MEETING THE SYSTEMATICITY CHALLENGE CHALLENGE: A NONLINGUISTIC ARGUMENT FOR A LANGUAGE OF THOUGHT

VÍCTOR M. VERDEJO
UNIVERSIDAD AUTÓNOMA DE MADRID

ABSTRACT: From Fodor and Pylyshyn's celebrated 1988 systematicity argument in favour of a language of thought (LOT), a challenge to connectionist models arises in the form of a dilemma: either these models do not explain systematicity or they are implementations of LOT. From consideration of this challenge and of systematicity in domains other than language, defenders of connectionism have mounted a parallel systematicity argument against LOT which results in a new self-defeating dilemma, what I call here the systematicity challenge challenge (SCC): either LOT does not explain nonlinguistic systematicity, or it is in no better position than its rivals to explain any systematicity, even linguistic systematicity. In this paper, first, I critically examine the SCC and some considerations that seem to support it. Second, I offer a response to the SCC by: (1) showing that LOT was never meant to be a cognitive model restricted only to linguistic systematicity, and (2) formulating a new argument in favour of LOT from nonlinguistic systematicity. Third, I argue that there is a central assumption underlying the SCC and maintain that it is mistaken. I conclude that the classical systematicity challenge continues to be fully valid for linguistic and nonlinguistic domains.

What is by and large the most discussed of the arguments in favour of so-called classical cognitive architecture, makes a case that the systematicity apparent in cognitive states and processes can only be explained by appeal to the existence of mental representations with a compositional syntax and semantics. That is to say that such systematicity can only be explained by appeal to a system of symbols in a language of thought (LOT henceforth). The celebrated systematicity challenge to connectionist models in cognitive science thus emerges from the classical systematicity argument in the form of a dilemma (Fodor and Pylyshyn 1988): either connectionist models cannot explain systematicity, or they only do so because they are implementations of a LOT. In the anti-LOT literature, there are various reactions to this classical systematicity challenge. Particularly deserving of attention is the connectionist response that consists of mounting a parallel systematicity argument to the effect that, if the defender of LOT is right in that linguistic systematicity involves representation in a LOT, then it must equally be the case that nonlinguistic systematicity involves representation that is of a kind other than that based on a LOT (Cummins 1996; Cummins, Blackmon, Byrd, Poirier, Roth, and schwarz 2001; Cummins, Blackmon, Byrd, Lee, and Roth 2005).[1] If we take into account the fact that a significant number of cognitive states and processes are independent of linguistic abilities—in humans, though more clearly in nonhumans—then we can see that the connectionist systematicity argument presents a challenge to the systematicity challenge, or what we can call the systematicity challenge challenge (sCC henceforth). The sCC thus questions the justification for seeing LOT as a *general* model of cognitive architecture.

As it stands, the sCC offers the opportunity to gain some much-needed clarification in our understanding of LOT architectures. In this paper, I will argue that, once we see the real commitments of LOT, the sCC can certainly be met; indeed, its proponents have not even started to show that such a challenge really exists. In the first section, I will carefully examine the sCC and some considerations coming from the classicist landscape that seem to support it. In the second section, I will present some textual evidence that, right from Fodor and Pylyshyn's original 1988 analysis of systematicity, in no relevant sense was LOT considered a model that is restricted only to *linguistic* systematic abilities. I will then, in the third section, present a new argument to the effect that, in the case of the systematicity exhibited in visual perception, LOT representation offers as good an explanation as it does in the linguistic case. A series of objections to that argument will be considered in section IV. The upshot will be that the sCC is just the result of a deep misunderstanding of LOT as applied to nonlinguistic cognitive domains, and hence, that the classical systematicity challenge, when correctly understood, maintains its full force. In section V, I will discuss what seems to me to be the crucial mistaken assumption that is behind the connectionist sCC, namely, that LOT representation must be interpreted as requiring that it shares its structure with (natural) language. In that section, I will argue that a weaker sense of LOT representation, without such a requirement, is enough for a LOT to be perfectly viable. A final conclusion will be offered in the sixth section.

I. THE SYSTEMATICITY CHALLENGE CHALLENGE

In the literature one can find a good number of arguments in favour of LOT.[2] However, the systematicity argument is by far the most ubiquitous. For clarity and for reasons related to the issues discussed in section III, it is helpful to briefly review that argument. A particularly succinct exposition is provided by Fodor:

> Ok, so here's the argument: Linguistic capacities are systematic, and that's because sentences have constituent structure. But cognitive capacities are

> systematic too, and that must be because *thoughts* have constituent structure. But if thoughts have constituent structure, then LOT is true. so I win and Aunty loses. Goody! (Fodor 1987, 150–151, emphasis his)

This needs some elucidation. Paradigmatically, the argument starts from the consideration of systematicity as it is observed in the exercise of linguistic abilities. The most illustrious example is that everybody capable of understanding/producing 'John loves Mary' is capable of understanding/producing 'Mary loves John'. The explanation of this capacity comes from the fact that sentences have constituent structure; that is, they consist of semantic and syntactic units—roughly words with a meaning and syntactic characterization—that can be compositionally arranged in different ways (in accordance with appropriate rules of grammar). The idea is that the systematicity of sentence production/comprehension comes from the subjects' recognizing or somehow being sensitive to such sentential structure. It is useful to observe that the argument, so far, is located at a high level of description; what we can roughly consider to be Marr's computational level 1 (Marr 1982, chap. 1). Thus, what we have so far is a description of what cognitive subjects actually do, which function they actually compute and the explanation that accounts for this fact at a high level—viz., the existence of sentential structure and the subjects' sensitivity to such structure.

The argument proceeds to explain systematicity at an algorithmic level—what we can take as approximately Marr's level 2—where the representation that accounts for the cognitive phenomenon identified at the high level is specified. Now, it is remarked that representational capacities—those underlying sentence production/comprehension—are systematic too; in particular, the capacity to mentally represent a given propositional content or thought is systematic. Again, the most familiar example is that anyone capable of thinking *John loves Mary* is, in point of nomological necessity, also capable of thinking *Mary loves John*. The explanation of this must parallel the case of sentences: thoughts are also semantically and syntactically structured. Finally, the conclusion is reached that the syntactic and semantic constituents of thoughts must involve, in effect, a system of symbols (roughly, neurophysiologic states with both semantic and syntactic properties) in a mental language at the implementational level (corresponding basically to Marr's level 3). In short, we can see the classical systematicity argument as consisting schematically of the following three steps:

(1) sensitivity to structure s (of sentences) explains a particular systematic (linguistic) behaviour B at a high level of description.

(2) structured (linguistic) representation R explains sensitivity to structure s (of sentences) at the algorithmic level.

(3) system of symbols ss explains structured (linguistic) representation R at the implementational level.

As is well known, from the systematicity argument the classical systematicity challenge to connectionist cognitive architecture emerges in the form of a dilemma: either connectionist models cannot explain systematicity—and hence they fail to be satisfactory as models of cognition—or they figure in explanations of systematicity as implementations of a classical cognitive architecture—and hence they fail to be an alternative to LOT (Fodor and Pylyshyn 1988). The reactions to this challenge vary from outright rejection to acceptance, in a range that passes through several degrees of clear annoyance and indifference. some authors have tried to deny the existence of systematicity phenomena as required by the argument (e.g., Johnson 2004). Others have presented a project for the general articulation of implementational connectionism (e.g., Marcus 2001). still others have advocated some dynamicist and anti-representational stance, so that the question of what kind of representation actually accounts for systematicity does not even arise in the first place (e.g., Calvo Garzón 2008). My concern here will be with those who accept the challenge and try to respond to it. More specifically, my concern will be with what I am here calling the sCC, as formulated by Cummins and his adherents (Cummins 1996; Cummins, Blackmon, Byrd, Poirier, Roth, and schwarz 2001; Cummins, Blackmon, Byrd, Lee, and Roth 2005).

Cummins et al. are by no means exclusively concerned with the sCC: it forms just part of a series of objections to classical cognitive architecture.[3] What follows is the essence of its formulation. The sCC begins with the assumption that the intuitive appeal of classical representation in linguistic systematicity has to do with a certain appearance of principleness or naturalness in classical explanations of systematicity with regards to the exercise of linguistic abilities. More precisely, the authors under consideration think that the naturalness in question comes from the fact that the kind of representation that explains linguistic systematicity in classical models *shares* its structure with (is a *structural* representation of) language (I will discuss this assumption in section V). However, it is not difficult to realize that there are a good number of different (specifically nonlinguistic) cognitive domains that also exhibit systematicity. The key idea of Cummins et al. is therefore that, when these different cognitive domains are considered—for instance, vision, colour, music, algebra—and compared with the domain constituted specifically by linguistic abilities, the intuitive appeal of classical explanations of systematicity automatically vanishes. The reason, they argue, is that nonlinguistic domains cannot be said to share their structure with classical LOT representation.

> Once we see that there is sensitivity to systematicity in nonlinguistic domains, an argument parallel to that given for the classical representation of language would be available for the nonclassical representation of these other domains. The perception of objects in space appears to be a case in point. No scheme can structurally represent both sentence structure and spatial structure, since these are not isomorphic to each other. Friends of the Fodorian argument that moves from sensitivity to linguistic systematicity to classical representation ought to be friendly to an analogous argument that moves from sensitivity to spatial systematicity to nonclassical representation, for the Fodorian argument is really just an argument from sensitivity to systematicity in [a domain] *D* to structural representation of *D*. (Cummins 1996, 604–605)

The parallel systematicity argument against LOT representation results more vividly in the formulation of a challenge that is also parallel to the classical systematicity challenge and is in the form of a new self-defeating dilemma for the defender of classical cognitive architecture. This challenge is the sCC: either admit that representation in nonlinguistic domains shares structure with the cognitive domain and is therefore not LOT representation—since it is not linguistically structured—or else accept that systematicity can be accounted for by means of representation that does not share structure with the domain it represents—and therefore accept that purely connectionist and classical explanations of linguistic systematicity are equally valid. The first horn of the dilemma is devastating if LOT is to be considered a general model of cognition for any domain, and furthermore, it involves the implausible thesis that each domain has its own scheme of structural mental representations (what the authors call *representational pluralism*). The second horn is even worse since, by admitting that systematicity explanations recognize that some representation does not share structure with the domain it represents (in the authors' vocabulary, by allowing *encodings*) the classicist withdraws from what has always been the strongest evidence for the defence of LOT. Here is how Cummins et al. present what I am here calling the sCC:

> since the inference from systematicity effects in domains other than language to nonclassical, structural representations is on a par with the inference from systematicity effects in language to classical representations, the systematicities in vision and audition are good evidence for some nonclassical, structural representations. Thus, if the FPM [Fodor, Pylyshyn, and McLaughlin] inference is sound, it constitutes a good argument for representational pluralism. The only way Fodor, Pylyshyn, and McLaughlin can retain the idea that mental representation is monistic [i.e., the idea that there is only a LOT scheme of mental representations for any domain] is by allowing that some systematicity effects can be adequately explained without recourse to a corresponding scheme of structural representations. some systematicity will have to be explained by appeal to mere *encodings* By allowing some encoding, however, friends of Fodor, Pylyshyn, and McLaughlin would forfeit the objection that connectionist data coverage is unprincipled, since encoding, on their view, forces us to be unprincipled somewhere. so it seems that either the objection must go, or Fodor, Pylyshyn, and McLaughlin are forced to accept representational pluralism. (Cummins, Blackmon, Byrd, Poirier, Roth, and schwarz 2001, 181, emphasis theirs)

In short, if the sCC is sound, LOT architecture is in a very bad position: either LOT is strictly limited and does not apply to nonlinguistic domains; or the possibility opens up that LOT does not apply at all, not even to linguistic domains.

strikingly, this new systematicity challenge seems to receive support from LOT advocates themselves. For instance, in a reply to Cummins et al., Davis (2005), while defending the LOT account in the linguistic case, is nonetheless happy to acknowledge that: “if the argument from linguistic systematicity to classical representations is sound, then so is an argument from visual systematicity to non-classical representation” (Davis 2005, 403). More importantly, Fodor's recently articulated view (Fodor 2008) that the compositionality of iconic representation must be differentiated from the compositionality of discursive or linguistic representation, apparently accords well with the connectionist contention that the mental representation involved in domains other than that of language must be of a distinctive, nonlinguistic kind. Iconic representation, which seems to be ubiquitous in human perception, is, according to Fodor, "a homogeneous kind of symbol from both the syntactic and the semantic point of view" (Fodor 2008, 174). Fodor seems to accept, then, the idea that at least some mental representation (specifically, in visual cases) is structurally different from mental representation in the linguistic case: the latter but not the former seems to require canonical constituents with distinctive semantic-plus-syntactic properties.[4] To mention yet another significant example, McLaughlin (2009), while defending the classical systematicity challenge from considerations raised by Cummins et al., does not think it convenient to critically respond to the sCC. This certainly suggests that either he does not know how to respond to it or else he is ready to accept its conclusion.

In my view, and contrary to the suggestions and explicit claims in the literature, the sCC can certainly be met. Indeed, it can be shown that, once we examine the main tenets of LOT, no real challenge for the defender of LOT can be derived from reflection on the existence of different cognitive domains or hence of structurally different systematicity phenomena. In the remainder of this paper I will be concerned with showing just that. A first step in doing so will be to present some textual evidence to the effect that LOT was meant to be, right from the start, a model of cognitive architecture to be applied *generally* in explanations of systematicity, and not only in explanations of *linguistic* systematicity.

II. SOME TEXTUAL EVIDENCE

Given the present debate, it may seem novel that LOT is not considered to be a model restricted *only* to linguistic systematicity. As it happens, it was never meant to be so restricted. Right from the start, in Fodor and Pylyshyn's original 1988 argumentation in favour of LOT, they commented that "linguistic capacity is a paradigm of systematic cognition, but it's wildly unlikely that it's the only example. On the contrary, there's every reason to believe that systematicity is a thoroughly pervasive feature of human and infrahuman mentation" (Fodor and Pylyshyn 1988, 37). Fodor and Pylyshyn do not acknowledge the existence of nonlinguistic systematicity just in passing. They use it as one of the points they make against purely connectionist cognitive architecture. The following quotation makes this apparent:

> It is not, however, plausible that only the minds of verbal organisms are systematic. Think what it would mean for this to be the case. It would have to be quite usual to find, for example, animals capable of representing the state of affairs ***a*R*b***, but incapable of representing the state of affairs ***b*R*a***. such animals would be, as it were, ***a*R*b*** sighted but ***b*R*a*** blind since, presumably, the representational capacities of its mind affect not just what an organism can think, but also what it can perceive. In consequence, such animals would be able to learn to respond selectively to ***a*R*b*** situations but quite *un*able to learn to respond selectively to ***b*R*a*** situations. (so that, though you could teach the creature to choose the picture with the square larger than the triangle, you couldn't for the life of you teach it to choose the picture with the triangle larger than the square.) . . .

That infraverbal cognition is pretty generally systematic seems, in short, to be about as secure as any empirical premise in this area can be. And, as we've just seen, it's a premise from which the inadequacy of Connectionist models as cognitive theories follows quite straightforwardly. (Fodor and Pylyshyn 1988, 40–41, emphasis theirs)

This is certainly surprising, since the proponents of the sCC simply seem to ignore the fact that, from the very beginning, nonlinguistic systematicity was specifically accepted by defenders of LOT. Indeed, nonlinguistic systematicity was taken to be further evidence against purely connectionist models of cognitive architecture. The present point does not apply exclusively to Fodor and Pylyshyn's original defence of LOT. Over the years, many authors have discussed the systematicity argument in favour of LOT. In no development that I am aware of is the systematicity in question meant to be only linguistic or only apparent in linguistic abilities; and what is certainly more important, some of them have considered definitions of systematicity that *explicitly* include nonlinguistic systematicity; in particular, systematicity exhibited in (human and infrahuman) perception. Thus, for instance, McLaughlin (1993) presents systematicity phenomena as capacities to have intentional states. Explicitly belonging to such a class of capacities we find "the capacity to see a visual stimulus as a square above a triangle *and* the capacity to see a visual stimulus as a triangle above a square" (McLaughlin 1993, 167, emphasis his).[5] More recently, although it is true that McLaughlin (2009) does not address (or even mention) the sCC, he remarks that he sides with Fodor and Pylyshyn "in thinking that having propositional attitudes such as belief, desire, intention, and the like by no means requires the ability to use natural language. ...... I think that pre-verbal children, deaf people that have never learned to sign, as well as primates, for instance, will non-vacuously satisfy various systematicity laws" (McLaughlin 2009, 259–260). Let me mention another example in this regard. In the course of an elucidation of the notion of systematicity, García-Carpintero (1995) also points out that, although he focuses on public-language linguistic abilities, "these are neither the only, nor even the core cases. The core cases involve abilities which are required to have public-language linguistic abilities." He goes on to add: "For instance, if an organism is able to perceive a red sphere one meter in front of him, he will typically be able to perceive also a green sphere of the same size one meter in front of him" (García-Carpintero 1995, 373n4).

Finally, if we focus on visual representation, we should stress the work carried out by Zenon Pylyshyn over the last thirty years or so. This work makes it apparent that one of the crucial questions when considering the structure of representations underlying visual perception, mental imagery and visual reasoning is precisely whether that structure is relevantly dissimilar from the structure of purely linguistic representation. In Pylyshyn's work, it is acknowledged that "any form of reasoning, including reasoning by visualizing, must meet the constraints of productivity and systematicity" and precisely this is taken to lead "to the inevitable conclusion that reasoning with mental imagery or reasoning by visualizing or 'visual thinking' requires a combinatorial system—a language of thought—that itself is not in any sense 'pictorial'" (Pylyshyn 2003, xv). More generally, Pylyshyn vigorously defends (e.g., Pylyshyn 2003 and 2007) that the idea of visual representation being depictive or picture-like is anchored in the (intuitively appealing) intentional fallacy of attributing to a mental representation the properties of what it represents. On reflection, however, the essentially symbolic nature of visual representation can naturally be brought into view.

> I often feel I have a vivid image of someone's face, but when asked whether the person wears glasses, I find that my image is silent on that question: it neither has nor lacks glasses, much as the blind spot neither provides information about the relevant portion of the visual field nor does it contain the information that something is missing. You might note that *sentences* (and other languagelike compositional encoding systems) have this sort of content indeterminacy, whereas pictures do not. You can say things in a language (including any language of thought) that fails to make certain commitments that any picture would have to make (e.g., your sentence can assert that *A* and *B* are beside one another while failing to say which is to right or left). (Pylyshyn 2003, 34)

The long and short of it is this: the sCC seems to completely disregard a series of implicit and explicit considerations made on the LOT side of this debate. Not only is there a fact of the matter about whether LOT was originally meant to apply to nonlinguistic domains, but there is certainly also a fact of the matter about whether representation involved in nonlinguistic domains is relevantly *non*linguistic. It is true that the evidence just presented does not amount to an argument to the effect that the sCC is unsound, but it clearly suggests that there is something deeply wrong with it; and the connectionist contender, qua proponent of the sCC, should be embarrassed by this shortcoming. Perhaps defenders of connectionism are right that LOT cannot account for nonlinguistic systematicity, but what is not acceptable is that they do not even mention that, prima facie, LOT purportedly *does* apply to nonlinguistic systematicity. Nonlinguistic systematicity may be news to some involved in this debate, but it certainly is not for the proponent of LOT cognitive architecture. It is true, nonetheless, that the pro-LOT literature has not been rigorous in illustrating the systematicity argument in favour of LOT for the case of nonlinguistic domains. In the majority of cases, as in what we have just seen, LOT explanations of nonlinguistic systematicity are merely sketched, rather simply. Furthermore, several questions need to be spelled out in relation to the "linguistic" nature of LOT architecture. To fill that gap, in the next section I present a careful articulation of LOT accounts for the case of systematicity in visual perception of objects and, in section IV, a series of objections that can, prima facie, be raised against that argument. The upshot will be that the sCC can certainly be met. More importantly, we will see that the sCC manifests a deep misunderstanding of the classical systematicity challenge.

III. YET ANOTHER ARGUMENT: FROM NONLINGUISTIC SYSTEMATICITY TO LOT

If we look for a cognitive field that exhibits nonlinguistic systematicity, the most promising and most studied candidate is, by far, vision. Needless to say, there is currently no definite computational theory of vision and, correspondingly, no definite account of the nature of visual representation. Nonetheless, we can confidently claim to know: (1) the kind of phenomena that we must take into account when considering systematicity as exhibited in visual perception and, more specifically, in visual recognition of objects; and (2) what computational possibilities may account for the systematicity of visual representation. I will address these two points in turn.

First, when we speak of systematicity in visual perception we refer to such facts as: any subject capable of seeing or perceiving a blue square and a red circle must, in point of nomological necessity, be able to see or perceive a red square and a blue circle. similarly, subjects are nomologically systematic in that, for any given pair of objects *a* and *b*, if they are capable of perceiving *a* being to the left of *b*, then they must be equally capable of perceiving *b* being to the left of *a*. These examples are entirely analogous to those Cummins mentions:

> sensitivity to nonlinguistic systematicities appears to abound in human psychology, however. Consider, for example, the perception of objects in space: anyone who can see (imagine) a scene involving objects o_1 and o_2 can see (imagine) a scene in which their locations are switched. (Cummins 1996, 604)

In effect, and contrary to what the developments followed by Cummins et al. suggest, systematicity phenomena have been widely discussed in cognitive research. More precisely, research models of human vision have long involved the study of systematicity phenomena as manifested in abilities to recognize objects: (i) from different viewpoints—and hence from different retinal inputs—(e.g., Hummel and Biederman 1992; Marr 1982; Ullman 1989); (ii) from different shapes—such as typically, in generalizations over classes of objects—(e.g., Edelman and Intrator 2003; Hummel 2001; Ullman 1998); and finally, (iii) in the perception of structurally similar objects (e.g., Edelman 1998; Edelman and Intrator 2003; Hummel 2000; Taylor and Hummel 2009). While disagreeing on the representational alternatives that would account for these systematicity phenomena in vision, cognitive researchers generally assume that these phenomena are an important feature of human and infrahuman vision and must therefore be explained by any serious cognitive model of visual perception.

Secondly, what are the computational possibilities that would explain visual systematicity?[6] The basic problem for an account of systematicity in vision is *how to represent the structure of objects* in a visual scene, and specifically, the objects themselves, the parts of those objects and their spatial relations. This is (part of) what, in the cognitive literature, has been called the binding problem—i.e., the problem of how the different features of objects are bound together in our experiences (Treisman 1996). For instance, the visual representation of *a circle above a square* involves the identification of two objects (or two parts of an object) standing in a particular relation to each other: *a circle, a square* and the relation of *being above*.[7] According to a certain proposal regarding systematicity, the binding problem requires, for each perceptual object, the representation of a *structural description* of the object that explicitly represents the relevant parts and the categorical spatial relations into which they enter. The most celebrated elaboration of this model is Biederman's (1987) *recognition by components*. Using Hummel's terminology (e.g., Hummel 2001), the kind of binding involved in this proposal is dynamic binding: the representational units corresponding to an object or part of an object—viz., a circle or a square—are combined in different ways in accordance with the relevant relation—viz., the relation of being above. According to an alternative proposal, the different configurations into which objects or parts of objects can be arranged correspond to different, explicitly represented, view-like or image-like representational units in which spatial relations are not explicitly represented and in which structural similarity is understood in terms of shared view-like properties; in the example, one unit would represent a circle above a square and another would represent a square above a circle so that the structural similarity of both representations is a function of basic shared shape features. The kind of binding involved in this alternative proposal is static binding: objects or parts of objects are bound via representations of whole figures (e.g., Edelman 1998; Riesenhuber and Poggio 1999; Ullman 1998).[8] still others have proposed a hybrid model with both dynamic and static binding so that both the effect of attention—in dynamic binding, which takes place over a period of time—and the automatic character of object recognition—in rapid, static binding—are accounted for (e.g., Hummel 2001; Thoma, Hummel, and Davidoff 2004).

We do not need to go into the details of each of these representational possibilities. In their roughest outlines, they all consist of neural networks that house distributed representation and perform a function that has stimulus in the receptive field as its input and delivers a representation of the relevant object as its output. What matters for present purposes is that, from a computational point of view, systematicity phenomena in vision require the postulation of explicit representation of the structure of objects. Thus, as I am about to argue, a systematicity argument in favour of LOT can also be formulated for the nonlinguistic case, which, in the present context, can be considered as a relatively new argument in favour of LOT from systematic visual abilities.

Here is the argument. There are certain perceptual capacities that are systematic. We can neutrally describe these capacities at the level of the function that is computed (Marr's level 1), that is, by simply specifying what subjects actually do. As mentioned, perception is known to be systematic regarding (at least) three kinds of capacities: (i) capacities about perception of objects from different and novel viewpoints (viewpoint invariance); (ii) perception of objects consisting of different shapes as being the same kind of object (shape generalization); and (iii) perception of structurally similar (and arguably of the structural similarity of) objects. As regards (i), subjects are systematic in that when they have the capacity to perceive an object, *a*, from viewpoint, *v*, and retinal image, *r*, then they must also be capable of perceiving that object from a different point of view, *v′*, and retinal image, *r′*. Meanwhile, (ii) is instantiated in the capacity to perceive an object, *a*, with shape *s* as belonging to class *c*, and the capacity to perceive a different object, *b*, with shape *s′* as belonging to the same class, *c*. An instance of (iii), to use Hummel's (2000) clear example, is that any subject capable of perceiving two objects, o_1 and o_2, with two respective sizes, s_1 and s_2, and being in a certain spatial relation *R*,—such as in the perception of a small circle atop a large square—must also be capable of perceiving those objects, o_1 and o_2, with their sizes and their role in the spatial relation *R* switched—such as in the perception of a small square atop a large circle. Now, it is crucial that the explanation of these systematic capacities, still at a high level of description (Marr's level 1), comes from the fact that a subject's capacity to perceive *x* involves the subject's identification of certain structure of *x*. Thus, it is because subjects identify or are sensitive to a particular structure of the perceived object *x*—a structure that remains constant over different views of *x*—that they show visual systematicity regarding (i). similarly, it is because subjects can identify a particular structure of a given class of objects, that they can show systematicity in relation to (ii). Finally, it is because subjects can structurally identify objects, properties of those objects, and spatial relations, that they show systematicity with respect to (iii). systematic abilities in the perception of an object *x*, therefore involve a specific structure of *x* in their explanation.

Now, as in the case of linguistic systematicity, the exercise of systematic visual abilities is understood as involving, at the algorithmic level (Marr's level 2), the corresponding kind of (visual) representation. The key point is that the capacity to visually represent involved in systematic visual abilities—whatever such visual representation consists of exactly—must itself be systematic. For instance, the capacity to visually represent object *a* from retinal image *r* must involve the capacity to visually represent that very object from retinal image *r′*. similarly, the capacity to visually represent a small circle atop a large square must involve the capacity to visually represent a small square atop a large circle. Now, the explanation of these systematic representational capacities—at the algorithmic level—must come from the fact that visual representation is explicitly structured. The visual representation of *x*, when it figures in the explanation of a systematic capacity, inevitably leads to the postulation of an explicit partition of *x*, the resulting parts of which can be combined in accordance with certain (psychological) rules.[9] Thus, for instance, the reason why the visual representation of *x* is viewpoint invariant is that it consists of the explicit representation of specific parts of *x*—usually a particular set of shape features such as contours, edges or surfaces—which are semantically evaluable—since they stand for parts of the real object—and are also arranged in the visual representation in a particular way—so that they result in a structural description or specification of the target object. viewpoint invariance therefore takes place because different views of and retinal inputs corresponding to a single object are computationally interpreted as being different views and retinal inputs *of the same structural parts* of the target object. something similar happens in the case of the systematic visual representation of a small circle atop a large square. Different semantically evaluable parts of the representation—for instance, the representational parts corresponding to *circle*, *square*, the relation of *being larger/smaller than* and the relation of *being above*—are explicitly represented and arranged in a particular way. The systematicity involved here is explained because different combinations of the parts explicitly represented in the visual representation can be constructed.

Ultimately then, at the level of implementation—that is, at Marr's level 3—the algorithmic identification of explicit structure leads to the postulation of rule-governed and semantically-evaluable representational posits that are physically instantiated (plausibly in the brain). More generally, this amounts to postulating a physical system of representations with rule-governed and semantically significant constituents: a system of neurophysiologic symbols with both semantic properties and properties that define their compositional nature.

The interpretation of visual representation as involving a system of symbols of the (LOT-) sort described should not be taken to be merely a loose way of talking. It is, from time to time, explicitly acknowledged in the literature on visual cognition. For instance, this is how Hummel analyses his own "circle/square" example from an algorithmic point of view in terms of distributed representations of objects and spatial relations:

> For example, each vector might code a separate set of bindings, with one vector binding *circle* to *above-agent* and *larger-patient* and another binding *square* to *above-patient* and *larger-agent* Augmented with a means for dynamically binding elements to relations (e.g., through synchrony of firing; Hummel and Biederman, 1992), a single vector can be part of a larger compositional structure (e.g., in the way that *circle* is part of *above (circle, square)*). To represent relational structures—that is, explicit relations—a representational system must be able to dynamically bind relational roles to their fillers. . . . The capacity to bind simple elements into relational structures is the hallmark of a structured, *symbolic* representation. (Hummel 2000, 165, emphasis his)

That is, once you have vectors corresponding to the elements (fillers) and relations (roles) of the constituents of the visual representation *small circle atop large square*, you can then easily combine them via suitable (dynamic) bindings of vectors. The possible set of combinations of the visual representation is therefore computationally tractable as a possible set of bindings of vectors corresponding to elements (fillers)—which in the example stand for *square* and *circle*—and vectors corresponding to explicit relational structures (roles)—which in the example represent the two-place relations *being larger than* and *being above*. As is apparent, the result is a full-blooded compositional system of symbols at the algorithmic level.

Hummel's algorithmic proposal specifically involves dynamic binding because of the need to explicitly represent spatial relations. However, it is important to stress that other algorithmic possibilities are also appropriate for a pro-LOT argument as long as they *explicitly* represent a certain structure of the representation. For instance, Edelman and Intrator (2003) offer a model with static binding and without explicit representation of relational structures. In their model, explicit representation of the parts of an object (in their experiment, digits situated above or below each other) is achieved through representational (image-like or view-like) units for fragments of the receptive field, which code shape and location (*what + where* units). These representational units explicitly represent the parts of the object, and allow us to evaluate systematicity (although not the spatial relations of the object via dynamic binding). For instance, in the case of the visual representation *eight above seven*, the visual representation explicitly represents *eight* and *seven* at a particular location—there is a view-like representational unit for the *eight* and the *seven* shape—but not the relation of being above.[10] Of course, dynamically bound representations differ greatly, from a theoretical point of view, from statically bound representations (and for that matter, from a given hybrid kind of representation). The present point is that, whatever exactly the kind of representation at stake, *any* theory must postulate a system of explicitly structured representations if it is to account for systematicity at all. In this respect, dynamically bound and statically bound representations are both clearly LOT representations: they only differ regarding the amount of structure explicitly represented.

In summary, the representation of structure in the explanation of visual systematicity is not alien, in fact it is wholly germane, to the representation of structure in linguistic systematicity. Language and vision are akin in the fundamental respect: systematicity demands structure, and the only way that such a demand can be fulfilled is by postulating representations that are themselves structured: symbols with both semantic and formal (syntactic or psychological) properties. Note that the steps followed in the nonlinguistic case are entirely analogous to those followed in the linguistic case:

(1) sensitivity to structure s (of perceptual objects) explains a particular systematic (perceptual) behaviour B at a high level of description.

(2) structured (visual) representation R explains sensitivity to structure s (of perceptual objects) at the algorithmic level.

(3) system of symbols ss explains structured (visual) representation R at the implementational level.

Nonlinguistic systematicity in vision therefore offers another argument for a "nonlinguistic" or visual LOT. This means that the sCC really has no force at all. Indeed, it demonstrates that the sCC is founded on a considerable misunderstanding of the fundamental tenets of LOT.

IV. SOME OBJECTIONS

so much for the presentation of the nonlinguistic argument for a LOT. The conclusion of the argument is substantive. If sound, it leads to the renewal of the classical systematicity challenge to nonlinguistic visual domains. To some authors that conclusion may seem hard to accept. Uneasiness can here take the form of, at least, three possible objections:

A. OBJECTION 1: NOT ONLY BASIC GEOMETRICAL OR SPATIAL INFORMATION

The first objection may proceed as follows. Paradigmatic instances of visual systematicity must rely on the capacity to represent properties of objects that are far from being merely spatial properties. For instance, human and infrahuman cognition must be capable of dealing systematically with visual properties of real objects such as colour, texture or motion. since the argument provided focuses only on the representation of spatial or shape structure, the conclusion that LOT representation explains visual systematicity is unwarranted.

I certainly agree that the notion of structure that is relevant in visual systematicity phenomena must be taken to include a large number of visual properties of objects that are not merely spatial or shape properties. However, the argument as I have put it forward has the resources to apply to any such nonspatial properties. For instance, the explicit representation of blueness in the capacity to visually perceive a blue square must be a component of the structure of the representation that is inherent to that perception, if that is a systematic capacity regarding the visual perception of blue objects. The incorrect impression that nonspatial properties—such as colour—are neglected in the present account comes from the fact that computational studies on the perception of objects usually (although by no means always) explore experimental settings that concentrate on shape or spatial information, leaving aside other features of real objects. In experimental research this is justified by the assumption that the representation of spatial information for the purposes of object identification is more basic. This assumption is very much the received view in the literature (e.g., Marr 1982; Biederman 1987; Hummel and Biederman 1992; Hummel 2001; Edelman and Intrator 2003) even if some authors have stressed in this context the importance of other kinds of information, such as information about stereoscopic depth (Burke 2005), the viewer's changing position (simons, Wang, and Roddenberry 2002), or the temporal correlation between object views (Wallis, Backus, Langer, Huebner, and Bülthoff 2009). The suggestion is, however, that all sorts of structural visual information must be explicitly represented if it is relevant in explanations of systematicity phenomena.

B. OBJECTION 2: REPRESENTATION OF SPECIFICALLY VISUAL INFORMATION

The second objection I would like to consider makes a case that visual representation involves information specific to the visual modality. systematic visual abilities must be abilities that involve representation of that information. But representation in a language cannot entirely account for all the visual information. The conclusion is therefore that visual representation cannot be LOT representation.

Nobody would deny that visual representation involves, as a matter of fact, modally specific information. For instance, Zenon Pylyshyn, defends a version of LOT as applied to the visual case according to which "there is no inner replica or picture of any kind in our head, neither literally nor in any other nonvacuous sense capable of explaining how we represent spatial information in perception and thought" (Pylyshyn 2007, 120). He nonetheless accurately delineates the features that are distinctive of spatial representation, such as the representation of 3D information, the continuity of space, or its links to sensory-motor engagement (cf. Pylyshyn 2007, chap. 5). Indeed, representation specificity is only to be expected granted a framework that conceives the mind as a modular system (Fodor 1983). Language comprehension and visual perception are perhaps the clearest examples of functionally distinct encapsulated systems and, to that extent, representation specificity is acknowledged upfront within classical models. Therefore, LOT must be thought of as being compatible with a wide variety of distinct, modality-specific representations.[11]

In effect, and contrary to positions such as Pylyshyn's, I wish to acknowledge representation specificity so as to include image-like, or analogue or sensory kinds of representation within LOT. I am not the first to accept, in general, that there may be explanatorily relevant iconic or imagistic representation in visual cognition that is consistent with LOT. Fodor's recent (Fodor 2008) and not so recent (Fodor 1981) work is a case in point (for discussion on the tension between Pylyshyn's and Fodor's views see Verdejo (forthcoming)). The rather new claim that I defend is that, in explanations of systematicity, even if visual representation is image-like or analogue, it must explicitly represent some structure so as to involve a system of symbols such that constituents can be combined in different ways. In other words, I defend that, in explanations of systematicity, even if visual representation is image-like or analogue it is representation in a LOT. The exact sense in which this claim has to be understood is spelled out in the response to the third objection.

C. OBJECTION 3: IMAGE-LIKE REPRESENTATION

This objection stresses that the postulation of image-like or view-like representations is central to most of the developments in visual cognition. This is particularly true in accounts of systematic visual abilities based on static binding or otherwise some process of image-matching or view-matching (e.g., Edelman 1998; Edelman and Intrator 2003; Riesenhuber and Poggio 1999; Ullman 1998). However, if representation is image-like or view-like, it follows that it cannot be part of a language: therefore, visual systematicity is not accounted for in terms of LOT.

The third objection is perhaps the most important for the present account and thus I will offer a detailed reply. Nonetheless, I would be content if my considerations simply showed that LOT is applicable in all the models that employ dynamic binding. Arguably, dynamic binding models are those that best fit the account of systematicity phenomena in general (see Hummel 2000 for a defense of this view). However, a broader conclusion is certainly plausible. Thus, I intend to show that, even if visual representation is view-like or image-like, the only way in which visual representations can figure in the explanation of systematicity phenomena is via the explicit representation of structure that is characteristic of a LOT, or in other words, via the postulation of semantically evaluable parts of the representation that can be combined in different ways. According to the view I recommend there *is* a crucial difference between dynamic binding models and static binding models in the account of systematicity phenomena: they differ in the kind of explicit representation of structure they postulate. However, both types of model involve a commitment to the explicit representation of structure, at least when dealing with systematicity explanations. It is uncontroversial that dynamic binding accounts, which in standard developments (e.g., Biederman 1987; Marr 1982; Ullman 1989) result in 3D object-centered structural descriptions, do explicitly represent structure in a system of combinatorial symbols. However, I want to argue that explicit representation of structure is also unavoidable in models that exclusively employ static binding or some process of image-matching or view-matching, if these models are to account for systematicity phenomena as well.

When dealing with systematicity phenomena (viz., viewpoint invariance, object generalization or object similarity), the basic computational procedure of image-like or static binding models consists of representing an object via a series of views. These views can be stored and retrieved when confronted with new views of objects so that the system can determine whether a new view corresponds to a previously encountered object, or class of object, or to an object similar to one encountered before. We should not be misled however, by the fact that these models use view-like or image-like representations. view-like or image-like representation does not mean, in this context, *unstructured* representation. Unsurprisingly, the paradigmatic way in which these models work is by making the representation sensitive to the shape of the object, or more precisely, to an abstract feature of the shape of the object—such as contours, edges or defining surfaces—as described in a parametric or functional space. My key point here is precisely that abstract structural features of the object (e.g., its shape) must be explicitly represented when systematicity is at stake, even in view-like or static binding models.

so, for instance, Ullman (1998) presents a model in which viewpoint invariance and object generalization is obtained in terms of (linear) combinations of object views. In Ullman's computational model, a set of 2D object views is originally taken to represent the object, and new views must then be evaluated as being or not being instances of the object represented. The system classifies the new view as being of the same object if the new view corresponds to a linear combination of the original set of views. It is important to note that 2D object views consist here of coordinates that define points that stand for the object's contours or similar basic shape features. Now we can see that talk of explicit representation of structure in Ullman's model is not a loose way of talking. The system does indeed explicitly represent the structure of the object in terms of coordinates that represent (points of) the shape of the object. The existence of structure here literally means the postulation of different explicitly represented (semantically evaluable) parts in the representation—different coordinates for points defining the shape of the object—which can be arranged in a number of different ways—in terms of different linear combinations of the coordinates that originally represent the object. True, Ullman's model differs greatly from models that construct an explicit 3D representation of the parts of the object and their spatial relations. The important point however, is that Ullman's postulated visual representation for dealing with systematicity in viewpoint invariance and class generalization "contains detailed information about the object's structure" (Ullman 1998, 41).

Another approach based upon image-like or analogue representation is Edelman's representation of similarity (Edelman 1998). In very rough terms, Edelman's approach delineates a computational mapping from distal shape to similarity points in a functional shape space. object recognition is therefore computationally interpreted as consisting of the proper location of a given stimulus in the proximal shape space. similarly, object categorization is achieved via a process of location of a stimulus in a demarcated region of the proximal shape space. Admittedly, the visual representation postulated by Edelman lacks explicit representation of structure in the way required by dynamic binding models. Nonetheless, the system explicitly represents the structural properties of shape of single objects: different points in the shape space explicitly represent different shape-structures of objects. This is clearly the case if we take into account the way in which systematic viewpoint invariance is accounted for and the problem of systematic object generalization is resolved in the model. Invariance to viewpoint is achieved by isolating components extraneous to the defining shape of the object—such as object orientation with respect to the observer or illumination—and removing them from the final representation in the shape space. Meanwhile, object categorization is achieved through defining a shape class of structurally similar objects in a region of the shape space. of course, this involves explicit representation of the structure of objects via abstract features of a basic shape dimension.[12] Again then, the existence of structure here literally means the postulation of different explicitly represented (and semantically evaluable) parts in the representation—different object-defining shape features—which can be arranged in a number of different ways—so as to define a particular location or region in the shape space.

From a more general point of view, the theoretical possibility of imagistic or somehow sensory representations being part of a combinatorial system of symbols is brought into view by Barsalou's forceful defence of perceptual symbol systems (Barsalou 1999). According to Barsalou, a perceptual symbol is a neural representation physically instantiated in sensory-motor areas of the brain which represents sensory properties of entities and events. Needless to say, visual sensory properties of an image-like or analogue sort are just such a case. Now, according to Barsalou, one of the most important features of perceptual symbols is their schematic character. schematicity is meant here to avoid perceptual symbols being interpreted as completely holistic or unstructured. A perceptual symbol is not just a piece of formless perceptual experience, but the result of a selection process by means of which a particular schematic aspect of sensory experience is represented. Unsurprisingly, schematicity is what allows Barsalou's perceptual symbols to accommodate productivity (and hence systematicity). With the aid of a simulation process, the system can combine different schematic aspects of perceptual experience. "Thus, schematicity makes productivity possible. For example, if a perceptual symbol for *ball* only represents its shape schematically after color and texture have been filtered out, then information about color and texture can later be added productively. For example, the simulation of a *ball* could evolve into a *blue ball* or a *smooth yellow ball*." (Barsalou 1999, 593) Barsalou's schematicity is therefore another way of considering, from a more general point of view, the presence of structure in image-like representation.

It seems to me that the analysis presented here shows that objection 3 is certainly wrong-headed. The lesson should not be very surprising. It is an innocent truism that images or pictures are not really words or sentences. It does not follow, however, that pictures cannot implement a symbolic system. Indeed, images or pictures can structurally represent all the relevant information needed to account for visual systematicity. It is important to note that the distinction between image-like and linguistic representation is *not* the distinction between unstructured and structured representation. Images, as we have seen, can have structural parts, whereas languages can be entirely constituted by unstructured sentences—as Evans showed in a celebrated example (Evans 1981). If the structure present in image-like or analogue representation is not the kind of structure authors expect from LOT representation, the problem is surely that they expect LOT to be structurally analogous to natural language systems. In the next section I offer some considerations that show that this expectation is unwarranted.

V. WHAT IS WRONG WITH THE SYSTEMATICITY CHALLENGE CHALLENGE

Without a doubt, there is a deep misconception of LOT behind the formulation of the sCC that crystallizes vividly in a single assumption. This is the assumption that it is essential for the structure of LOT representation that it be sufficiently similar to, indeed entirely analogous to, the structure of natural language (usually English). In order to see that this is actually the key assumption in Cummins et al.'s developments we only need to go back to the way they characterize classical or LOT representation in the explanation of systematicity phenomena. As advanced, Cummins et al. characterize classical representation as being a structural representation of the linguistic domain. The idea of being a structural representation of the linguistic domain is that "in order to be sensitive to the combinatorial syntax and associated semantics of a language, there must be a system of internal representations that has the same (or corresponding) syntactic and semantic features" (Cummins, Blackmon, Byrd, Poirier, Roth, and schwarz 2001, 180). since Cummins et al. focus on the explanation of typical systematicity cases of the 'John loves Mary'/'Mary loves John' sort, the target syntax and semantics of the language must correspond roughly to those found in basic grammatical characterizations of the associated natural language sentences. Now, a paradigm case of a system of internal representations that has "the same (or corresponding) syntactic and semantic features" as those found in basic grammatical descriptions of natural language sentences would then be a programming language, such as the List Processing programming language (LIsP). For instance, in LIsP, it is the case that there is just one algorithmic counterpart, call it *Mary*, corresponding to the natural language word 'Mary,' which is explicitly encoded in a lexical entry as meaning Mary and as being of syntactic category NP. In LIsP, explicit rules of the grammar that define well-formed formulas—that state, say that [NP[VP[NP]]] is well-formed—and all lexical units are also explicitly defined—such as the lexical unit *John* corresponding to the word 'John' and the lexical unit *love* corresponding to the word 'love.' All this makes possible a systematic parse-dependent function for which, if the system registers *Mary loves John*, it can—via *car*, *cdr* and *cons* operations—construct the systematic counterpart *John loves Mary*. since, on Cummins et al.'s account, LIsP is a paradigm case of LOT representation, it is suggested that the intuitive advantage of LOT representation over connectionist representation is the appearance of principleness or naturalness that comes from the similarity between structures in formulas of the programming language and in natural language sentences.

Thus, Cummins et al.'s reading of LOT results in a fundamental assumption that is in effect twofold: not only must the structure of LOT representation be shared with the structure of a given cognitive domain, in addition, the cognitive domain in question must be (roughly) natural language sentences. It follows quite naturally from this (double) assumption that we only need to find a differently structured domain, a nonlinguistic domain, in order to defeat the proponents of LOT cognitive architecture. such a strategy would indeed be promising if only the assumption were true. As it happens, and as I will now try to show, this assumption marks a highly implausible and utterly unnecessarily strong reading of the demand that LOT representation must be "linguistic." In what follows, I will argue that a much weaker understanding of the linguistic character of representation is enough for there to be LOT representation.

so, the question I am now concerned with is: how structurally similar to a natural language must a given representation be for it to be LOT representation? The answer is: not very similar at all. The only, albeit substantive, thing that is required for LOT representation is a system of symbols, that is, a system of (algorithmically and ultimately neurophysiologically defined) representations with both semantic and formal properties that allow combination and recombination of their constituents. More specifically, when systematicity is under consideration, the only kind of representation needed is one that accounts for such systematicity in a causally relevant way. Thus, the reason why *Mary* in the thought *Mary loves John* must be explicitly represented is not, as Cummins et al. have suggested, because that would make the representation seem or be principled, or seem or be natural. The reason is that the postulation of *Mary* as a symbol is explanatorily relevant to the systematicity involved in representing a huge series of thoughts: *Mary loves John*, *John loves Mary*, *Mary runs*, *everybody is fond of Mary*, and a long, etcetera. similarly, the reason why *square* must be explicitly represented in the visual perception of *large square atop a small circle* is that the visual representation *square* can figure in the explanation of a subject's capacity to perceptually represent the systematic variants that have *square* as a constituent: *small square atop a large circle*, *red square to the left of a blue circle*, etc. From this, it does not even begin to follow, and it is certainly false, that LOT, as a representational scheme, requires a direct correspondence with the representational categories and structure of any natural language whatsoever. I propose to reflect on the following three points of increasing significance to make this clear.

First, as nobody would deny, LOT does not require explicit representation of all the linguistic information. At the very minimum, it is an old point that LOT need not explicitly represent grammatical rules or whatever rules account for linguistic behaviour. It is true that programming languages, such as LIsP, which have been taken to be paradigmatic examples of LOT representation, do explicitly represent all the grammatical information and in particular, grammatical rules. However, that is not a requirement of LOT:

> [I]t is universally assumed by Connectionists that Classical [that is, LOT] models are committed to claiming that regular behaviours must arise from explicitly encoded rules. But this is simply untrue. Not only is there no reason

> why Classical models are required to be rule-explicit but—as a matter of fact—arguments over which, if any, rules are explicitly mentally represented have raged for decades *within* the Classicist camp. The one thing that Classical theorists do agree about is that it *can't* be that *all* behavioural regularities are determined by explicit rules; at least some of the causal determinants of compliant behaviour *must* be *im*plicit. (Fodor and Pylyshyn 1988, 60, emphasis theirs)

In general, the question of which information, if any, is explicitly represented is certainly an empirical matter for which natural language can indeed be a very misleading guide. For instance, there are lots of examples of natural language idioms whose structure is almost certainly not explicitly represented in thought. Furthermore, it is to be expected that the postulation of any amount of explicit representation must be fully accepted only after empirical confirmation. This is so independently of what exactly the preferred notion of explicit representation turns out to be.[13] For instance, it is an empirical possibility that the structure of such thoughts as *Mary loves John* or *John loves Mary* involves the explicit representation of only some of the corresponding constituents in natural language. For instance, the two-place relation denoted by the word 'loves' in 'Mary loves John' could be implicitly encoded via some hardware specific to the system. To be sure, the conceivable possibilities here are countless. Of course, systematicity phenomena impose serious constraints on those possibilities. We should not however, rule out the existence of LOT representation if such representation does not fit exactly with the representation we can expect from usual grammatical analysis in natural or public language. After all, grammatical structure of natural or public language *is* apt for empirical confirmation. Why should we not think of the structure of LOT representation in the same way?

Let me go on to the second sort of considerations. The key point now is that a minimally plausible conception of LOT must ensure that it is consistent with a notion of constituency that is not the kind of spatio-temporal constituency found in natural or public languages. Thus, adjacent representations in natural language are spatially (in written language) or temporally (in spoken language) adjacent. However, the relevant kind of compositional adjacency that LOT demands is just functional and certainly not spatio-temporal. The basic idea of functional adjacency or functional compositionality—a notion vindicated by illustrious defenders of connectionist models (see van Gelder 1990)—comes from reflection on the proper relation between spatio-temporal adjacency relations found at a high level of description and lower-level implementations of that relation. All that is required for there to be adjacency relations at low levels is a functional or causally appropriate relation that mimics the one found at a high level. For instance, if *A&B* is identified as a complex representation that has to be algorithmically and physically implemented, it would be required that a token of the algorithmic and physical states corresponding to *A&B* be causally linked to tokens of the algorithmic and physical states corresponding to their constituents (say, *A*, *&* and *B*). Consistent with this, algorithmic and physical adjacency can be instantiated between algorithmic and physical states that are far from being spatio-temporally adjacent. This point has been emphasized by several authors (e.g., García-Carpintero 1996; Aydede 1997) and it should be clear right from the start if only we pay sufficient attention to Fodor and Pylyshyn's remarks about the physical realization of functional adjacency:

> For example, conventional [that is, LOT] architecture requires that there be distinct symbolic expressions for each state of affairs that it can represent. since such expressions often have a structure consisting of concatenated parts, the adjacency relation must be instantiated by *some* physical relation when the architecture is implemented. However, since the relation to be physically realized is *functional* adjacency, there is no necessity that physical instantiations of adjacent symbols be *spatially* adjacent. similarly, although complex expressions are made out of atomic elements, and the distinction between atomic and complex symbols must somehow be physically instantiated, there is no necessity that a token of an atomic symbol be assigned a smaller region in space than a token of a complex symbol; even a token of a complex symbol of which it is a constituent. (Fodor and Pylyshyn 1988, 57, emphasis theirs)

Finally, I would like to focus on a point that perhaps is not emphasized in the literature (although see García-Carpintero 1995): a minimally charitable reading of LOT must make it possible for a given high-level representation—such as *Mary* in *Mary loves John* or *square* in *large square atop a small circle*—to have more than one possible algorithmic and physical implementation. The reason is not only that such a possibility would respect the classical Marrian doctrine that there exists a one-to-many relationship between levels of explanation. More importantly, it is to be expected that, as is the rule in the special sciences, one and the same high-level category can have more than one low-level (algorithmic and ultimately physical) implementation. This is what is known in philosophical circles as the multiple realizability thesis. Let me elaborate the point.

To take an example from (rather humble) economics, the high-level category *unemployment* is algorithmically defined by a particular formula and physically realized by a certain (unlucky) part of the population in a given economy/country at a given time. However, the fact that in any given particular economy and time, just one algorithmic and physical implementation of *unemployment* must hold, does not mean that *unemployment*, qua scientific category, must on each particular occasion have one and the same algorithmic and physical implementation. Let us consider, for the sake of argument, that, at the high level, the category *unemployment* is individuated by a given unemployment rate. Thus, *5 percent of the labour force* would then be one possible individuation of *unemployment*. For the individuated high-level category *5 percent of the labour force* to have scientific credibility—or otherwise, for it to be a real category according to scientific standards—it is, of course, not necessary for it to be implemented by just one algorithm. As we know, the formula at stake—(unemployed workers/total labour force) x 100—is a statistical formula for which different statistical criteria, and hence, different algorithms, can be given. Granted this, there are indefinitely many rival algorithms that correspond to a given individuated *unemployment* category, for instance, indefinitely many algorithms corresponding to *5 percent of the labour force*. Note further that, even if one particular statistical formula is fixed, there would still be indefinitely many particular algorithmic implementations corresponding to *5 percent of the* *labour force* on different occasions, that is, if the number of unemployed workers and the number of the population constituting the labour force vary in the same proportion on different particular occasions. Now, as regards physical realization, scientific rigour does not require either that *5 percent of the labour force* be always physically realized by a series of particular subjects, let alone by a series of particular types of subject. For all that, it is conceivable that the relevant population that physically realizes *5 percent of the labour force* is constituted of intelligent plastic robots[14] or Martians of weird physical configuration. All that (scientifically respectable) physical implementations of an individuated high-level category require is that they are explanatorily correlated with the individuated high-level category. In summary, an individuated high-level category—viz., *5 percent of the labour force*—only requires first, an algorithmic definition—roughly, a statistical formula which can vary more or less at the politicians' mercy; and second, the existence of an explanatory correlation between the high-level category and its physical realization—say, between unemployment rates and actual people without jobs. In particular, for the category *unemployment* to be scientifically acceptable, it is not necessary that there be only one algorithm for each unemployment rate or that there is just one set of subjects that realizes the category at the physical level. This is so even if it *is* required that, in any given particular case, just one algorithm and just one physical realization holds.

Now, if these points are uncontroversial regarding categories in economics, they should be equally uncontroversial with regard to categories in a postulated LOT. Thus, it should also be on the (LOT) cards that one and the same high-level representation—say *Mary* or *square*—may, on the one hand, have rival algorithmic candidates (in the economics example, different statistical formulas), while, on the other hand, it is expected *within* LOT that, once a certain algorithmic system is fixed, one and the same high-level representation may have more than one algorithmic implementation *on different particular occasions* (in the example, this would be the case if the number of unemployed population and of the total labour force varies in different applications of the *unemployment*-formula with the same result). If we take this into account, we begin to appreciate clearly that for there to be a LOT, it is preposterous to require that, at the algorithmic (and ultimately at the neurophysiological) level there exist a very direct correspondence with, say, words or sentences in a natural language. Perhaps that could be the case if the algorithmic representation consisted of representations in a programming language. However, it certainly is not mandatory. First, consistent with LOT, a given thought or high-level representation can have rival algorithmic implementations: a LIsP formula or a filler/role binding vector or a uniquely factorable number—such as in Gödel numbering. second, once we fix a particular algorithmic system of representation, the same high-level categories—say, *Mary* or *square*—may have (in the case of vectorial or Gödel numbering representations) different (numerical) implementations on different occasions (depending on syntactical or spatial position).[15]

similar considerations hold when dealing with the lowest physical realization level. The physical realization of a given high-level representation—such as *Mary* or *square*—does not require that just one particular physical state is tokened, let alone one particular type of physical state. What is required is that, for any given instantiation of a high-level category in any particular case, just one explanatorily correlated neurophysiological realization is tokened. This parallels the requirement in economics that for any given particular instantiation of the individuated high-level category *5 percent of the labour force*, just one particular set of individuals in the economy is actually without jobs. It does not follow that LOT, as a matter of principle, needs the postulation of a direct correspondence between realizing states and words or sentences in a public or natural language. The same high-level representation may, consistent with a plausible version of LOT, be physically realized in many different neurological structures.

The long and short: what is crucial for LOT is that algorithmic (and ultimately neurophysiological) representations, whatever they consist of exactly, have relevant and definite formal and semantic properties that correlate with high-level explanatory categories; indeed that these representations constitute an algorithmically and neurophysiologically determined system of symbols. In explanations of systematicity, the relevant and definite formal and semantic properties are quite simply those that permit us to account for systematicity phenomena.

The three considerations just elucidated might seem speculative at points. They are certainly not incontestable. Taken together however, they clearly support the conclusion that LOT does not require (although it may allow) mental representation that shares structure with natural language, or a representational scheme structurally analogous to natural or public language. If this is sound, then the basic assumption that supports the sCC is unwarranted.[16]

VI. CONCLUSION

There certainly are many interesting things to be learned about LOT representation and, in particular, about the ways that connectionist models can implement a LOT architecture and improve our explanations of systematicity phenomena. It is a merit of LOT-sceptical approaches such as those considered here that they draw our attention to the problem of the application of LOT to cognitive domains other than language and bring pressure to bear on over-simplistic conceptions of LOT as applied to those domains. In addition, connectionist developments, and research in perceptual psychology in general, have shown that the kind of representation needed to account for systematicity in different cognitive domains may vary across different modalities and, to that extent, be modally specific. Finally, the considerations in this paper are compatible with Cummins et al.'s favoured account of systematicity—in terms of smolensky's tensor-product networks—being correct, as long as it is not claimed that it is a real alternative to LOT models.

That said, in this paper I have argued that the only way in which the application of LOT to nonlinguistic domains is a problem is by considering the "linguistic" nature of LOT in a very implausibly strong sense, one that demands that LOT shares its structure with natural language. A much weaker sense is enough for a LOT to actually be in place, a sense according to which LOT consists merely of the postulation of explicitly structured symbols at high, algorithmic and physical levels (independently of whether LOT shares structure with any cognitive domain whatsoever). My view strongly relies on a pro-LOT argument that focuses on computational studies of visual representation. However, the suggestion is that LOT should also be applicable to systematicity phenomena at large, including systematicity phenomena in other sensory modalities—such as audition, proprioception or olfaction—, the precise nature of which is a field still awaiting exploration.

LOT is certainly not the whole story to be told about systematicity phenomena (at most it is the beginning of such a story). What should be clear so far is that LOT representation need not be interpreted in the strong sense of sharing its structure with natural language. That means, inter alia, that no sound argument can be mounted against LOT structure simply from the consideration of nonlinguistic systematicity. A fortiori, that means in turn that the sCC is as much of a red herring as anything can be. Just as there is no requirement for LOT representation to share its structure with natural language, neither is the application of LOT to nonlinguistic domains a challenge. What we have seen is in effect a renewal of the classical systematicity challenge for the case of nonlinguistic, and specifically visual, representation. This is not to say that all kinds of nonlinguistic representation must be LOT representation. After all, it is certainly *not* the case that each and every cognitive domain must be systematic. However, as things stand, if a given (nonlinguistic) cognitive capacity is systematic, then that capacity must be, at least until someone produces a sound argument to the contrary, accounted for within the representational powers of a LOT.[17]

ENDNOTES

1. Here, as everywhere else in this paper, the expression 'linguistic/nonlinguistic systematicity' means systematicity exhibited in linguistic or nonlinguistic abilities as shown in behaviour. so understood, systematic abilities in nonverbal creatures are paradigmatic instances of nonlinguistic systematicity. As will become apparent, to acknowledge linguistic or nonlinguistic systematicity does not, in and of itself, amount to an acknowledgement of underlying LOT or non-LOT representation, respectively.

2. In Fodor and Pylyshyn 1988 alone, apart from the argument that starts from systematicity, there are also arguments from productivity, inferential coherence and compositionality. Especially salient in this respect are the eight consequences of intentional realism (or "essential mentalism" as he calls it) that Rey presents as eight possible starting points for a LOT argument. see Rey 1991.

3. I must confess that, although the sCC clearly figures among the dialectical resources of Cummins et al., it is not at all clear to me what their privileged line of argument against LOT in explanations of systematicity really is. In particular, and apart from that challenge, it is not clear to me whether they ultimately want to argue that, unlike connectionist explanations, classical explanations are principled and "[p]rincipled explanations are easier to test, but they are no more likely to be true" (Cummins 1996, 608); or that "classical explanation and the tensor-product [connectionist] explanation explain the systematicity of thought equally well" (Cummins, Blackmon, Byrd, Poirier, Roth, and schwarz 2001, 170); or quite differently that "until sT [the systematicity of thought] can be demonstrated, there is no need to explain

it" (Cummins, Blackmon, Byrd, Lee, and Roth 2005, 407). Probably, one might think, a mixture of them all. But note that as the years go by, radicalism seems to be growing, and the difference matters: if those authors' ultimate aim is to cast doubts on systematicity itself, it certainly seems rather pointless to look for any explanations of systematicity whatsoever and, a fortiori, for *connectionist* explanations of systematicity, be it linguistic or otherwise. For the purposes of my discussion, I will only focus on the formulation of the sCC given by Cummins et al., which involves acknowledgement of systematicity phenomena and of the need to respond to the classical systematicity challenge.

4. I do not mean to suggest, of course, that Fodor would agree with the moral of the sCC. However, it seems to me that the only way in which Fodor may resist the sCC is by acknowledging that icons are precisely *not* the kind of representation that can figure in explanations of visual systematicity. The reason is that systematicity certainly demands semantic and syntactic heterogeneity of the representation, while icons are, by definition, semantically and syntactically homogeneous. see verdejo forthcoming for further elaboration of this point.

5. As an anonymous referee points out, McLaughlin seems to be assuming that, even in the perceptual case, systematic capacities are thoroughly conceptual capacities since they "involve possession of the same concepts and faculties" (McLaughlin 1993, 169). That being so, the present point is simply to show that classicist developments have also focused on cases that do not involve, in any obvious sense, linguistic abilities. It is worth emphasizing however that, according to the view that I am defending, there is no commitment to the claim that the application of LOT in systematic nonlinguistic abilities requires the exercise of concepts on the part of the subject or, anyhow, a psychologically real translation of the content of experience into conceptual content (this will be clearer in section III below). see also n7.

6. In my analysis of the computational possibilities that may account for visual systematicity, I focus on models of (visual) information processing mechanisms, that is, on models that suit approaches that posit operations on representations of visual information. I conspicuously ignore those models that do not or may not fit this mould: for instance, so-called enactive, embodied or ecological models. For an overview of these "heterodox views" see Noë and Thompson 2002.

7. It is important to note that, in the literature of visual cognition under consideration, although it mostly focuses on *human* visual cognition, there is no explicit commitment to the "identification of objects" being a conceptual phenomenon. This should be clear if, as one assumes, the capacity to identify or recognize objects, the parts of an object, and their spatial relations, is a capacity that we humans must share with a large number of nonverbal animals. According to the present account, the identification of objects need not (although certainly may) involve classification of the object as belonging to a given conceptual category or the use of a singular-term or demonstrative concept. This is in accordance with the familiar view that information processing systems in vision involve subpersonal states with nonconceptual content. see Bermúdez 1995. Finally, the conceptually-free use of the term 'identification' or 'recognition' in this context is consonant with Pylyshyn's postulation of primitive perceptual indexes (FINsT indexes), whose function is to individuate and keep track of distal objects in a preconceptual way. see Pylyshyn 2007. That said, it is open to a theorist to mount an argument to the effect that, at least in humans, object identification or recognition necessarily involves conceptual capacities. The present point is that nothing in my approach relies on such a line of argument.

8. It bears emphasis that the difference between static binding and dynamic binding is not the difference between object recognition that involves, and that does not involve, structured representation, at least not in explanations of systematicity. The capacity to recognize objects or parts of objects via some process of image-matching or view-matching, where this is a systematic capacity, must always require a certain structure of the representation, even if this representation is image-like, or analogue, or somehow anchored in sensory properties of experience. I will extensively discuss this point in the next section.

9. In the visual case, psychological rules play the role that (psychologically real) grammatical rules play in the case of linguistic abilities. For instance, as regards the systematic capacities of viewpoint invariance, it is known that object perception is not invariant with respect to object rotation in the picture plane—such as upside-down inversion. Thus psychological rules constrain the range of systematic visual abilities just as (psychologically real) grammatical rules constrain the range of linguistic abilities.

10. In the next section I will address the question of how view-like or image-like representations in static binding generally involve explicit representation of structure. For now, the fundamental point is that Edelman and Intrator's model would here be analogous to a linguistic case in which, although an explicit representation of *Mary* and of *John* is postulated in the representation of *Mary loves John*, no such symbol is postulated for the relation *loves*, because it is somehow hardwired into the system. Though perhaps implausible (due to the decrease in representational capacities), it would nonetheless be a possible system of LOT symbols; it is just that the system would not entirely correspond to the one we intuit from natural language. In the visual case, Edelman and Intrator remark that a given spatial relation can be represented by a set of measurement functions of shape patterns at a given location. The set of measurement functions that identifies objects u and v at location t_1 and at location t_2 can be seen as implicitly representing a certain two-place spatial relation (such as the relation of *being above*). To that extent, as Edelman and Intrator stress, their computational model is not compositional *if* we take as our paradigm of compositional relations those found in natural or formal languages—where (spatial) relations *are* explicitly represented. However, the important point is that, even if spatial relations are not explicitly represented, the objects that stand in a particular spatial relation are explicitly represented via measurement functions. Thus, "[t]he role of a measurement function is to provide a *perceptual symbol*. . . , which stands for a particular visual event" (Edelman and Intrator 2003, 104, emphasis theirs). The different ways in which perceptual or view-like symbols can be composed would, in this case, correspond to the different locations at which they can be displayed in a parametric space.

11. As an anonymous referee points out, this result of the propounded account makes it compatible with the postulation of different kinds of representation (e.g., strictly linguistic/logical vis-à-vis image-like or map-like representation). It still does not follow however, that there is something to Cummins et al.'s claim that LOT leads to representational pluralism, that is, to the thesis that "for every differently structured domain for which we show a systematicity effect, we employ a scheme of mental representation that shares structure with that domain" (Cummins, Blackmon, Byrd, Poirier, Roth, and schwarz 2001, 181). In fact, I will argue in the next section that the requirement that LOT must share structure with any domain whatsoever—and in particular with language—is unwarranted. According to my account, there is just one scheme of mental representation (i.e., that of LOT architectures, correctly understood), which is applied to specific domains and therefore exhibits representation-specificity. Nonetheless, the wide choice available within LOT models certainly has limits. For instance, Fodor, when dealing with the nature of perceptual mechanisms in

terms of subsidiary systems that make environmental information available to computational processes, mentions a format-condition, according to which, “if we think of the perceptual mechanisms as analogous to such devices [such subsidiary systems], then we are saying that *what perception must do is to so represent the world as to make it accessible to thought*. The condition on appropriateness of format is by way of emphasizing that not every representation of the world will do for this purpose” (Fodor 1983, 40, emphasis his).

12. Edelman defends the claim that his propounded scheme “can be adapted to attend selectively to different dimensions of variation of the stimuli in several ways” (Edelman 1998, 486) so as, for instance, to also accommodate a colour and texture dimension (not just a shape dimension). That, of course, would add to the point: Edelman’s holistic representation can then be taken to offer structural descriptions of objects across a variety of dimensions.

13. The options in the interpretation of the notion of explicit representation vary between more technical definitions in terms of, for instance, processing time—as in Kirsh 1990—and more intuitive ones based upon ordinary use—see Hadley 1995.

14. The possibility of robotic implementations of social categories should not strike us as surprising, given the increasing interest in the robotic modeling of collective intelligence. In the literature, one can even find preliminary attempts to apply robotics to categories of institutional economics. see for instance silva and Lima 2007.

15. For example, in Gödel numbering, different algorithmic implementations—that is, different numbers—may correspond to the high-level category *Mary*, depending on whether *Mary* is in the subject syntactic position—as in *Mary loves John*—or in the object syntactic position—as in *John loves Mary*—see Cummins, Blackmon, Byrd, Poirier, Roth, and schwarz 2001 for details. It does not even begin to follow that the structure at the algorithmic level is here non-LOT structure. A minimally plausible conception of LOT must allow for high-level representations to have different algorithmic implementations on different occasions, just as, mutatis mutandis, the high-level category *5 percent of the labour force* may have different numerical algorithmic implementations on different occasions—as in e.g., (1,000,000 unemployed workers/20,000,000 total labour force) x 100 and (1,500,000 unemployed workers/30,000,000 total labour force) x 100.

16. Considerations similar to the ones presented in this section are also relevant when critically responding to Cummins et al.’s view that (connectionist) encodings offer a real alternative to LOT representation. see Verdejo 2009, chap. 5.

17. I am sincerely grateful to David Pineda, Daniel Quesada, Martin Roth, Josefa Toribio, Christofer Evans and an anonymous referee for this journal for helpful comments and suggestions on earlier drafts of this paper. support for this work has been provided in part by the spanish government (MICINN: research project FFI2008-06164-C02-02 and CONsOLIDER INGENIO 2010 Program CsD2009-0056), and the Catalan government (consolidated research group GRECC sGR2009-1528).

BIBLIOGRAPHY

Aydede, Murat. 1997. “Language of Thought. The Connectionist Contribution.” *Minds and Machines* 7: 57–101.

Barsalou, Lawrence. 1999. “Perceptual symbol systems.” *Behavioural and Brain Sciences* 22: 577–609.

Bermúdez, José Luis. 1995. “Nonconceptual Content: From Perceptual Experience to subpersonal Computational states.” *Mind and Language* 10: 333–369.

Biederman, Irving. 1987. “Recognition by Components: A Theory of Human Image Understanding.” *Psychological Review* 94: 115–147.

Burke, Darren. 2005. “Combining Disparate Views of Objects: Viewpoints Costs Are Reduced by stereopsis.” *Visual Cognition* 12: 705–719.

Calvo Garzón, Francisco. 2008. “Towards a General Theory of Antirepresentationalism.” *The British Journal for the Philosophy of Science* 59: 259–292.

Cummins, Robert. 1996. “systematicity.” *The Journal of Philosophy* 93: 591–614.

Cummins, Robert, Jim Blackmon, David Byrd, Alexa Lee, and Martin Roth. 2005. “What systematicity Isn’t: Reply to Davis.” *The Journal of Philosophical Research* 30: 404–408.

Cummins, Robert, Jim Blackmon, David Byrd, Pierre Poirier, Martin Roth, and Georg schwarz. 2001. “systematicity and the Cognition of structured Domains.” *The Journal of Philosophy* 98: 167–185.

Davis, Wayne. 2005. “On Begging the systematicity Question.” *The Journal of Philosophical Research* 30: 399–408.

Edelman, shimon. 1998. “Representation Is Representation of similarities.” *Behavioral and Brain Sciences* 21: 449–498.

Edelman, shimon, and Nathan Intrator. 2003. “Towards structural systematicity in Distributed, statically Bound Visual Representations.” *Cognitive Science* 27: 73–109.

Evans, Gareth. 1981. “semantic Theory and Tacit Knowledge.” In *Collected Papers,* ed. Antonia Phillips, 322–342. Oxford: Oxford University Press 1985.

Fodor, J. A. 1981. “Imagistic Representation.” In *Readings of Philosophy of Psychology,* vol. 2, ed. Ned Block, 135–149. London: Methuen.

———. 1983. *The Modularity of Mind*. Cambridge: MIT Press.

———. 1987. *Psychosemantics. The Problem of Meaning in the Philosophy of Mind*. Cambridge: MIT Press.

———. 2008. *LOT2: The Language of Thought Revisited*. Oxford: Oxford University Press.

Fodor, J. A., and Zenon Pylyshyn. 1988. “Connectionism and Cognitive Architecture: A Critical Analysis.” *Cognition* 28: 3–71.

García-Carpintero, Manuel. 1995. “The Philosophical Import of Connectionism: A Critical Notice of Andy Clark’s Associative Engines.” *Mind and Language* 10: 370–401.

———. 1996. “Two spurious Varieties of Compositionality.” *Minds and Machines* 6:159–172.

Hadley, Robert. 1995. “The ‘Explicit-Implicit’ Distinction.” *Minds and Machines* 5: 219–242.

Hummel, J. E. 2000. “Where View-based Theories of Human Object Recognition Break Down: The Role of structure in Human shape Perception.” In *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines,* ed. E. Dietrich and A. Markman, chap. 7. Hillsdale, NJ: Lawrence Erlbaum.

———. 2001. “Complementary solutions to the Binding Problem in Vision: Implications for shape Perception and Object Recognition.” *Visual Cognition* 8: 489–517.

Hummel, J. E., and Irving Biederman. 1992. “Dynamic Binding in a Neural Network for shape Recognition.” *Psychological Review* 98: 352–376.

Johnson, Kent. 2004. “On the systematicity of Language and Thought.” *The Journal of Philosophy* 3: 111–139.

Kirsh, David. 1990. “When is Information Explicitly Represented?” In *Information, Language and Cognition*, ed. Philip Hanson, 340–365. Oxford: Oxford University Press 1992.

Marcus, Gary. 2001. *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. Cambridge: MIT Press.

Marr, David. 1982. *Vision.* san Francisco: Freeman.

McLaughlin, B. P. 1993. “The Connectionism/Classicism Battle to Win souls.” *Philosophical Studies* 71: 163–190.

———. 2009. “systematicity Redux.” *Synthese* 170: 251–274.

Noë, Alva, and Evan Thompson. 2002. “Introduction.” In *Vision and Mind. Selected Readings in the Philosophy of Perception*, ed. Alva Noë and Evan Thompson. Cambridge: MIT Press.

Pylyshyn, Zenon. 2003. *Seeing and Visualizing: It's Not What You Think*. Cambridge: MIT Press.

———. 2007. *Things and Places*. Cambridge: MIT Press.

Rey, George. 1991. “An Explanatory Budget for Connectionism and Eliminativism.” In *Connectionism and the Philosophy of Mind*, ed. Terence Horgan and John Tienson, 219–240. Dordrecht: Kluwer Academic Publishers.

Riesenhuber, Maximilian, and Tomaso Poggio. 1999. “Hierarchical Models of Object Recognition in Cortex.” *Nature Neuroscience* 2: 1019–1025.

silva, Porfirio, and P. U. Lima. 2007. “Institutional Robotics.” In *Advances in Artificial Life. Proceedings of the 9th European Conference, ECAL 2007*, ed. Fernando Almeida, Luis Mateus, Ernesto Costa, Inman Harvey, and António Coutinho, 595–604. Berlin: springer.

simons, Daniel, R. F. Wang, and David Roddenberry. 2002. “Object Recognition Is Mediated by Extraretinal Information.” *Perception and Psychophysics* 64: 521–530.

Taylor, Eric, and J. E. Hummel. 2009. “Finding similarity in a Model of Relational Reasoning.” *Cognitive Systems Research* 10: 229–239.

Thoma, Volker, J. E. Hummel, and Jules Davidoff. 2004. “Evidence for Holistic Representations of Ignored Images and Analytic Representations of Attended Images.” *Journal of Experimental Psychology: Human Perception and Performance* 30: 257–267.

Treisman, Anne. 1996. “The Binding Problem.” *Current Opinion in Neurobiology* 5: 520–526.

Ullman, shimon. 1989. “Aligning Pictorial Descriptions: An Approach to Object Recognition.” *Cognition* 32: 193–254.

———. 1998. “Three-dimensional Object Recognition Based on the Combination of Views.” *Cognition* 67: 21–44.

Van Gelder, Tim. 1990. “Compositionality: A Connectionist Variation on a Classical Theme.” *Cognitive Science* 14: 355–384.

Verdejo, V. M. 2009. *The Rational and Natural Mind: From Concepts to the Language of Thought.* saarbrücken: VDM Verlag.

———. Forthcoming. “The Visual Language of Thought: Fodor vs. Pylyshyn.” *Teorema*.

Wallis, Guy, B. T. Backus, Michael Langer, Gesche Huebner, and Heinrich Bülthoff. 2009. “Learning Illumination- and Orientation- Invariant Representations of Objects through temporal Association.” *Journal of Vision* 9: 1–8.