle and explanatorily irrelevant sense in relation to such capacities.

Relatedly, inappropriate reading of the target characterization seems to rule out paradigmatic instances of systematicity. Mentally performing a piece of mathematical or logical reasoning is certainly a case in point. There is no obvious sense in which we should consider the steps in the mental calculus as different sorts of behaviors—indeed, it is not obvious why we should consider them as overt behaviors at all—in which case the propounded characterization may be called into question. Such cases are characterized by the absence of specific observable behaviors associated with them. Sufficiently complex analysis, however, might construe cases of logical or mathematical reasoning as involving relevant sorts of (intelligent) overt behavior such as successfully performing calculations or providing correct answers to logical problems. The details of such a story may be hard to fill. The presumption, however, is that such a story would provide candidate behaviors that would render abstract reasoning and similar processes of high-level cognition systematic in our sense.

Finally, there are two further conditions that the relevant behaviors must reasonably meet. First, the nomological relation of behaviors that are candidates for systematic-

ity must plausibly be symmetric, that is, the conditionals of the form ‘if behavior A, then (potentially) behavior B’ must be reversible to ‘if behavior B, then (potentially) behavior A’. If the nomological relation is not reversible in this way, this would be quite strong evidence that there is no common behavioral capacity involved in their production. Secondly, the kind of behaviors at stake must not be purely reactive behaviors. The amoeba moving around in water may be taken to satisfy some instances of BS. Granted that the behavior of the amoeba is indeed purely reactive or automatic or involves purely reactive capacities, this is an undesired result. This result is avoided if we restrict our characterization to *bona fide*, non-purely reactive cognitive behaviors.

3 Systematic behavior in the honey bee

In a completely different context, Carruthers (2004) has argued that honey bee behavior involves a bee’s mind in the quite demanding sense of a belief/desire psychology. Nothing of the sort will be defended here. However, bee behavior is a case of systematicity especially salient in this context. To put it mildly, bee behavior is a significant kind of sensorimotor, embedded, embodied and completely practical behavior. Bluntly put, bee behavior is not a chess game or some other sort of logical or abstract cognitive phenomenon apt for easy computational or representational interpretation. My focus on such a case involves a clear suggestion: if anti-representational dynamicism cannot explain all there is, from a purely behavioral standpoint, to bee systematic behavior, that must be because systematic behavior generally does not fit the anti-representational dynamicist mould. Here is the broad outline of the argument that follows:

- (1) Systematic behaviors are exercises of the same behavioral capacities (as per BS).
- (2) Anti-representational dynamicism offers (eventually very complex) specifications of behavior in dynamic fields, but no account of the underlying behavioral capacities of such specifications, and hence no account of systematicity.
- (3) Representational theories do offer an account of behavioral capacities underlying systematic behavior in terms of representational schemes.

Conclusion: Anti-representational dynamicism cannot adequately explain systematicity if it is not by dynamically implementing a representational account.

To establish (2), it will be useful to investigate the explanatory import of dynamical accounts regarding (merely) nomological behaviors (Subsect. 3.1) and behaviors that are furthermore systematic (Subsect. 3.2). This (abductive) argument will be completed in Subsect. 3.3 with a defense of (3) via an analysis of the representational alternative.

3.1 Merely nomologically related behavior: honey bees flying behavior

It is common wisdom that honey bees—*Apis mellifera*—spend their lives foraging nectar and pollen for the hive’s colony. They look for a source of nectar and then go back to the nest over and over again from birth until death. Detailed observation has shown that honey bee spatial behavior is richer than in many other insects, involving a variety of flight strategies such as straight flight trajectories, landmark exploitation and

short cutting. For present purposes, we can select the following extremely simplified conditional regarding bee cognitive behavior: if a honey bee is capable of flying from the hive to sources of nectar, then, as a matter of empirical and contingent truth, it is also capable of flying back from the sources of nectar to the hive. We have then the following pair of nomologically related flying behaviors:

FB1: flying from the hive to sources of nectar.

FB2: flying from sources of nectar to the hive.

Scholars have postulated quite complex kinds of representations in order to account for bee flying and navigating behavior. They include both egocentric view-based dead reckoning and allocentric and general map-like spatial representation (see e.g., [Menzel et al. 2006](#)). But our aim is to consider whether an account in terms of anti-representational dynamicism could be given that can dispense with representations altogether and, still, be a satisfactory account of bees flying behavior. On the assumption that FB1 and FB2 are not systematic, but merely nomologically and causally connected cognitive behaviors,⁷ there is no obvious reason why we should doubt that such an explanation is possible.

Let us illustrate this point with a toy example. Consider a dynamical function where FB1 and FB2 target behaviors are accounted for in terms of the bee's flying position as continuously changing over time with respect to the hive's position. Let us introduce, then, the position of the bee in the hive's range (x_b) and the position of the hive (x_0). Bee flying behavior can be seen as responding to distance with respect to the hive ($x_b - x_0$). When a maximum distance in a foraging flight has been reached (and a quantity goal of nectar achieved), bee dynamics make it return back home. That is, the effect of distance would be a continuously increasing negative force towards the hive's position. Once at hive, however, a certain 'foraging inertia' produces another foraging flight, which again, would be associated with a continuously increasing back-to-hive flying response. Thus, to a very rough approximation, bee flying behavior (FB1, FB2) can be accounted for in terms of the following second order linear differential equation:

$$m_b(d^2x_b/dt^2) = -k_b(x_b - x_0),$$

where m_b is interpreted as a bee foraging constant, d^2x_b/dt^2 stands for the continuously changing bee foraging impetus, and k_b is a homing constant. To be sure, this serves at most as a rough and even metaphorical approximation to a dynamical account of bee flying behavior, but it is enough to illustrate how easily we can begin to provide a dynamical account of such behavior. Indeed, so interpreted, bee flying behavior can be assimilated to the behavior of a simple harmonic oscillator—such as a spring or a pendulum—where m_b is the inertial mass, d^2x_b/dt^2 is the acceleration and k_b is the

⁷ The point of the assumption about the non-systematic character of the target behaviors is both illustrative and concessive, regarding the kinds of accounts that anti-representational dynamicism may be taken to provide. In addition, there seems to be *prima facie* nothing really puzzling about the idea that the capacities involved in the production of FB1 ($c_1^{fb1}, c_2^{fb1}, \dots, c_n^{fb1}$) and FB2 ($c_1^{fb2}, c_2^{fb2}, \dots, c_n^{fb2}$) are such that $\neg\exists c_i^{fb1} c_u^{fb2} (c_i^{fb1} = c_u^{fb2})$. It is only needed that the capacities involved in feeding-flights are distinct from the capacities involved in returning flights.

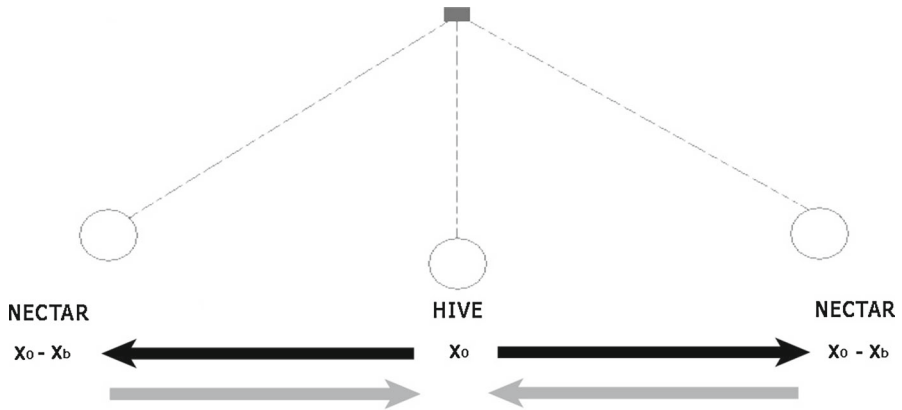


Fig. 1 Bee flying behavior can be initially modeled as an undamped pendulum or some other form of simple harmonic motion. Foraging displacements (*black arrows*) are associated to a restoring force towards equilibrium position x_0 (the hive) which produces homing displacements (*grey arrows*) at the maximum distance. Once bees arrive at hive, a certain foraging inertia make them go past equilibrium to complete the cycle

stiffness constant: a paradigmatic instance of behavior accountable in dynamical terms (see Fig. 1).⁸

As wished, the dynamics here are completely free from representational posits and offer an account of why FB1 behavior is connected in nature to FB2 behavior: the bee flies in an oscillatory fashion in such a way that it seeks for nectar when it flies from the hive until it reaches a maximum distance (FB1) and, granted the restoring force associated with the maximum distance, it goes back to the hive to unload its charge (FB2). The bee repeats this operation because of a certain ‘foraging inertia’ that makes it go past the hive’s equilibrium position and repeats the operation back and forth through the hive. In short, with sufficient aid of mathematical sophistication, there is no obvious reason for doubting that a satisfactory account could be given of the nomological relation between FB1 and FB2 in non-representational dynamics.

3.2 Systematic behavior: honey bees communicating behavior

As is known, honey bees do something more than just flying around foraging nectar for the hive’s colony. They also possess a unique and considerably rich system of communication to report to other bees the presence and position of nectar, pollen or water relative to the position of the hive and the position of the sun in the sky. They dance in a ‘drawing eight’ pattern in the hive so as to signal with straight movements crossing the center of the ‘eight figure’ the direction of the source of nectar (determined by the angle described with respect to the sun) and the distance (determined by the number of waggles performed) (Gould and Gould 1988). For present purposes, let

⁸ The suggested model is a toy example but not a mere speculation as a model in cognitive research. Oscillators have been the benchmark for dynamical accounts of limb movement and coordination at least since Haken et al. (1985).

us concentrate on the following simplified conditional regarding bees communicating behavior. If honey bees are capable of communicating, via an appropriately performed waggle dance, the position of nectar at 500 m north of the hive, then they are, as a matter of empirical and contingent truth, also capable of communicating, via an appropriately performed waggle dance, the position of nectar at 500 m south of the hive. The relevant pair of nomologically connected communicating behaviors is then CB1 and CB2:

CB1: perform waggle dance so as to signal nectar 500 m north of hive.

CB2: perform waggle dance so as to signal nectar 500 m south of hive.

This is, as required, a completely neutral description of bees communicating behavior. Now, can this communicating behavior be explained in non-representational dynamical terms?

Let us have a look at a possible strategy. Thelen and colleagues first suggested (Thelen and Smith 1994) and then carefully articulated (Thelen et al. 2001) a dynamical account of a classic case of infant perseverative reaching behavior, the so-called ‘A-not-B error’ originally reported by Piaget (1954). Infants between 7 and 12 months of age exhibit reaching-object behavior as if objects presented to them and then hid had lasting existence where they first appeared (location A), even if they observe how the object is switched to another hiding location (location B). In broadest outline, Thelen et al. model the infant’s reaching behavior as a “coupled dynamics of looking, planning, reaching and remembering within the particular context of the task” (Thelen et al. 2001, p. 5). Thus, reaching A and reaching B behaviors are the result of the dynamic activation of a motor planning field, which changes over time and is itself the result of the integration of the activation of visual and memory parameters (namely, the task input concerning the target locations in the experimental setting, the specific input relative to the effect of a cued location and the memory input regarding previous reaching decisions) together with a developmentally constrained cooperativity parameter. When the dynamics yield activation in the motor field above threshold, the reaching movement thereby specified gets generated by the infant (see Thelen et al. 2001 for details).

Now, Thelen et al.’s (2001) non-representational model works very well in dynamically accounting for the different aspects that affect child development (including notably child’s memory) and that determine the ‘A-not-B’ error. This and similar models that account for the dynamics of motor planning (e.g., Erlhagen and Schöner 2002; Schöner et al. 1997) may therefore constitute a substantial anti-representational basis for the explanation of honeybees communicating behavior.

Consider the following sketch of a model by way of illustration. Analogously to the way in which Thelen et al. model the infant’s (reaching A, and reaching B) behavior, communicating bees behavior can be modeled in terms of a motor planning field—in this case a communicating planning field—which shows the activation of a continuous communicating parameter, x , over time. Different levels of activation of the communicating parameter, $u(x)$, would correspond to different specifications of communicating behavior. If the level of activation corresponding to a specification of communicating behavior is above threshold, then the bee would actually perform the corresponding

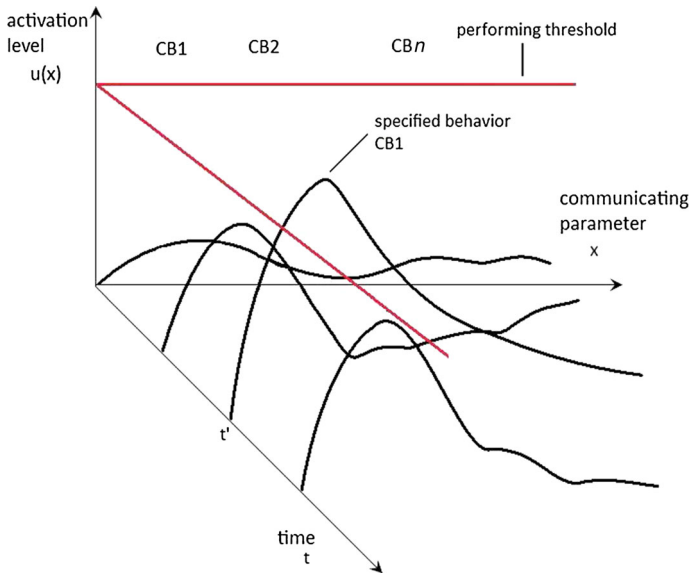


Fig. 2 The dynamic field for bee communicating behavior could be seen as a parameter, x , whose activation, $u(x)$, specifies a range of different behaviors (CB1, CB2, ..., CB n) and changes continuously over time. At t' , the level of activation for CB1 is above threshold and the bee would actually perform the corresponding waggle dance. These dynamics can be understood as a function of a variety of input parameters, such as, say, the one corresponding to the quality of nectar detected: if the quality of nectar is higher at a given location, the activation of communicating parameter for that location would also be higher

waggle dance. The target communicating behaviors are CB1 and CB2 above, but the model could be extended to any communicating behavior CB n . The dynamic field may thus be a function of the communicating parameter and time, $u(x, t)$. The system, in addition, can be taken to be a function of a great deal of different input parameters (such as the ones corresponding to sources of nectar detected, quantity/quality of nectar, distance, moment of the day, etc.). These input parameters would be the counterpart of Thelen et al.'s task input, specific input and memory input. Now, if the input parameters and the evolution of the system in time are appropriate, a given level of activation of the communicating parameter would be above the performing threshold and the bee would actually perform the corresponding waggle dance (see Fig. 2).

Arguably, if fully laid out, one such account would be promising for the explanation of the nomological connection of communicating behaviors CB1 and CB2. However, CB1 and CB2 are not only nomologically related from a strictly behavioral point of view. In addition, bee communicating behavior is, to all intents and purposes, a clear instance of *systematic* behavior.

To repeat, systematic behaviors are (causally and nomologically) connected behaviors which are, as a matter of empirically testable fact, explained as involving exercises of *the same* behavioral capacities (see BS in Sect. 2). CB1 and CB2 are a case in point. CB1 clearly differs from CB2: it is not the same, from a behavioral point of view to signal nectar 500 m north of the hive than signaling nectar 500 m south of the hive. However, and this is the crucial point, no one would reasonably deny that CB1 and

CB2 are exercises of the same behavioral capacities, namely, the capacities involved in signaling nectar found in foraging tasks somewhere in the hive's range.⁹

The non-representational dynamicist account seems meager as an account of systematicity in this behavioral sense. Even if we could formulate a mathematical model with the assumptions and language of dynamic systems in this case, and even if all the variables and complexity of real bee behavior were taken into account in this mathematical model, the nontrivial question would still arise as of why bee communicating behavior CB1 is at all *systematically* linked with bee communicating behavior CB2, that is, why CB1 and CB2 involve exercises of the same behavioral capacities. In the appropriate system of differential equations, CB1 would be just a part (e.g., a particular point, a set of points, a basin of attraction, an arrangement of basins of attraction) of the continuous dynamic field, whereas CB2 would be just another part of the dynamic field. It would remain mysterious why CB1 and CB2 are exercises of the same bee communicating capacities. The systematicity question is thus how one can tell apart, just by analyzing the differential equations or parameters of a given dynamics, which behaviors are nomologically related and which are *furthermore* systematic. No easy answer to this question seems to be available because any behavior (independently of whether it is systematic or not) would be, from a dynamical perspective, just a part of the dynamic field. If this is correct, anti-representational dynamicism is, as a matter of principle, unable to properly identify, let alone explain, systematicity phenomena.

Anti-representationalists may wish to reply to the foregoing considerations in a number of ways. First, they may straightforwardly object, that, in dynamical models, CB1 and CB2 can be considered as exercises of the same behavioral capacity insofar as they are specified by the same set of nonlinear differential equations.

This suggestion, however, clearly underestimates the problem under consideration, namely, the problem of properly distinguishing nomologically and causally connected behaviors from systematic ones. Let us assume that we have a correct dynamical account for actual CB1 and CB2 in terms of differential equations. Now, let us imagine a counterfactual situation in which CB1 and CB2 are not exercises of the same behavioral capacity (say, that CB1 is the exercise of a signaling-north communicating capacity and that CB2 is the exercise of a signaling-south communicating capacity). The question is: would the difference between actual (systematic) and counterfactual (non-systematic) scenarios involve a difference in the corresponding dynamic systems of differential equations? Well, it is hard to see how given that the target behaviors (as opposed to the underlying capacities) would still be the same and, hence, the dynamic activation values specifying the (occurrence of the) behaviors must also be the same in the actual and counterfactual cases. And what else can a set of differential equations provide apart from a dynamic activation field which specifies (the occurrence of) the target behaviors? Since the same specifications of bee behavior may correspond

⁹ Unlike perhaps the case of flying behavior (for which one may conceive FB1 and FB2 as arising out of different capacities), it seems very hard to deny that one the same kind of communicating behavioral capacities are involved in FB1 and FB2 and therefore that these are not punctate kinds of behavior. These communicating capacities might be of a very complex cognitive sort such as the capacity to identify the position of the sun in the sky and the angle with respect to the target source or the capacity to respond selectively to distances, quantity and quality of the target.

to different underlying capacities, these specifications are clearly insensitive to the capacities underlying such behavior. Therefore, “same set of differential equations” is not plausibly interpreted as “same underlying behavioral capacities”.

The anti-representational dynamicist may also try to reply that discrimination between behaviors that are exercises of the same capacities and behaviors that are not can be made, in dynamical approaches, via empirical data somehow captured or predicted in the models. Relevant developmental and counterfactual data about a given pair of nomologically connected behaviors would provide evidence about whether the pair involves the same or distinct behavioral capacities. What the dynamicist does is to provide a dynamical account of those empirically confirmed systematicities.

This reply would just be a plain acknowledgement of the main thesis here defended, namely, that anti-representational dynamicism cannot, in and of itself, discriminate and let alone provide a satisfactory account of systematic behaviors. The dynamicist needs to appeal to other sorts of considerations, beyond dynamical approaches as such, in order to offer principled criteria for distinguishing merely nomologically connected behaviors and behaviors that are, in addition, systematic. Once the discrimination is made via alternative means, the dynamicist has nothing on offer that would explain or ground such discrimination.

The problem is not compellingly addressed by appealing to future improvements in dynamical accounts and concepts. Forthcoming dynamical analyses, one would be tempted to argue, might identify in state space or dynamic field which points, sets of points, attractors or bifurcations would correspond to a given behavioral capacity so as to identify in turn combinatorial structures involved in properly systematic behaviors. This strategy is, to my knowledge, quite speculative, but it probably illustrates the polemical situation once we have arrived at this point. A new systematicity loop would seem to begin. Granted that we have a dynamic system with a definite dynamic field, it is hard to see what could ground the required identification of systematic behaviors—viz. the required selection of points or sets of points or basins of attraction in the dynamic field—apart from sheer stipulation. The dynamicist contender would try to provide improved and extremely complex mathematical criteria to address this worry. Still, the decision as to what parts of the dynamic field correspond to systematic behaviors would arguably be based upon considerations other than purely dynamic considerations. The suggestion is that what would do the job in satisfactory dynamicist explanations of systematicity is not the dynamics of the system *per se*, but our best available theory about the target behavioral capacities. Which one is such a theory? No matter which one you choose exactly, it is to all appearances a representational theory of cognition.

Another possible way of approaching the problem at stake is in terms of the distinction between covering-law and mechanistic explanations. It is often claimed that dynamical models distinctively provide covering-law explanations, that is to say, explanations that proceed via subsumption of a target phenomenon under a natural law from which it can be predicted when relevant conditions are stated (Bechtel 1998; Walmsley 2008; Chemero 2009). The problem with anti-representational dynamicism regarding systematicity can be seen as a consequence of the fact that no causal processes or mechanisms would seem to be specified in the account. Covering-law explanations for any pair of behaviors A and B are clearly silent regarding underly-

ing behavioral capacities and, therefore, seem compatible with these behaviors being either systematic or unsystematic. This is so over and above the predictive power of the alleged dynamical model. But the point against dynamical accounts of systematicity is not merely, as it might be, that dynamical explanations are not mechanistic and are, for this reason, at fault (cf. [Eliasmith 1996](#)). The point is more exactly that, insofar as they are not mechanistic, no criterion appears to be in view for the discrimination between merely nomologically connected and genuinely systematic behaviors.

True, some authors would be ready to argue that dynamical explanations are also mechanistic or causal (e.g., [Gervais and Weber 2011](#); [Zednik 2011](#)). But conceding this point does not lead the anti-representational dynamicist too far. For the problem would then simply turn to whether the dynamicist could really invoke non-representational mechanisms and causes in this context. Indeed, it is hard to see that the alleged mechanisms or causal agents, if any, figuring in would-be satisfactory dynamical explanations of systematicity are other than, precisely, the corresponding mechanisms and causal agents postulated in schemes of representational explanation. This point is in agreement, for instance, with Carlos Zednik's observation that advocates of mechanistic explanations in dynamical approaches "may be steering toward reconciliation with proponents of *representationalism*" ([Zednik 2011](#), p. 261, emphasis his). Thus, the dilemma between covering-law and mechanistic-causal explanation in dynamicism as regards systematicity is plausibly seen as just a version of the dilemma between either failing to explain or else implementing representational explanations of systematicity phenomena. Let us examine, therefore, the latter kind of explanation.

3.3 The representational solution

Compare the above scenario with the following: let us postulate a representational system or scheme for the honey bee, say, a system that involves representation of nectar position relative to the hive. For present purposes, we can state the representational scheme as a function R that takes direction, d , distance with respect to the hive, dis , and target, t , as inputs and delivers the corresponding waggle dance as output. Let us assume that the function $R(d, dis, t)$ determines a possible range of contents (of the form 'Target t , at direction d and distance dis ') expressible in a waggle dance. Let us assume further that this function is subject to malfunctioning and misrepresentation. Since the psychological reality of R would ultimately require its physical implementation in bee organisms, it follows that R would involve the sorts of states, sets of states and processes meeting our previous characterization of representation (see Sect. 2 above).

Although described in roughest outline, this representational scheme constitutes what we can call a nectar-location representational system for the honey bee. With some such representational system in hand, we begin to see what would explain the fact that CB1 and CB2 are exercises of the same behavioral capacity: if a system such as R is used by the bee then we can consider it as being the basis of the different manifestations of bee communicating behavior. Thus, if bees exploit R in order to signal ' $t = \text{NECTAR}$, at $d_1 = \text{NORTH}$ and $dis = 500 \text{ m}$ ', then it is only to be expected that they exploit the same representational scheme in order to signal ' $t = \text{NECTAR}$ at

$d_2 = \text{SOUTH}$ and $dis = 500 \text{ m}'$. The exercise of one and the same behavioral capacity is seen, under this view, as the use of one and the same representational scheme R.

Filling in the details of the representational story involves all sorts of complex and substantial matters. The correct statement of a representational system must include fundamental claims about the exact content and structure of representations, biological and developmental plausibility, actual availability for bee organisms and much else. Note however that the postulation of representational schemes of this sort is not an ad hoc explanation. A representational system R would be empirically grounded and made precise in the light of the bee's testable signaling practices.¹⁰ Its postulation need not be supported by prior, uncritically assumed behavioral capacities free from empirical commitments.

The explanatory import of representational systems seems to generalize wildly. It is plausible that such systems provide an account of behavioral systematicity across many different cognitive domains and across many different kinds of organism. The relevant question in this context is of a familiar form: what, if not representational theories, would account for systematicity? The suggestion is that representational theories of cognition (whatever exactly their specific nature) are actually our best estimates of systematic behaviors. To all appearances then, the only way anti-representational dynamicism can account for systematicity is by appealing (however tacitly) to some such theory. There is of course nothing theoretically objectionable about this. There is a price to pay nonetheless. The price is that dynamicism is no longer anti-representational and becomes an implementation of representational theories. This is known territory, it is just the old systematicity challenge reformulated against anti-representational dynamicism.

4 Conclusion: a renewed systematicity challenge

This is the argument so far. Once equipped with neutral notions of representation and behavioral systematicity, we can analyze a substantial case of systematicity in nature, namely, honey bee communicating behavior. Unlike perhaps honey bee flying behavior, honey bee communicating behavior is blatantly behavior that arises from the same behavioral capacities. Dynamic fields cannot account for this fact because dynamic fields would include all sorts of behaviors and would offer no way of distinguishing, from among those behaviors, the ones that are exercises of the same behavioral capacities and the ones that are not. Things change dramatically, however, if we appeal to an appropriate scheme of representations. The moral is simple, as it is familiar. Anti-representational dynamicism has one of two options: either dismiss systematicity in any neutrally described terms (that is, dismiss systematicity *tout court*), or else be conceived as an implementation of representational systems. This would seem fatal

¹⁰ Indeed the exact nature of honey bee representational system for communication and navigation is an empirically open question at least since [Von Frisch \(1967\)](#). Recent data using harmonic radar tracking strongly suggest that bee waggle dance spatial representation may not only involve estimates of direction and distance but complex sorts of navigational and motivational information stored in a map-like memory (see [Menzel et al. 2006](#); [Menzel 2011](#)).

for anti-representational dynamicism, considered as an alternative and general theory of high cognition. We should try to do better, if we can.

The foregoing argument would perhaps not persuade anti-representationalists of the failure of their anti-representationalism. Even so, it is clear that the promise of anti-representational dynamicism for a radically new kind of explanation in the systematicity debate is simply ungrounded. In this context, anti-representational dynamicism might at best lead us to a slightly modified version of Fodor and Pylyshyn's well-known systematicity challenge.

The here defended line of reasoning also provides support for a more general and far-reaching conclusion regarding the benefits of collaborative research which may be succinctly expressed in terms of levels of explanation. Marr (1982) famously delineated a three-level distinction as regards explanation in cognitive research (i.e., the distinction of the computational, the algorithmic and the implementation levels). Even if Marr's own developments may be polemical in several ways, the existence of different levels of explanation is perhaps the only unchallenged idea of old computational research (see Verdejo and Quesada 2011 for discussion). However, the dynamicist's anti-representational turn actually involves an unjustified underestimation of the importance of the different levels, and especially in this context, of the algorithmic, representational level. The claim that there are no representations or that we should dispense with representations is, nonetheless, a substantial and extreme claim concerning Marr's algorithmic level 2: it amounts, in fact, to a radical dismissal of research at that level as a genuine source for scientific progress. By contrast, it would seem that we can accommodate everything that dynamicism has on offer within an overall framework for representational computationalism, writ large, in which explanatory (Marrian) levels are better demarcated and integrated. This would probably require the task of investigating, for each particular dynamical model on offer, at which level or levels it operates.¹¹

This does not mean that dynamical approaches would be secondary in this scenario. Here, as always, it is useful to keep in mind Marr's own dictum: levels of explanation must be differentiated but *all* these 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). Since information-processing devices can also be taken to include *dynamical* information-processing devices, nothing at first sight seems to tell against the idea of a dynamical account of systematicity that is part of an integrated account at various levels of explanation, including, notably, the representational or algorithmic level.¹² This strategy for a unified cognitive science would put us far from the pose-challenge/respond-to-challenge dialectics and is, plausibly, the

¹¹ The lack of consensus amongst scholars on the proper location of dynamical explanations at a particular level suggests that this task is largely overdue. Horgan and Tienson (1994), for instance, ironically locate dynamical models at the algorithmic level as proposing a change from the discrete mathematics of classicism to the continuous mathematics of dynamical systems. The more recent dynamicist focus on embodiment (Calvo Garzón 2008; Chemero 2009) would suggest that they are located mostly at Marr's implementation level 3.

¹² This proposal is of course not new. Many authors have shown sympathy to a collaborative view of the different schools in the cognitive sciences (e.g. Clark 1997; Bechtel 1998; Grush 2003; Beer 2014; Zednik 2011; Fresco 2012; Dawson 2013).

most promising framework for reaching a powerful and rich account of systematicity phenomena.

Acknowledgments Daniel Quesada, Jesús Vega, Paco Calvo, Toni Gomila and an anonymous referee for this journal contributed with insightful comments to improve earlier versions of this material. I would like to thank all of them and the audience of the workshop “Systematicity and the post-connectionist era” held in San José (Spain) in May 2011, where the germ of this paper was presented. This work has received the support of the Secretary for Universities and Research of the Ministry of Economy and Knowledge of the Government of Catalonia, the Spanish Ministry of Economy and Competitiveness (research projects FFI2013-41415-P and FFI2012-35153) and the consolidated research group GRECC (SGR2014-406).

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Aizawa, K. (2003). *The systematicity arguments*. Dordrecht: Kluwer Academic Publishers.
- Aydede, M. (1997). Language of thought. *The connectionist contribution. Minds and Machines*, 7, 57–101.
- Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist’s challenge in cognitive science. *Cognitive Science*, 22, 295–318.
- Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 173–215.
- Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11, 209–243.
- Beer, R. (2014). Dynamical systems and embedded cognition. In K. Frankish & W. Ramsey (Eds.), *The Cambridge handbook of artificial intelligence* (pp. 128–148). Cambridge: Cambridge University Press.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Burge, T. (2010). *The origins of objectivity*. Oxford: Oxford University Press.
- 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, CA: AAAI Press.
- Calvo Garzón, F. (2008). Towards a general theory of anti-representationalism. *British Journal for the Philosophy of Science*, 59, 259–292.
- Carruthers, P. (2004). On being simple minded. *American Philosophical Quarterly*, 41, 205–220.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Chiel, H. J., & Beer, R. (1997). The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 20, 553–557.
- Clark, A. (1997). *Being there: Putting brain, body and world together again*. Cambridge, MA: MIT Press.
- Clark, A., & Toribio, P. (1994). Doing without representing? *Synthese*, 101, 401–431.
- Cummins, R. (1996). Systematicity. *Journal of Philosophy*, 93, 591–614.
- 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.
- Davies, M. (1991). Concepts, connectionism, and the language of thought. In W. Ramsey, S. Stich, & D. Rumelhart (Eds.), *Philosophy and connectionist theory* (pp. 229–257). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Dawson, M. R. W. (2013). *Mind, body, world: Foundations of cognitive science*. Edmonton: Athabasca University Press.
- Dorffner, G. (1999). The connectionist route to embodiment and dynamicism. In A. Riegler, M. Peschl, & A. von Stein (Eds.), *Understanding representation in the cognitive sciences* (pp. 23–32). Dordrecht: Kluwer Academic Publishers.
- Eliasmith, C. (1996). The third contender: A critical examination of the dynamicist theory of cognition. *Philosophical Psychology*, 9, 441–463.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, 109, 545–572.

- Fodor, J. A. (1987). *Psychosemantics. The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1997). Connectionism and the problem of systematicity (continued): Why Smolensky's solution still doesn't work. *Cognition*, *62*, 109–119.
- Fodor, J. A., & McLaughlin, B. P. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, *35*, 183–204.
- Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71.
- Fresco, N. (2012). The explanatory role of computation in cognitive science. *Minds and Machines*, *22*, 353–380.
- Fresco, N. (2013). Information processing as an account of concrete digital computation. *Philosophy and Technology*, *26*, 31–60.
- García-Carpintero, M. (1995). The philosophical import of connectionism: a critical notice of Andy Clark's Associative Engines. *Mind and Language*, *10*, 370–401.
- Gervais, R., & Weber, E. (2011). The covering law model applied to dynamical cognitive science: A comment on Joel Walmsley. *Minds & Machines*, *21*, 33–39.
- Gomila, A., Travieso, D., & Lobo, L. (2012). Wherein is human cognition systematic? *Minds & Machines*, *22*, 101–115.
- Gould, J., & Gould, C. (1988). *The honey bee*. New York: W.H. Freeman.
- Grush, R. (2003). In defense of some 'Cartesian' assumptions concerning the brain and its operation. *Biology and Philosophy*, *18*, 53–93.
- Hadley, R. (2004). On the proper treatment of semantic systematicity. *Minds and Machines*, *14*, 145–172.
- Haken, H., Kelso, J. A. S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, *51*, 347–356.
- Horgan, T., & Tienson, J. (1994). A nonclassical framework for cognitive science. *Synthese*, *101*, 305–345.
- Haugeland, J. (1991). Representational genera. In W. Ramsey, S. Stich, & D. Rumelhart (Eds.), *Philosophy and connectionist theory* (pp. 61–89). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Johnson, K. (2004). On the systematicity of language and thought. *Journal of Philosophy*, *3*, 111–139.
- Kelso, J. A. S. (1995). *Dynamic patterns*. Cambridge, MA: MIT Press.
- Marcus, G. F. (2001). *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. Cambridge, MA: MIT Press.
- Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.
- Matthews, R. J. (1994). Three-concept monte: explanation, implementation and systematicity. *Synthese*, *101*, 347–363.
- McLaughlin, B. P. (1993). The connectionism/classicism battle to win souls. *Philosophical Studies*, *71*, 163–190.
- McLaughlin, B. P. (2009). Systematicity redux. *Synthese*, *170*, 251–274.
- Menzel, R., De Marco, R. J., & Greggers, U. (2006). Spatial memory, navigation and dance behavior in *Apis mellifera*. *Journal of Comparative Physiology A*, *192*, 889–903.
- Menzel, R. (2011). Navigation and communication in honeybees. In R. Menzel & J. Fischer (Eds.), *Animal thinking: Contemporary issues in comparative cognition* (pp. 9–22). Cambridge, MA: MIT Press.
- Petitot, J. (1995). Morphodynamics and attractor syntax: Constituency in visual perception and cognitive grammar. In R. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 227–283). Cambridge, MA: MIT Press.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.
- Piccinini, G., & Scarantino, A. (2011). Information processing, computation, and cognition. *Journal of Biological Physics*, *37*, 1–38.
- Port, R., & van Gelder, T. (Eds.) (1995). *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.
- Schöner, G., Kopecz, K., & Erlhagen, W. (1997). The dynamic neural field theory of motor programming: Arm and eye movements. In P. G. Morasso & V. Sanguineti (Eds.), *Self-organization, computational maps, and motor control* (pp. 271–310). Amsterdam: Elsevier-North Holland.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures 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 (Tech. Rep. 92–08). Boulder, CO: University of Colorado, Institute of Cognitive Science.

- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, *24*, 1–34.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Thompson, E. (2007). *Mind in life: Biology, phenomenology and the sciences of mind*. Cambridge, MA: Harvard University Press.
- 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. *Behavioral and Brain Sciences*, *21*, 615–665.
- Van Gelder, T., & Port, R. (1995). It's about time: An overview of the dynamical approach to cognition. In R. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 1–43). Cambridge, MA: MIT Press.
- Van Rooij, I., Bongers, R. M., & Haselager, W. F. G. (2002). A non-representational approach to imagined action. *Cognitive Science*, *26*, 345–375.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Verdejo, V. M. (2012). Meeting the systematicity 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.
- Von Frisch, K. (1967). *The dance language and orientation of bees*. Cambridge, MA: Harvard University Press.
- Walmsley, J. (2008). Explanation in dynamical cognitive science. *Minds & Machines*, *18*, 331–348.
- Zednik, C. (2011). The nature of dynamical explanation. *Philosophy of Science*, *78*, 238–263.