Systematicity and Conceptual Pluralism

Fernando Martínez-Manrique
University of Granada
fmmanriq@ugr.es

1. Introduction

The systematicity argument (henceforth SA), offered by Fodor and Pylyshyn (1988) against the plausibility of connectionism as an alternative theory of cognition, can be characterized in terms of three different claims –an *empirical* claim, an *explanatory* claim, and a *definitional* claim– from which a dilemma for connectionism arises. Let me present the four elements in a sketchy form before saying a little more about each of them:

SA

(i) Empirical claim: systematicity is a pervasive property of cognition

(ii) Explanatory claim: the only plausible explanation for systematicity is to posit a compositional system of representations

(iii) Definitional claim: compositionality is a defining property of classical representational systems

(iv) Dilemma: if connectionism is not compositional then it cannot account for systematicity and so it does not provide a full account of cognition (from i & ii); if connectionism can account for systematicity then it is actually implementing a classical system (from ii & iii)

SA has been haunting connectionist approaches ever since, and main responses to it can be classified depending on whether they focus on (i), (ii) or (iii).[1] Much can be said about the relative success of each such response, yet one important common point is that Fodor and Pylyshyn's argument would work as a global refutation of connectionist explanations only if systematicity were regarded as a property of cognition in general. However, SA *per se* does not include the latter commitment. Truly, Fodor and Pylyshyn stated that "there's every reason to believe that systematicity is a thoroughly pervasive

[1] See McLaughlin (1993) for a different and earlier–and consequently less complete– classification of connectionist replies to argument.

feature of human and infrahuman mentation" (1988, p. 37). Yet, unless further arguments for the universality of systematicity are provided, the statement can be read simply as claiming that it is an important phenomenon that needs explanation, and practically everybody agrees that much. This leaves open the issue whether there are cognitive domains or processes that are not systematic in the way intended by SA, and one may conjecture, as many connectionist authors do, that some non-classicist model could just account for them.

Still, the fact that some cognitive processes were not systematic in the way intended by SA would not be enough for non-classical models to carry the day. To this end they not only must show that their models are capable to deal with such cognitive processes but that they are in a better position than their classical competitors to do so. In other words, what they would need is something like a SA for themselves –let me call it the Non-classical Systematicity Argument– that would run roughly as follows:

NSA

(i') Empirical claim: X is a pervasive property of cognition

(ii') Explanatory claim: the only plausible explanation for X is property Y

(iii') Definitional claim: Y is a defining property of such and such non-classical systems

(iv') Dilemma: if classicism cannot account for property Y then it does not provide a full account of cognition (from i & ii); if classicism can account for Y then it is actually implementing a non-classical system (from ii & iii)

My aim in this paper is to provide a path to construct such an argument. I want to stress that my main focus is not NSA itself, but the elements that may allow us to get at NSA. First, I offer an overlook of the connectionist answers to SA, classified as focusing on (i), (ii) or (iii), followed by a quick assessment of the debate. This assessment is negative for the connectionist side, in the sense that it never managed to substantiate an alternative explanation of the phenomenon pointed out by Fodor and Pylyshyn. Of course, I lack the space to go into details, so connectionist fans of this or that particular reply may think that I am being unfair to it. Yet, apart from the general considerations that I will provide to back my negative assessment, it seems to me that it is reinforced by the sheer fact that there is no agreement with respect to which reply to SA works best. My aim in this section, thus, is just to motivate the view that classical models still

stand as the *most plausible* explanation for classical systematicity. Second, I will deal with the question whether systematicity is actually a general property of cognition. I will argue that the best chances to support such a view come from regarding Evans's well-known Generality Constraint as a constraint on the architecture of conceptual creatures –a constraint that only concepts that exhibit classical systematicity seem to satisfy. Then I will show a different way of understanding the constraint, in terms of *attributions of belief*, that opens the door to architectures with concepts that do not exhibit classical systematicity. Third I will present and motivate the thesis of conceptual pluralism, arguing that concepts split into subkinds that share two fundamental properties: they are central and they grant belief-attributions. I will draw on Camp's (2009) analysis to make the case that there are actually two kinds of concepts. Finally, I will rely on Dual Systems Theory and on Penn *et al*'s (2008) recent review of differences between animal and human cognition to motivate a plausible scenario of two different processing systems that work on different kinds of concepts with properties that give raise to two different sorts of systematicity. I will sketch then a way in which NSA could be filled but my goal is not to endorse a particular non-classical approach as a filler for the argument. To repeat, my aim is not to try to reply to SA for the umpteenth time, but simply to show that while non-classical approaches lack the resources to meet SA, the elements for an alternative NSA argument can be provided.

2. The elements of the systematicity argument

2.1. The empirical claim

Fodor and Pylyshyn plainly took their claim about the systematicity of cognition as an empirical one. Systematicity can be characterized as the property of having the ability to think systematically related thoughts. It is a matter of fact that creatures that have the ability to think aRb have also the ability to think bRa. Apparently, some critics have failed to see this point. For instance, early in the debate Clark (1989) argued that the relation between the abilities to think aRb and bRa is not an empirical but a conceptual fact. It is not that we cannot find organisms with punctate thoughts but that the fact that they are punctate is enough to deny that they are thoughts. It is part of our concept of what it takes to have thoughts that they be systematically related. McLaughlin (1993) replied that if systematicity is a conceptual property then the challenge posed by Fodor

and Pylyshyn would be strengthened given that we would get an *a priori* constraint for the constitutive basis of cognition. More recently, Chemero (2009) also complained 1) that SA is a conceptual argument (or, as he calls it, a Hegelian one) against an empirical claim, and 2) that Fodor and Pylyshyn provide almost no empirical evidence to support premise (i). Actually, Chemero is wrong about both complaints. First, having poor empirical evidence for one's argument does not make it a conceptual argument –it makes it a poor argument. Second, their empirical evidence is not so poor as Chemero intends us to believe. It is based on a parallelism with language understanding, the most famous example being that just as you do not find anyone who can understand 'John loves Mary' and cannot understand 'Mary loves John', you do not find anyone who can think that John loves Mary and cannot think that Mary loves John. Fodor and Pylyshyn think that examples like this come on the cheap so it is no wonder that they do not feel the need to provide plenty of them. In other words, they assume that the extent of their empirical evidence is as large as the extent of language itself.

Other critics accepted the claim as an empirical one but they rejected it as false. Some of them focused on the idea of systematicity as "a thoroughly pervasive feature of human and *infrahuman* mentation" (Fodor and Pylyshyn 1988: 37, my emphasis), and alleged that non-human animals do not exhibit the sort of systematicity exemplified by the *J loves M* case (Sterelny, 1990: 182-83; Dennett, 1991; Kaye, 1995). More recently, Gomila *et al* (2012) reject the claim that systematicity is a general property of cognition. In their view, it is only related to those cognitive abilities that are possible by the acquisition of language, and it is derived precisely from the systematicity of linguistic structure. As I will argue later, I concur with Gomila *et al* that there are grounds to deny that systematicity is a general property of cognition. Yet, this does not entail a rejection of the classical explanation. On the one hand, even if SA only applied to human cognition, or to language-related cognition, it would still be a *significant* property. On the other hand, the best explanation of this property is still classical. For instance, even if the explanation of systematicity lied in the properties of language, as Gomila *et al* (2012) contend, the way of fleshing out such an explanation is still by regarding language as a classical system itself –i.e., systematicity is still explained in terms of language's alleged compositional structures and processes that are sensitive to those structures.[2] So inasmuch as connectionism could not avail itself of this explanation, it

[2] On the other hand, there are reasons to doubt that language is fully compositional in the required

would be in trouble to account for cognition, and this is how many authors viewed the issue. In other words, connectionist attempts at rejecting systematicity as a *general* property of cognition do not entail, even if they were successful, rejecting classicism as the architecture of at least *part* of cognition.

2.2. The explanatory claim

The second claim in SA is an instance of a "best explanation" argument. The idea is that a straightforward and plausible way of explaining systematic relations of the *J loves M* type is to posit a compositional semantics, i.e., a system of context-free, recombinable semantic pieces in which the semantics of the composed whole depends in a systematic way on the semantic values of the pieces. Many critics focused on this explanatory relation. Some of them complained that the *explanandum* –i.e., systematicity– had been poorly characterized and consequently devoted their efforts to reformulate it in a way that could be explained by non-classical systems. For instance, Clark (1989: 149) insisted that what has to be explained "is not the systematicity of thoughts but the systematicity of the behavior, which grants thought ascription"; Goschke and Koppelberg (1991) or Bechtel (1994) regarded systematicity not as a property of thoughts but of an external symbolic language; Niklasson and van Gelder (1994), and Cummins (1996; Cummins *et al*, 2001) examined forms of systematicity different from the language-based cases; Johnson (2004), on the other hand, addresses systematicity from the linguistic perspective and provides a definition of systematicity so as to contend that language is not systematic after all.[3]

Other critics focused instead on the *explanans* –i.e., compositionality– and tried to offer distinctions that helped connectionism to meet the explanatory challenge. The most notable of them was due to van Gelder (1990), who made a distinction between concatenative and functional compositionality.[4] In concatenative composition,

---

sense. (See Vicente and Martínez-Manrique (2005) for a rejection of the claim that semantics can provide fully determined compositional thoughts and its consequences on the views that regard language as a cognitive vehicle). Language may be simply a combinatorial system, and thus the picture presented by Gomila *et al* would be of a classical compositional system getting installed thanks to the combinatorial properties of language. But notice that SA is neutral about how systematicity is acquired, its claim being about how it is explained, and its explanation in such an acquisition model is still a classical one.

[3] See McLaughlin (2009) for an extensive analysis of Cummins's and Johnson's claims, in which he contends that they miss the point about what has to be accounted for –which, in his view, are the lawful psychological patterns revealed in systematic relations between thoughts.

[4] Van Gelder and Port (1994) extended the analysis by proposing six different parameters –

tokenings of constituents of an expression (and the sequential relation between them) are literally preserved in the expression itself. In functional compositionality there are general, reliable processes to decompose an expression in their constituents, and to produce it again from them, but it is not necessary that the expressions contain literally their constituents. Van Gelder argued that even if connectionist networks only exhibit the latter kind, this is enough to account for systematicity.

The trouble with reformulations of compositionality is that they failed to provide a global alternative explanation of systematicity, e.g., one that relied on functional compositionality as a fundamental property of non-classical cognition, in the same sense as compositionality plays the central role in classical conceptions. Even though many people, myself included, acknowledge the relevance of the distinction, I know of no overarching connectionism conception in which it plays that pivotal role. So regarding van Gelder's prediction (1991) that functional compositionality would be one of the central aspects for connectionism to become a truly alternative paradigm, one must say that it is a prediction yet to be fulfilled. Indeed, we will see later that recent approaches that dwell on van Gelder's distinction use it to characterize the features of two *different* systems, so functional compositionality could be seen as playing an explanatory role only in part of cognition.

The trouble with reformulations of systematicity, on the other hand, is that they easily change the subject matter. The facts that behavior or language are also systematic, or that there are non-linguistic instances of systematicity, do not deny the systematicity of thought that is the basis for SA. It is good to say that there are other things to explain apart from the systematicity in SA, but unless one wants to say that the latter property is unreal, SA itself remains untouched. Indeed, the line that I am going to follow in this paper is an instance of the "change subject matter" strategy but not to defeat SA, only to create a different argument that leaves room for non-classical systems as an account of *part* of cognition.

1.3. The definitional claim

Having a combinatorial syntax and semantics for mental representations, and having processes that are sensitive to the structure of the representations so constructed are

---

properties of primitive tokens, and properties of modes of combination– in terms of which to distinguish varieties of compositionality. However, concatenative vs. functional still seems to be the crucial one.

defining properties of classical models, according to Fodor and Pylyshyn (1988: 13). There are two sides to this claim: One is that given the principles of classical computationalism explaining systematicity comes as a necessary consequence, i.e., it is not possible to have a classical system that is not systematic in the demanded sense. The challenge can be thus reformulated as a demand that the opponent should provide models based on different principles, in which systematicity appears as a consequence of those principles (Fodor and McLaughlin, 1990). In terms of Aizawa (1997; 2003), the challenge is not to exhibit systematicity –i.e., to show that it is possible to have a systematic connectionist model– but to explain it –i.e., to show that systematicity follows necessarily from the principles of the theory. It is the latter challenge that connectionists fail to meet. My view is that even if systematicity is not strictly entailed by the principles of classical models, as Aizawa contends,[5] it is still the case that these models have a much more robust explanation of the phenomenon than their connectionist counterparts.

The second side of the definitional claim is that if compositionality plus structure-sensitive processes are defining properties of classical systems, then any system that resorts to them will count *ipso facto* as a classical one. The early debate between Smolensky (1988, 1991a, 1991b) and Fodor *et al* (Fodor and McLaughlin, 1990; Fodor, 1997) can be understood in those terms, and the gist of the dilemma posed by Fodor *et al* comes to this: if Smolensky is capable of showing that his models do have a constituent structure, then they are implementations of a classical system –given that they are based on the same relevant explanatory principles; if they do not have a constituent structure, then they cannot account for systematicity. The countless subsequent connectionist attempts of proving that this or that network has systematic capabilities –I will save space referring to Hadley (1994) for a review and criticism of early attempts, and to Frank *et al* (2009) for later ones– are subject, despite their differences, to basically the same sort of objection.

2.4. Quick assessment of the debate

I think that connectionist attempts never provided a satisfactory answer to SA, and I think that this applies both to those that tried to reformulate systematicity or

[5] Incidentally, Aizawa thinks that neither connectionist nor classical models can explain systematicity without the aid of further additional hypotheses.

compositionality, and to those that tried to provide practical refutations of the classicist challenge. The problem with the former, as I said above, is that they easily changed the subject matter without really meeting the challenge. The problem with the latter is that they easily fell prey to the classicist dilemma.

Someone could object that this assessment is too quick and unfair with some of the connectionist contenders, and it is possible to point towards this or that particular model to argue that it offers better chances to deal with the classical challenge.[6] I do not deny that some models work better than others and that the process of trying to cope with SA has unraveled many interesting aspects of the properties of both classical and non-classical systems. What I deny is that there is, as of today, an answer that satisfies most authors on the connectionist side, and this is enough to be at least suspicious that the challenge has been met. To put but one recent example, Frank *et al* (2009) review previous connectionist attempts to provide a model with semantic systematicity (Hadley, 1994) without implementing a classical system. They find all of them wanting only to propose their own model that, allegedly, succeeds in the task. One gets the impression that it is only a matter of time before someone comes up with a similar criticism of their model and a similar optimistic claim.

Indeed, I think that the problem with connectionist attempts can be put in different terms: what Fodor *et al* were demanding was not a new family of computational models but a new family of explanatory principles. Even though connectionists claimed to be providing just this when they talked about vector representations, learning algorithms, activation propagation, and the like, the thing is that they did not have an easy day when it came to explain how those principles connected with explaining the relation between the ability to think *J loves M* and *M loves J*. It seemed that in order to do so it was necessary to appeal to how those relations emerged from the network's behavior. Yet all the explanatory load seemed to remain *on what emerged* –the elements *J*, *M*, and *love* and their relations– and not on the goings-on of the system from which it emerged. The latter was, to use the classical parlance, implementation detail. To put it bluntly, what connectionism had to provide, and failed to do, is a new theory of mind.

3. How to view systematicity as a general property of cognition

[6] I owe this objection to a referee that wanted to know what was wrong with a specific model. Obviously, answering questions like this exceeds the limits and goals of this paper.

As I said, SA rests simply on the claim that a lot of cognition is systematic, not necessarily all of it. However, there are two claims that, when taken jointly, may sustain the view that systematicity is a general property of cognition. The claims, which are part and parcel of Fodor's view of mind, are:

(1) Cognition as concept involving. As Fodor says (1998: vii), "the heart of a cognitive science is its theory of concepts". What distinguishes cognition from, say, perception is that cognitive processes work on concepts. Hence processes that work on non-conceptual representations are of relatively little interest for the central claims about the nature of cognition.

(2) *Compositionality as a non-negotiable property of concepts.*[7] Whatever concepts are, they are compositional, i.e., they can be combined with other concepts to form larger conceptual structures in such a way that the content of the compound is a function of the contents of the concepts it contains and their mode of combination.

Taken together (1) and (2) entail that the constitutive elements of cognition –concepts– have a fundamental property –compositionality– that is the source of systematicity –i.e., a conceptual system is *ipso facto* a systematic system. In other words, systematicity is a general property of cognition that derives from the nature of the cognitive elements.

Do we have good grounds to maintain (1) and (2)? I am going to assume that (1) is true and I will take issue with (2). I am not going to provide an argument for (1) but let me say briefly that it is an assumption that, tacitly or explicitly, is widely endorsed in cognitive science. Even in those accounts that try to blur the distinction between cognition and perception, such as Prinz's neoempiricist theory of concepts (Prinz, 2002), there is something that distinguishes concepts from other mental representations and, therefore, that distinguishes cognition from perception. For instance, in Prinz's view even if concepts are copies of percepts the former have the distinctive property of being under internal control.

Let me thus focus on (2). The question of compositionality has been in the

[7] See Fodor (1998, ch. 2). Fodor's idea of a non-negotiable condition for a theory of concepts is that the condition is fallible but abandoning it entails abandoning the representational theory of mind itself.

agenda for years, especially due to Fodor's insistence on using it against non-atomistic theories of concepts (Fodor, 1998; Fodor and Lepore, 2002). His argument, in a nutshell, goes like this: concepts are the basic elements of thought; compositionality is a "non-negotiable" property of concepts; but non-atomistic theories of concepts –i.e., those that contend that concepts are structured representations such as prototypes– are incapable to meet compositionality demands; hence non-atomic concepts are ill-suited to figure as the basic elements of thought.

The problem of compositionality was already detected by early proponents of prototype theory (Osherson and Smith, 1981), and some technical solutions have been attempted (e.g., Kamp and Partee, 1995). There is recently a defence against the compositionality argument –endorsed by Prinz (2002, 2012), Robbins (2002), or Weiskopf (2009a)– that relies on the idea that it is a modal property. The idea is that concepts *can* combine compositionally but they do not necessarily have to do it all the time. Prinz (2012) contends that this weaker requirement allows us to regard prototypes as compositional given that there are cases in which they behave compositionally (i.e., the semantics of the compound is fully determined by the semantics of its parts), and there are others in which the compositional mechanism may not be used, or it may be regularly supplemented with other combination mechanisms.

I think that this defense is weak. First, notice that the "can" involved in it demands that there is something in the nature of concepts that allows them to be compositionally combined. So the defense assumes that compositionality is a general *constitutive* property of concepts, and it seems to demand that there are general compositional mechanisms that can work on concepts, even if sometimes they are not used. If this is the case, then it still follows that systematicity is a general property of cognition, even if sometimes it does not show up. Second, to show that prototypes are compositional, the relevant thing is to show that they are combined *as prototypes*. Yet it seems that instances of prototype combination are compositional inasmuch as their prototypical features are simply dropped away.

Although I do not wish to address the debate on compositionality in the limited space of this paper, I dare to say that Fodor's criticisms have never been properly rebutted. Compositionality is still a problem for prototypes and other structured concepts. However, the compositionality of concepts cannot be used to support the view that systematicity is a general property of cognition. The reason is that Fodor's argument for the compositionality of concepts hinges precisely on the systematicity of cognition –

i.e., if cognition is systematic, the better explanation is a compositional system– so the extent to which concepts are compositional will be given by the extent to which cognition is systematic. But you still need an argument to show that cognition is *generally* systematic in the way classicism demands. Otherwise one can hypothesize that a part of cognition is systematic in the required sense –hence works on compositional concepts, hence poses a problem for prototype-like explanations– and another is not –hence does not work on compositional concepts, hence might be accounted for by prototypes or other structured concepts. This hypothesis entails defending a version of conceptual pluralism, which I will provide in the next section. Before doing so, I want to consider a different, although related, argument that may offer independent reasons to hold that cognition is systematic and compositional.

The argument arises from Evans's well-known Generality Constraint. The constraint can be succinctly put thus:

> "[I]f a subject can be credited with the thought that *a* is *F*, then he must have the conceptual resources for entertaining the thought that *a* is *G*, for every property of being *G* of which he has a conception." (Evans, 1982: 104)

Weiskopf (2010) argues that the constraint can be understood as an *architectural* constraint, that is, "as a constraint on the sorts of representation combining capacities a creature must have in order to possess concepts" (109, fn. 1). The constraint acts as a closure principle for the conceptual system so that "[n]othing could be a concept unless it was capable of entering into this kind of system of relations, and nothing could count as possessing a conceptual system unless it had a system of representations that were organized in such a way" (109). Notice that this is the sort of claim that turns systematicity into a non-empirical property, in the sense I referred to in section 1.1. In other words, systematicity would be a demand on mental architecture derived not from our theories on how concepts actually are but from deep intuitions on what concepts *have to be*.

I agree that the Generality Constraint arises from deep intuitions about thought. However, I contend that it is possible to interpret it in a way that does not pose the strong architectural constraint that Weiskopf suggests. If one looks closely to the formulation by Evans, the constraint can be seen primarily in terms of *how to credit* a subject with a thought. In other words, it is a constraint on how to attribute beliefs: it is

not possible to attribute a creature the belief that *a is F* and the belief that *b is G* without allowing the *possibility* of attributing it the belief that *a is G* and the belief that *b is F*. We need an extra assumption to turn the Generality Constraint into an architectural constraint that demands full combinability of concepts in the creature's internal system of representation. This is the assumption that concepts are components of thoughts that have to be combinable in ways that mirror the structure of the beliefs attributable to the creature. Yet, as I am going to argue, there is room to resist this view as a general relation between beliefs and concepts. There will be cases in which concepts will combine in complexes whose structure mirrors the structure of the corresponding beliefs but there will be also others in which there will not be such mirroring. In the latter case, a creature can be credited with the belief, and the credit is grounded in its representational abilities, but the elements in its representations will not correspond part-to-part to the elements in the attributed belief.

In short, what I am going to defend is a version of *conceptual pluralism* that allows us to resist the line of reasoning that leads from the intuitions of the Generality Constraint to the conclusion that systematicity is a general property of cognition. The point is that the conclusion is warranted only for concepts that have the property of being combinable in ways that mirror the structure of the beliefs. If there are other elements of cognition that can be still regarded as concepts but that do not have such a property, then they will not be systematic in the way required by SA. Thus there are two things that I must do to support this line of defense: the first one is to show that conceptual pluralism is a cogent notion, i.e., that it is possible to find elements in cognition that share fundamental properties that characterize them as concepts yet split into different subkinds; the second one is to show that there are subkinds that differ precisely with respect to the property that is the source of systematicity, namely, compositionality.[8]

4. Conceptual pluralism and compositionality

[8] A referee complains that this looks like an unnecessarily circuitous route. Should not be enough for the purposes of the paper to show the second, i.e, that there are elements in cognition that are not systematic in the way required by SA? I don't think so. The point is that one has to motivate first the view that they are precisely elements *in cognition*, that is, conceptual elements. Otherwise one might brush aside the suggestion that there is a different kind of systematicity by saying that it has to do with perceptual or other less-than-cognitive elements. The point of the next section, thus, is to show that there is a general way of characterizing concepts so that they comply with the Generality Constraint, understood as a constraint on belief-granting capabilities, while at the same time they split into subkinds that differ in important respects.

Conceptual pluralism is the thesis that concepts constitute a kind that splits into a number of different subkinds. The notion appeared in the context of the debate against Machery's claim that concepts are not a genuine natural kind, and hence they are not fit to figure in psychological theories (Machery, 2009). The basis for this eliminativist claim is that what psychologists call concept is served by an assorted collection of representations, such as prototypes, exemplars, or mini-theories, that have very little in common, either in terms of their structure, or of the processes that operate on them. So Machery contends that there are not many useful generalizations that can be made about them.

In contrast, pluralistic approaches to concepts (Weiskopf, 2009a, 2009b) hold that there are different kinds of mental representations that can be rightfully regarded as concepts. Psychological literature shows, indeed, that prototypes, exemplars or theory-like structures appear to have a role to play in dissimilar cognitive tasks.[9] Yet the conclusion to draw is that minds have the three kinds of representational structures at their disposal, and they make a selective use of each of them depending on the type of task in which they are engaged. Still, those different kinds of representations have enough in common to be regarded as subkinds of a more inclusive, superordinate kind – the kind of concepts.

What are those common properties that unify concepts as a kind? They have to be properties picked at a different level than those that unify each subkind of concepts. In other words, in order to show that concepts are a kind you cannot use criteria that split themselves, i.e., criteria that are applied differently to the different hypothesized subkinds. What is needed is some middle point at which one can find common high-level properties that are robust enough to block the eliminativist conclusion but still permit a plurality of kinds that possess them. In other words, one needs to show, first, that there are properties that qualify concepts as a class and, second, that there are different subkinds that share those properties and yet differ in other significant properties. Among the properties of concepts suggested in the literature, there are two that stand out as the most prominent ones: their *centrality*, and their role in *attributions*

[9] Although I do not wish to enter the debate on conceptual atomism, I would like to point out that conceptual pluralism allows for the possibility of atomic concepts as one more among the subkinds of concepts. Weiskopf (2009a) seems to forget this possibility when he opposes atomism to pluralism. As I pointed out in Martínez-Manrique (2010), the relevant opposition is between pluralism and monism, and the former can admit atoms in the repertoire as long as they are not mistaken as the whole class of concepts.

of belief. Consequently, I contend that they pose the minimal common requirements that qualify concepts as a class. On the other hand, I will argue that a significant property in which subkinds of concepts differ is compositionality and, hence, systematicity. Let me elaborate a little on the common properties of concepts in the next subsection and leave the question of the differences in systematicity for the following one.

4.1. Common properties of concepts: centrality and belief-attribution

Centrality is the idea that concepts are *central* mental representations, as opposed to *peripheral* ones. By 'peripheral' I mean mental representations that are closer to the stimuli or input. This distinction has been used in different ways in theories of concepts. For instance, to point out a couple of recent examples, Camp (2009) singles out stimulus independence as one of the crucial factors that mark conceptuality, while Prinz (2002) appeals to internal control as the distinctive property between concepts and percepts –which in his view are undistinguishable with respect to its modality-specific constitution. The distinction between central and peripheral also plays a pivotal role in classical modularist views of mind (Fodor, 1983), where peripheral representations correspond to the proprietary bases of input modules, and central representations are typically the concepts handled by the central processor. Indeed, even massive modularist views of mind (Carruthers, 2006) make a distinction between conceptual and perceptual modules, which depends on architectural considerations regarding the distance to the input.

The second prominent property of concepts is that they are the representations whose possession allows the possibility of *attributing belief-like states* (as well as other kinds of propositional attitude states) to an individual. I intend this property to be neutral between those theories that hold that beliefs must be actually composed of concepts (Fodor, 1998), and more instrumentally-inclined theories that hold that beliefs can be ascribed to creatures with representational capabilities without necessarily holding that the tokened representational structures are literally composed by parts that correspond to those of the attributed belief (Dennett, 1987). The point I want to make is that it is possible to make compatible, on the one hand, the rejection of the notion of beliefs as actually composed by concepts as smaller pieces with, on the other, a representationalist stance on concepts. Concepts would be the sort of mental representations whose possession allows an organism the possibility of exhibiting

behaviors that grant attributions of belief.

Let me illustrate this with a toy example from the literature on animal cognition. Consider birds, such as jays (Clayton *et al*, 2003), that are capable of remembering the location where they stored food some time ago. One can describe the bird's performance by saying that the jay remembered where it stored the food, which involves attributing it the belief *that there is food at location l*. I think that there are two claims about this description that is necessary to reconcile. One is the claim that it is a genuinely explanatory statement: it provides a description that allows one to make generalizations that are useful, perspicuous and predictive. The other is the claim that it possibly strains the capability of birds (more on this later) to say that they are capable of combining concepts such as FOOD and LOCATION so as to form beliefs like the one I mentioned. Following the first claim, someone would like to contend that the bird does literally possess the structured belief that is composed by those concepts. Following the second claim, someone would like to contend that belief attribution is a wholly pragmatic affair that does not reflect the innards of the creature. However, there is a middle ground between both contentions: given the bird's food-tracking abilities, it is possible that it deploys actual mental representations for the attributed concepts FOOD and LOCATION, without deploying anything like a structured representation for the attributed belief *there is food at location l*. In other words, one can be (approximately) a *realist* about concepts and, at the same time, (approximately) an *instrumentalist* about beliefs. Belief attributions like this would not be merely instrumental and observer-dependent but would be supported by certain representational abilities that some organisms possess and others do not. Concepts would thus be *those mental representations that it is necessary to possess so as to be the kind of organism to which one can attribute beliefs*.

Nothing prohibits, however, that in certain cases the structure of the attributed belief could be actually mirrored by the structure of the representational structure that grants the attribution. Yet this does not split the notion of belief into two different kinds –one for beliefs that are representationally mirrored and another for beliefs that are not so. Attributing beliefs has principally to do with the possibility of making predictions and generalizations regarding the organism's behavior, and this possibility can be served whether the representational states that underlie the behavior mirror those beliefs or not. This opens the door to the possibility of having two kinds of concepts, managed by two kinds of mechanisms, that underlie attributions of belief.

The point to consider now is whether there are elements that can be rightfully

regarded as concepts, inasmuch as they exhibit the properties of being central and being the representations that underlie attributions of belief, and yet split into subkinds that differ with respect to properties that are the source of systematicity. The relevant property in this respect, of course, is compositionality.

4.2. Compositional and non-compositional concepts

Let me take stock: I said above that the Systematicity Argument works on the premise that systematicity is a significant property of cognition yet it does not contain itself the stronger notion that systematicity is a general property of cognition. To support the latter one may appeal to the claim that concepts are non-negotiably compositional and back this claim with intuitions from Evans's Generality Constraint. I tried to debunk the idea that the constraint mandates a certain architecture, so as to show that there may be different kinds of representations that possess the minimal requirements for concepthood and that satisfy the constraint. Now it is time to argue that those kinds of concepts differ in some respect that does not allow us to regard systematicity as a general property of them. I want to claim that there are mental representations that qualify as concepts, in terms of being central and involved in belief attributions, and yet are not compositional, and hence systematic, in the way SA contends. The upshot is that we would have two different kinds of concepts that differ in their compositional properties.

Let me back that claim adapting some ideas from Camp (2009), who provides a careful analysis of the concept of 'concept' that takes into account evidence from animal abilities. She begins by noting that notions of 'concept' typically oscillate between two extremes: *concept minimalism,* in which for a cognitive ability to be regarded as conceptual it simply has to be systematically recombinable; and *concept intellectualism*, which links conceptuality to linguistic abilities, so that language, or some capacity that is only possible by means of language –e.g., the capacity for thinking about one's thoughts– becomes necessary for conceptual thought. Both extremes would delimit a continuum in which Camp thinks it is possible to distinguish three notions of concept:

> "a minimalist ‘‘concept$_1$,’’ denoting cognitive, representational abilities that are causally counterfactually recombinable; a moderate ‘‘concept$_2$,’’ denoting cognitive, representational abilities that are systematically recombinable in an

> actively self-generated, stimulus-independent way; and an intellectualist ‘‘concept$_3$,’’ denoting concept$_2$-type representational abilities whose epistemic status the thinker can reflect upon, where we assume that this latter ability is possible only in the context of language." (2009: 302)

Concept$_1$ is involved in activities that demand little more that passive triggering and marks the lower limit of the notion. Concept$_2$ is typically associated to cognitive abilities engaged in instrumental reasoning. This cognitive activity, which we find in a number of non-human animals, demands from the creature the capacity to represent states of affairs that are not directly provided by the environment, namely, the goal-states that the creature wants to achieve and the mean-states that bring it closer to that goal in a number of stages. Finally, concept$_3$ marks the upper limit and it is here, she contends, where Evans's Generality Constraint can be actually met because only concept$_3$ grants full recombinability, i.e., the capacity to combine arbitrarily any *a* and *b* with any *F* or *G* of which the creature has a conception. Concept$_2$ cannot grant this capacity because, even if its representational power is removed from the immediate environmental stimulation, its deployment is still tied to the creature's immediate needs. To put it in Camp's terms, a chimpanzee would never entertain any of the potential thoughts that Evans's constraint refers to “because they are utterly useless for solving any problems that it actually confronts” (2009, 297). In contrast, creatures with language and the ability for epistemic reflection –the requirements for concept$_3$– can find some use for the most arbitrary combinations once they have certain epistemic drives, such as curiosity and imagination.

Appealing as I find this analysis, there are two important points that I find unconvincing. First, Camp states that concept$_1$ is less theoretically useful to provide an account of conceptual thought. In fact, I think that it is doubtful that this notion even meets the minimal requirements for concepthood. Camp relies on some capacity for recombination as a minimal requirement to count as conceptual. However, the fact that this capacity can be found in systems that are directly triggered by perceptual stimulation ought to make one suspicious of the proposal. As I pointed out in section 3, one wants an account of cognition as concept-involving in a way that lets one distinguish it from perception. Centrality, I argued in section 4.1, is a way to mark such a distinction. Yet the notion of concept$_1$ is clearly tied to non-central capacities, so it does not meet the minimal criteria for concepthood. Recombinability is a red herring

because one can find it in non-conceptual structures.

The second unconvincing point is Camp's treatment of the Generality Constraint. Camp endorses the view, that I resisted above, that it is an architectural constraint. At the same time, she contends that it works as an ideal rather than as a necessary constraint to grant conceptual thought. To meet the constraint one needs the fully systematic recombinability that permits arbitrary combinations to occur. Yet, in her view, even creatures with concept$_3$ capacities would often not meet the constraint given that many times they would be reluctant to form the arbitrary thoughts that, according to the constraint, they must be capable to form.[10] This way she makes room for a way to accept the Generality Constraint that, at the same time, allows to regard conceptuality as a matter of degree. In other words, the constraint is ok but it is too strong to be met in full for most practical concerns. Now, the reasoning behind this conclusion seems to me close to the reasoning behind the modal defense of compositionality that I discussed in section 3, and thus committing the same sort of mistake but in the opposite direction. Let me explain.

Recall that the reasoning of the modal defense was that representations that can *sometimes* combine compositionally count as compositional, even if other times they cannot so combine. This was used to support the compositionality of representations such as prototypes. Camp's reasoning is that creatures with concept$_3$ capacities *sometimes* are not capable to entertain certain combinations for practical purposes. This is used to deny that they meet the Generality Constraint "in full". The mistake in both cases is the same: what it takes for representations to count as compositional, and to meet the Generality Constraint, is that they are capable to be arbitrarily recombined *as a matter of how they are constituted* (and given certain processes sensitive to this constitution). It is irrelevant whether as a matter of fact they sometimes do or do not combine. The upshot is that, despite what Camp contends, her notion of concept$_3$ *does* meet the Generality Constraint. But if this is the case, and one still wants to maintain that the constraint restricts the suitable conceptual architecture, now one may object to her pluralist gradable analysis of the notion of concept. One could say that, as we have only a class of representations that meet the constraint –concept$_3$– we'd better regard this class as the *genuine* notion of concept, and the other two notions as varieties of non-

[10] To put it in Camp's words, "we also fall short of full generality: precisely because certain potential thoughts are so absurd, it's unlikely that anyone would ever think them or utter sentences expressing them in any practical context" (2009, 306).

conceptual representations.

However, notice that in section 3 I offered an alternative reading of Evans's constraint that poses much more lax restrictions on the representations that a creature must posses in order to satisfy it. So it does not matter much whether a class of concepts includes representations with limited compositionality. What is crucial is that they are central representations whose possession is required in order to grant systematic attributions of belief. In this respect the notion of $concept_2$ appears as a suitable candidate for concepthood, unlike the peripheral, perceptually-bound representations in $concept_1$.

Where do these considerations leave us? I think that Camp's analysis allows for the existence of just two kinds of concepts. One of them, roughly corresponding to her $concept_2$, has the minimal common requirements to be regarded as conceptual but does not appear to be compositional in the classical sense[11]; hence, it is incapable of giving raise to the sort of systematicity referred in SA. The other kind, roughly corresponding to $concept_3$, is compositional and supports systematicity in SA.

Now, there are two final related issues that I wish to address to end paving the way to an alternative non-classical SA. One is: even if it were possible to tell two notions of concept apart, systematicity could still be a general property of cognition. The reason is that each notion could be applicable to a different type of creature. For instance, Camp's analysis suggests a scenario in which $concept_2$ is simply the basis of non-human animal thought, while human thought is exclusively constituted by $concept_3$. If this is the case, then one may contend that systematicity is a general property of *human* cognition, which is still a strong claim. To debunk this claim one ought to show that both kinds of concepts have a place in human cognition.

The second issue is that even if humans possess both kinds of concept, it still may be the case that classical systems can account for them. In other words, one must show not only that SA applies just to a part of cognition –the one that deals with $concept_3$– but that it does not apply to the other part –the one that deals with $concept_2$. I address these two issues in the next and final section.

---

[11] As we will see in the next section, there is the question whether they are compositional in a different sense. Now, I do not wish to fight for the term 'compositional'. I am ready to leave it as the property that characterizes the class of concepts present in classical systems ($concept_3$) and accept that the other class of concepts is just non-compositional. What matters for this paper is that these concepts give rise to different systematic properties not accounted for by classical systems.

5. Two kinds of systematicity

The aim of this section, then, is to motivate the view that there are two kinds of systematicity in human minds, each of them related to a different kind of concept. It is more than mere wordplay to say that a kind of systematicity involves a kind of system. It is because classical symbol systems have the defining properties that they have *as systems* that they exhibit the sort of systematicity of SA. So to argue for two kinds of systematicity one must search for reasons that back the existence of two kinds of systems, each of them working on a conceptual kind. Dual systems theory (DST) is an obvious candidate to provide the backbone of such an approach.[12]

DST is the view that human minds are constituted by two distinct kinds of cognitive processing systems (Evans and Frankish, 2009). Although their detailed characterizations and properties vary depending on the specific theory, in general terms one is typically characterized as fast, automatic, holistic, inflexible, difficult to verbalize, evolutionarily old and nonconscious, and the other as slow, controlled, analytic, flexible, more easy to verbalize, evolutionarily recent and conscious. Following standard usage, I will refer to those systems as S1 and S2, respectively. Even though there are differences about how to articulate this general view (Evans 2009; Stanovich 2009), I will not take them into account. DST has been mainly applied to explain reasoning and social cognition but in their most ambitious forms it purports to provide a general vision of mental life, in which the basic distinction between two kinds of systems is the fundamental architectural design of human cognition that helps to account for a range of mental phenomena (Carruthers, 2006; Samuels, 2009).

The first thing to note is that both S1 and S2 have to be conceptual systems: they are involved in paradigmatic central cognitive processes, such as reasoning, not in perception or other input-controlled processes; and the sorts of behaviors that any of them controls, such as decision-taking, give rise to belief-attributions. The question now is whether each system can be conceived of as working on a different conceptual kind. In other words, whether creatures with a dual system architecture are endowed *both*

[12] The approach by Gomila *et al* (2012) has elements that are congenial to the proposal I am making in this paper. For instance, they also resort to dual systems theory as the overarching architecture of mind. Yet I do not agree with their claim that "this duality also corresponds to the divide between non-systematic and systematic processes". There is much systematicity in S1 and of a kind that demands conceptual processing, even if not the kind of concepts that are processed by S2. So I doubt that dynamic, embodied approaches, as the one they endorse for S1, provide a good account for this system either, at least if they are couched in non-representational terms.

with concept$_2$ and concept$_3$ abilities. The scenario to be considered in terms of DST would be one in which the first kind of concept is handled by S1 and the second by S2. The reason is that the properties exhibited by S1 resemble those of the conceptual capabilities associated to instrumental reasoning that, as we saw above, are arguably present in some non-human animals. This also fits the idea the S1 is evolutionarily older than S2, and that S2 is likely to be exclusive of humans. Both systems would be capable of performing typical conceptual functions, such as categorization, reasoning and meaning extraction, yet in different ways and with different limits in the kinds of thoughts that they are capable of delivering. In particular, S2 would be capable of satisfying the Generality Constraint understood as an architectural constraint but S1 would not.

Is this a plausible scenario? Support for a positive answer can be found for the recent extensive review by Penn *et al* (2008) comparing human and animal cognition. Their aim is to show that there is a "profound functional discontinuity between human and nonhuman mind" (2008: 110). The discontinuity is revealed in a wide range of domains, such as the ability to cope with relational (as opposed to perceptual) similarity, to make analogical relations, to generalize novel rules, to make transitive inferences, to handle hierarchical or causal relations, or to develop a theory of mind. Penn *et al*'s point is basically to show that the discontinuity between human and non-human minds can be cashed out in terms of the presence of a capacity for systematically reinterpreting first-order perceptual relations in terms of higher-order relational structures akin to those found in a physical symbol system (PSS) –the archetypal classical system.[13] This is the sort of capacity that, in Camp's analysis, required something like concept$_3$.

Non-linguistic creatures do not exhibit such kind of systematicity. Instead, they manifest a different kind of systematicity that "is limited to perceptually based relations in which the values that each argument can take on in the relation are constrained only by observable features of the constituents in question" (2008: 127). Borrowing Bermúdez's (2008) term, I will call it *featural systematicity*. They think that this systematicity would be accounted for by compositional properties[14] different from those

[13] In order to cope with critics of the PSS hypothesis, Penn *et al* borrow a milder version from Smolensky (1999), the "Symbolic Approximation" hypothesis. The distinction is irrelevant for the purposes of this paper, given that the point is still that symbolic approximators require an architecture that is different from the one that supports animal capacities.

[14] As I said in fn. 11, it is irrelevant whether one does not want to call them 'compositional' and prefers to reserve the term for classical compositionality. What matter is that there is a different kind of systematicity accounted for by different properties of the system.

that characterize a PSS. Penn *et al* resort to van Gelder's (1990) notion of functional compositionality to account for the kind of compositionality present in animals. Unlike van Gelder, however, Penn *et al* do not regard functional compositionality as capable of underlying the sort of systematicity exhibited by humans –i.e., as capable of satisfying SA. Animal compositional capacities would be limited to "some generally reliable and productive mechanism for encoding the relation between particular constituents" that would account for "the well-documented ability of nonhuman animals to keep tracks of means-ends contingencies and predicate argument relationships in a combinatorial fashion" (2008: 125). The animal abilities referred are basically of the same kind as those that, according to Camp, grant the attribution of $concept_2$.

As DST –a theory that Penn *et al* regard as related to their view– the thesis is that both kinds of systematicity appear in humans, so it is necessary to explain how. Penn *et al* propose that the representational system unique to humans "has been grafted onto the cognitive architecture we inherited from our nonhuman ancestors" (2008: 111). In search of an explanation of how such "grafting" might be possible, they resort to computational models. Nonclassical connectionist models might explain the kind of systematicity that we find in animals, whereas recent connectionist-symbolic models might account for the grafting of human new representational abilities to the preexisting representational machinery. Even though they back their proposal with computational models of their own (e.g., Hummel and Holyoak, 1997, 2003) one might object that strong evidence for it is still lacking. However, I want to consider a different kind of objection that is more relevant for the purposes of this paper: accepting that there are two different processing systems, why could not one resort to a classical explanation for *both* of them? In other words, one could insist on the possibility that animal systematic capabilities were underlied by a classical compositional symbolic system, perhaps limited with respect to the range of represented contents that it can deal with but still working on the same principles of concatenative recombination. If this were the case, one could contend that the difference between both systems –or between the concepts on which they operate, or between the systematicity they exhibit– was not one of *kind*. Classical systems could then still constitute the keystone of cognition in general, just as SA contends.

I think that there are good grounds to reject this possibility. Its problem, in a nutshell, is that symbol systems are *too strong* for that. Recall that from the classicist perspective it is impossible to have a classical system that is not systematic in the sense

posited by SA. So if animal minds included in some way a classical system, *then they would ipso facto be endowed with standard full systematicity*. The extensive evidence reviewed by Penn *et al* shows precisely that this is not the case. As they contend, if there were no differences in kind, then one could expect that the observed discontinuities would be erased under appropriate conditions. For instance, animals under a "special training regime", which let them access to a larger range of contents and relations, would at least approximate human behavior. Yet the evidence shows that even those animals have a poor performance.

It seems to me that we have finally reached the elements that would allow us to construct a Non-classical Systematicity Argument. Recall the general form that such an argument would have:

NSA

(i') Empirical claim: X is a pervasive property of cognition

(ii') Explanatory claim: the only plausible explanation for X is property Y

(iii') Definitional claim: Y is a defining property of such and such non-classical systems

(iv') Dilemma: Dilemma: if classicism cannot account for property Y then it does not provide a full account of cognition (from i & ii); if classicism can account for Y then it is actually implementing a non-classical system (from ii & iii)

Now we have ways of seeing how the different claims in the argument could be substantiated. First, the X that we have to explain is the kind of systematicity exhibited by nonhuman animals in terms of their limited recombination abilities –limited by their perceptual repertoire even if not bound to the immediate environment, and limited in the kinds of relations that they allow. Moreover, it is also a pervasive property of human cognition given that it belongs to the inherited part of our cognitive machinery.

Second, the Y that constitutes the best explanation of this systematicity is some property of non-classical systems. It cannot be a product of classical systems because, as I have just argued, this would endow animals with human systematic capabilities. A plausible candidate for Y comes from the set of properties characteristic of *distributed representations.* Perhaps, as Penn *et al* observe, distributed systems as we currently envision them may need to be supplemented to account for animal minds. Yet it would suffice for NSA that distributed representations are essentially involved in the

explanation of featural systematicity, and that whatever supplement they require *cannot be* classical.

This would also satisfy the definitional claim, given that it simply says that whatever property is Y it is constitutive and characteristic of some non-classical system. In this respect, distribution is a defining property of distributed non-classical systems, from which it follows that it is simply not possible to be such a system and not to have distributed representations. To conclude, the dilemma for the classicist position comes to this: if it cannot account for the sort of systematicity exhibited by animals and by part of human cognition, then it does not provide a full account of cognition; and if they offer a model that exhibits non-classical property Y –for instance, distributed representation– then given that Y is defining of non-classical systems the model would count immediately as an implementation of a non-classical system.

6. Conclusion

After all these years the Sytematicity Argument still poses a powerful challenge to any attempt at explaining cognition. Part of its force resides in its simplicity: "here is this notorious property of cognition; here is a conspicuous explanation of this property; does anyone have an explanation that does not collapse into ours?". In this paper I claimed that the answer to the latter question is negative. Despite the attempts, nobody has come with a better explanation for the sort of systematicity that the argument alludes to than a compositional system of representations. And nobody has a complete account of cognition unless one is able of explaining properties of that sort.[15] However, I also contend that this is not the end of the story: there are other cognitive properties to explain, and classicism is not in a better position to do so. Just because one has a powerful explanation of an important mental property, it does not mean that one can transfer this explanation to every other mental property. If, as the evidence is increasingly supporting, the human mind includes two fundamentally different kinds of systems, and each system exhibits a different way of being systematic, then classical symbol systems cannot account for both of them.

[15] So attempts at providing a whole alternative framework to computational-representational cognitive science, such as Chemero (2009), seem to be flawed inasmuch as they simply ignore those properties. For instance, there is no single clue in his book about how radical embodied cognitive science would deal with language comprehension or with reasoning processes, just to mention two paradigmatic domains where resort to classical representations is more natural.

The bottom line can be put thus: while nonclassical systems are *too weak* to account for human-like compositionality-based systematicity, classical systems are *too strong* to account for non-compositionality-based systematicity. The reason is precisely that any system that has a classical computational-representational architecture *necessarily* exhibits compositional systematicity as a consequence of architectural design. Yet I have argued that the evidence suggests that, both in animals and in humans, there are genuinely cognitive processes that fail to exhibit such kind of systematicity. They are genuinely cognitive because they are concept-involving: they are not tied to immediate perceptual stimuli, and they control behaviors that are complex enough so as to merit attributions of belief. If nonclassical approaches are able to explain such processes –and not only, as their critics often complain, early perceptual processing– then they will have an account of part of our mental life, even if not of all of it.

To sum up, the picture of cognition that I tried to motivate in this paper comes to this: an architecture that supports at least two distinct subkinds of concepts with different kinds of systematicity, neither of which is assimilable to the other. This picture sets a whole new agenda of problems to solve, particularly regarding the relation between both systems. In particular, one may wonder whether non-classical systematicity is exactly the same in humans and in those animals that exhibit analogous properties, or perhaps it is affected by its coexistence with compositional systematicity; one may wonder whether it is possible to integrate both kinds of concepts in some respect, perhaps to form a sort of hybrid structure; one may wonder whether compositional systematicity is exclusively related to linguistic cognition. These are the sorts of questions that I think it will be interesting to address in future research.

References

Aizawa, K. (1997) "Explaining systematicity". *Mind and Language, 12*: 115-136.

Aizawa, K. (2003) *The Systematicity Arguments*. Dordrecht: Kluwer.

Bechtel, W. (1994) "Natural deduction in connectionist systems". *Synthese,* 101: 433-463.

Bermúdez, J. L. (2008) "The reinterpretation hypothesis: Explanation or redescription?" *Behavioral and Brain Sciences*, 31: 131-132

Camp, E. (2009) "Putting thoughts to work". *Philosophy and Phenomenological*

Research, 78: 275-311.

Carruthers, P. (2006) *The Architecture of Mind*. Oxford: Oxford University Press.

Chemero, T. (2009) *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.

Clark, A. (1989) *Microcognition* Cambridge, MA: MIT Press.

Clayton, N. S., Yu, K. S., and Dickinson, A. (2003) "Interacting cache memories: Evidence for flexible memory use by Western scrub-jays (Aphelocoma californica)". *Journal of Experimental Psychology: Animal Behavior Processes*, 29(1): 14-22.

Cummins, R. (1996) "Systematicity". *Journal of Philosophy*, 93: 591-614.

Cummins, R., Blackmon, J., Byrd, D., Poirier, P., Roth, M., and Schwarz, G. (2001) "Systematicity and the cognition of structured domains". *Journal of Philosophy*, 98: 1-19.

Dennett, D. C. (1987) *The Intentional Stance*. Cambridge, MA: MIT Press.

Dennett, D. C. (1991) "Mother Nature versus the Walking Encyclopedia: A Western drama". In Ramsey *et al*, pp. 21-30.

Evans, G. (1982) *The Varieties of Reference*. Oxford: Clarendon Press.

Evans, J. St. B. T. (2009) "How many dual-process theories do we need? One, two, or many?". In Evans and Frankish, 2009, pp. 33-54.

Evans, J. St. B. T., and Frankish, K. (eds.) (2009). *In Two Minds: Dual Processes and Beyond*. Oxford, Oxford University Press.

Fodor, J. (1983) *The Modularity of Mind*, Cambridge, MA: MIT Press

Fodor, J. A. (1998) *Concepts*. 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., and Lepore, E. (2002) *The Compositionality Papers*. Oxford: Oxford University Press.

Fodor, J. A., and McLaughlin, B. P. (1990) "Connectionism and the problem of systematicity: Why Smolensky's solution soesn't work". *Cognition,* 35: 183-204.

Fodor, J. A., and Pylyshyn, Z. W. (1988) "Connectionism and cognitive architecture: A critical analysis". *Cognition,* 28: 3-71.

Frank, S. L, Haselager, W. F. G., and van Rooij, I. (2009) "Connectionist semantic systematicity". *Cognition*, 110: 358-379.

Gomila, A., Travieso, D., and Lobo, L. (2012) "Wherein is human cognition systematic?". *Minds and Machines* doi 10.1007/s11023-012-9277-z

Goschke, T., and Koppelberg, D. (1991) “The concept of representation and the representation of concepts in connectionist models”. In Ramsey *et al* 1991.

Hadley, R. F. (1994) “Systematicity in connectionist language learning”. *Mind and Language,* 9: 247-272.

Hummel, J. E. & Holyoak, K. J. (1997) “Distributed representations of structure: A theory of analogical access and mapping”. *Psychological Review, 104*: 427-66.

Hummel, J. E. & Holyoak, K. J. (2003) “A symbolic-connectionist theory of relational inference and generalization”. *Psychological Review, 110*: 220-64.

Johnson, K. (2004) “On the systematicity of language and thought”. *Journal of Philosophy* 3: 111-139.

Kaye, L. J. (1995) "The languages of thought". *Philosophy of Science,* 62: 92-110.

Kamp, H., and Partee, B. (1995) "Prototype theory and compositionality". *Cognition*,. 57: 129-91.

Machery, E. (2009) *Doing without Concepts.* Oxford: Oxford University Press.

Martínez Manrique, F. (2010) “On the distinction between semantic and conceptual representation”. *Dialectica,* 64 (1): 57-78

McLaughlin, B. (1993) "Systematicity, conceptual truth, and evolution". *Philosophy, Supplement,* 34: 217-234.

McLaughlin, B. (2009) "Systematicity redux". *Synthese,* 170: 251-274.

Niklasson, L. F., and van Gelder, T. (1994) "On being systematically connectionist". *Mind and Language,* 9: 288-302

Osherson and Smith “On the adequacy of prototype theory as a theory of concepts”. *Cognition*, 9(1): 35-58.

Penn, D. C., Holyoak, K. J., and Povinelli, D. J. (2008) "Darwin’s mistake: Explaining the discontinuity between human and nonhuman minds". *Behavioral and Brain Sciences*, 31: 109-178.

Prinz, J. (2002) *Furnishing the Mind*, Cambridge, MA: MIT Press.

Prinz, J. (2012) "Regaining composure: A defense of prototype compositionality". In M. Werning, W. Hinzen, and E. Machery, eds. (2012) *The Oxford Handbook of Compositionality*. Oxford: Oxford University Press.

Ramsey, W., Stich, S. P., and Rumelhart, D. E., eds. (1991) *Philosophy and Connectionist Theory,* Hillsdale, NJ: Lawrence Erlbaum.

Robbins, P. (2002) “How to blunt the sword of compositionality”. *Nous* 36: 313-334.

Samuels, R. (2009) “The magical number two, plus or minus: Dual Process Theory as a

theory of cognitive kinds". In Evans and Frankish, 2009, pp. 129-146.

Smolensky, P. (1988) "On the proper treatment of connectionism". *Behavioral and Brain Sciences,* 11: 1-23.

Smolensky, P. (1991a) "The constituent structure of connectionist mental states: A reply to Fodor and Pylyshyn". In T. Horgan and J. Tienson, eds. (1991) *Connectionism and the Philosophy of Mind* Dordrecht: Kluwer, pp. 281-308.

Smolensky, P. (1991b) "Connectionism, constituency, and the language of thought". In G. Rey, ed. (1991), *Fodor and his critics*, Oxford: Blackwell, pp. 201-227.

Smolensky, P. (1999) "Grammar-based connectionist approaches to language". *Cognitive Science*, 23: 589-613.

Stanovich, K. E. (2009) "Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory?". In Evans and Frankish, 2009, pp. 55-88.

Sterelny, K. (1990) *The Representational Theory of Mind.* Oxford: Blackwell.

van Gelder, T. (1990) "Compositionality: A connectionist variation on a classical theme". *Cognitive Science,* 14*:* 355-384.

van Gelder, T. (1991) "What is the "D" in "PDP"? A Survey of the Concept of Distribution". In Ramsey *et al*, pp. 33-60.

van Gelder, T., and Port, R. (1994) "Beyond symbolic: Toward a Kama-Sutra of compositionality". In V. Honavar and L. Uhr, eds. *Artificial Intelligence and Neural Networks: Steps toward Principled Integration* London: Academic Press, pp. 107-126.

Vicente, A., and Martínez-Manrique, F. (2005) "Semantic underdetermination and the cognitive uses of language". *Mind and Language*, 20: 537-558.

Weiskopf, D. (2009a) "Atomism, pluralism and conceptual content". *Philosophy and Phenomenological Research* 79, pp. 131-163

Weiskopf, D. (2009b) "The plurality of concepts". *Synthese,* 169: 145-173

Weiskopf, D. (2010) "Concepts and the modularity of thought". *Dialectica,* 64: 107-130.