

ON LANGUAGE AND THOUGHT: A QUESTION OF FORM

DAVID J. LOBINA

dj.lobina@gmail.com

LOGOS/Faculty of Philosophy,
University of Barcelona

ABSTRACT. The literature on the relationship between language and thought is immense, but it is not always clear what phenomenon exactly most publications are in fact examining – or whether the various studies are even compatible. This is mainly due to two interrelated reasons: the interested scholars seldom offer clear definitions for language and thought, and the terms of comparison between these two phenomena are hardly ever specified. The aim of this paper is twofold: a) to showcase the perils of this state of affairs by analyzing a particular but little considered type of mismatch between linguistic and conceptual representations, thus providing the primer for a competence-like analysis of thought, and in so doing b) defend a specific way to approach the language and thought nexus, prior to and above all others. In particular, I argue that the study of *the relationship* should give priority to a comparative analysis of the representational vehicles language and thought exhibit in order to work out what thought representations are actually like.

Keywords: language; thought; representational vehicle; level of explanation; question

How to cite: Lobina, David J. (2019). “On Language and Thought: A Question of Form,” *Linguistic and Philosophical Investigations* 18: 37–63.

Received 6 November 2017 • Received in revised form 23 January 2018

Accepted 24 January 2018 • Available online 15 February 2018

1. The Question Now Being Put

The question of how language and thought relate is possibly time immemorial, and the extant literature too large to review satisfactorily in a single paper. Such a description of the general state of affairs won't surprise anyone; what may raise a few eyebrows, perhaps, is the claim that large tracts of the modern literature are afflicted by two rather important problems. Firstly, scholars focused on this relationship are hardly ever clear as to what they take language and thought to be exactly, as definitions are rarely put forward; and complementing this, the very same scholars are usually rather vague regarding how language and thought are supposed to relate at all.

Myriad possibilities can be imagined in either respect, but identifying the different prospects necessitates significant exegetical effort and the rewards are to be found outside the language-and-thought literature. Natural language, to begin here, may be defined in many different ways: as a computational system of the mind, along with its connection to other mental systems (Chomsky, 2013); as a set of mental representations, perhaps in terms of *propositional attitudes* (Fodor, 1979, 2001a); as a purely communicative system, thereby closely connected to *mind reading* and *joint attention* abilities (Tomasello, 2003); as an exclusively real-time processing phenomenon (Frank, Bod, & Christiansen, 2012); as a symbolic system (Deacon, 1997); etc.

Whatever perspective one adopts, the particular choice will surely have a significant effect on what sort of issues are considered in the study of how language and thought are actually related. In fact, it is not an uncommon feature of the language-and-thought literature that the eventual terms of comparison employed in many a study are in fact the result of the conclusion to an argument on how to relate the two phenomena rather than a principled way to approach the problem at hand. Among the different possibilities, we find the following, *prima facie* incompatible options: linguistic (viz., syntactic) representations may be the actual vehicles of the main *medium of thought* humans employ (Hinzen, 2013); thought may be in fact impossible without language, this capacity necessitating the representation of the propositional attitudes, an ability some scholars have argued to be only achievable by employing language (Davidson, 1997, 2001; Millikan, 2001); language may connect different conceptual systems of the mind via the (syntactic) level of what used to be called *logical form*, these systems otherwise unconnected in the absence of a fully formed language faculty (Carruthers, 2002); the employment of language in speech (either inner or outer) may augment one's computational powers during explicit mental processes, language therefore at least enhancing (but probably not replacing) thought (see Clark's contribution in Carruthers & Boucher, 1998); natural language may be in fact inadequate as a system of thought, its main function being the *communication* of thoughts rather than its *representation*, thought therefore constituting a (perhaps slightly) different mental phenomenon altogether (Fodor, 1975, 2001b, among others); etc.

The result is that scholars are seemingly talking about very different phenomena when discussing the reputed relationship, from representational and architectural issues to processing phenomena. What's more, these scholars also differ in terms of how central they take language to be for thought, from those who defend that the most sophisticated forms of thought are conducted in natural language (Davidson, Carruthers), in some cases concluding that there is no *medium of thought* other than language, thus cutting the distinction very loose indeed (Hinzen), to those who believe that thought is instead conducted in a medium *other than* natural language, language and thought thus (again, perhaps slightly) different domains of the mind (Fodor).

It would certainly be a good thing to catalogue the different arguments (and data) to make sense of what they tell us about human thought *tout court*, but my objectives are focused elsewhere here. I am most interested in analyzing the *format* properties of linguistic and conceptual representations in order to work out whether language can be regarded as being constitutive of thought – to work out, that is, if language exhibits the right features to be regarded as the (main) medium of thought. To that end, I will analyze a specific type of mismatch between language and thought, and in so doing, provide the internal structure of a specific conceptual representation: the thought representation underlying an interestingly deviant English sentence. Naturally, it will be necessary to be rather clear as to what *language*, *thought*, and indeed *being a medium of thought* are supposed to mean in all this, and ample space shall be devoted to that in what follows. Assuming I have done all this right, I should have developed the following objectives: a) to argue in favor of both a certain way of defining language and thought and a certain way of relating them; b) to provide a primer as to what sort of structure, or shape, thought representations exhibit; and, c) to show that when one focuses on the deceptively simple question of figuring out what thought underlies a given sentence, a number of rather different issues unexpectedly surface, oftentimes pointing to different aspects of mental architecture.

I shall proceed in two stages. Firstly, I shall put together the machinery needed to relate language and thought; that is, I will define what I will take language and thought to be and I will claim that these two phenomena apparently involve slightly different mental systems. After that, I will propose a specific way to relate language and thought directly and I will exemplify the overall framework by discussing the aforementioned English sentence. According to this layout, the essay is thematically divided into two sections, neither more central than the other, but they certainly form a progression: the next section sets the stage for the second half. The essay ends by mentioning some additional linguistic data that further suggest that language and thought are indeed underlain by slightly different mental systems, and a reference is made to future work along these lines.

2. On Language and Thought

By a natural language I will here understand, following the line defended by the generative grammar enterprise in the last 60 or so years, a *state* of the faculty of language (FoL, in short), an architectural system of the mind composed of the following components: a combinatorial operation (recently termed *merge*); a set of lexical items (bundles of syntactic, phonological/morpho-phonemic, and, perhaps, semantic features); and two interfaces that connect the faculty with other systems of the mind, namely the sensorimotor (SM; roughly, the sound system) and the conceptual-intentional (C/I; even more roughly, the thought/meaning systems). The FoL is said to be in an *initial state* at birth, only to change during the language acquisition process in reaction to the input, ending in a *mature state*, the result of a

one-time selection of lexical items and how this particular selection interacts with a number of computational principles and interface conditions (see Chomsky, 2008, for a succinct description, and Yang, 2002, for how language acquisition might actually proceed).¹

Thus, to know a natural language, say English, is to be in possession of a particular state of the FoL that more or less matches the mature FoL states of other English speakers. It is possibly the case that no two English speakers possess the exact same system of rules and representations (another way to refer to the *states* of the FoL), and the faculty clearly does not cease to change during one's lifetime – by mature we certainly do not mean *final*. Nevertheless, the description just given appropriately characterizes what is usually meant by the phrase “knowing a language,” and that will suffice for the purposes of this paper (see, for some discussion and relevant references, Collins, 2008).

Such a conception, moreover, captures the fact that there are different aspects to language – syntactic, phonological, or at the interface with external systems – and whilst the full richness of language may require we look at the faculty in its entirety, many linguistic properties will be evident in some of its sub-systems. In fact, the different definitions of language advanced in the previous section – as a communicative system, as a set of propositional attitudes, etc. – implicate some of the sub-systems and not always necessarily the overall language faculty. In this sense, what we call Language, the overall linguistic capacity, may simply be the combination of its various components.

Language so described, what are the products of such a mental system? The language faculty generates a number of rather different sort of representations, two being the most central: one for the SM system, termed a PHONetic representation (a flat and linear object ready for externalisation, be in speech or any other modality), and one for the C/I systems, this one called SEMantic (a hierarchical object encompassing the containment relations inherent in a syntactic object; despite its name, a SEM representation is a syntactic object and thus it is not to be confused with the sort of representations semanticists postulate and work with; see footnote 2). These two representations are compiled from different sets of lexical features (morpho-phonemic features in one case, syntactic in the other), thereby yielding different types of representations and computations (or derivations, as linguists typically call them), the latter possibly the result of the different constraints the SM and C/I systems impose upon the underlying generative procedure. What are these constraints precisely? The very nature of the physical channel in which we communicate accounts for why the SM systems require a flat and linear representation for externalization (in either inner or outer speech, and in fact in other modalities too), whilst the C/I systems demand a structured object, given that syntactic structure goes a long way towards establishing the meaning and interpretation of a sentence (in addition, one would need to take into account the surrounding context, speakers intentions, and other such things, but I won't get into any of this here).²

In describing the language faculty in such manner, one is naturally pointing to the computational-representational paradigm so dominant in studies of the mind. According to this perspective, most of cognition involves computational processes over mental representations (or mental symbols), a view of the mind that linguists have actively elaborated, albeit in their own way. Indeed, the computations linguists have postulated are not meant to take place in real time, the focus of much of cognitive science; rather, linguistic derivations constitute abstract specifications of the principles and constraints underlying linguistic structure. The linguist's is a study of *competence*, whereas most of cognitive science presumably focuses on the actual processing mechanisms operative in behavior, the study of *performance* (these terms, as is well known, were introduced by Chomsky, 1965, but I will employ them somewhat differently here; Chomsky himself has moved on from the particular description given in his 1965 book). Fodor (1983), a foundational work of the computational and representational theories of mind, is a study of some of the processing mechanisms the mind employs in behavior, and has little to say about architectural/representational systems such as the language faculty (the relevant interconnections notwithstanding) and many a textbook in cognitive science paint a similar picture (e.g., Gallistel & King, 2009).

The distinction between abstract computations and real-time processes can be found in the formal sciences as well (viz., in mathematical logic and computer science); it is not an exotic curiosity of cognitive science. The computations the language faculty is said to effect are closer in nature to a logical proof, or to what some mathematicians call a *theory of the computation* (McCarthy, 1963), an “abstract formulation of *what* is being computed and *why*” (Marr, 1977: 37), to quote a scholar who was greatly influenced by the formal sciences (much like Chomsky, in fact; see Tomalin, 2006, and Lobina, 2017, for some details). Such a perspective would appear to fit well with the study of linguistic competence, with its focus on *what* language qua systematic body of knowledge is like, a study that is conducted in generative (i.e., computational) terms in order to capture language's open-ended nature. The processes that engage most of cognitive science are instead closer to what computer scientists call *models of computations* (Moschovakis, 2001), step-by-step computations that typically proceed in real time, as when a particular process is executed. As I shall explain below, I will keep to the more abstract computational theory to connect language and thought, even if some processing/behavioral factors will have to be discussed too.

The notion of a representation is also treated inconsistently among the disciplines encompassing cognitive science; in particular, this term means something rather different to a philosopher than it does to a linguist. For the former, a mental representation depicts an aspect of the world, and thus it must be veridical and accurate if it is to play a role in thinking – indeed, if it is to be considered a mental reality at all. Linguists, however, treat representations as nothing more than internal data structures that play specific roles within linguistic derivations. A linguistic representation is simply either the input or the output of a linguistic operation and

they are composed of lexical features that cannot be analyzed relationally to externalia (Collins, 2004). Indeed, the lexical features, or combinations thereof, that linguistic derivations manipulate cannot be said to represent any external reality; there is no *wh* anything out there, to point to the ubiquitous *wh*-feature of linguists; thus, there can be no question as to a lexical feature's accuracy or veridicality. In this case too, I shall adopt the linguist's usage when talking of representations, but the actual points to be discussed will be very central to the philosopher.

Representational and computational theories of mind have much going on in their favor. It has been a cornerstone of cognitive science that human behavior is to be explained in terms of the internal representations that mediate between environmental stimuli and behavioral responses (including how these internal representations combine with each other), a sort of mentalism that is meant to account for the degree of freedom there often is between stimuli and responses (Fodor, 1975: 157; Laurence and Margolis, 1997: 65). According to this reasoning, the key to explaining human cognition lies on the internal properties of the mind's medium of computation and not on the stimulus-response pairs themselves, the units of explanation that once-upon-a-time behaviorism would have us focused on. Crucially, this sort of story provides the appropriate basis for a definition of thought (and thinking); and in turn for relating language and thought.

I wouldn't want to suggest, however, that all of cognition can be explained in these terms. In particular, I am conscious of Fodor's argument, which I take to go some way, that if a system of the mind is not modular, then it is very likely that we will not be able to scientifically study it at all (Fodor, 1983). This is especially true for what I take the *activity* of thought to be about primarily – namely, the fixation of belief that constantly goes on in mental life – the very process Fodor has argued to be non-modular and non-computational (Fodor, 2000). That may be so; my interest, however, does not lie on the *activity* of thinking per se (a study of *performance*, effectively), which is what I take Fodor to have been talking about anyway, but on the sort of representations that support thinking (in a way, a *competence* matter). That is, I'm interested in the sort of *format* and *form* that thought representations must have in order to support the fixation of belief. Thus, I'm proposing to characterize *thought* in terms of its representational vehicles along with the contents these vehicles typify, this in itself a line Fodor has argued for extensively. The famed relationship, therefore, will be here framed in terms of the types of representations language and thought exhibit, and thus the focus will fall upon the format, or shape, of mental representations (the actual details will be provided in the next section).

Thought, then, is to be regarded as being closely connected to what philosophers usually call content, or a proposition, the kind of objects over which propositional attitudes range, and the objects of beliefs themselves. To have a thought, according to this view, is to entertain a proposition (or a group of propositions), while to think is to combine propositions in various ways, from embeddings and combinations of various kinds (via connectives, for instance) to the premises-and-conclusion organi-

zation typical of reasoned thought. To this end, a proposition must be fully explicit and complete, bear truth values, and (to my mind) exhibit a constituent structure (viz., a predicate with its arguments).³

A proposition's predicate-argument structure will be the crucial issue here; consider the following two examples in order to clarify (the first typical of introductions to formal logic, the second due to Steintal, as quoted in Seuren, 2009).

- (1) Caesar invaded Gaul.
- (2) Coffee grows in Africa.

From the point of view of what proposition underlies (1), the sentence ascribes the property *having been invaded to Gaul*, and thus the predicate would be *having invaded* (or something like that), with *Cesar* and *Gaul* as the arguments – in schematic form: < having invaded < caesar, gaul >> (I use angle brackets to represent propositions and Polish notation to mark the predication). It is quite immaterial how this fact is grammatically expressed (in the active or passive voice, with a cleft construction, etc.), as the state of affairs would be the same in every case: the proposition *Caesar's invasion of Gaul*. Following Seuren (2009), but paraphrasing a bit, a proposition is the act of ascribing a property to a mental representation. Regarding (2), and again, from the point of view of the underlying proposition, we find the property of *being in Africa*, the predicate, being ascribed to the argument *the growing of coffee*. This certainly contrasts with the sentence's syntactic structure, where *coffee* would be the subject of the sentence and *grows in Africa* the predicate along with its argument. Instead, the structure of the proposition would have to be something like this: < being in Africa < the growing of coffee >>. Thus, predicates and their arguments (or subjects) would apply rather differently in each domain.⁴

A rather apposite way to highlight the mismatch is in terms of the two-layered structural analysis described in Seuren (2009). According to Seuren, sentences exhibit two layers of structure, corresponding, I add, to two different types of analysis. The analysis of the first layer would be focused on outlining the grammatical subject-predicate structure of a sentence, whilst the analysis of the second layer would instead identify the *logic-like* subject-predicate structure of the underlying proposition. The first analysis would be presumably syntactic in nature and the second more conceptual and thus logic-like; a sort of progression from a sentence's syntactic structure to the corresponding (logical) structure in thought. We would have to add a semantic level to this picture in order to mediate between the two layers, especially for the purposes of this paper. Indeed, it is customary in semantics to run meaning composition off the syntactic structure of sentences, and I submit that the information yielded at this particular juncture – the syntax-semantics interface – ought to be particularly useful in order to characterize conceptual representations properly. What I am suggesting is that differences in meaning point to specific differences in conceptual representations, and this is precisely the theoretical space I am keen to explore here, which I shall approach by paying particular

attention to the different structural relations each domain exhibits – syntax and semantics, on the one hand, the conceptual and logic-like on the other (the study of pragmatics is also relevant, and I will say something about it later on).⁵

Accordingly, the example in (2) suggests that the theoretical modeling required to account for thought representations seems to be quite different from the model that is required to account for linguistic representations – or more accurately, that the structure you have to postulate in one case is rather different from the structure you postulate in the other. One could use a different notation to describe thought representations, but the difference in structure would have to be accounted for either way. Note, also, that the fact that in (1) there is a post-verbal noun phrase complement and in (2) a post-verbal prepositional phrase adjunct is quite immaterial to the point I am making – different grammatical structures for different thought representations does not establish a constitutive relationship between language and thought, given that thought representations are such regardless of much grammatical detail.

Crucially, these two examples point to some of the mismatches between linguistic and, what I take to be, conceptual representations, with the internal structure of propositions qua mental representations constituting the key issue here. Propositional constituents must allow for the flexibility and creativity involved in belief fixation, this flexibility no more evident than in the very common phenomenon in which different types of perceptual inputs can be combined with each other and with pre-existent beliefs during the construction of a thought. An explanation for cognitive flexibility may be approached by taking thought constituents to be what some philosophers and psychologists call concepts, the mental particulars that underlie propositions (Fodor, 1998). Concepts are abstract thus amodal, but also pretty stable and as a result re-usable, two general properties that ought to allow for the combination of mental representations of various kinds into ever more complex representations of a general type.⁶

The sort of mental reality I envision for propositions and concepts is a steady and permanent one; namely, they must constitute some sort of structure in long-term memory, a *type* of mental representation rather than a *token*, even though these representations would certainly be causally tokened in mental processes, belief fixation being the most general case. In this sense, by the phrase mental particular I do not mean *particular* tokenings of concepts, but simply the much more general fact that the concept SAINT, for instance, is a different mental particular to the concept SAGE. Naturally enough, these concepts can be combined into more complex mental representations, such as THE ISLAND OF SAGES AND SAINTS, but the *particular* tokening of the complex concept THE ISLAND OF SAGES AND SAINTS would remain distinct from the mental particular of the corresponding concept *type*.⁷ I also do not mean to suggest that all possible concepts and propositions are stored in long-term memory, even though many of them would certainly be; all we need here is a set of concepts and some sort of combinatory/compositional system that creates complex concepts out of primitive concepts, much as is the case

in linguistic cognition. Thus, conceptual structure is a bit like natural language – thought has a syntactic structure and a compositional semantics – and it is for this very reason that the question of whether natural language can be regarded as the main medium of thought arises at all. Also, it will have been noticed that the overall picture is very similar to Fodor’s (1975) *language of thought* story (LoT, *postea*), and I shall indeed keep referring to the LoT in what follows, at least in the narrower sense of it being a conceptual, representational system of some kind. So what sort of structure must thought representations exhibit?

Consider Gareth Evans’s Generality Constraint, according to which thoughts must be structured (Evans, 1982: 100), not in terms of their internal elements, something Evans in fact rejects, but in terms of “their being a complex of the exercise of several distinct conceptual *abilities*” (101; his emphasis). Evans is here drawing attention to the apparent fact that if one can entertain a thought in which a given property, call it *F*, can be ascribed to one individual, *a*, this is the result of two abilities: understanding *F*, understanding *a*, and applying the predicate *F* to the argument *a*. Consequently, if one also understands property *G* and individual *b* – that is, one understands the two sentences/thoughts, *Fa* and *Gb* – then there are “no conceptual barriers” (101) to entertain the sentences/thoughts *Fb* and *Ga*.

Ultimately, this is an ability “to think of an object in a series of indefinitely many thoughts” (104). The Generality Constraint is related to, but is not quite the same as, what Fodor has come to call the systematicity of thought (Fodor, 1987), the claim that our ability to entertain some thoughts is intrinsically connected to our ability to entertain similar thoughts (Fodor & Pylyshyn, 1988). According to Fodor and McLaughlin (1990), this property is a reflection of constituent structure, given that the stated similarity amongst thoughts is a matter of the form these thoughts have, their internal structure (this is precisely the aspect of systematicity Evans would not have accepted). Thus, if one can entertain the thought that $P \rightarrow (Q \rightarrow R)$, one should also *ipso facto* be able to entertain the thought that $(P \rightarrow Q) \rightarrow R$; clustered thoughts, as it were (McLaughlin, 2009: 253).⁸

From Evans’s structured cognitive abilities to Fodor’s structured propositions, but what does *all this* yield exactly? The above description provides the desiderata that a medium of thought must meet; namely, a *language of thought* must be able to: a) appropriately represent the contents of thoughts; b) accurately distinguish the contents of different thoughts; c) faithfully represent the propositional attitudes; and d) play a causal role in mental processes.⁹ Critically, these four requirements can provide the yardstick against which linguistic representations can be measured as a potential medium of thought, and in various ways. Indeed, the first three desiderata are related to matters of competence, whereas the last one involves the study of performance, and as I shall show in the next section, what one says about language-and-thought at the level of competence may well be very different from what one says about it at the level of performance.

The analyses I shall provide in the following section ought to bring this point home, but before sallying forth, it is important to note that the conception of thought

I have put together is based on a logic-like notion of what thought must be like, as discussed in reference to the sentences in (1–2). That is, I have offered a characterization of thought that is meant to be independent from what we know about language. After all, I have intimated that the most central issue of thought representations is predication – i.e., the ascription of a property to a mental object (a concept or mental representation) – and this property is certainly independent of much grammatical detail. This is rather important, for at the very least shows that it is possible to characterize thought independently of language, allowing us to then evaluate whether natural language can be regarded as *the* language of thought – an empirical issue, surely.¹⁰

I now turn to outlining how language and thought are to be connected in order to approach this question. I will analyze a deviant English sentence and focus on the first two desiderata I have listed (content representation and differentiation), the overall discussion framed in terms of representations and computations as the linguist understands these notions. As will be increasingly clear, I will eventually come to frame *the relationship* in terms of how the *language of thought* (LoT) and the *faculty of language* (FoL) interrelate qua architectural systems of the mind, but some other factors will have to be clarified along the way, as announced at the beginning.

3. On Relating Language and Thought

According to the two conceptions of a computation outlined in the previous section, there are at first sight two main ways in which to relate language and thought: at the level of competence and at the level of performance. Much as has been the case in linguistics, I will argue that the first of these can be regarded as taking precedence over the other, and I shall present them in this very order below (complicating things somewhat, the two perspectives point to yet other ways of relating language and thought, as I will show).

This more primordial way involves what Pylyshyn (1984) calls the symbolic (or syntactic) level, a level of description of cognition that mandates focusing on two interconnected issues: a) the shape, or form, of mental representations, thereby abstracting away from such questions as intentionality or truth, quintessential philosophical issues to be studied at the semantic level, as Pylyshyn calls that other plane of cognitive studies; and b) the (formal) properties that specify how symbolic primitives combine into more complex mental representations (the rules of composition, say). At Pylyshyn's syntactic level, then, the only features that matter are the abstract and formal properties of computations and representations.

Pylyshyn's syntactic level accords well with what generative linguists themselves claim to be studying; Chomsky (1986) explicitly states that it is precisely at Pylyshyn's symbolic level that the linguist proposes “a system of rules, representations, and computations” (262), the perspective that Chomsky called competence and which locates the language faculty within what Pylyshyn (1984) denominates

the “functional (or cognitive) architecture” of the mind, the conglomerate of the most basic and primitive systems of cognition. As stated in the previous section, the FoL may simply be a specific complex of some of the invariant components of mental architecture, one that would not include memory systems, the parser, and other components (these being part of a study of performance), but lexical features, *merge*, and the interfaces instead. Given the centrality of both language and thought in human cognition, it is precisely at the level of the cognitive architecture where language and thought would relate to begin with, and that shall be my remit here – an FoL-LoT interrelation. Adopting this perspective allow us to study the way language relates to other systems of the mind, including those subsuming thought, from the point of view of whether these other cognitive systems share the exact same components as language, a subset of these, or instead none at all.

The language-and-thought literature has so far neglected this viewpoint, and yet it would seem to be the principal way to relate the two phenomena. Every product of cognition is the result of how the internal mental machinery functions, and it would be just natural to focus on such machinery as our object of study. Further, focusing on mental architecture is a fairly straightforward strategy to follow: figure out what sort of mental structure supports language and then compare it to the sort of mental structure that supports thought. If they are not more or less the same, then we are likely to be talking about different mental phenomena. Naturally, the requisite mental structures cannot be directly observed in the brain and cognitive psychology certainly does not provide a clear picture (see Lobina & García-Albea, 2017, for some examples), but we could do worse than to concentrate on the outputs of each phenomenon – i.e., the representations language and thought generate – in order to tease out what type of mental structure supports each.

The focus will fall on the shape of the representations supporting language and thought as well as on the computational theories that bring about such representations. To conduct, that is, a competence-like analysis of language-and-thought as a way to unearth the relevant parts of mental architecture. To this end, I shall understand the locution *being a medium of thought* to refer to the mental code (representational vehicle) required for belief fixation, its study a question of working out what sort of format this internal code has, and not an investigation of the real-time processes at play during the activity of belief fixation. From this perspective, then, in studying the language-and-thought nexus one would want to determine, first of all, what sort of representations thought subsumes in order to then move on to an evaluation of whether language exhibits the appropriate features – the latter being the *sine qua non* condition on natural language being *the* language of thought.

The second main way to relate the LoT and the FoL is at the level of what Chomsky calls performance, which for our purposes is the level of explanation that concerns the real-time processes that a good part of cognitive science, and certainly most of cognitive psychology, focus on. This is a level of analysis that goes beyond the mental components outlined for a competence-like study, a peculiarity that is

true for both language and thought. Language behavior, including comprehension and production, would plausibly involve, *inter alia*, the FoL (at the very least as a knowledge base specifying what sort of representations are licit), a parser that chunks the input into constituents of various kinds (for language comprehension), a formulator that puts together the words and phrases to be outputted (for language production), and whatever perceptual and other cognitive systems participate in either process (such as attention and memory resources, the pragmatic component, joint attention, etc.). The act of thinking, on the other hand, would engage mental processes such as categorization and reasoning – but more generally belief fixation – and mental states such as propositions and propositional attitudes, an overall (mental) activity that is obviously also a real-time phenomenon and thus would also make use of memory, attention, etc.

Importantly, it is only when the different elements underlying competence and performance are laid out that it is possible to evaluate the supporting mental architecture for each phenomenon (the strategy I have defended we should follow). That is, if language and thought are studied at the level of competence, then the investigation will implicate a very different mental reality from what one might expect in a study of performance – *a fortiori*, rather different cognitive theories stemming from different methodologies would be constructed at each level. It seems to me, in fact, that most of the literature on the relationship between language and thought centers on the level of performance, as evidenced by the numerous accounts on offer on how language influences thought during actual behavior (or real-time processing), some of which were referenced earlier on. That this is so muddies the issues at hand a great deal, for many a datum from performance may not bear a transparent relationship to architectural/representational properties (see *infra* in relation to spatial reorientation data).

More importantly, the question of how language and thought relate at the level of competence must take precedence over how they interact during actual behavior. We need to first resolve the question of whether language really is constitutive of thought – that is, whether the mental representations thought uses are provided by the language faculty – before considering how language is employed during actual bouts of thinking, as the behavioral data will look very different depending on how this prior question is answered. This general point applies to the architectural perspective as well as to some of the other ways in which language and thought can relate, such as the general effect acquiring a language has on thought and cognition, the physical representation and implementation of language and thought in the brain, or the evolution of language and thought. Much as is the case with the level of performance, I would argue that these other ways of relating language and thought are also dependent on the results of a competence-like study of *the relationship*.

In order to bring this point home, I now finally turn to analyzing a type of mismatch between linguistic and conceptual representations that has received little attention. I aim to accomplish two things with this analysis: a) highlight that it is imperative to employ appropriate definitions for language and thought, and appro-

priate ways to relate these two phenomena; and b) (re-)establish the position that language and thought *do* constitute different phenomena, against the recent efforts of Hinzen, Carruthers, & Pietroski to conclude otherwise (we would have to add Chomsky to this bunch, at least in some of his recent moods).

Let us discuss the sentence in (3) below, which was once employed by a journalist to describe a particular event in the British Parliament:¹¹

(3) Pete Wishart, the SNP MP, says he wants to put the question that the question not be put be now be put.

This sentence, a convoluted example no doubt, but somewhat typical of British political discourse, may pose great difficulty upon first encounter, and probably because of its repetitive nature: one too many *question*, clearly too many *bes* and *puts*. An account for what the sentence means (and what thought it expresses) can be sought both at the level of competence and at the level of performance, and each will yield a different sort of explanation (or set of concerns). The importance of this particular example, however, rests on the mismatch it highlights between the way things need to be phrased in Parliament and the corresponding thoughts speakers intend to communicate under and, as we shall see, beyond such constraints. The mismatch, moreover, provides some criteria as to what a competence theory of thought will look like vis-à-vis a fully developed competence theory of language, and that is precisely what I am after in this paper (at least in the form of a primer).

Let us consider the level of performance first; this is not an unnatural place to start from, considering that most people are first exposed to such sentences in settings that involve some sort of behavior (e.g., during a verbal exchange, listening to a parliamentary debate or, indeed, reading a philosophy paper). I'll then move on to the question of what thought underlies the sentence, at which point I shall argue that there is a rather interesting mismatch between the sentence's linguistic structure and the structure of the corresponding thought. Such a take on things will give me the opportunity to provide a competence-like analysis of the underlying thought representation, and I will end the discussion of this sentence by offering a few remarks regarding the first two representational desiderata established earlier on for a *language of thought*.

From the perspective of language comprehension, upon reading or hearing a sentence the linguistic parser is receiving one element after another in a linear fashion, and as it does so builds a syntactic structure for the sentence, which may or may not be the actual structure intended by the speaker/writer. At the same time (or rather, concurrently to that), the overall comprehension system is attempting to assign a meaning to the sentence – that is, building a semantic representation – and does so at every opportunity it has (say, after every word or phrase); much as in the case of syntactic processing, and very commonly because of the operations of it, this interpretation may not be the intended one.

Keeping to the syntactic side of things, the parser operates in an incremental fashion, both building a hierarchical structure *and* predicting the next type of

linguistic item to be parsed at each stage of the comprehension process. Thus, upon having processed, say, a noun, the processor would both build up a noun phrase structure for it and predict/expect, typically, a verb phrase to come next, the sentence's predicate for the sentence's subject (all this *ceteris paribus*, of course). In the case at hand, the parser would process *Pete Wishart* as the subject of the sentence and *says* as the main verb, and then predict/expect a direct object argument. The parser is instead inputted the pronoun *he* after *says*, indicating an internal, embedded clause, forcing the parser to re-evaluate its operations. The pronoun is followed by a verb, in this case an infinitival construction, *wants to*, another verb, *put*, and finally a direct object, the *question*. So far, so straightforward for the parser, but what follows introduces another internal phrase (this time marked by the complementiser *that*) and much repetition: another *question*, three *bes*, and two *puts*, the whole eventuality possibly driving the parser into a bottleneck.

There is, of course, more to language comprehension than the operations of the parser, as the overall system can access and make use of a number of different sources of information to process a sentence, sometimes to the rescue of the parser. Central amongst these is prosodic information, as it provides information on the sentence's syntactic structure, but how do the syntactic and prosodic phrases of the sentence in (3) match up? It may take a couple of passes before the reader is able to link the right syntactic structure to the right prosodic structure – one of the crucial steps to understand the sentence in – but once this is done, computing the sentence's intended syntactic structure should pose no great effort, even if the actual meaning could still be elusive (as we shall see). That is, if one were to utter the sentence in (3), and did so correctly – that is, with the right intonation – that would prove to be a very strong cue for a hearer, allowing them to work out the underlying syntactic structure with not too much difficulty.

As it happens, working out the right prosody and the corresponding syntactic structure required quizzing the SNP MP, but he complied. At first sight, the Wishart sentence can be understood in one of two ways. According to the first possibility, Wishart has a particular question in mind, namely *the question that the question not be put*, and with that question in mind, he then moves that this question be now be put (whatever this means; I will come to this soon enough). According to the second possibility, Wishart has a different question in mind, namely *the question that the question not be put be now be put*, and with that question in mind he then moves this question (again, and for the time being, whatever this means exactly). So which one is the right interpretation?

(4) below uses the | symbol to mark where the prosodic phrases would roughly end, according to the reputed MP, with (5) showing the sentence's simplified and lightly annotated syntactic structure. The syntactic structure in (5) establishes that the string *that the question not be put* is a relative clause (RC) modifying the previous noun phrase (NP) *the question*, and that after this internal clause there is a verbal phrase (VP) followed by what looks like another verbal phrase. According

to the prosodic and syntactic structures in (4–5), then, Wishart had the first meaning in mind when uttering (3).

(4) Pete Wishart says he wants to put the question | that the question not be put | be now | be put.

(5) Pete Wishart says he wants to put [NP the question [RC that the question not be put]] [VP be now [VP? be put]].

I expect that this will not have helped matters all that much, though. Indeed, the prosodic and syntactic structures may be well matched-up now, but the sentence would still be rejected as meaningless or incomprehensible, and not only because its repetitive nature may prove to be too great a burden. That is, the sentence may be regarded as being well-formed once the intended structure has been ascertained (maybe), but remain meaningless nonetheless. Further parses and some concentrated reflection upon what the sentence is trying to convey could help, but proper comprehension necessitates knowledge of what some of the sentence's units actually mean (or stand for). Pronouncing and perceiving the sentence correctly does not settle the issue; what we need to do is pronounce and perceive the sentence with the right *competence*.

A rather central factor here is the idiomatic or formulaic nature of British parliamentary discourse. As it happens, the phrase *that the question be now put* is usually employed in Parliament to request a “closure motion,” that is, a closing of the debate at hand so that the question under consideration can *be put* to a vote. Thus, in order to know what the sentence means one would have to be privy of this knowledge; one would have to be acquainted with this particular idiom. What the sentence amounts to, then, is a situation in which there is at present a question under discussion that is not currently being considered for a vote (i.e., the question not put, or in parliamentary terms, the question not be put) and Mr Wishart wants to put this very question to a vote (viz., that the question now be put, or in Wishart's terms, that the question be now be put). The linguistic processor can't patently return the meaning of the sentence if the parliamentary idiom isn't part of one's linguistic knowledge.

Interestingly, Wishart seems to have gone beyond the parliamentary formula, creating his own idiom in (4–5). In particular, Wishart is playing with the standard wording a little bit and appears to have taken the phrase *be put* to be sort of detachable from the whole, *that the question be now put* expression. That is, he is closing the motion in a slightly different way; namely, in putting the question to a vote, Wishart takes it that one might want to request that the question *be put* to a vote. Accordingly, the *be put* within the *that the question be now put* would be part of Wishart's own idiom in addition to being the subjunctive verbal form within the more conventional parliamentary idiom. This would explain the presence of what seems to be an extra *be* (and thus an extra VP), the second last in (4). Indeed, the sentence sounds much better without this extra *be*, as shown in (6) below, and I

gather that most people would consider the sentence unquestionably grammatical in this case.¹²

(6) Pete Wishart says he wants to put the question [that the question not be put [now [be put]]].

(7) *? Pete Wishart says he wants to put the question [that the question not be put [now [be [be put]]]].

In the original example, then, Wishart seems to have thought that, in addition to the subjunctive *be*, which the sentence must obviously have, the sentence must also exhibit a *be put* as a unit, given that, for Wishart at least, in order to close a motion in parliament one must *be putting* the question to do so. If so, the word *be* appears in two different guises at the end of (3): as part of a common parliamentary expression – the original idiom – and as a subjunctive. Thus, Wishart could have conceivably uttered (7) instead, considering that the adverb *now* can appear on either side of the subjunctive. Needless to say, it would be very hard to convince anyone that the latter version ought to be regarded as grammatical too.¹³

The extra *be* is key to the point I aim to make here, but we need to consider the question of how the productivity of natural language syntax relates to the productivity of thought in order to show why. In particular, we need to evaluate whether the match between the productivity of language and the productivity of thought is one-to-one. The productivity of each domain has been stressed before in the literature (and in this very paper), and so has the point that they must be somehow related. Searle (1969), for instance, has argued in favor of a *principle of expressibility*, according to which whatever can be meant (thought, in my terms) can be said (16), a principle that is closely related to what Katz (1978) has called *effability*, namely, the inter-translatibility of whatever content one might be able entertain (in a conceptual system, in my imitable terms). I think that both principles, properly hinged, must be true, but we must nonetheless allow for greater flexibility at the “what can be said” side of the relationship. After all, there are many different ways in which one may wish to linguistically communicate the very same content, and these different ways are not always the result of conceptual differences.

In fact, the match between linguistic and conceptual representations is very unlikely to ever be, strictly speaking, one-to-one. This is clearest in the case of (3), where the idiomatic nature of the sentence results in the matching-up of rather disparate units. After all, an idiom involves, by definition, a match between a constructed syntactic structure and an unequivocal meaning (or conceptual representation). More importantly, a closer look at the internal structure of the sentence in (3) will show how transparent the relationship between language and thought is in this case.

The field of generative grammar is particularly useful in this respect, with its focus on the hierarchical relations among different lexical items within a sentence and how these are formed (or derived). As explained earlier, a generative grammar employs a combinatorial operation (i.e., *merge*) that manipulates the syntactic

features of lexical items in order to generate representations that interact with the sound and meaning/thought systems, and it does so in stages, or *phases* (these are points in a derivation where specific sets of operations are completed and the resultant representations are closed off from further manipulation). According to this theoretical perspective, *merge* takes two lexical items at each time and combines them into a unit, starting the overall derivation with what looks like the most embedded material. In the case of (3), repeated once again below, *merge* would presumably start by taking the final *be* and *put* and form the unit *be put* (as we shall see, this is not quite what happens, though).

(8) Pete Wishart, the SNP MP, says he wants to put the question *that the question not be put be now be put*.

I will not quite run the entire derivation of the sentence; some steps will only be described and I will also omit the very first step I have just mentioned and simply assume that *be put* is a single unit throughout (I do this to ease the exposition). We need to concentrate on the derivation of *that the question not be put be now be put*, the *content* that needs to be accounted for, and thus I will ignore the rest of the sentence (note that I have highlighted this content in (8)). The sentence actually exhibits an internal, embedded clause, and this results in a rather intricate derivation, which may create some confusion. (9) outlines such a derivation, which I shall describe in some detail.¹⁴

- (9) (1) [be put] [the question]
 (2) [not] [be put the question]
 (3) [the question]_i [not be put _____i]
 (4) [be put] [[the question]_i not be put _____i]
 (5) [now] [be put [the question]_i not be put _____i]
 (6) [(be)] [now be put [the question]_i not be put _____i]
 (7) [[[the question]_i not be put] j _____i] [(be) now be put _____j]

The derivation outlines the formal properties of the sentence and manifests a number of peculiarities. One of these is that the initial stages of a derivation involve the argument structure of the overall sentence. Thus, instead than generating a sentence from left to right, as linguistic production and comprehension would presumably mandate, modern generative grammar proceeds from the bottom upwards. The argument structure of the fragment I am analyzing is composed of the predicate (*be*) *now be put* and the argument *the question not be put*, and thus we start with the derivation of the latter, a complex syntactic structure in its own right. Steps 1–3 outline the derivation of *the question not be put*; an argument structure itself, the derivation starts by combining the predicate *be put* and its argument *the question* side by side (step 1), the negation is then introduced in step 2, and the phrase the question moves in step 3 to its final position (a case of displacement, in this case fronting). The underlining and the subscripts mark the relevant co-dependencies

between the starting position of the question and the position where it moves to; an old-fashioned, and not entirely correct, way of marking the movement operation, but let it stand for now. The phrase *the question* is introduced adjacent to *be put* in step 1 in order to receive thematic interpretation from the predicate (*be* theta-marked, that is) and so that when *not* is introduced in step 2 the negation takes scope over the whole *be put the question*, accounting for the sentence's syntactic and semantic facts. Steps 4–6 show the introduction of the sentence's main predicate (along with the theta-marking) and in step 7 the whole phrase *the question not be put* moves from its original position vis-à-vis (*be*) *now be put* to its final position, where it is in fact pronounced (note the new underlining and subscripts, and the fact that the scope of the negation remains within the internal, embedded phrase).

The one major peculiarity in all this for our purposes is the derivation of Wishart's own idiomatic creation, as shown in the presence of the extra *be* in step 6 (thus also introducing an extra VP, as mentioned), which the parliamentary idiom would not have (it would have been noticed that I have marked this contrast with a parenthesis throughout). The result is the self-contained unit that is put together in steps 4–6: *be now be put*. Naturally, both the *be now* and the *be put* refer to *the question not be put*, even if this is not entirely clear in Wishart's formulation, this partly on account of Wishart's apparent word play, where *be put* is taken to be detachable from the whole idiom, as stressed earlier. Significantly, even if we were to regard Wishart's sentence as ungrammatical, its derivation actually employs the same sort of operations and constraints than the grammatical, parliamentary version would – all that is needed in this case is the dispensation of an extra *be* along with the corresponding steps in the derivation.¹⁵

Be that as the syntactic facts may be, in simple terms the meaning of the sentence is that there is a question that has not been put to a vote and Wishart wants to now put it to a vote (as noted, I'm ignoring the rest of the sentence, including the indirect speech verb *says* and the propositional attitude verb *wants*, but this matters not). In parliamentary terms, that there is *a question not be put* and that someone wants to have it *now be put*. Or, in Wishart's terms, that the question *be now be put*. So what does this all mean for the conceptual representation underlying the sentence?

The differences in meaning, which point to the relevant conceptual differences, as argued earlier, are really key to understanding how the linguistic and conceptual representations match up. The phrase *be put* is understood to be short for “put something to a vote”; or better, it means something like “to ask people to vote on a proposal.” Thus, we can take PUT THE QUESTION TO A VOTE as the relevant conceptual representation here (I will revert to using angle brackets for the underlying predicate-argument proposition), with *the question be put* a paraphrase in parliamentary language. Thus, moreover, *the question not be put* would have to be translated into conceptual representation as something like THERE IS NO VOTE ON THE QUESTION, and both the standard *now be put* and the non-standard *be now be put* as PUT THE (UNCONSIDERED) QUESTION TO A VOTE NOW (or

something like that). I'm playing fast and loose with the conceptual representations – how could it be otherwise? – but I think the point I'm trying to make is clear enough: there is some distance between how you put things in language – parliamentary or not, in Wishartian terms or not – and how this is in fact represented in conceptual terms.

More to *the point*, the content-bearing segment I am analyzing – *that the question not be put (be) now be put* – reveals two specific predications in the conceptual representation (again, I'm ignoring the rest of the sentence). The first relates to the question not being put to a vote – an unconsidered question, that is. This is the first internal structure the conceptual system must put together: the predication of < not being considered > to the argument < the question >. That is, the proposition < not being considered << the question >> (or simply *Fa*, say). This particular thought representation would interface with steps 1–2 in the linguistic derivation, as current generative theory dictates. However, and rather crucially, the conceptual representation would not interact with step 3, as the representation that obtains at such step plays no role in accounting for the requisite conceptual differences – that is, it has no analogue in thought. It would seem that the derivation in step 3 would be the result of some internal, linguistic requirement, and it is hard to see what need there might be for movement operations in thought at all. Displacement seems to have an exclusively linguistic function.

What function in particular is an interesting topic in itself, and I am but alluding to a rather contentious issue, actually; namely, whether movement phenomena are due to either the C/I interface or the SM interface. If movement operations were the result of the SM interface, my position would be strengthened, and this is the position defended in Moro (2000), in fact. I am more or less non-committal about this, though. My intention is simply to highlight the mismatch between linguistic representations (and operations) and the corresponding conceptual ones, as stressed throughout.¹⁶

After the initial predication, the conceptual system would have to construct a predication of the so-far constructed representation; namely, it would have to assign the predicate < being put to a vote > to the proposition < not being considered << the question >>, resulting in, simplifying somewhat, the conceptual representation < being put to a vote << unconsidered question >> (or simply *P(Fa)*). Again, the generation of this particular conceptual representation would interface with the steps 4–6 in the linguistic derivation (more about this soon enough), but not with step 7, where there is further movement. In any case, the derivation (or competence-like theory) of the conceptual representation would be, mostly, a two-stage affair, as shown below.

- (10) (1) predicating < not being considered > of < the question >
- (2) predicating < being put to a vote > of < not being considered << the question >>

So how does all this bode for the relationship between the linguistic and conceptual representations underlying (3)? A number of factors are worth discussing. First of all, even though I am using a similar vocabulary to the sentence's actual words to describe the conceptual representations (being put to a vote, the question, etc.), I very much doubt that the requisite concepts are all that similar to the relevant lexical items. That is, these concepts are likely to be very different indeed, and it is worth noting that I have had no need to employ so many *bes* and *be puts* in the conceptual representations and the corresponding derivations. Nevertheless, the employment of the copula *to be* plus a predicate is rather close to the conceptual facts, and that suffices for my purposes (recall the predicate-argument(s) structures provided in (1)–(2)).

In any case, the differences in vocabulary point to the different types of atoms at play in conceptual representations, above and beyond the fact that we are dealing with an idiom, as reflected in the rather different derivations I have proposed for each domain – different competence analyses, that is (cf. Higginbotham (1998) on what he calls conceptual competence, which is close to what I have in mind here). The two representations must obviously connect at some point if what is meant (or thought) is to be said, but it is not the case that the two derivations match up step-by-step. This, if anything, is simply a reflection of the mismatch I am delineating.

More importantly, there is the fact that the parliamentary idiom is manifestly more transparent as to its meaning – and, thus, as to how it relates to the underlying conceptual representation – than Wishart's own idiom. That is, given the relevant knowledge of parliamentary procedure, the sentence in (6), which does not have an extra *be*, is readily accepted as grammatical and its meaning easily computed in a way that the sentence in (8) clearly is not. Careful reflection may well help work out the actual meaning of (8), this being an interesting issue in itself, but the sentence could very easily be rejected as ungrammatical. The extra *be* in (8) is having an inordinate effect on the comprehension of the sentence, as anticipated earlier on, and in how the overall linguistic representation connects to the corresponding conceptual representation. There really aren't any extra *bes* and VPs in the conceptual representation, these are additions at the "what is said" side of the equation; the thought itself is something else altogether. This, in turn, highlights the actual mismatch between the linguistic and conceptual representations, which I suggest generalizes to many other sentences, the non-idiomatic included.

Putting it all together, then, the analysis of (8) indicates that its underlying syntactic representation does not respect the two main desiderata a language of thought must meet: content representation and content differentiation. Among some of the reasons adduced here, the various *bes* and *be puts* express, or are related to, rather different concepts – in at least one case *be put* forms a complex expression (or idiom), in another it is simply a verb form – and these *conceptual* differences are not explicitly codified in either the PHON or the SEM representation for the sentence. The underlying morphophonemic and syntactic features certainly license the respective linguistic representations, but they do not capture the differences in

meaning I have outlined. This is more dramatically demonstrated in the case of the extra *be*, which is entirely absent in the thought representation. Thus, the conceptual differences would appear to come from outside of language and in this case at least language appears to be *reflecting* rather than *constituting* the relevant conceptual differences.

Before closing this discussion, I should stress that what the linguist has to say about the sentence in (3), at the level of competence, is rather different from what the psycholinguist would have to say about it at the level of performance. In particular, the sort of mental architecture the linguist would have to postulate is very different from the sort of mental architecture that the psychology must, and this is likely to confuse matters a great deal. A case in point is the data on spatial reorientation, and the role of language therein, reported in Hermer-Vázquez, Spelke, & Katsnelson (1999) and Hermer-Vázquez, Moffet, & Munkholm (2001), as discussed by, *inter alia*, Carruthers (2002, 2012), Samuels (2002), and Lobina & García-Albea (2017). As argued in the latter two sources, those data point to the different ways in which various mental systems interact during a spatial cognition task, but they do not establish that natural language is in fact employed to connect different sources of information during such a task. In fact, there are reasons to believe that natural language does not play such a role at all (Lobina & García-Albea, 2017). As a matter of fact, psychology has no precedence in the study of thought (or thinking), even if psychologists usually have a more interesting and complicated story to tell, as is the case for (8). Instead, an analysis of the format properties of linguistic and conceptual representations must be undertaken before we venture into studying, *inter alia*, what effect acquiring a specific language has on cognition, how language is used in reasoning tasks, etc.

4. The Final Thought-in-Language

In this essay I have showcased some of the problems afflicting the literature on the relationship between language and thought, identifying two main problems. On the one hand, language and thought can be defined in various ways, and depending on which definition is employed the relationship will result in different and clearly incompatible sort of studies. On the other hand, the two phenomena can be related in various ways, and here too which way is selected will result in rather different but also ultimately incompatible studies. I have proposed well-established definitions for each domain and argued that priority should be given to an analysis of the representational vehicles each domain exhibits in order to work out what sort of structure thought representations have. To that end, I have offered but a primer of a competence-like analysis of the LoT as it pertains to a particular mismatch between linguistic and conceptual representations. Crucially, I have claimed, it is only when this question is more or less resolved that we can move on to such issues as the role of language in actual behavior, the interaction between language acquisition and cognitive development, the evolution of language and thought, and others.

Having said that, I do not wish to suggest that thoughts are generally not as convoluted as the syntactic structures they are related to, even if there may be some truth to that; nor would I want to claim that thought in general is structurally simpler than language. In the particular case I have discussed, I would argue that the sentence's convoluted structure points to the great *expressive* flexibility language confers to speakers, but this is a specific characteristic of linguistic cognition and there is no reason to believe that thought is like that too – after all, what may be meant need not ever be said. What I *do* claim is that there is a *particular* but widespread mismatch between the structure of a sentence's proposition and the syntactic structure of such a sentence. Namely, it is not only that the structures of thoughts and sentences are often not identical; my actual point is that they involve fundamentally distinct representational vehicles.

Other mismatches between linguistic and conceptual representations have been noted in the literature before, as mentioned earlier, but not quite in these terms. Fodor (2001b), for example, talks of language not being explicit enough to represent thought, as in a sentence such as *It's raining* – where, when, Fodor would ask. In such “Fodor cases,” the relevant sentences do not seem to codify the thought one is having in full. Collins (2007), for his part, has pointed out that, when it comes to determining a sentence's proposition, language oftentimes offers either too much structure, as witnessed by the linguistic phenomena of raising/control pairs, ambiguity, passivization, etc., which find no analogue in thought representations, or too little, as the meaning of any given sentence typically needs to be enriched by specific pragmatic processes. My analysis exemplifies a different phenomenon altogether; it is not that there is too much or too little linguistic material in sentences such as (3) from the perspective of working out the sentence's proposition; rather, the point is that the *primitives* of linguistic and conceptual representations appear to be *fundamentally* different, and so are their respective structures and uses (cf. Segal, 2001).

What can we conclude regarding the underlying mental architectures of language and thought from all this, then? I submit the following final thoughts. An account of linguistic competence necessitates a computational operation (*merge*), lexical features, and two mechanisms to connect PHONs and SEMs to their respective interfaces, the sensorimotor and the conceptual/intentional. Thought, on the other hand, whilst also requiring a computational mechanism of some kind (perhaps even one like *merge*), would manipulate *concepts* rather than lexical items and a priori there would be no need for its mental representations to meet any of the interface conditions language has to respect. Consequently, *language* and *thought* appear to share fewer mental elements than what might at first be supposed. This is of course a tentative position, and the resolution of this issue is certainly an empirical question. What I have provided in this essay is a) a framework with which to approach such an investigation; and b) a discussion of the particular sort of mismatches that this framework brings to the fore.

There are many more cases of linguistic/conceptual misalignments – among others, grammatical illusions, ungrammatical but thinkable sentences, ellipsis phenomena, etc. – data that can be employed to firmly establish that language and thought *do* fall under different purviews, thereby requiring different types of theories. At the same time, however, we know a lot more about language than we know about thought, and this is reflected in the sophistication of our theories of language vis-à-vis our theories of thought (barren as the latter may currently be). This needs to be remedied if we are to settle the question at hand and I am hoping to provide some details anon, but first things first.

Acknowledgements

I would like to thank Prof. Martin Davies and Dr. Mark Brenchley for various conversations and comments on some of the material discussed here. This paper constitutes a continuation to Lobina & García-Albea (2017).

NOTES

1. This section draws from Lobina and García-Albea (2017), which it supersedes in some ways.

2. What are the other, less central representations the language faculty generates? Phonological representations, composed of phonological features and comprising hierarchical structures such as metrical and prosodic tiers, are plausibly compiled in tandem with, or right after, the generation of syntactic structures, but none of the phonological tiers would in principle be constitutive of thought (PHONs would emerge from these phonological representations, though, and since inner/outer speech is based on PHONs, they would supposedly play a role in the sort of story Clark, cited earlier on, envisions). The language faculty also generates properly semantic representations at the C/I interface in order to compute a sentence's interpretation (Pietroski, 2007), this time *paru passi* with syntactic derivations (Chomsky, 2008). Semantic representations are probably more relevant in a study of the language-thought relationship, but it is rarely argued that semantic representations are literally the representations employed in thought (Jackendoff, 1992, might be an exception), unlike the case of syntactic representations (see the references to Hinzen and Carruthers mentioned earlier).

3. I should add that I do not wish to hold the view that propositions are atemporal externalia, or possible worlds, or else. I here simply take a proposition to strictly refer to the *format* properties of mental representations, and I'm assuming that mental representations *are* indeed structured.

4. Despite the different structural relations, the actual notions of subjects (or arguments) and predicates would remain roughly the same in each domain (cf. Strawson, 1974, on this issue).

5. Seuren (2009) argues that logical subject-predicate structures are better described as topic-comment structures, but I shall employ the old-fashioned Aristotelian way of describing propositions as predicate-argument structures here.

6. That there is often a mismatch between linguistic and conceptual representations has certainly been noted in the literature before (e.g., in Collins, 2007; Fodor, 2001b). However, previous studies have not analyzed any mismatch in any detail, let alone from the

perspective of what predicate-argument structure underlie particular sentences/thoughts, my approach here, as we shall see. Moreover, the discussion I shall conduct will expand upon previous studies in specific ways, thus providing a much more comprehensive take.

7. To my mind, entertaining a complex concept such as THE ISLAND OF SAGES AND SAINTS constitutes a case of having a thought too, but I shall focus on full propositions here.

8. Systematicity does not involve, as Johnson (2004) and Pullum and Scholz (2007) have it, the permutability of constituents across syntactic templates, which in many cases in fact fail (cf. *Alice showed Martha the book* vs. **Alice described Martha the book*, taken from Johnson's paper). As mentioned, systematicity is a matter of form and constituent structure, and certainly not of substituting constituents across sentence (or thought) types (see McLaughlin, 2009: 260–6, for further comments).

9. I am basing some of this characterization on Deirdre Wilson's 2004 lectures on the subject at UCL.

10. Even though I have used the works of Evans and Fodor to outline the internal structure of thought, my take on what thought must be like is not entirely compatible with their own. I won't spend any space drawing out the distinctions, though, and it is more productive to emphasize the similarities. In particular, the methodology they employed is certainly not foreign to my story; even though both Evans and Fodor employed linguistic examples to make their points, the phenomena they characterized is very often described in contrast to, or despite of, specific properties of language (Fodor, 2001b, is clearest in this respect, especially regarding the *compositionality* of thought vis-à-vis the compositionality of language).

11. A debate that was narrated by said journalist in the following web-link, where this sentence may be found: <http://tinyurl.com/qgt6qa9>. SNP and MP stand for Scottish National Party and Member of Parliament, respectively, but that's thankfully immaterial.

12. I am referring to the second last *be* as the extra *be* for convenience only, as I could instead refer to the last *be* as the extra *be*; indeed, both *be now put* and *now be put* are fine. We will have to keep this choice in mind in the discussion, especially as it relates to the question of which VP phrase exactly is actually the extra VP. The linguistic derivation of (3), which I will introduce presently, suggests that the extra VP is the second last instead than the last (as marked in 5), but I shall keep this specific, and somewhat secondary, issue to footnotes.

13. It is worth pointing out that the noted MP avoided the extra *be* when describing the event on his twitter account; *I tried to move – the question be not now put – now be put*, he tweeted at roughly the same time as he tried to close the motion in parliament (this can be found in the web-link provided in footnote 11). That is, in that communication Wishart doesn't seem to be taking *be put* to be idiomatic in itself, only the whole phrase, with no need for an extra *be*. The parliamentary report for the day (i.e., the Hansard) doesn't record the actual words of Pete Wishart MP in "putting the question," only that he tried to close the motion twice (parliamentary protocol apparently establishes how such things are recorded in the official record). In any case, I assume the aforementioned journalist recorded the words of the MP faithfully; if not, it is the journalist who took the phrase *be put* to be a possible idiom within a common expression in British parliamentary discourse. It matters not; it was the output of someone's language faculty, either *a viva voce* or in writing. The standard phrase alongside the relevant procedure employed to close a motion can be found here: <http://www.parliament.uk/site-information/glossary/closure-motion>.

14. I depart from normal linguistic practice in that I do not employ the notation of set theory to describe the operations of *merge*. Instead, I use square brackets to mark the relevant syntactic units, two per line, this simply signaling that *merge* combines two such units each time it operates (each product will be evident in the very next line).

15. Given that a linguistic derivation establishes that predicates and arguments are generated side-by-side, the second last *be* in (3) is the extra *be* under my reading, but the correct extra VP would be the phrase headed by this *be*, contrary to what I had actually indicated in (5). Do note that the extra *be* and its accompanying VP phrase are both introduced in step 6. This is just one of the many mismatches between strings (or sentences) and their underlying syntactic structures, which I had to respect in the presentation as to not anticipate anything (and thus creating some more confusion).

16. This is entirely independent of the apparent fact that *merge* operates in basically the same way regardless of whether it puts two new lexical items together or moves an item from an already-built structure to a new position. And likewise regarding the more recent, and more accurate, description of movement operations as involving structures that contain two copies of the moved element; that is, [the question]_i [not be put [the question]_i] rather than [the question]_i [not be put _____i].

REFERENCES

- Carruthers, P. (2002). "The Cognitive Functions of Language," *Behavioral and Brain Sciences* 25(6): 657–726.
- Carruthers, P. (2012). "Language in Cognition," in E. Margolis, R. Samuels, & S. Stich (eds.), *The Oxford Handbook of Philosophy of Cognitive Science*. Oxford: Oxford University Press, 381–401.
- Carruthers, P., & J. Boucher (eds.) (1998). *Language and Thought*. Cambridge: Cambridge University Press.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: The MIT Press.
- Chomsky, N. (1986). *Knowledge of Language*. New York, NY: Praeger Press.
- Chomsky, N. (2008). "On Phases," in R. Freidin, C. P. Otero, & M. L. Zubizarreta (eds.), *Foundational Issues in Linguistic Theory*. Cambridge, MA: The MIT Press, 133–166.
- Chomsky, N. (2013). "What Kind of Creatures Are We? Lecture I: What Is Language?," *The Journal of Philosophy* 110(12): 645–662.
- Collins, J. (2004). "Faculty Disputes," *Mind and Language* 19(5): 503–533.
- Collins, J. (2007). "Syntax, More or Less," *Mind* 116(464): 805–850.
- Collins, J. (2008). "Knowledge of Language Redux," *Croatian Journal of Philosophy* 8(1): 3–43.
- Davidson, D. (1997). "Seeing through Language," in J. Preston (ed.), *Thought and Language*. Cambridge: Cambridge University Press, 15–27.
- Davidson, D. (2001). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- Deacon, T. (1997). *The Symbolic Species: The Co-evolution of Language and the Brain*. New York: Norton.
- Evans, G. (1982). *The Varieties of Reference*. J. McDowell (ed.). Oxford: Clarendon Press.
- Fodor, J. A. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.

- Fodor, J. A. (1979). "Fixation of Belief and Concept Acquisition," in M. Piatelli-Palmarini (ed.), *Language and Learning: The Debate between Jean Piaget and Noam Chomsky*. Cambridge, MA: Routledge, 142–161.
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: Bradford Books/The MIT Press.
- Fodor, J. A. (1987). *Psychosemantics*. Cambridge, MA: The MIT Press.
- Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J. A. (2000). *The Mind Doesn't Work That Way*. Cambridge, MA: The MIT Press.
- Fodor, J. A. (2001a). "Doing without What's Within," *Mind* 110(437): 99–148.
- Fodor, J. A. (2001b). "Language, Thought, and Compositionality," *Mind and Language* 16(1): 1–15.
- Fodor, J. A., & B. McLaughlin (1990). "Connectionism and the Problem of Systematicity: Why Smolensky's Solution Doesn't Work," *Cognition* 35(2): 183–204.
- Fodor, J. A., & Z. Pylyshyn (1988). "Connectionism and Cognitive Architecture: A Critical Analysis," *Cognition* 28(1/2): 3–71.
- Frank, S. L., R. Bod, & M. H. Christiansen (2012). "How Hierarchical Is Language Use?," *Proceedings of the Royal Society B: Biological Sciences* 279(1747): 4522–4531.
- Gallistel, C. R., & A. P. King (2009). *Memory and the Computational Brain*. Malden, MA: Wiley-Blackwell.
- Hermer-Vázquez, L., A. Moffet, & P. Munkholm (2001). "Language, Space, and the Development of Cognitive Flexibility in Humans: The Case of Two Spatial Memory Tasks," *Cognition* 79(3): 263–299.
- Hermer-Vázquez, L., E. Spelke, & A. S. Katsnelson (1999). "Sources of Flexibility in Human Cognition: Dual-task Studies of Space and Language," *Cognitive Psychology* 39(1): 3–36.
- Higginbotham, J. (1998). "Conceptual Competence," *Philosophical Issues* 9: 149–162.
- Hinzen, W. (2013). "Narrow Syntax and the Language of Thought," *Philosophical Psychology* 26(1): 1–23.
- Jackendoff, R. (1992). *Semantic Structures*. Cambridge, MA: The MIT Press.
- Johnson, K. (2004). "On the Systematicity of Language and Thought," *The Journal of Philosophy* 101(3): 111–139.
- Katz, J. J. (1978). "Effability and Translation," in F. Guenther & M. Guenther-Reuter (eds.), *Meaning and Translation: Philosophical and Linguistic Approaches*. New York, NY: New York University Press, 157–189.
- Laurence, S., & E. Margolis (1997). "Regress Arguments against the Language of Thought," *Analysis* 57(1): 60–66.
- Lobina, D. J. (2017). *Recursion*. Oxford: Oxford University Press.
- Lobina, D. J., & J. E. García-Albea (2017). "On Language and Thought: A Question of Format," in R. G. De Almeida & L. Gleitman (eds.), *On Concepts, Modules, and Language: Cognitive Science at Its Core*. Oxford: Oxford University Press, 249–274.
- Marr, D. (1977). "Artificial Intelligence – A Personal View," *Artificial Intelligence* 9(1): 37–48.
- McCarthy, J. (1963). "A Basis for a Mathematical Theory of Computation," in P. Braffort & D. Hirshberg (eds.), *Computer Programming and Formal Systems*. Amsterdam: North-Holland, 33–70.
- McLaughlin, B. P. (2009). "Systematicity Redux," *Synthese* 170(2): 251–274.

- Millikan, R. G. (2001). "The Language–Thought Partnership: A Bird’s Eye View," *Language and Communication* 21: 157–166.
- Moro, A. (2000). *Dynamic Antisymmetry*. Cambridge, MA: The MIT Press.
- Moschovakis, Y. N. (2001). "What Is an Algorithm?," in B. Engquist & W. Schmid (eds.), *Mathematics Unlimited: 2001 and Beyond*. Berlin: Springer, 919–936.
- Pietroski, P. M. (2007). "Systematicity via Monadicity," *Croatian Journal of Philosophy* 7(21): 343–374.
- Pullum, G. K., & B. C. Scholz (2007). "Systematicity and Natural Language Syntax," *Croatian Journal of Philosophy* 7(21): 375–402.
- Pylyshyn, Z. W. (1984). *Computation and Cognition*. Cambridge, MA: The MIT Press.
- Samuels, R. (2002). "The Spatial Reorientation Data Do Not Support the Thesis that Language Is the Medium of Cross-modular Thought," *Behavioral and Brain Sciences* 25(6): 697–698.
- Searle, J. R. (1969). *Speech Acts*. Cambridge: Cambridge University Press.
- Segal, G. M. A. (2001). "On a Difference between Language and Thought," *Linguistics and Philosophy* 124: 125–129.
- Seuren, P. A. M. (2009). *Language in Cognition: Language from Within*. Vol. 1. Oxford: Oxford University Press.
- Strawson, P. F. (1974). *Subject and Predicate in Logic and Grammar*. Farnham: Ashgate.
- Tomalin, M. (2006). *Linguistics and the Formal Sciences*. Cambridge: Cambridge University Press.
- Tomasello, M. (2003). *Constructing a Language: A Usage-based Theory of Language Acquisition*. Cambridge, MA: Harvard University Press.
- Yang, C. (2002). *Knowledge and Learning in Natural Language*. Oxford: Oxford University Press.