

The Commitment to LOT

VÍCTOR M. VERDEJO *University College London*

ABSTRACT: I argue that acceptance of realist intentional explanations of cognitive behaviour inescapably lead to a commitment to the language of thought (LOT) and that this is, therefore, a widely held commitment of philosophers of mind. In the course of the discussion, I offer a succinct and precise statement of the hypothesis and analyze a representative series of examples of pro-LOT argumentation. After examining two cases of resistance to this line of reasoning, I show, by way of conclusion, that the commitment to LOT is an empirically substantial one in spite of the flexibility and incomplete character of the hypothesis.

RÉSUMÉ : Je soutiens qu'accepter les explications réalistes intentionnelles du comportement cognitif conduit inévitablement à endosser l'hypothèse du langage de la pensée («language of thought», LOT), et que cette position théorique est, par conséquent, largement répandue chez les philosophes de l'esprit. Au cours de la discussion, je propose un exposé succinct et précis de cette hypothèse et j'analyse une série d'exemples représentatifs de l'argumentation pro-LOT. Après avoir examiné deux cas de résistance à ce type de raisonnement, je conclus en montrant que le soutien accordé à la LOT est empiriquement substantiel, en dépit de la flexibilité et du caractère incomplet de cette hypothèse.

Keywords: intentional explanation, language of thought, realism, representation, supervenience, computation, systematicity

Consider the following conditional claim regarding the language of thought (LOT) hypothesis, namely and to a first approximation, the hypothesis that cognition involves language-like representation (C):

(C)

If one accepts intentional explanations of cognitive behaviour and realism about the categories posited in those explanations, then one is committed to the truth of LOT.

Dialogue 55 (2016), 313–341.

© Canadian Philosophical Association/Association canadienne de philosophie 2016

doi:10.1017/S0012217316000408



If true, (C) would seem a direct consequence of the attitudinally free (I):

(I)

If intentional explanations of cognitive behaviour and realism about the categories posited in those explanations are true, then LOT is true.

In this paper, I endeavour to defend (C) through the careful analysis and defence of (I). The presumption connecting (C) and (I) is, of course, that one is committed to the (reasonably direct) entailments of the theses one accepts to be true. (C) highlights a substantial commitment of cognitive theorizing. The reason is that, presumably, the antecedent of (C) is fulfilled by a vast majority of contemporary philosophers of mind and cognitive scientists of all stripes. Thus, if (I) is true, the commitment to LOT is probably the rule rather than the exception in spite of a certain prevailing unpopularity and explicit indecision towards the hypothesis.

This is the structure of what follows. After demarcating the sceptical background surrounding (I) in Section 1, I shall state precisely and illustrate the meaning of the relevant terms in Section 2. I will then proceed to argue for the truth of (I) in Section 3 by means of an analysis of the fundamental features of pro-LOT argumentation. Since (I) is a conditional claim, it is important to stress that I will very much not argue in favour of the detached version of the claim. I do not wish to argue that LOT is actually true, let alone true *a priori*—even if *a priori* and *a posteriori* endorsements of (I) are indeed available.

Even though resistance to LOT is probably widespread, and (I) certainly not a matter of course among philosophers, there are not many explicit dismissals of (I) in the literature. In Section 4, I will examine two representative cases of disavowal, namely, Christopher Peacocke's developments regarding level of explanation 1.5 and Jonathan Knowles's upfront denial of the thesis.¹ I will defend that failure to embrace the commitment of LOT in these cases stems from attributing to LOT inessential theoretical features. Finally, in Section 5, I will show, by way of conclusion, that the commitment to LOT is an empirically substantial one in spite of its admittedly flexible and incomplete character.

1. Sceptical Background

(C) and (I) must be sharply distinguished from the Fodorian well-known contention that LOT would involve a vindication of intentional psychology.² First, this contention concerns a vindicating relation from LOT to intentional or folk psychology (i.e., from the truth or scientific merits of LOT, to the truth or vindication of intentional psychology), one that suggests right the opposite direction of entailment in relation to (I). Second, as is known, the Fodorian interpretation

¹ Peacocke (1986a, 1986b, 1989), Knowles (2001, 2002).

² Fodor (1987, Chapter 1), Devitt (1990).

of the vindication does not result in an entailment of the sort (I) expresses but, quite differently, in an argument to the best explanation. By Fodorian lights, LOT is our best (perhaps the only) explanation available of realism about intentional psychology. On this account, it does not follow that LOT is implied by a realist intentional psychology.

In a sense, (I) is not new, however. It can fairly be taken to signal one central theme in the existing literature about LOT. Notably, (I) has been the backdrop of developments due to *both* defenders and opponents of the hypothesis. Advocates of LOT have relied on something like (I) to argue in its favour, and even do so on *a priori* grounds, granted the plausibility of various forms of intentional realism.³ Furthermore, detractors of the hypothesis have taken versions of (I) to advance a reductio of folk psychological theory, in the light of approaches that impugn the truth of LOT.⁴

In spite of its significance, (I)—and its attitudinal interpretation (C)—is not seriously considered and certainly heard of less often these days. A main reason is a ubiquitous and strong scepticism about LOT. For many researchers, LOT has largely fallen into disrepute as an outmoded and limited standpoint in cognitive theorizing. This idea was popular some time ago with the emergence of connectionism, which challenged the symbolic kind of representation characteristic of LOT.⁵ But LOT seems to be much further away from newest interests in ongoing cognitive research if we give any credence to recent approaches that depart radically from the computational paradigm and even deny the existence of representations in cognition.⁶ Indeed, it is not hard to find characterizations of the available positions in cognitive science that relegate LOT or ‘classicist computationalism’ to a rather restricted locus among many less stringent views in the cognitive and even computational spectrum.⁷

Another source of scepticism about LOT is, ironically enough, also found in the work of its chief defender, Jerry Fodor, who has vehemently argued for the existence of insuperable problems such as what is known as the relevance or frame problem—the problem of determining what is computationally relevant—or the globality problem—the problem of capturing the mind’s sensitivity to global and contextual properties of representation.⁸ Although alternatives to Fodor’s pessimism on this score have been proposed,⁹ the impression remains that it is still an open question whether LOT has the resources to deal with

³ Davies (1991), Lycan (1993), García-Carpintero (1995).

⁴ Churchland (1989), Churchland and Sejnowski (1989), Ramsey, Stich and Garon (1991).

⁵ Rumelhart et al. (1986).

⁶ Port and van Gelder (1995), Chemero (2009).

⁷ E.g., Piccinini (2009), Piccinini and Scarantino (2011), Egan (2012), Fresco (2012).

⁸ Fodor (2000, 2008).

⁹ E.g., Carruthers (2006), Schneider (2011).

problems that run deep into the very nature of computational explanation. It is, therefore, clear even among staunch sympathizers that LOT is, if true, far from being the whole truth about cognition. There is more to it, in fact, since there are sure to be mental processes and types of representations that escape the LOT mould in a range that includes vision and iconic representation,¹⁰ mental processes in other perceptual modalities¹¹ or mathematical thinking and thinking about magnitudes,¹² to name some examples.

So far, the reasons highlighted for the lack of interest regarding (I) concern its consequent. However, there is also pervading mistrust concerning its antecedent. For instance, a good number of theorists would cast doubts on the earlier-considered most clear evidence and paradigmatic cases of intentional explanations, namely, intentional explanations of systematicity phenomena.¹³ These theorists would tend to think that (the fleshing out of) the antecedent of (I) is also unwarranted in fundamental cases. Some would even suggest that we cannot make sense of intentional realism in the way required by LOT proponents.¹⁴

The foregoing considerations intimate that (I) enjoys at present a thin philosophical credence and a prevalent lack of attention. It is my view, however, that even if the antecedent and consequent of (I) are controversial in various ways—indeed even if both the antecedent and consequent turned out to be indisputably false—(I) is true nonetheless. To show this, however, we must first analyze and clarify the thesis in greater detail.

2. Setting the Stage

The assessment of (I) requires of us to make explicit the exact meaning of the terms involved. First, the antecedent of (I) makes use of the notion of intentional explanation. By ‘intentional explanation’ in this context I mean simply explanations that appeal to intentional states and processes, that is to say, states and processes endowed with distinctive representational or informational content or semantic properties so that they stand in for or carry information about other states or events.

The term ‘distinctive’ is here meant to exclude states and processes that are intentional in some general or homogeneous way that does not unambiguously individuate the state or process in question. This condition rules out, for instance, picture-like and other (purely) reproductive sorts of representation at the intentional level. For instance, since pictures or photographs represent homogeneously (i.e., every part of a picture of X is the picture of a part of X), a picture-like

¹⁰ Fodor (2008, Chapter 6).

¹¹ Kaye (1995).

¹² Schneider and Katz (2012).

¹³ E.g., Johnson (2004), Gomila, Travieso and Lobo (2012).

¹⁴ Garson (1997).

state P (say, my mobile photo of Big Ben) lacks distinctive or unambiguously individuating intentional features: whatever picture-like features P has are potentially the features of indefinitely many other distinct states $P_1, P_2, \dots P_n$ (say, my mobile photo of the crowd around Big Ben, of the Thames, of Westminster Bridge, of a sunny day in London, and so on).¹⁵ Note that to put aside homogeneous representation in this context is not to privilege representation in linguistic domains. Distinctive representation of the sort apt to figure in a LOT is arguably also present in vision, navigation and other non-linguistic domains.¹⁶

There is nothing necessarily folk about intentional explanations so understood. These explanations may be, and frequently are, simply surmised in folk psychology (e.g., as when one explains a man's taking an umbrella on the basis of his belief that it is raining) but they may also be carefully articulated in philosophical theories of mind (e.g., by appealing to a particular theory of concepts or mental content), or else, precisely specified and empirically confirmed by scientific theories (e.g., by postulating egocentric and allocentric representational capacities in navigation for different organisms). They need not be specified at the personal level either. Explanations that qualify as intentional, given the present context, are explanations that appeal to content-carrying states, regardless of whether those states are personally (such as in self-conscious belief and intentional action) or subpersonally characterized (such as in internal grammar states involved in language learning or states carrying retinal information in early visual processing).

The target explanations may concern a large number of explananda and, as we will see below, include a wide variety of areas of study concerning human and infrahuman cognitive capacities. The antecedent of (I) does not presuppose that a particularly narrow understanding of explanation must be correct. Importantly, however, explanations in the relevant sense cannot be mere descriptions of phenomena, or covering-law explanations that proceed via subsumption of phenomena under natural law as in the classical deductive-nomological model.¹⁷ Although, strictly speaking, explanation and causality are indeed separable notions,¹⁸ for present purposes it will be helpful to assume that intentional explanations are causal in the specific sense that they must involve the postulation of the states or processes that are causally responsible for the phenomenon under study. The relevant kind of intentional explanation in the antecedent of (I) can be seen, therefore, as a form of mechanistic explanation.¹⁹

¹⁵ Cf. Fodor (2008, Chapter 6).

¹⁶ See Camp (2007), Rescorla (2009), Verdejo (2012a, 2012b) for extensions of LOT to non-linguistic or non-sentential domains. I offer a discussion of the notion of homogeneity in representation in Verdejo (manuscript).

¹⁷ Hempel and Oppenheim (1948).

¹⁸ Cartwright (2006).

¹⁹ Craver and Bechtel (2006).

This is so independently of the sophistication, empirically complex or folk character of the mechanisms proposed.

Now, the antecedent of (I) also expresses a condition about the reality of the categories that make a causal explanatory role. The reality in question is understood, as in any other scientific discipline outside psychology, in the very specific sense that the causal-explanatory components found in intentional explanations must ultimately have a physical realization. Here we follow Fodor in believing that “mind/brain supervenience (and/or mind/brain identity) is, after all, the best idea that anyone has had so far about how mental causation is possible.”²⁰ Nonetheless, I would like to remain neutral about the specific kind of physical realization needed. For present illustrative purposes, it is enough that we assume a form of supervenience of the intentional categories on the categories relevant at lower physical levels. The understanding of intentional ‘reality’ in (I) is therefore explicitly physicalistic.²¹ The basic idea, which I will not develop in any detail here, is that intentional causation is not different from other forms of causation found in science in respecting this physicalistic approach. Alternative, non-physicalistic interpretations of realism are also available but will be put to one side in what follows.²²

It is time to move on to the consequent of (I). How is LOT understood in this context? LOT is often spelled out as a rather involved confederation of theses.²³ For present purposes, it will be useful to consider a capsule-form characterization. We may state the hypothesis in terms of (L):

(L)

LOT is true of an organism O regarding cognitive task T if and only if O physically implements language-like representational computation in order to accomplish T.

(L) explicitly discards merely abstract, heuristic or descriptive readings of the hypothesis. In this sense, (L) has the advantage of being clearer than other characterizations regarding what is required, at the concrete physical or neuro-physiological level, in order for LOT to be true of an organism.

In (L) ‘representational computation’ stands for the manipulation or processing of representations in accordance to rules. Digital computers, and other information-processing devices are paradigmatic examples of representational computation. Representational computation in this sense may well be of a

²⁰ Fodor (1987, 30).

²¹ García-Carpintero (1995, 1996).

²² See e.g., Crane and Mellor (1990), Dennett (1992). I am, therefore, assuming, in contrast with these developments, that the antecedent of (I) is fulfilled only when a physicalistic interpretation of ‘reality’ is at play. Physicalism does not, however, entail in this context physical reductionism, as we will see in more detail below.

²³ Fodor (1987), Schneider (2011), Aydede (2010).

connectionist variety, including various models of neural networks.²⁴ The claim that the target representation is language-like means that (i) the representation is digital or discrete (as opposed to analog or continuous) and (ii) the representation has specific semantic and syntactic properties, where syntactic properties are understood in Fodor's sense as causal properties correlated with semantic properties.²⁵ As it stands, the characterization in (L) remains neutral about whether semantic properties are essential to LOT computation and representation²⁶ or whether only syntactic properties are.²⁷

There is no restriction as to the kind of semantics and syntax that language-like or LOT representation allows beyond (i) and (ii). In particular, language-like representation is not literally meant to involve a (natural) language in any straightforward way. Consider a concrete digital computer *C* that calculates additions through computations and compositional principles defined over a finite body of strings of 0s and 1s representing a finite set of numbers. Thus, for instance, having as input the string of digits representing the number 100 together with the string of digits representing the number 1, would cause *C* to deliver as output the string of digits representing 101. Were *O* to perform addition on anything like this model, *O* would instantiate LOT computation in the required sense, namely, in the sense of computation defined over (number) representations with semantic and syntactic properties.

Under alternative construals,²⁸ representation in a LOT must not only respect (i) and (ii) but, furthermore, necessarily involve a compositional syntax and semantics, i.e., a system of (de)composable semantic and syntactic units. Compositionality, however, is neither a necessary nor a sufficient condition for language-likeness in our sense. First, as Fodor has pointed out, compositionality in mental representation may also be a feature of picture-like representation.²⁹ Second, a language entirely constituted by a list of atomic and unconnected formulae is, albeit extremely limited and biologically implausible, a language nonetheless.³⁰ If physically realized in a (simple) organism and endowed with a specific causal role, such a system of formulae would involve states with semantic and syntactic properties in Fodor's sense. Thus, although compositionality is a possible (chief) feature of language-likeness in our sense, it is not a necessary or sufficient one. This is so even if the most cogent arguments in favour of LOT exploit the fact that only LOT would provide a compositional scheme of representations valid for handling systematicity.

²⁴ Piccinini and Scarantino (2011), Fresco (2012).

²⁵ Fodor (1987, 16-21).

²⁶ Crane (1990), Shagrir (2001).

²⁷ Stich (1983), Piccinini (2015).

²⁸ E.g., Fodor and Pylyshyn (1988, 12-13), Fodor and McLaughlin (1990, 198).

²⁹ Fodor (2008, Chapter 6).

³⁰ Cf. Evans's (1981) analysis of unstructured semantic theories.

Finally, in (L) ‘implementation’ is intended to mean physical and possibly multiple realization in an organism or organism’s physical parts. In short, the truth of LOT for an organism O concerns the physical realization in O of a particular kind of computation, namely, computation over digital language-like representations. According to (L), it suffices for the truth of LOT that some of the workings of a creature’s mind regarding a certain cognitive task T—as opposed to most or most fundamental such workings—involve the implementation of language-like representational computation. Note also that, insofar as the computational process *is* physically implemented, it must obey the laws of physics and, indeed, plausibly involve a dynamical system that continuously changes over time at a certain level of description.

It must be noted that (L) is compatible with several readings and possible emphases. For instance, (L) may be taken to express primarily a claim about psychological states (*viz.* representational states figuring in a computation) or psychological processes (*viz.* computational processes sensitive to semantic/syntactic features of the representations). Some would see (L) as a condition specifically concerning creatures endowed with a propositional-attitude psychology.³¹ Others would argue that it concerns capacities dealing with tasks that belong to structured cognitive domains.³² In short, and although I lack the space to substantiate this exegetical remark with a thorough discussion, the presumption is that (L) captures succinctly the fundamental thesis that the many currently available versions of LOT actually share.³³

3. The Commitment to LOT

Once (I) is clarified, we are in a position to assess its plausibility. In order to show the force and the widespread character of (I), I will invoke a comprehensive series of exemplary arguments in favour of LOT that exist in the literature. As will become apparent, they are all attempts to establish the truth of LOT out of specific realist intentional explanations. The validity of this series of arguments supports the view that, generally, acceptance of such explanations suffices for the commitment to LOT. Consider an intentional state or property of a state, *i*, a cognitive state or process, Δ , and some physical supervenience base, Σ . The examples to follow provide variations of the following simplest formulation (SF):

(SF)

If *i* causes Δ , then there is a supervenience base Σ that physically realizes *i* in a language-like computational process.

³¹ Fodor (1987), Aydede (2010).

³² Cummins (1996), Cummins et al. (2001).

³³ E.g., Harman (1973), Fodor (1975, 2008), Field (1978), Maloney (1984), Von Eckardt (1993), Devitt (1996), Rey (1997), Carruthers (2006), Schneider (2009, 2011).

We may see (SF) as the simplest schema for (I). (SF) will help us to swiftly capture the essential features of the various arguments under consideration below. The antecedent of (SF) brings out a causal relation attributable to a discrete state i in virtue of its intentional or semantic properties. This is the minimal expression of a causal explanation at the intentional level. Since, on the assumption of physicalism, it is the supervenience base Σ that has causal powers at the realization level, and since it is meant literally that i causes Δ in the antecedent of (SF) (intentional realism), it follows that Σ must realize (and, therefore, have unambiguously associated) specific semantic properties (*viz.* the semantic properties of i) as well as syntactic properties (*viz.* the causal properties appropriately correlated with such semantic properties). Therefore, the target causal process involves language-like representational computation: a process defined over specific discrete states with semantic and syntactic properties realized by Σ in which intentional relations or rules are preserved.

To illustrate, let us suppose that the target intentional discrete state (i) is my seeing that it is raining outside. Let us grant, for the sake of the argument, that our preferred intentional theory takes it that this visual state, in normal circumstances, causes me to believe that it is raining outside (and to exhibit a suitably related behaviour such as my putting on a raincoat) (Δ). Now, since causation is taken literally, there must be a neurophysiological state, or a combination of such states (Σ), that actually realizes the causal impact of my seeing that it is raining outside. Since, by assumption, the causal impact of my seeing is wedded to a distinctive content (*viz.* the content that it is raining outside), the neurophysiological realizing basis must itself be wedded to that distinctive content and be unambiguously and strictly distinguished from other bases that would realize the causal impact of indefinitely many other intentional states with distinct contents, such as my seeing that it is sunny outside, that it is snowing outside, that it is windy outside, and so on. This in turn means that the realizing state(s) must (a) have syntactic properties or causal properties correlated with semantic properties, and (b) conform to a general intentional rule of roughly the form: in normal circumstances, S's seeing that p causes S to believe that p . But a causal process which is sensitive to the semantic properties of their intervening (discrete) states through properties of their syntax in accordance to rules just *is* a language-like representational computational process, indeed the computational process of a LOT as spelled out in (L).

(SF) predicts the falsity of its antecedent whenever rival (non-LOT) forms of computation are assumed, such as subsymbolic and analog computation. In general, and as we will see in more detail below (Sections 4.1 and 5), the antecedent of (SF) will not be compatible with forms of computation that are not themselves implementations of LOT computation in the sense described.

3.1. Tacit Knowledge

A particularly salient case that follows precisely the line of reasoning laid out in (SF) is found in the work of Martin Davies.³⁴ For many years now, Davies has argued that neo-Fregean accounts of concepts and systematic thought—a paradigmatic case of intentional explanation in our sense—involve an *a priori* commitment to LOT. At some points, Davies articulates this commitment in terms of a non-reductive interactive relationship between personal and subpersonal levels of description.³⁵ The derivation of LOT from neo-Fregean theories of conceptualized thought³⁶ is, in Davies's terms, a case of 'downward inference' from *a priori* philosophical theories to empirical psychological requirements of information-processing mechanisms. The downward inference proceeds in two stages. First, the neo-Fregean notion of possession of a concept is elucidated as involving causally efficacious tacit knowledge of a rule of inference, where tacit knowledge is understood as "a state that figures as a common factor in content-involving causal explanations of certain transitions between representations or states of information."³⁷ In the second stage, it is shown that, under the assumption of intentional realism, there must be an intrinsic connection between inferential transitions—explained in terms of tacit knowledge—and the required syntactic properties of the physical configurations that figure as inputs in those inferential transitions. Thus, inputs of a cognitive function involving tacit knowledge (of a rule of inference) must be causally adequate to engage the tacit knowledge as a common causal factor in the input-output pattern, and hence the semantic properties of the input state must correlate with its causal properties. The conclusion is that genuinely causal tacit knowledge must be underpinned by sentences in a language of thought: causally efficacious physical states with semantically correlated (hence, syntactic) properties figuring in a computational process that realizes the target inferential transitions. Davies's considerations offer a clear-cut illustration of (I), that is to say, a case in which intentional, causal explanations (of systematicity and conceptualized thought) in terms of tacit knowledge lead to the truth of LOT.³⁸

³⁴ Davies (1991, 2000a, 2000b, 2004, 2015).

³⁵ Davies (2000a, 2000b, 2004).

³⁶ Evans (1982), Peacocke (1992).

³⁷ Davies (2015, 77).

³⁸ These considerations bear no commitment, as it does in Davies's accounts (1991, 2000a, 2000b, 2004) to the view that LOT can be established purely *a priori* (see also Rey 1995 and Lycan 1993). For Davies, intentional realism is an assumption to which we are committed "by our everyday practices of personal-level description and explanation" (Davies 2000b, 97). According to the here recommended view, however, the prospects of intentional realism are better seen as those of (a form of) physicalism as generally found in empirically-confirmed causal explanation.

Davies's line of argument has not gone uncontested. For instance, Jürgen Schröder³⁹ invokes the idea of a common categorical base for having inferential dispositions in order to question the one-to-one correspondence Davies assumes between inferential rules and mechanisms of tacit knowledge. However, while this point threatens the evidence we have for rendering states of tacit knowledge as causally explanatory states (that is to say, as physically realized at all), it leaves untouched the view that states of tacit knowledge must be computationally realized by means of causal properties of the input states *if* they actually were causally explanatory states—which is what would affect (I) in this particular case. I will address further the philosophical resilience to (I) in Section 4 below.

I have detoured to present in some detail Davies's developments as a paradigmatic instance of (I). However, Davies's style of argument resonates well beyond neo-Fregean theories of conceptualized thought. Indeed, the practical totality of pro-LOT arguments in the literature, including arguments that are a long distance from Davies's *a priori* argument based on tacit knowledge, can be seen as specific articulations of (I).

3.2. Systematicity

Many agree that systematicity, however exactly understood, is a defining feature of central cognitive capacities. Several prominent arguments for LOT capitalize on this feature. Although this class of arguments may vary greatly, they all involve the identification of certain explanatory intentional categories or features as real or causally effective, in such a way that the requirement of a computational process manipulating language-like representations at the physical level ensues. If this is correct, these developments are better seen as instances of (I). With the aim of sacrificing detailed exegesis for the sake of a far reaching overall picture, I will put detailed discussion of the views to one side in the survey that follows.

As is known, Fodor and allies have offered the most influential systematicity arguments in favour of LOT. Fodorian developments take LOT to constitute the best (causal) explanation of systematicity understood as the existence of structurally related intentional capacities or capacities to have intentional states.⁴⁰ For these authors, only a compositional system of representations can explain systematicity, e.g., the fact that whenever subjects have the capacity to possess a given belief (such as the belief that John loves Mary) they have the capacity to possess beliefs that share the same intentional parts (such as the belief that Mary loves John). Thus, the compositional structure of an organism's intentional states causally explains the organism's ability to have structurally

³⁹ Schröder (1998, §2).

⁴⁰ Fodor (1987), Fodor and Pylyshyn (1988), Fodor and McLaughlin (1990), McLaughlin (1993, 2009), Aydede (1997). See Calvo and Symons (2014) for recent discussion.

related intentional states. If true, this explanation very obviously requires the organism to realize the processing or manipulation of language-like complex representations that capture the compositional semantics identified at the intentional level. This is a clear case of (I): an inference from (the best or only) intentional explanation of systematicity to the truth of LOT.

There are a number of variations to this line of reasoning. Robert Hadley⁴¹ takes as an explanatory target so-called ‘strong (semantic) systematicity,’ considered to be a feature of learning-based processes of correct generalization of syntactic and semantic properties to novel cases. On Hadley’s account, part of the challenge posed by systematicity phenomena to LOT-opponents consists of accounting for such syntactic and semantic generalizations in a way that does not involve causal constituent structure, and hence LOT.

Ken Aizawa,⁴² following Fodorian developments, distinguishes systematicity arguments that have the systematicity of inference, the systematicity of cognitive representation and the content-relatedness of thoughts or compositionality itself as explananda.⁴³ These related phenomena all concern capacities among intentional states or features of intentional states themselves whose explanation, the classicist claims, rely on the causal properties realized in language-like computation. For instance, Aizawa takes classicism to offer a superior explanation of the fact that whenever thoughts are counterfactually dependent—i.e., whenever a subject’s capacity for thought A is nomologically necessary and sufficient for the subject’s capacity for thought B—they also are semantically related. Roughly, the reason is that the classicist causal account of the counterfactual dependency of A and B is grounded in the sharing of language-like representational structures that involve computational operations where the sharing of such structures already entails that A and B are semantically related.⁴⁴

In an insightful discussion, Manuel García-Carpintero⁴⁵ provides a general framework for the analysis of systematicity arguments in favour of LOT. In a nutshell, according to García-Carpintero, the commitment to LOT stems, generally, from systematicity understood in terms of intentionally complex explanations of behaviour, that is to say, in terms of causal explanations of behaviour that invoke complex representational states.⁴⁶ Under the assumption of intentional realism, the physical correlates of any such complex representational states figuring in intentional causal explanations of systematicity must inevitably lead to acknowledging the existence of LOT.

⁴¹ Hadley (1994, 2004).

⁴² Aizawa (2003).

⁴³ See Aydede (2010) for a similar classification of these arguments.

⁴⁴ Cf. Aizawa (2003, Chapters 6-8).

⁴⁵ García-Carpintero (1995, 1996).

⁴⁶ See Verdejo (2012a, 2015) for a similar analysis.

The above considerations involve no restriction as to the source and nature of complex intentional explanations in the antecedent of (I). Indeed, we should expect to find intentional explanations of systematicity belonging to a wide range of cognitive abilities of “human and infrahuman mentation.”⁴⁷ The possibilities go well beyond the various forms of learning, perception and production characteristic of linguistic abilities. For instance, the explanatory need for LOT may also arise from intentional explanations of vision and visual processes⁴⁸ or navigation.⁴⁹

3.3. Beyond Systematicity

The number of possibilities for a pro-LOT argument that instantiates (I) transcends systematicity arguments, properly so-called. Examples include the well-known arguments that start from productivity⁵⁰ or inferential coherence in thinking,⁵¹ but also thought,⁵² the tracking of objects,⁵³ broad content⁵⁴ or the complexity of behaviour.⁵⁵ Georges Rey⁵⁶ speaks of eight consequences of intentional realism (or ‘essential mentalism,’ as he calls it) that may be better explained (if at all) by LOT rather than by connectionist models, and to that extent, eight possible starting points for a pro-LOT argument. These include rational and irrational relations and errors among attitudes, conceptual stability, the fine-grainedness, the multiple roles and the causal efficacy of attitudes. Finally, the actual research conducted in cognitive science, which also appeals to intentional explanations in the relevant sense, is also a fundamental and widely known starting point for an argument in favour of LOT.⁵⁷

It is because there are intentional causal explanations of the target phenomena (*viz.* productivity, inferential coherence, content relations, etc.) that we are entitled to conclude that the workings of our minds at the intentional level are realized by a language-like computational process. For instance, if an organism’s inference from *A&B* to *A* (conjunction elimination) is explained, at the intentional level, by the fact that a series of transitions are causally enabled for ...&... structures, then there must be a physical state or mechanism defined over the

⁴⁷ Fodor and Pylyshyn (1988, 37).

⁴⁸ Pylyshyn (2003), Verdejo (2012a).

⁴⁹ Camp (2007), Rescorla (2009).

⁵⁰ E.g., Fodor and Pylyshyn (1988, 33-37), Fodor (1990, 16-19), Aizawa (2003, Chapter 3).

⁵¹ E.g., Fodor (1987, 143-147), Fodor and Pylyshyn (1988, 46-48), Fodor (1990, 19-24).

⁵² Rey (1995).

⁵³ Horgan and Tienson (1996, §5.2).

⁵⁴ Crane (1992).

⁵⁵ Fodor (1987, 141-143).

⁵⁶ Rey (1991, 224-225; 1997, §8.8 and §9.4).

⁵⁷ Fodor (1975, 2008).

manipulation or processing of states realizing ...&... structures. This would be, again, a case of language-like representational computation in the required sense.

The list of this and the preceding sections could be enlarged further. What we have seen already should suffice to establish that there are indefinitely many starting points for a pro-LOT argument conforming to (I), indeed, indefinitely many ways of stating intentional causal explanations in cognition that lead to the postulation of language-like computation. This makes a strong case for the truth of (I).

4. Rejecting the Commitment to LOT

The points of the foregoing section are not a matter of course among philosophers. Sceptics include a ragbag of theorists (e.g., Davidson, Dennett, McDowell, Brandom, Bealer or the second Wittgenstein under certain interpretations) who, in one way or another, would defend the autonomy or lack of continuity of *a priori* or intentional explanations with respect to scientific explanations, properly so-called. This may also involve different notions of ‘reality’ for intentional realism. I would like to put aside such ‘autonomy’ developments. For present purposes, we may gloss this philosophical position, conceived broadly, as simply rejecting intentional realism in the relevant sense and, hence, as being orthogonal to the truth of (I).

The real challenge to (I) comes rather from the work of philosophers who, while officially accepting both intentional explanations *and* realism in our sense, still refuse the commitment to LOT. Although a certain renegade spirit pervades among intentional realists, it is not so easy to find works that explicitly take up this line of thought. In this section, I will consider two significant and especially clear cases of this kind of LOT desertion. Confinement to these representative and temporally distant developments will also help to keep the discussion focused. The first one concerns Peacocke’s discussion about levels of explanation. Secondly, I shall focus on Knowles’s upfront denial of (I).

4.1. Peacocke on Levels of Explanation

It should strike readers as surprising that Peacocke’s own developments some time ago are at odds with the view just highlighted in Section 3 (and especially 3.1). Peacocke’s main concerns in those developments are, on the one hand, to defend the specificity of certain explanations to be located at a new level—so-called ‘level 1.5’—between David Marr’s level 1 and Marr’s level 2, and, on the other hand, to show that this specific level of explanation can yield a criterion for deciding psychological reality at least for the case of semantic theories and grammars.⁵⁸ For the purposes of this piece, Peacocke’s illuminating discussion

⁵⁸ Peacocke (1986a, 1986b, 1989). For starters, whereas Marr’s level 1 states what the system does (and why)—i.e., the input-output function—Marr’s level 2 concerns how the system does it—i.e., the algorithm and representation. Finally, Marr’s level 3 deals with the physical realization in hardware structures (Marr 1982, Chapter 1). See Verdejo and Quesada (2011) for a defence of the importance of the Marrian distinction.

as regards those issues may remain untouched. However, in the course of these developments, Peacocke explicitly endorses the view that statements at level 1.5 are compatible with connectionist algorithmic implementations at Marr's level 2, and hence, that explanations at that level are neutral with respect to the question of whether LOT or connectionist models are preferable in cognitive research. Since 1.5 explanations, if true, clearly amount to intentional realist explanations, Peacocke's neutrality regarding LOT is, therefore, an important case of denial of (I).

Peacocke introduced explanations at level 1.5 by way of three kinds of examples: finite languages, perception of non-linguistic objects and perception of kinds. For simplicity's sake, let us focus exclusively on the first of these. Once this case is appraised, extension of the analysis to the other cases will be easy to figure out. Level 1.5 explanations state the information upon which the algorithm draws. In the case of a finite language with 10 monadic predicates (F, G, \dots) and 10 proper names (a, b, \dots), we can explain a subject's (S) understanding of each sentence of the language as being systematic because we can cite common informational elements whenever S understands sentences with common linguistic items. For instance, we can explain S's understanding of every sentence containing a as drawing upon the information that a denotes John. However, according to Peacocke, this kind of explanation does not fall into Marr's level 1. At this level, it is only relevant the statement of the function in extension—in this particular case, the function from whole sentences to their meanings. This is why, at Marr's level 1, no distinction can be made between S's understanding being systematic or unsystematic. But these explanations do not belong to Marr's level 2 because, Peacocke observes, the specified information drawn upon is compatible with the information being drawn upon by different algorithms specified at that level. This is why connectionist and LOT algorithmic implementations are claimed to be equally valid.⁵⁹ Consequently, the commitment to LOT out of explanations at level 1.5 is explicitly denied:

It may be queried whether we can really make sense of an algorithm or mechanism drawing on information unless we accept some kind of hypothesis that there is a language of thought in which the information used is written out. I would dispute that there is any such commitment.⁶⁰

Peacocke maintains his neutrality in a similar way when considering the features of a possible Informational Criterion for the psychological reality of grammars based on explanations at level 1.5. In outline, according to such a criterion, for a set of grammatical rules R_1, \dots, R_n (that state p_1, \dots, p_n , respectively) to be psychologically real is for any of its derivable statements (q) to

⁵⁹ Peacocke 1986a, 104.

⁶⁰ Peacocke (1986a, 111).

have a corresponding explanation at level 1.5 that specifies the information drawn upon by a given algorithm (such information consisting in the statements p_1, \dots, p_n of the set of rules R_1, \dots, R_n in question).⁶¹ Granted that some such criterion is correct, Peacocke rejects the idea that it yields a commitment to LOT.⁶²

There is no doubt that 1.5 explanations are intentionally real explanations. They invoke informational states, whether or not such informational states are specified at the subpersonal level. Furthermore, the target explanations are intended to be real in our sense since (a) they are taken to be a form of causal explanation⁶³ and (b) are used to provide a criterion for the psychological reality of semantic theories and grammars.⁶⁴ But can one's 1.5 level explanation be intentional and causal and, still, fail to involve computation in the sense of (L) above?

Although Peacocke's writings do not help us at this point, we may try to address this question by supposing, for the sake of the argument, that one's 1.5 explanation involves subsymbolic or analog computation instead. Subsymbolic computation consists of the manipulation of discrete states without syntactic properties in Fodor's sense (*viz.* causal, realizing properties correlated with semantic properties).⁶⁵ Therefore, the question from this angle is: how can the information identified at level 1.5 literally play the causal explanatory role it is supposed to play if it is deprived of a causally efficacious basis underpinning that very information? For instance, how can the information that *a* denotes John causally explain the systematicity involved in understanding *Fa*, *Ga*, *Ha*, ... if there is no causally efficacious physical state correlated with such information? If causal properties are not correlated with the relevant information, it is mysterious how that information has a causal role in (as opposed to being epiphenomenally related to) the explanation of, e.g., systematic understanding.

On the other hand, analog computation involves the manipulation of continuous variables, or variables that vary continuously over time and whose real values can only be measured within a margin of error.⁶⁶ But if the 1.5 account identifies a discrete state i^* (e.g., the state with the information that *a* denotes John) as the one that is explanatory regarding algorithm α , how can i^* be realized by an analog magnitude ambiguously connected, from the point of view

⁶¹ For a detailed exposition of the criterion, of which I have only given a simple version, see Peacocke (1989). A detailed exposition of the criterion is not needed to follow the present considerations.

⁶² Peacocke (1989, 118).

⁶³ Peacocke (1986a, 102; 1989, 113).

⁶⁴ Peacocke (1986a, 115), Peacocke (1989, 114).

⁶⁵ Cf. Smolensky (1988).

⁶⁶ Pour-El (1974), Piccinini and Scarantino (2011, 11), Piccinini (2015, Chapter 12).

of realization, with indefinitely many other states i_1, i_2, \dots, i_n (e.g., the state with the information that a denotes Mary, or that a denotes Peter, and so on). Since continuous variables of this sort are only ambiguously determined within certain boundaries, the very idea of analog computation thwarts the contention that i^* is a state with distinctive causal-explanatory intentional features.

In sum, these computational possibilities are just huge revisions if not upfront dismissals of the view that 1.5 level information or information on which an algorithm draws is explanatory in the relevant realist sense. Peacocke's neutral position regarding LOT, therefore, suffers from inconsistency.

This rather straightforward criticism might not be entirely fair to Peacocke's standpoint in these papers. Failure to notice the aforementioned inconsistency in the case at hand can be explained by a conception of LOT as involving both more and less than what (L) actually requires. On the one hand, Peacocke's position is based on the assumption that LOT involves more than it does, since he describes it as bearing a commitment to the axioms derivable from explanations of systematic capacities being explicitly represented in the language of thought.⁶⁷ However, this is not a necessary feature of LOT models.⁶⁸ On the other hand, Peacocke also suggests that LOT requires less than (L), since he describes it as a thesis restricted to Marr's level 2. However, LOT is more accurately viewed as an empirical hypothesis about the proper inter-level explanatory relations spelled out in terms of (L), from the highest intentional level to the lowest physical level.

It would be incorrect to suggest that Peacocke's neutrality has been invariably maintained. Although there are neutral avowals in other writings,⁶⁹ Peacocke's more recent developments would seem to agree with the points advanced here in accepting that only LOT can back content-involving explanations, such as the ones related to grasp of a canonical concept:

It is very plausible that part of the subpersonal explanation of how it is that a thinker is able to enjoy such content-sensitivity in his grasp of the canonical concept of a concept is that the subpersonal realizations of the relevant mental states to which he is sensitive contain some representations in a subpersonal, Fodorian language of thought, structured representations that have the Thought in question as its assigned concept. I agree with those who say that it is hard to see how there can be any other explanation of all the phenomena in which mental states with content are implicated.⁷⁰

Fresher Peacockean views would, therefore, seem to duck the attribution of inessential theses to LOT in favour of a less vacillating embracement of the

⁶⁷ Peacocke (1986a, 104).

⁶⁸ E.g., Fodor (1987, 25).

⁶⁹ E.g., Peacocke 1992, 185; Peacocke 1994, 310.

⁷⁰ Peacocke 2008, 294-295. See also Peacocke 2004, 97.

hypothesis. Be this exegetical point as it may, the present considerations are not so much concerned with Peacocke's developments as they are with a clarification of LOT commitments. The upshot is that realist intentional explanations, such as those found in level 1.5 theories, exhibit the fundamental commitments of LOT.

4.2. Knowles's Upfront Rejection of (I)

Knowles⁷¹ agrees that intentional explanations, and more specifically, intentional folk explanations, are of a realist or genuinely scientific kind.

I believe, like Fodor, that F[olk] P[psychology] is essentially continuous with scientific psychology, and hence in the idea of I[ntentional] P[psychology] as a nomological scheme of explanation whose posits should be fully realistically interpreted (i.e., as realistically as other scientific posits).⁷²

However, he explicitly rejects any implications of (psychologically real) explanations regarding the representational theory of mind or LOT. Knowles discusses several ways in which intentional explanations can yield an argument in favour of LOT and finds reasons to cancel the commitment in every case. Two lines of argument stand out.

A first line of argument takes into account a number of intentional folk categories, namely, propositional content, broad content and Fregean modes of presentation from which an argument for LOT can be crafted.⁷³ For present purposes, we may put aside Knowles's treatment of Fregean modes of presentation insofar as it ultimately leads him to deny or doubt their existence.⁷⁴ Since intentional realism is thus suspended in this case, this particular set of considerations does not really threaten (I) after all. The discussion of propositions and broad content is different in this respect. If Knowles is right that to assume the existence of causally efficacious content does not amount to LOT, then we may say with Popperian eyes, (I) would be falsified right away. Let us thus consider the argument from broad content [B], which is of greater generality:

First premise: Content is broad. Beliefs are, as such, about the environment of the believer.

Second premise: Relations to environmental objects are not causally efficacious in the production of behaviour.

Conclusion: Local causal surrogates corresponding to the environmental objects explain how beliefs can have causal powers in line with their contents.⁷⁵

⁷¹ Knowles (2001, 2002).

⁷² Knowles (2001, 350).

⁷³ Knowles (2001, 350-360).

⁷⁴ Knowles (2001, §3.2).

⁷⁵ Cf. Knowles (2001, §3.1).

Knowles uses the expression ‘local causal surrogate’ as equivalent to ‘mental representation,’ that is to say, states of the organism with semantic and syntactic properties. The bulk of Knowles’s argument against [B] is to accept both premises but then reject the conclusion on the grounds that we need not posit such local causal surrogates given that the environmental objects themselves are the ones responsible for the causal powers of belief:

My being related to these things [the environmental objects] may not be causally efficacious, but that doesn’t seem to stop the objects being efficacious themselves in my behaviour, nor the thinker’s dispositions in relations to these objects conforming to those we would expect from a belief.⁷⁶

Knowles is surely right that [B] does not show that “the environmental features implicated in broad contents can’t themselves be causally efficacious in intentional causation.”⁷⁷ All the same, in order to make sense of this suggestion, we must note that environmental features, as such, would lack the kind of specificity and distinctiveness characteristic of intentional contents. By ‘specificity’ I do not mean simply the fact that content is usually understood in intensional, rather than extensional terms—something that Knowles may be willing to impugn along with modes of presentation generally.⁷⁸ Rather, environmental objects as such lack specificity in the sense that they lack the distinctive and unambiguous character needed in order to make sense of intentional causation in the first place. Roughly, it is surely not water (H₂O) itself what causes me to drink; it is the fact that I represent water as being drinkable, together with certain physiological and practical vicissitudes, which does. But granted that intentional causation relies on such specificity of representation, it is hard to see how the resulting picture Knowles proposes is still an anti-LOT picture, short of denying the intentional realist stance explicitly embraced by him. And this is so no matter what amount of extendedness about content one’s theory of mind allows and, in particular, no matter how non-local the corresponding causal surrogates are. “No Intentional Causation without Explicit Representation” used to be Fodor’s pro-LOT battle cry.⁷⁹ Knowles does not provide any reasons to suspend the battle cry in the case in which representation is however non-local.

The second main line of argument considered and rejected in Knowles’s developments takes naturalization and cognitive science as the starting point. Here is a reconstruction of this kind of argument [CS]:

First premise: There are intentional regularities—such as productivity, systematicity or inferential coherence—that must obtain in a physical world.

⁷⁶ Knowles (2001, 354).

⁷⁷ Knowles (2001, 354).

⁷⁸ Knowles (2001, §3.2).

⁷⁹ Fodor (1987, 25).

Second premise: LOT provides a neat (perhaps only) explanation of these regularities by recourse to a physical language-like medium of computation.

Conclusion: LOT is true.⁸⁰

Knowles's worry about [CS] is announced to be about the alleged need to explain the target intentional regularities in straightforwardly physicalistic terms. He acknowledges some pressure to explain how the high/capacity level and the physical/lower level are related but contends that:

... it is important to stress that accepting this does not commit one to the idea that the kind of implementation constraints the neurophysiological details may impose on the theory of a cognitive capacity will exhaust the content of, let alone supplant the need for, the higher level theory. The felicity of the latter is not dependent on uncovering any neat mapping between it and neurophysiology.⁸¹

Knowles does not restrict his discussion to the quoted considerations but, as far as I can make out, other points and suggestions are largely inessential for the assessment of [CS]. Knowles's fundamental worry in this respect is that, as a consequence of the denial of a straightforward or reductive physicalism, there is no sound argument from (intentional explanations in) cognitive science to the truth of LOT.

The cogency of Knowles's objection to [CS] depends, therefore, on the assumption that LOT requires a very direct physical relation between high level, intentional theory, and physical or implementational theory. In the quoted paragraph, he also suggests that the truth of LOT would somehow undermine high level, intentional theories to the extent that they would be supplanted by or reduced to implementational LOT theories. However, as noted above (Sections 2 and 3), prominent defences of LOT have endorsed non-reductive forms of physicalism.⁸² This kind of physicalism is also contemplated in García-Carpintero's developments, cited by Knowles.⁸³

More generally, and independently of the explicit claims in the literature, there is no obvious reason that we should accept that LOT demands such a direct explanatory relationship between physical correlates and intentional categories. It is enough—and indeed, a substantial finding were it to be true—that there is some range of states that could be identified as the supervenience base and hence the (multiple) realization of the causal explanatory categories at the intentional level. From all this, it seems fair to conclude that resistance to [CS] is unjustified on the ground of a too-demanding inter-level explanatory relationship.

⁸⁰ Cf. Knowles (2001, 360-372).

⁸¹ Knowles (2001, 367).

⁸² Fodor (1987), Davies (2000a), (2000b), (2004).

⁸³ Cf. García-Carpintero 1995, 371 and 379.

5. Conclusion: A Flexible but Substantial Commitment

It is time to assess the impact of the detachable commitment to LOT (I) and (C) highlight. As shown, both Peacocke's and Knowles's resistance to (I) can be seen as stemming from attributing to LOT features or theses that go well beyond those brought out in (L) (such as explicit representation of rules, local character of representation or physical reductionism). This may leave the impression that the commitment of LOT, as spelled out in (I) and (L), is no heavy burden or of no substantial empirical import. There are, in general, at least two sources for this impression.

First, LOT is, at best, only part of the story about cognition and is, for all we have seen, a commitment concerning a (perhaps highly) restricted range of mental phenomena, namely, those for which we need the recourse to intentional explanations. It is quite plausible that a LOT research programme can only yield a minimally satisfactory account of cognition in combination with alternative research programmes, such as connectionism⁸⁴ or computational neuroscience.⁸⁵ Second, the many arguments conforming to (I) are all existential quantification arguments that deliberately avoid entering into the details of the hypothesis. This generality of LOT crystallizes in a great flexibility. In particular, there are many specific computational frameworks concerning radically different cognitive domains that would constitute concrete and highly disparate specifications of a language of thought.

To acknowledge both the incomplete character and flexibility of LOT in this sense links up with the task of freeing the hypothesis from claims and theses that do not belong to the essence of the explanations it proposes (as illustrated in Section 4). However, incomplete character and flexibility need not mean lack of empirical significance. Once inessential features are factored out, LOT is still a substantial empirical hypothesis, one commitment to which theorists must be wary not to ignore.

Since my focus is on the implementation (rather than the abstract/general statement) of language-like computation relative to specific intentional explanations concerning organisms (O) and cognitive tasks (T), the relevance and substantial character of LOT can be fully appreciated once we lay out several competing empirical hypothesis that would be discarded, were such commitment to hold good in a particular case.

Most obviously, if LOT is true of O regarding T, then extreme anti-computationalist approaches in cognitive science would be false for O and T. This would result in the rejection of a family of 'post-cognitivist' theories whose number and growing significance is hard to overemphasize.⁸⁶ They include

⁸⁴ Marcus (2001).

⁸⁵ Eliasmith and Anderson (2003).

⁸⁶ Wallace et al. (2007).

radical dynamicism, embodied and enactive approaches which, in a variety of ways, are united by the dismissal of computationalism and representation.

Besides, and as often emphasized at the roots of the systematicity debate, if LOT is true of O regarding T, then cognition cannot be *purely* connectionist computation for O and T. Connectionist computation in this respect would at best be part of the implementation of language-like digital computation and thus, from among the many varieties of the connectionist paradigm, the connectionist models to be seriously considered for O and T would only be the ones that can accommodate symbolic as opposed to subsymbolic computation. The same would go for the models in computational neuroscience or neurophysiologically and neuroanatomically constrained neural networks.⁸⁷

Progress in cognitive science and our improved understanding of the notion of computation allows us to point out further LOT incompatibilities *within* the computational paradigm. For instance, LOT cannot consistently be combined with views that opt for a non-representational or non-information-processing kind of computation. As several authors have pointed out, computation does not require information processing or representation.⁸⁸ However, if LOT is true of O regarding T, we would have to regard non-representational or non-informational notions of computation as irrelevant for O and T because LOT consists of the processing of language-like representational information.

Finally, to the extent that LOT is true of O regarding T, it cannot be the case that O implements analog computation regarding T, where analog computation is understood as the manipulation of continuous variables that take any real values in systems of differential equations (see also Section 4.1). If O implements language-like digital computation in relation to T, then it cannot implement analog computation for T.

In conclusion, the truth of LOT regarding a particular cognitive task of an organism would rule out a significant number of views about cognition and computation in spite of its acknowledged incomplete character and flexibility. This is so even if we free LOT of all inessential commitments that result in stricter and misleading versions of the hypothesis. If this is sound, it would seem that (I) highlights a substantial and widespread empirical commitment of contemporary theorizing.

Acknowledgements: Thanks are due to Gualtiero Piccinini and four anonymous referees for this journal for their comments on this paper. Particular thanks are owed to Daniel Quesada and Manuel García-Carpintero, whose thinking first put me on to the problems explored here. This work has been supported by the Secretary for Universities and Research of the Ministry of Economy and Knowledge of the Government of Catalonia, the Spanish Ministry of Economy

⁸⁷ O'Reilly and Munakata (2000), Piccinini and Bahar (2012).

⁸⁸ Piccinini and Scarantino (2011), Piccinini (2009, 2015).

and Competitiveness, via the research projects FFI2013-41415-P and FFI2015-63892-P (MINECO/FEDER, UE), and the consolidated research group GRECC SGR2014-406.

References

- Aizawa, Ken
2003 *The Systematicity Arguments*. Dordrecht: Kluwer Academic Press.
- Aydede, Murat
1997 "Language of Thought: The Connectionist Contribution." *Minds and Machines* 7: 57–101.
- Aydede, Murat
2010 "The Language of Thought Hypothesis." In E.N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. URL=<<http://plato.stanford.edu/archives/fall2010/entries/language-thought/>>
- Calvo, Paco and John Symons (eds.)
2014 *The Architecture of Cognition: Rethinking Fodor and Pylyshyn's Systematicity Challenge*. Cambridge, MA: MIT Press.
- Camp, Elisabeth
2007 "Thinking with Maps." *Philosophical Perspectives* 21: 145–182.
- Carruthers, Peter
2006 *The Architecture of the Mind: Massive Modularity and Flexibility of Thought*. Oxford: Oxford University Press.
- Cartwright, Nancy D.
2006 "From Causation to Explanation and Back." In B. Leiter (ed.) *The Future of Philosophy*. Oxford Clarendon Press, 230–245. Also as: Causality: Metaphysics and Methods. Technical Report CTR 09-03, CPNSS, LSE.
- Chemero, Anthony
2009 *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.
- Churchland, Paul M.
1989 *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge, MA: MIT Press.
- Churchland, Patricia S. and Terrence Sejnowski
1989 "Neural Representation and Neural Computation." In L. Nadel, L.A. Cooper, P. Culicover and R.M. Harnish (eds.) *Neural Connections, Mental Computation*. Cambridge, MA: MIT Press, 15–48.
- Crane, Tim
1990 "The Language of Thought: No Syntax Without Semantics." *Mind and Language* 5: 187–212.
- Crane, Tim
1992 "Mental Causation and Mental Reality." *Proceedings of the Aristotelian Society* 92: 185–202.
- Crane, Tim and D.H. Mellor
1990 "There is No Question of Physicalism." *Mind* 99: 185–206.

- Craver, Carl F. and William Bechtel
 2006 “Mechanism.” In S. Sarkar and J. Pfeifer (eds.) *Philosophy of Science: An Encyclopedia*. New York: Routledge, 469–478.
- Cummins, Robert
 1996 “Systematicity.” *Journal of Philosophy* 93: 591–614.
- Cummins, Robert, James Blackmon, David Byrd, Pierre Poirier, Martin Roth and Georg Schwarz
 2001 “Systematicity and the Cognition of Structured Domains.” *Journal of Philosophy* 98: 167–185.
- Davies, Martin
 1991 “Concepts, Connectionism, and the Language of Thought.” In W. Ramsey, S. Stich and D. Rumelhart (eds.) *Philosophy and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates, 229–257.
- Davies, Martin
 2000a “Persons and their Underpinnings.” *Philosophical Explorations* 3: 43–62.
- Davies, Martin
 2000b “Interaction without Reduction: The Relationship between Personal and Sub-personal levels of Description.” *Mind and Society* 2, 87–105.
- Davies, Martin
 2004 “Aunty’s Argument and Armchair Knowledge.” In J.M. Larrazabal and L.A. Pérez Miranda (eds.) *Language, Knowledge, and Representation*. Dordrecht: Kluwer Academic Publishers, 19–37.
- Davies, Martin
 2015 “Knowledge—Explicit, Implicit and Tacit: Philosophical Aspects.” In J.D. Wright (ed.) *International Encyclopedia of the Social and Behavioural Sciences. 2nd Edition*. Oxford: Elsevier, 74–90.
- Dennett, Daniel
 1992 “The Self as a Center of Narrative Gravity.” In F.S. Kessel, P.M. Cole and D.L. Johnson (eds.) *Self and Consciousness: Multiple Perspectives*. Hillsdale, NJ: Erlbaum Associates, 103–115.
- Devitt, Michael
 1990 “A Narrow Representational Theory of the Mind.” In W.G. Lycan (ed.) *Mind and Cognition*. Oxford: Basil Blackwell, 369–402.
- Devitt, Michael
 1996 *Coming to Our Senses*. Cambridge: Cambridge University Press.
- Egan, Frances
 2012 “Representationalism.” In E. Margolis, R. Samuels and S.P. Stich (eds.) *The Oxford Handbook of Philosophy of Cognitive Science*, Chapter 11, 250–272.
- Eliasmith, Chris and Charles H. Anderson
 2003 *Neural Engineering: Computation, Representation and Dynamics in Neurobiological Systems*. Cambridge, MA: MIT Press.

- Evans, Gareth
1981 "Semantic Theory and Tacit Knowledge." In A. Phillips (ed.) *Collected Papers*. Oxford: Oxford University Press, 118–137.
- Evans, Gareth
1982 *The Varieties of Reference*. Oxford: Oxford University Press.
- Field, Hartry
1978 "Mental Representation." *Erkenntnis* 13: 9–61.
- Fodor, Jerry A.
1975 *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Fodor, Jerry A.
1987 *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Fodor, Jerry A.
1990 *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- Fodor, Jerry A.
2000 *The Mind Doesn't Work that Way: The Scope and Limits of Computational Psychology*. Cambridge, MA: MIT Press.
- Fodor, Jerry A.
2008 *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.
- Fodor, Jerry A. and Brian P. McLaughlin
1990 "Connectionism and the Problem of Systematicity: Why Smolensky's Solution Doesn't Work." *Cognition* 35: 183–204.
- Fodor, J.A. and Zenon Pylyshyn
1988 "Connectionism and Cognitive Architecture: A Critical Analysis." *Cognition* 28: 3–71.
- Fresco, Nir
2012 "The Explanatory Role of Computation in Cognitive Science." *Minds and Machines* 22: 353–380.
- García-Carpintero, Manuel
1995 "The Philosophical Import of Connectionism: A Critical Notice of Andy Clark's *Associative Engines*." *Mind and Language* 10: 370–401.
- García-Carpintero, Manuel
1996 "Two Spurious Varieties of Compositionality." *Minds and Machines* 6: 159–172.
- Garson, James W.
1997 "Syntax in a Dynamic Brain." *Synthese* 110: 343–355.
- Gomila, Antoni, David Travieso and Lorena Lobo
2012 "Wherein is Human Cognition Systematic?" *Minds and Machines* 22: 101–115.
- Hadley, Robert F.
1994 "Systematicity in Connectionist Language Learning." *Mind and Language* 9: 247–272.

Hadley, Robert F.

- 2004 "On the Proper Treatment of Semantic Systematicity." *Minds and Machines* 14: 145–172.

Harman, Gilbert

- 1973 *Thought*. Princeton, NJ: Princeton University Press.

Hempel, Carl G. and Paul Oppenheim

- 1948 "Studies in the Logic of Explanation." *Philosophy of Science* 15: 135–175.

Horgan, Terence E. and John L. Tienson

- 1996 *Connectionism and the Philosophy of Psychology*. Cambridge, MA: MIT Press.

Johnson, Kent

- 2004 "On the Systematicity of Language and Thought." *Journal of Philosophy* 101: 111–139.

Kaye, Lawrence J.

- 1995 "The Languages of Thought." *Philosophy of Science* 62, 92–110.

Knowles, Jonathan

- 2001 "Does Intentional Psychology Need Vindicating by Cognitive Science?" *Minds and Machines* 11: 347–377.

Knowles, Jonathan

- 2002 "Is Folk Psychology Different?" *Erkenntnis* 57: 199–230.

Lycan, William G.

- 1993 "A Deductive Argument for the Representational Theory of Thinking." *Mind and Language* 8: 404–422.

Maloney, Christopher J.

- 1984 "The Mundane Mental Language: How to Do Words with Things." *Synthese* 59: 251–294.

Marcus, Gary F.

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

Marr, David

- 1982 *Vision*. San Francisco, CA: Freeman.

McLaughlin, Brian P.

- 1993 "The Connectionism/Classicism Battle to Win Souls." *Philosophical Studies* 71: 163–190.

McLaughlin, Brian P.

- 2009 "Systematicity Redux." *Synthese* 170: 251–274.

O'Reilly, Randall C. and Yuko Munakata

- 2000 *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press.

Peacocke, Christopher

- 1986a "Explanation in Computational Psychology: Language, Perception and Level 1.5." *Mind and Language* 1: 101–123.

Peacocke, Christopher

- 1986b "Replies to Commentators." *Mind and Language* 1: 388–402.

- Peacocke, Christopher
1989 "When Is a Grammar Psychologically Real?" In G. Alexander (ed.) *Reflections on Chomsky*. Oxford: Basil Blackwell, 111–130.
- Peacocke, Christopher
1992 *A Study of Concepts*. Cambridge, MA: MIT Press.
- Peacocke, Christopher
1994 "Content, Computation and Externalism." *Mind and Language* 9: 301–335.
- Peacocke, Christopher
2004 "Interrelations: Concepts, Knowledge, Reference and Structure." *Mind and Language* 19: 85–98.
- Peacocke, Christopher
2008 *Truly Understood*. Oxford: Oxford University Press.
- Piccinini, Gualtiero
2009 "Computationalism in the Philosophy of Mind." *Philosophy Compass* 4: 515–532.
- Piccinini, Gualtiero
2015 *Physical Computation*. Oxford: Oxford University Press.
- Piccinini, Gualtiero and Sonya Bahar
2012 "Neural Computation and the Computational Theory of Cognition." *Cognitive Science* 34: 453–488.
- Piccinini, Gualtiero and Andrea Scarantino
2011 "Information Processing, Computation, and Cognition." *Journal of Biological Physics* 37: 1–38.
- Port, Robert F. and Tim van Gelder
1995 *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.
- Pour-El, Marian Boykan
1974 "Abstract Computability and its Relation to the General Purpose Analog Computer." *Transactions of the American Mathematical Society* 199: 1–28.
- Pylyshyn, Zenon
2003 *Seeing and Visualizing: It's Not What You Think*. Cambridge, MA: MIT Press.
- Ramsey, William, Stephen Stich and Joseph Garon
1991 "Connectionism, Eliminativism, and the Future of Folk Psychology." In W. Ramsey, S. Stich and D. Rumelhart (eds.) *Philosophy and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates, 199–228.
- Rescorla, Michael
2009 "Cognitive Maps and the Language of Thought." *British Journal for the Philosophy of Science* 60: 377–407.
- Rey, Georges
1991 "An Explanatory Budget for Connectionism and Eliminativism." In T. Horgan and J. Tienson (eds.) *Connectionism and the Philosophy of Mind*. Dordrecht: Kluwer Academic Publishers, 219–240.

- Rey, Georges
1995 "A Not Merely Empirical Argument for a Language of Thought." *Philosophical Perspectives* 9, 201–222.
- Rey, Georges
1997 *Contemporary Philosophy of Mind: A Contentiously Classical Approach*. Oxford: Blackwell.
- Rumelhart, David E., James M. McClelland and the PDP Research Group
1986 *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT Press.
- Schneider, Susan
2009 "LOT, CTM, and the Elephant in the Room." *Synthese* 170: 235–250.
- Schneider, Susan
2011 *The Language of Thought. A New Philosophical Direction*. Cambridge, MA: MIT Press.
- Schneider, Susan and Matthew Katz
2012 "Rethinking the Language of Thought." *WIREs Cognitive Science* 3: 153–162.
- Schröder, Jürgen
1998 "Knowledge of Rules, Causal Systematicity, and the Language of Thought." *Synthese* 117: 313–330.
- Shagrir, Oron
2001 "Content, Computation and Externalism." *Mind* 110: 369–400.
- Smolensky, Paul
1988 "On the Proper Treatment of Connectionism." *Behavioural and Brain Sciences* 11: 1–23.
- Stich, Stephen
1983 *From Folk Psychology to Cognitive Science*. Cambridge, MA: MIT Press.
- Verdejo, Victor M.
2012a "Meeting the Systematicity Challenge Challenge: A Nonlinguistic Argument for a Language of Thought." *Journal of Philosophical Research* 37: 155–183.
- Verdejo, Victor M.
2012b "The Visual Language of Thought: Fodor vs. Pylyshyn." *Teorema* 31: 59–74.
- Verdejo, Victor M.
2015 "The Systematicity Challenge to Antirepresentational Dynamicism." *Synthese* 192: 701–722.
- Verdejo, Victor M.
Manuscript. "Determinability of Perception as Homogeneity of Representation."
- Verdejo, Victor M. and Daniel Quesada
2011 "Levels of Explanation Vindicated." *Review of Philosophy and Psychology* 2, 77–88.

Von Eckardt, Barbara

1993 *What is Cognitive Science?* Cambridge, MA: MIT Press.

Wallace, Brendan, Alastair Ross, John Davies and Tony Anderson

2007 *The Mind, the Body and the World: Psychology after Cognitivism?*
Exeter: Imprint Academic.