Glossa *a journal of general linguistics*

RESEARCH

Language as an instrument of thought

Eran Asoulin

Macquarie University, NSW, Australia
eran.asoulin@mq.edu.au

I show that there are good arguments and evidence to boot that support the language as an instrument of thought hypothesis. The underlying mechanisms of language, comprising of expressions structured hierarchically and recursively, provide a perspective (in the form of a conceptual structure) on the world, for it is only via language that certain perspectives are available to us and to our thought processes. These mechanisms provide us with a uniquely human way of thinking and talking about the world that is different to the sort of thinking we share with other animals. If the primary function of language were communication then one would expect that the underlying mechanisms of language will be structured in a way that favours successful communication. I show that not only is this not the case, but that the underlying mechanisms of language are in fact structured in a way to maximise computational efficiency, even if it means causing communicative problems. Moreover, I discuss evidence from comparative, neuropathological, developmental, and neuroscientific evidence that supports the claim that language is an instrument of thought.

1 Introduction

In introductory textbooks, as well as elsewhere in the literature, whenever human natural language is discussed it is often contrasted with animal communication, implying that the particular author sees language as a more complex instrument for symbolic communication as compared with animal communication systems. This is the case in linguistics, philosophy, psychology, and cognitive science. According to Mangum (2010: 257), "[m]ost linguists consider human language a unique type of communication system". Williams (1993: 91) concurs when she argues that "[m]uch modern linguistic theory is based on the assumption that the primary and fundamental function of language is communication." Millikan (2005: 25) argues that "a primary function of the human language faculty is to support linguistic conventions, and that these have an essentially communicative function." Deacon (1997: 11–12, 50) refers to language as "our unique and complex mode of communication" and argues that it is a "fact that language is an unprecedented form of naturally evolved communication". Carruthers (2002: 657–658) notes that "most members of the cognitive science community" endorse what he calls "the (purely) communicative conception of language". And Jackendoff (2002: 123) takes as a basic assumption that "language arose primarily in the interests of enhancing communication, and only secondarily in the interests of enhancing thought."

There is, however, an alternative construal of language that sees it as an instrument of thought. This is the view of the rationalist tradition, most notably as it is manifested in philosophy of language and in linguistics. In philosophy, we have Frege, who "insists that thought-content is prior to matters of use" (Moravcsik 1981: 106). Frege

saw the communicative function of language as "merely peripheral" and argued that the "expression of thought must figure centrally in explanations of syntactic and semantic facts" (Moravcsik 1981: 106). Fodor argues that language has no semantics per se as distinct from the content of the thoughts it expresses. "Learning English," he says, "isn't learning a theory about what its sentences mean, it's learning how to associate its sentences with the corresponding thoughts" (Fodor 1998: 9). That is, to "know English is to know, for example, that the form of words 'there are cats' is standardly used to express the thought that there are cats" (cf. also Fodor 1975; 2001; 2008). The rationalist tradition in linguistics was most famously articulated by von Humboldt (2000 [1836]), and in modern times it is most clearly expressed in generative linguistics, a fundamental aspect of which is that natural language is an independent domain of inquiry, subject to its own principles that are distinct of thought but, amongst other uses, expressive of it. It is also expressed in the broader, generative-oriented biolinguistics program, where language is regarded as an internal computational system, a recursive mechanism that produces an infinite set of hierarchically structured expressions that are employed by the conceptual-intentional systems (systems of thought) and the sensorimotor systems to yield language production and comprehension. This particular functional design is strongly shaped by its interface with the systems of thought, rather than by the peripheral process of externalisation inherent in the link with the sensorimotor systems (for an overview of biolinguistics cf., amongst others, Chomsky 2007; Boeckx 2011; Di Sciullo & Boeckx 2011).

Language, of course, can be used to *express* thought, but often this claim is made with the implicit assumption that the structure of language is designed for the communication of thoughts. A rather different claim is that language is an *instrument* of thought. I think that this distinction is too easily overlooked. For example, Pinker & Bloom (1990: 714–715) remark that "the facts of grammar make it difficult to argue that language shows design for 'the expression of thought' in any sense that is substantially distinct from 'communication'". They paraphrase Chomsky's claim as emphasizing that "people's *use* of language does not tightly serve utilitarian goals of communication but is an autonomous competence to express thought" (Pinker & Bloom 1990: 719, emphasis in original). However, the counter to language's function being communication is that it is an *instrument* of thought, not merely that it is a tool for the *expression* of thought. I think that what is at issue is whether language does more than merely express pre-formed thoughts.

There are two ways to construe the claim that language is an instrument of thought: a weak and a strong claim. The weaker claim is that language is used primarily for the expression of thought, whereas the stronger claim is that language to some extent *structures* thought (or at least a subset or particular types of thought). The stronger claim is not a Whorfian one; thought is certainly independent of language, and what can be expressed or thought by a speaker of one language can certainly be expressed or thought by a speaker of a very different language. Moreover, this is not a claim that a natural language is the medium of thought, nor that language is necessary (conceptually or empirically) for thought.[1] This is clearly far too strong: animals have a rich mental life that involves thoughts of many kinds but no natural language. The work of Charles Randy Gallistel, for example, has shown the complexity and richness of animal cognition (Gallistel 1990; 1991; 2009; 2011; cf. also Hauser 2000; Povinelli 2000; de Waal 2001; Shettleworth 2010). Gallistel (2009) reviews the literature in experimental psychology and experimental zoology demonstrating that birds and insects, species with whom humans last shared

---

[1] For the claim that all thought requires language, see Davidson (1975; 1982), Dummett (1981; 1989), and McDowell (1994). Their arguments are not the same but the conclusion they reach is (cf. Carruthers 2002: 660ff. for further discussion). Cf. Ryle (1968) and Slezak (2002) for criticism of this view.

an ancestor several hundred million years ago, make complex computations to learn the time of day at which events such as daily feedings happen, they learn the approximate duration of such events and can calculate the intervals between them. Moreover, they can assess number and rate, and they can create a cognitive map of their environment in order to compute their current location by integrating their velocity with respect to time (Gallistel 2009: 61ff.). Complex thought processes range across the animal world, from bees, who navigate by computing the local solar ephemeris (Dyer & Dickinson 1994), to the higher primates, who have many impressive cognitive abilities (cf. Hurford 2007 for an overview that compares humans and, among others, primates).

This impressive array of animal cognition, however, is missing a specific kind of thinking that appears to be unique to humans. That is, while there is no doubt that animals think, there is little evidence that their thoughts display productivity and systematicity (cf. Hauser et al. 2014 for a review of the literature showing the dearth of evidence in this regard; cf. also Berwick et al. 2011). So there is a discontinuity, a partial overlap, between animal thought and human thought that must be accounted for. I argue that what accounts for this discontinuity and allows humans to think these particular types of thoughts are the *underlying mechanisms* of language that structure these thoughts in a particular way.[2] Note the stress on the underlying mechanisms of language and not on any particular natural language – this is a crucial distinction that will be detailed below. To that end, unless stated otherwise, the term *language* here refers to the underlying mechanisms in virtue of which the production and comprehension of natural languages is made possible. This is in contrast to the use of the term *language* to mean a particular natural language such as English or Italian.

With the above in mind, I show in what follows that there are good arguments and evidence to boot that support the language as an instrument of thought hypothesis.

2 Language: Its use and function

Before discussing the different views on the function of language, a brief discussion of the notion of *function* is in order. What does it mean to claim that language (or anything) has a certain function? There is an intuitive distinction between, say, the operation of a machine having certain effects and one of these effects being the machine's function. This is because functional attributions are inherently teleological – we know what the machine is *for* because of the intentions of its designer. Everyone can agree that the operation of a refrigerator has many effects (its motor makes a noise, say) but that only one (or very few) of these effects is its function. However, when one moves away from human-made objects, which were constructed with a certain purpose in mind, and into the biological realm this pre-theoretical intuition becomes problematic: it is not clear whether and to what extent the notion of function applies to biological entities. And so a naturalised teleology is needed in order to bridge the gap between the purported explanatory role of functional ascriptions in biology and naturalistic inquiry (Walsh 2008). In other words, justifying functional attributions in the case of human-made objects is, all things being equal, simple – its function is what its designer intended it to be used for. In the case of biological entities, however, justifying functional attributions is much more complex.

One way in which to justify attributions of function to biological entities – the account I favour – is the systematic account of functional attribution (Cummins & Roth 2009). This account utilises functional analysis as its explanatory strategy, where the operation

---

[2] The empirical question of just what this discontinuity consists in is still open for debate. A current hypothesis, which I favour, is that of Hauser et al. (2002), according to which one of the crucial differences, perhaps the only one, is recursion. For more on recursion in this regard, see the edited collections of Sauerland & Gärtner (2007) and Roeper & Speas (2014).

of systems is explained by the operation of their constituent parts. So complex systems are explained in terms of their – usually simpler – constituent parts as, for example, "amplification gets analyzed into the capacities of resistors, conductors, capacitors, power supplies, etc." (Cummins & Roth 2009: 74). The same is true in biology where explanations of organisms are given in terms of constituent systems, for example, the immune system, which in turn is analysed into constituent organs and structures. This strategy can be pursued until pure physiology takes over (Cummins 1975). Under this explanatory strategy – sometimes called the analytical strategy – a function is defined in terms of an exercise of an analysed capacity *within a particular explanatory theory*. So certain systems can have many effects, but only some – or one – of these effects can properly count as the system's function. This function, however, only makes sense as a concept within an explanatory theory that appeals to some specific effects and disregards others.

Take the heart, for example: relative to the need to explain the circulatory system's capacity to transport food and oxygen, the appeal to the heart's function being a pump makes explanatory sense. However, relative to a system of medical diagnosis, the function of the heart may be to make thumping noises. As Cummins & Roth (2009: 75) note, this may be one of the things one needs to know in order to understand how a certain sort of medical diagnosis works (cf. also Cummins 1975: 762). This is as it should be, for nothing in biology has a function in an objective or abstract sense. Biological entities have any number of functions that are inherent in their internal structure – they can be used for anything for which their structure allows. But the question "What is *the* function of biological entity X?" does not make sense in the abstract. Talk of functions is only appropriate relative to a certain explanatory theory. The analytical strategy, then, relativises the notion of function to a particular explanatory context and thus rids functional attributions of "a vestige of an unscientific teleology" (Cummins & Roth 2009: 75).

I want to suggest that the same is true in the case of language. Systems and their constituent parts have many effects, but the effect that counts as the system's function does so because of the explanatory role it plays in a theory. Analogously, the language faculty has an effect – that of allowing us to communicate – but the question is what explanatory role, if any, this function plays within the theory. How does the claim that the function of language is communication fit in with an explanatory theory of language? Moreover, since the notion of function makes the best sense within a systematic account of functional attribution, a good way to decipher the function of language is to look at its underlying mechanisms, to look at the way they are structured and the way in which they operate. My claim is that the nature of the underlying mechanisms of language indicates that its primary function is that of an instrument of thought.

3 Language and communication
3.1 What is communication?

Let us first be clear about what is meant when it is claimed that the function of language is primarily to support communication. Communication is usually understood as involving a transfer of information, though a communicative act involves much more (or much less; often a communicative act will involve no transfer of information or propositional content, when, say, social grooming is involved or formalities are exchanged). Communication is standardly construed in two ways: the encoding-decoding model, and the inferential model.

The encoding-decoding model, also known as the message model (Akmajian et al. 1980), involves the speaker encoding a message and transmitting it via sound/sign to the hearer, who then decodes the message. This is the common-sense and folk psychological notion of language where language is used as a conduit for ideas (Reddy 1979). An analysis

of metaphorical expressions used in English, for example, shows that English speakers conceptualise the way they communicate in terms of the conduit metaphor (Vanparys 1995; Semino 2006). This is of course not only a folk psychological concept, but one also taken seriously by many linguists and philosophers.

There are, however, problems with the message model of communication. It cannot adequately account for the way in which we successfully and predictively disambiguate utterances. It also cannot deal with cases of non-literal uses of language, in which the hearer does not decode but rather infer the meaning of an utterance by using various cues, only one of which is the literal meaning of the utterance. There is also the problem of the reference of utterances: how does the hearer know what the speaker is referring to when they produce utterances such as *this happy child*? Does the message that is supposed to be decoded by the hearer contain the information required to determine the reference of the utterance? One could answer this question by proposing some form of referential semantics; however, a better answer seems to be that the hearer infers the reference of the utterance by making use of several types of evidence, including but not limited to the literal meaning of the utterance. Despite these problems, though, the message model of communication was the basis for most explanations of communication well into the late twentieth century (cf. Origgi & Sperber 1998 for discussion). But it is now clear that the message model must be supplemented by other principles if one is to have a satisfactory explanation of the complex nature of human communication. This is where the inferential model, originally developed by Grice (1957; 1975), comes in.

According to the inferential model, communication involves the hearer identifying the intention of the speaker. In producing an utterance, a speaker communicates by giving evidence of what they intend to communicate. The hearer then uses this evidence, only part of which is the linguistic meaning of the utterance, in order to infer the message that the speaker intended to communicate. Grice famously distinguished between linguistic meaning and speaker's meaning. Linguistic meanings are merely one part of a larger set of data a hearer makes use of in order to infer what the speaker intended to communicate, which is the speaker's meaning. This larger set includes a set of shared beliefs and presumptions that speakers and hearers have of each other. It also includes a set of inferential strategies, predictable patterns of inference from linguistic meaning to speaker's meaning. This model of communication has been developed, albeit in different ways, by Bach & Harnish (1979) and Sperber & Wilson (1986).

It should be noted, however, that even though communication has been standardly construed in these two ways, they are not mutually exclusive; some combination of the two is taken for granted in current theorising about human communication.

3.2 A reason for claiming that language is for communication

As already mentioned, in much of the theoretical and empirical work into language it is assumed that the function of language is communication. Often this is the starting point of the discussion; it is an unquestioned working assumption that is not thought to be problematic. Ellis is a further example: he speaks of the "the obvious fact that the entire nature, structure, and development of language is a vehicle for communication or meaning transmission" (Ellis 1999: 53). And Gauker (2002: 687) states that "[t]radition and the contemporary majority hold that language serves communication by allowing speakers to reveal to hearers the conceptual contents of underlying thoughts".[3] Given the nature

---

[3] Cf. also Tomasello (2006: 258): "[…] linguistic constructions are meaningful linguistic symbols. They are nothing other than the patterns in which meaningful linguistic symbols are used in communication". Searle, who is critical of generative linguistics in this regard, sees language as essentially a system for communication – he says that the latter is the common-sense picture of human language and that the former is "peculiar and eccentric" (Searle 1974: 16).

of this working assumption, then, and the fact that many find it obviously true, one needs to dig a little deeper in order to find an explicit argument. I think that the main argument that attempts to ground the claim that language is for communication is the constellation of related evolutionary arguments according to which the adaptive value of language use is its communicative function: language fitness is said to correspond to communicative success. Since the most popular way in which to ground functional attributions in general has been via evolutionary considerations, it is no surprise that arguments in favour of the function of language being communication also take this route.

One of the classic texts that synthesizes this is Pinker & Bloom (1990).[4] They argue – indeed, they see it as entirely uncontroversial – that the structure of language shows evidence of a complex design for the communication of propositional structures. Since natural selection is the best scientific explanation for the emergence of complex structures in organisms, they naturally use evolutionary considerations to ground their claim that the function of language is communication. They do so by arguing that there was a selective advantage in human evolutionary history for using language for communication, and so its primary function must therefore be communication. However, I think that there are two issues here that need to be kept separate: the reasons for why a certain structure remains in the species are not necessarily the same as the reasons that make that structure what it is.

In other words, suppose that we are studying an organism with structure S. We can study the internal structure of S and how that organism uses S in its current environment. But such a study is separate to the further question of why S has remained in the species to which that organism belongs. For example, there could have been a change in the environment that made S useful for survival and thus it remained in the species. Or S could be part of a larger set of structures or mechanisms that together – but not individually – served a useful evolutionary purpose. Or S could have remained in the species for internal biological reasons that had little to do with their interaction with the environment or with natural selection. We can imagine several other scenarios according to which S may or may not be useful to the organism in its particular environment. Now, such (empirical) questions are valid and interesting and by no means trivial, but they are different to the question of what S is. To take a concrete example, knowing the biology of the mammalian lungs tells us very little about why they remained in the species in their current form. It is only by looking at the relation between the environment and the internal biology that we can figure out why the lungs are the way they are. Conversely, knowing that mammalian lungs remained in the species because their function was to facilitate breathing and thus keep humans alive tells us little about the structure of the lungs. We want to know *how* they achieve this (for there is more than one way). Being told that the lungs facilitate breathing given the current state of our environment sets up the problem to be solved for biology. Again, the two questions are both valid and interesting, and both can be pursued in parallel and illuminate one another, but they are separate questions.

I think the same is true of generative linguistics and evolutionary theory. The latter argues that the language faculty remained in the species due to its selective advantage in fostering better communication and co-operation,[5] but this tells us little about the structure of the faculty itself. In fact, as I argue below, looking at the structure of the language

---

[4] Much has been written on the topic since: cf., for example, the more recent Pinker & Jackendoff (2005), which is a reply to the scepticism of Hauser et al. (2002) in regard to the claim that language evolved for (and hence its main function being) communication. The debate continued in Fitch et al. (2005) and the reply to it in Jackendoff & Pinker (2005). Cf. also Jackendoff (1999).

[5] Cf. Millikan (2005: 39): "Coordinations effected through language conventions typically involve communication. [… A] primary function of the human language faculty is to support linguistic conventions, and these have an essentially communicative function. Therefore, a primary function of the human language faculty is to support communication."

faculty suggests that communication is a secondary aspect of language use. Another way to put the matter is that it is near impossible to derive the properties of the underlying computational mechanisms of language from functional accounts of language use. This is because communicative systems are consistent with more than one sort of language faculty, but the question for the biolinguist is why we have this particular language faculty and not some other (Reinhart 2006).

The upshot of the above is that, as per the analytical strategy detailed above, functional attributions only make sense within an explanatory theory that makes use of such attributions in its explanations. Generative linguistics and evolutionary theory are different theories giving different explanations of different phenomena. Functions are not essential properties of organs or of biological mechanisms, they are attributed in a way that best fits the explanatory purpose at hand. It should be stressed that this is not the case for man-made objects, where the function is defined by reference to the intentions of the designer. You can use a bread knife in any way you like (as a paperweight, as a shoehorn, or as a murder weapon) but there is no sense in which its main function and the reason for its existence is not for cutting bread. This is merely a definitional matter established by reference to the intentions of the designer of the object in question. The object would not exist if it were not designed and constructed with a specific function in mind. But in regard to biological entities, like the language faculty in the mind, this pre-theoretical intuition does not help us. There is no analogue in the biological realm to a designer of a man-made object: natural selection is both blind and without foresight. In other words, man-made objects have their functions essentially, whereas biological entities have their functions ascribed to suit the explanatory purposes at hand. Thus, I argue below that the functional attribution that best fits the explanatory purposes of generative linguistics is that of language being an instrument of thought, and this linguistic theory is unswayed by the claim that within evolutionary theory the function of language is communication.

I should note here that I do not wish to rehearse or to adjudicate on the debate in regard to the evolution of language. Rather, I wish to point to what the motivations may be for making the claim that the function of language is communication. It seems to me that the main motivation (perhaps the only fully articulated one) for making this claim is that since there was a selective advantage in human evolutionary history for using language for communication, its primary function must therefore be communication.[6] And besides, the purpose of this article is to flesh out what might be the reasons for seeing the function of language as that of an instrument of thought. Of course, merely pointing to the fact that, as opposed to evolutionary theory, the functional attribution that best fits the explanatory purposes of generative linguistics is that of language being an instrument of thought does not mean that this attribution is correct or even helpful in improving linguistics as a scientifically fecund explanatory theory. The question that must be faced is what explanatory scope does arguing that language is an instrument of thought buy us. To this end we now turn.

4 Language as an instrument of thought

My claim is that the underlying mechanisms of language do not merely express pre-formed thoughts but rather that they also allow humans to think particular types of thoughts that are unavailable to beings who do not have these mechanisms. This is the strong claim in regard to language being an instrument of thought. This is of course not to deny that animals can think without language, and this should not be taken to imply that all thought

---

[6] This is the assumption behind much evolutionary modelling of the emergence of language in humans. Cf. Baronchelli et al. (2012).

is due to the underlying mechanisms of language. As mentioned above, animal cognition is impressive indeed but it is missing a specific kind of thinking that appears to be unique to humans. In order to fully understand the nature of these uniquely human thoughts, let us see what productivity and systematicity consist in.

4.1 Productivity and systematicity

Linguistic productivity is part of what is known as the creative aspect of language use (Asoulin 2013). It is the ability to produce and understand an unlimited number of sentences. This feature of language was noticed by Descartes, who viewed productivity in all domains — language, mathematics, vision, etc. — as deriving from a single source. Modern cognitive science, however, has taken a modular approach, insisting that each domain has its own productivity engine (Brattico & Liikkanen 2009). In order for the underlying mechanisms of language to be able to produce from the set of finite primitive elements an infinite set of expressions they must allow for recursion. For present purposes assume that recursion involves embedding a structural object within another instance of itself — as when a noun phrase is embedded within another noun phrase (cf. Parker 2006; Tomalin 2007; Zwart 2011). The productivity of language means that there is no non-arbitrary limit on the length of a natural language sentence – a sentence (say, S) can always be made longer by embedding it in yet another sentence ("she said that S"), ad infinitum.

Notice that an iterative procedure can of course also produce infinite expressions from a finite set, but iteration is not the same as recursion. While the two procedures are similar in that both can yield structural repetition and thus a potential infinite set, they differ in the way in which they do this and thus in the sorts of expressions they can produce. A procedure is recursive if it builds structures by increasing embedding depth, whereas an iterative procedure can only yield flat structures that have no depth of this kind (cf. Karlsson 2010). Recursive procedures can therefore produce linguistic expressions that are, say, centre-embedded, and that lead to long-distance dependencies. Iterative procedures, on the other hand, cannot produce such expressions. In other words, the indefinite repetition or concatenation of elements (iteration) is not the same as the indefinite embedding of elements within other elements of the same type (recursion). Productivity, then, helps explain how we can deal with novel linguistic contexts and how we can produce and understand sentences that we have not previously encountered.

Linguistic systematicity, on the other hand, refers to the fact that our ability to produce and comprehend expressions of a certain kind guarantees that we can produce or comprehend other systematically related expressions.[7] The classic illustration of systematicity holds that anyone that can understand the sentence *Mary loves John* can also understand the sentence *John loves Mary* – indeed, it is impossible to understand one without also understanding the other. What accounts for this systematicity are our abstract linguistic structures, or, more specifically, our ability to construct structural representations of sentences – looking at the syntax, we have here the abstract syntactic structure [NP [V NP]]. Such abstract structures explain the systematic relation between expressions. Thus, productivity and systematicity are perhaps the best indicators of the creative and open-ended nature of human language.[8]

---

[7] Though I do not agree with some of what he argues, cf. Cummins (1996) for a good discussion of systematicity. Cf. also Cummins et al. (2001) and McLaughlin (2009).

[8] The open-ended nature of language is of course an idealisation, for there are practical limits such as short-term memory limitations or attention span that constrain the length of sentences that people may produce or comprehend. But such limitations are due to systems outside of the language faculty (but internal to the brain) and thus do not change the nature of the language faculty itself.

As Fodor (1975) famously argued, the above is also applicable to human thought. That is, just like language is productive and systematic, so is thought. There is no non-arbitrary limit constraining the length of thoughts; like sentences, the number of different thoughts we can have is infinite. And just as sentences are related to each other in a systematic way, thoughts too are related to each other systematically. Though much has changed in linguistics and cognitive science since the 1970s, Fodor's main argument for why both language and thought are productive and systematic remains unchanged: language and thought both employ a generative procedure that allows the creation of an unbounded set of structured expressions.[9] This procedure of course cannot be responsible for all of our thought processes, for much of what we share with animal cognition is clearly rich and complex but does not involve language nor its underlying mechanisms. So what does allow humans to have a specific type of thought that we don't share with any other species? Note again that this is a difference in kind, not just a difference in degree. Some human thoughts are not just more complicated than animal thoughts, they are structured in a productive and systematic way that is unavailable to non-human animals.

It might seem paradoxical to try to express in language the sort of thoughts that would or would not be possible without the underlying mechanisms of language that generate them. But these underlying mechanisms are of course not linguistic in nature – for if they were they would be explanatorily vacuous in regard to how language and thought work. As will be detailed in the next section, the underlying mechanisms of language include the process that creates recursive and hierarchically structured expressions – this process takes place *before* expressions are given a phonological or semantic interpretation in a particular natural language. With that in mind, consider the thought experiment from Reinhart (2006: 2ff.), which fleshes out one of Chomsky's thought experiments. She imagines a primate that has acquired by some mystery of genetic development the full set of human cognitive abilities but that does not have the language faculty. This fictitious primate would have, in addition to its cognitive abilities that allow it to think like its fellow primates, a set of concepts that is the same as that of humans and a set of sensorimotor systems that enable it to perceive and code information in sounds. Moreover, Reinhart imagines that this primate would also have the human system of logic, the abstract formal system that contains an inventory of abstract symbols, connectives, functions, and definitions necessary for inference. Given the nature of this primate, then, what would it be able to do with these systems? That is, given all these additions but lacking the language faculty, can this fictitious primate add to its thinking abilities the sorts of thoughts that display productivity and systematicity and that at present appear to be unique to humans? Reinhart argues that it could not.

At first blush this seems to be a surprising claim. If the primate has acquired the rich conceptual system of humans then presumably its preexisting inference system should allow it to use these newly acquired systems to construct more sophisticated theories that it can then use to, say, better navigate a complex terrain or make better and more complex inferences about its world. But this is not the case. What prevents this fictitious primate from making use of the new systems and concepts it has acquired is the fact that our inference system operates on *propositions*, not on concepts. The primate can of course communicate its preexisting concepts to its fellows, and it can make inferences typical of primates, but since it does not have the ability to create recursive and hierarchically structured expressions it cannot construct or comprehend propositions necessary for *higher-order* inference.

---

[9] Fodor's mentalese is but one attempt at elucidating what this underlying procedure is. As I detail below, however, I think that narrow syntax as construed within biolinguistics is a better explanation for the specifically human type of thought.

In other words, this imagined primate has concepts and knowledge of first-order logic, which it can use and comprehend, but that is not enough to produce and comprehend propositions nor to make second-order and higher-order inferences. In order to be able to do the latter, the primate in the thought experiment must – but does not – possess recursion. A fortiori, this primate cannot comprehend the entailment relations between propositions – it cannot think those sorts of thoughts. Now compare this fictitious primate to real world humans: we *can* think those sorts of thoughts. This is because the way in which the underlying mechanisms of language work in humans is by providing us with higher-order logic (cf. Crain 2012 on the relation between natural language and classical logic), by providing us with a computational system that creates recursive and hierarchically structured expressions that display productivity and systematicity and that we use to, amongst other uses, talk and think about the world (cf. also Hinzen 2006; 2013).

In what follows, then, I want to pursue the stronger claim in regard to language being an instrument of thought, and the evidence that may be adduced in its favour. The type of evidence and sorts of arguments can be divided into two kinds: the first is the *argument from linguistics*, according to which the externalisation of language – in, say, verbal communication – is a peripheral phenomenon because the phonological features of expressions in linguistic computations are secondary (and perhaps irrelevant) to the conceptual-intentional features of the expressions. The second is the *design-features argument*, according to which the design features of language, especially when seen from the perspective of their internal structure, suggest that language developed and functions for purposes that are not primarily those of communication.

4.2 The argument from linguistics

A strong argument in favour of language being primarily an instrument of thought has to do with the phonological properties of lexical items. Within biolinguistics,[10] lexical items, and all expressions generated from them, have properties that must be interpretable at both the language faculty's interface with the conceptual-intentional systems and at the interface with the sensorimotor systems. Briefly, the idea is that the internal computational processes of the language faculty (syntax in a broad sense) generate linguistic objects that are employed by the conceptual-intentional systems (systems of thought) and the sensorimotor systems to yield language production and comprehension. Notice that on this view the language faculty is embedded within, but separate from, the performance systems. So we have a device that generates structured expressions of the form $Exp = \langle Phon, Sem \rangle$, where *Phon* provides the sound instructions of which the sensorimotor systems make use, and *Sem* provides the meaning instructions of which the systems of thought make use. *Phon* contains information in a form interpretable by the sensorimotor systems, including linear precedence, stress, temporal order, prosodic and syllable structure, and other articulatory features. *Sem* contains information interpretable by the systems of thought, including event and quantification structure, and certain arrays of semantic features.

The expression *Exp* is generated by the operation Merge, which takes objects already constructed and constructs from them a new object. If two objects are merged, and principles of efficient computation hold, then neither will be changed – this is indeed the result of the recursive operation that generates *Exp*. Such expressions are not the same as linguistic utterances but rather provide the information required for the sensorimotor systems and the systems of thought to function, largely in language-independent ways. In other words, the sensorimotor systems and the systems of thought operate independently

---

[10] What follows is a summary of the current state of the art in biolinguistics. Cf., amongst others, Chomsky (1995; 2007), and the edited collections by Di Sciullo & Aguero (2008) and Piattelli-Palmarini et al. (2009).

of (but at times in close interaction with) the faculty of language. A mapping to two interfaces is necessary because the systems have different and often conflicting requirements. That is, the systems of thought require a particular sort of hierarchical structure in order to, for example, calculate relations such as scope; the sensorimotor systems, on the other hand, often require the elimination of this hierarchy because, for example, pronunciation must take place serially.

The instructions at the *Sem* interface that are interpreted by the performance systems are used in acts of talking and thinking about the world – in, say, reasoning or organising action. These semantic properties "focus attention on selected aspects of the world as it is taken to be by other cognitive systems, and provide intricate and highly specialised perspectives from which to view them, crucially involving human interests and concerns even in the simplest cases" (Chomsky 2000: 125). On this view, then, linguistic expressions provide a perspective (in the form of a conceptual structure) on the world, for it is only via language that certain perspectives are available to us and to our thought processes. This is the sense in which I take language to be an instrument of thought. Language does not structure human thought in a Whorfian way, nor does it merely express pre-formed thoughts; rather, language (with its expressions arranged hierarchically and recursively) provides us with a unique way of thinking and talking about the world.

Lexical items, then, and all expressions generated from them, are linguistic objects with a double interface property: they have phonological and semantic features through which the linguistic computations can interact with other cognitive systems – indeed, the only principles allowed under the minimalist program are those that can function at the interfaces. Thus, if one were to imagine an order of operations, the process would be as follows: first a lexical item is created with syntactic, phonological, and semantic features. Then, in the process known as Spell Out, the phonological features are sent to the sensorimotor interface, leaving the syntactic and semantic features together to be sent to the conceptual-intentional interface (cf. Burton-Roberts 2011). The upshot of this model is that the structure and content of lexical meanings are composed independently of how they are expressed phonologically in sound/sign. This is strong evidence in favour of the thesis that language is an instrument of thought, for the central computations in which lexical meanings are produced are carried out independently of any consideration as to how or whether they are to be communicated. These computations create recursive and hierarchically structured expressions that are used internally – as we'll see in the next section, the design features of language indicate that these expressions are optimised for computational as opposed to communicative efficiency. Thus, the externalisation of language is a peripheral phenomenon in the sense that the phonological features of expressions in linguistic computations are peripheral to the syntactic and semantic features of these expressions. This is in line with Sigurðsson's (2003; 2004) Silence Principle, according to which a great deal of what languages have in common are silent categories that are present in narrow syntax but silent at the *Phon* interface. Language thus has "innate elements and structures irrespective of whether or how they are overtly expressed" (Sigurðsson 2004: 235).

In addition to the above, we have independent evidence from comparative, neuropathological, developmental, and neuroscientific research that supports the existence of an asymmetry between the interfaces in favour of the semantic side, pushing externalisation to the periphery. I'm thinking specifically of research that shows the modality-independence of the externalisation of language. The work of Laura-Ann Petitto, for example, has shown that speech per se is not critical to the human language acquisition process. That is, the acquisition of language occurs in the same way in all healthy children, irrespective of the modality in which the child is exposed to language (speech in hearing children, sign in

deaf children, and even the tactile modality). This suggests that the brain is hardwired to tune in to the structure and meaning of what is expressed, but that the modality through which this is transmitted is irrelevant (Petitto 2005). In other words, the syntax and semantics of language are processed in the same brain site regardless of the modality in which they are expressed and perceived. Such evidence gives weight to the biolinguistic argument that syntax and semantics are computed together without recourse to the way in which (if at all) the product of this computation (say, lexical meanings) is to be externalised. There is further evidence of this sort: it appears that the neural specialization for processing language structure is not modifiable, whereas the neural pathways for externalising language are highly modifiable (Petitto et al. 2000). This again suggests that the language areas of the brain are optimized for processing linguistic structures and meaning, and that their externalisation is not only secondary but also that their type is not fixed – any modality would do as long as the brain can interpret the required linguistic patterns in the input.

Recent work by Ding et al. (2016) also points in this direction. Their experiments showed that "the neural tracking of hierarchical linguistic structures was dissociated from the encoding of acoustic cues and from the predictability of incoming words" (Ding et al. 2016: 158). They found that the neural tracking of multiple levels of linguistic structure (such as the phrasal or sentential levels) was dissociated from the encoding of acoustic cues, and concluded that "there are cortical circuits specifically encoding larger, abstract linguistic structures without responding to syllabic-level acoustic features of speech" (Ding et al. 2016: 162). These cortical circuits track abstract linguistic structures that are internally constructed and that are based on syntax. Further evidence of the modality independence of language, indeed the condition under which it is most acute, comes from cases where there is practically no externalisation (perhaps only the ability to say a few phonemes) but where the receptive language ability is completely intact. Eric Lenneberg discusses the case of a child "who is typical of a large group of children with deficits in their motor execution of language skills but who can learn to understand language even in the total absence of articulation" (Lenneberg 1962: 419). This form of developmental speech dyspraxia suggests that the ability to comprehend language and make normal grammaticality judgments does not depend on normal language production (Stromswold 1999). In other words, as the work of Caplan et al. (2007) suggests, the language difficulties of aphasics is due to their intermittent or complete failure to externalise language. That is, the linguistic competence at the syntactic and semantic levels remains intact but these patients have difficulty in linking this competence with the performance systems – they have difficulty in externalising the internally constructed expressions.

It should be stressed that the above does not merely point to the modality independence of language, which was of course known well before these recent neuroscientific studies (Yamada 1990; Smith & Tsimpli 1995). The above is direct evidence in support of the claim that there exists a separation in the underlying mechanisms of language between, on the one hand, the processing of structure and meaning, and, on the other hand, their externalisation. That is, not only is the processing of non-language information dissociated from the processing of information used in language, but also that the processing of the language information itself is separated into *Phon* and *Sem,* just as biolinguistics predicts. Note that this asymmetry regards the underlying mechanisms of language and thus does not apply in the same way to natural languages. So whilst it makes sense to separate *Phon* from *Sem* when one studies the underlying mechanisms of language, specific natural languages are a different matter. That is, a natural language encapsulates the use of the *Phon* and *Sem* interfaces – in conjunction with other modules – in the act of communication via sound or sign, and so the *Phon* interface is inseparable from what

a natural language is and the way it is used. In contrast to this, the claim that language is an instrument of thought regards the part of the underlying mechanisms of natural languages that creates the hierarchical and recursive expressions that provide humans with a unique way of thinking about the world. This part on its own is of course not yet a particular natural language, for it is not yet in a form in which it can be externalised. In order to become a natural language it needs to be paired with the *Phon* interface and then, together with other systems, be used in the act of communication.

Returning to the double interface object, one might wonder why the asymmetry between the interfaces is in favour the semantic side, pushing externalisation to the periphery. Why isn't it the other way around? I think the answer to this comes in the form of the *design-features argument*.

4.3 The design-features argument

If one does not share the general framework of biolinguistics, then they will perhaps be unconvinced by the *argument from linguistics* above. *The design-features argument*, on the other hand, has much wider scope and is not entirely dependent upon a particular linguistics school of thought. By design features I mean the kind of features one discovers upon investigating language as a system in its own right. Such features include, amongst many others, displacement, linear order, agreement, and anaphora. One may then investigate the communicative and computational efficiency of these features as they relate to language as a whole system, and ask whether these features are better optimised for communication or for computation.[11] The closer the design features of language are to being optimised for computation, the stronger the case for language being an instrument of thought. Of course, many comparisons of this sort can be made, and some particular selection that depicts a conflict between communicative efficiency and computational efficiency might seem tendentious, but I think that the conflicts of the sort highlighted below, in which computational efficiency wins out, represent one of several chinks in the armour of the orthodoxy that assumes that the function of language is communication. Let us now consider the case of the explanation of the linear order of expressions.

The linear order imposed on verbal expressions is not a language-specific constraint: it is not a consequence of the structure of the language faculty. Rather, it is a necessary consequence of the structure of the sensorimotor systems and the obvious fact that expressions cannot be produced or comprehended in parallel. The sensorimotor systems are constrained by the "hardware" available in the brain for sound/sign production and perception – for example, the time, short-term memory, and linear order constraints of real-time parsing. Assuming this is the case, then, what is the effect of such constraints on, say, the computations involved in parsing sound inputs into linguistic representations? If language is optimised for communication and if sound is our main source of externalisation, then one would predict that many of the features of language would respect linear order and favour operations that support it even if they conflict with computational efficiency. Closer investigation, however, suggests that this is not the case. Consider, for example, how co-reference is interpreted in sentences such as *In her study, Jane is mostly productive,* where *her* and *Jane* are interpreted as being co-referential. It was initially thought (Langacker 1969; Jackendoff 1972; Lasnik 1976) that in order to explain the difference between, say, (1) and (2) below, a linear relationship of precede-and-command was needed, according to which the pronoun cannot both precede and command its antecedent.

---

[11] By computation here I do not mean anything other than what is assumed in the computational theory of mind, which is the central doctrine of cognitive science in regard to computational systems such as language or vision (cf. Gallistel & King 2010 for an excellent discussion of computation).

(1)   *She$_i$ denied that Jane$_i$ met the minister.

(2)   The man who travelled with her$_i$ denied that Jane$_i$ met the minister.

The explanation used to be that in (1) the pronoun precedes and commands the full noun phrase and therefore the co-referential interpretation is blocked. In (2), conversely, it was claimed that the pronoun precedes but does not command the full noun phrase and therefore a co-referential interpretation is permitted.

   However, as Reinhart (1983) shows, the domains over which the precede-and-command operations are defined are quite arbitrary; the parts of the expressions that are preceded or commanded by other parts often do not correspond to independently characterisable syntactic units. On independent grounds, then, it would be surprising if such an arbitrary linear relationship would turn out to be the operative co-referential explanation. This is clear in (3) and (4) below, which cannot be explained by precede-and-command operations (cf. Reinhart 1983: 36ff.). In (3a) the pronoun cannot refer to *Mary*, whereas in (3b) the co-referential interpretation is permitted. However, when we consider (4), which is the pre-preposed version of the sentences in (3), the co-referential interpretation is blocked in both (4a) and (4b).

(3)   a. *In John's picture of Mary$_i$, she$_i$ found a scratch.
        b.   In John's picture of Mary$_i$, she$_i$ looks sick.

(4)   a. *She$_i$ found a scratch in John's picture of Mary$_i$.
        b. *She$_i$ looks sick in John's picture of Mary$_i$.

Thus, no ordering explanation such as precede-and-command can account for the difference between (3a) and (3b). Or compare (5a) and (5b), both of which are allowed by the relation of precede-and-command but only one of which has an acceptable co-referential reading.

(5)   a. *Ben's$_i$ problems, he$_i$ won't talk about.
        b.   Ben's$_i$ problems, you can't talk to him$_i$ about.

As Reinhart shows with a range of other examples, there is good reason to think that, instead of a linear order operation, the explanation of co-reference has to do with the structural properties of the expressions. According to the structure-dependent analysis, coreferential interpretations are only permitted when anaphors are bound by another nominal. This binding is a structure sensitive and asymmetric relation according to which a subject can bind an object, but an object cannot bind a subject. In regard to the above examples, there is an asymmetry between the coreference options of subjects and those of objects (or non-subjects), for in cases with preposed constituents forward pronominalisation is impossible where the pronoun is the subject – as in (3a) and (5a) – but possible where the pronoun is not the subject – as in (3b) and (5b). Thus, the hierarchical relation of binding, involving both c-command and coindexation, supplies us with a more encompassing and much improved explanation of the phenomena – it explains not only what the relation of precede-and-command explains, but also the cases that cannot be explained by invoking ordering relationships.

   The structural relation of c-command has been shown to be a fundamental relation in syntax that underlies many diverse linguistic phenomena (cf. Barker & Pullum 1990 for a detailed discussion of command relations), and the above is one example of a general phenomenon in which computationally simpler operations, such as linear distance,

are ignored in favour of computationally complex structure-dependent operations.[12] Moreover, there is by now a great deal of psycholinguistic evidence in favour of this claim (cf. Phillips & Wagers 2007 for a recent overview of the psycholinguistic research on structure dependency; cf. also the collection edited by Sprouse & Hornstein 2013). It should be noted that this holds for a specific kind of computation instantiated in the human brain – biological computation, if you will. If modern computers could consistently parse natural language expressions by methods that assume that the expressions are based on linear distance or statistical regularities, that would be an interesting and valuable outcome that could be put to numerous practical uses. However, whether or not computers can or would be able to do this is not relevant to this discussion because our brains do not work in that way: language appears to use structure-dependent operations almost entirely (cf. Moro 2013; 2014).

These operations are often irrelevant to externalisation and in many cases are in direct conflict with the efficient operation or needs of the sensorimotor systems as they are used in communicating. Linear distance is more efficient, arguably less taxing for the parser, and simpler from the point of view of communication but is largely absent in the crucial cases where one would expect it. As Chomsky (2013) argues, one explanation for this phenomenon is that linear distance is simply not available to the child during language acquisition; they are instead guided by a principle that dictates that there is no such thing as linear order and that only structure-dependent operations are to be considered. But structure-dependent operations cause problems for communication that would not arise if, say, linear distance was used instead. Linguistic expressions seem to be optimised for computational efficiency, they are not structured in a way that favours ease of communication.

A key source of evidence in favour of the claim that language is optimised for computational efficiency is that computational efficiency appears to be a feature of biological systems, which of course include the human brain and the language faculty within it. In other words, evidence for the computational efficiency of the human brain is also evidence for the computational efficiency of the language faculty because the latter is part of the former. Christopher Cherniak's work on computational neuroanatomy applies combinatorial network optimisation theory (Garey & Johnson 1979) to brain structure. His neuroanatomical studies deal with neural component placement optimisation, and show that "when anatomical positioning is treated like a microchip layout wire-minimisation problem, a 'best of all possible brains' hypothesis predicts actual placement of brains, their ganglia, and even their nerve cells" (Cherniak 1994: 89). In other words, layout optimisation predicts anatomy in that the so-called "save wire" principle appears to be an organising principle of brain structure (Chklovskii et al. 2002). The neural connections in the brain are a highly constrained and finite resource, especially the longer range ones that are subject to constraints due to volume and signal-propagation times. In order to carry out the wide range of tasks it is capable of, the brain must deal with these constraints in a finite time, and it does so in a way that is not merely "good enough" or just satisfying (Simon 1956). There are innumerable local maxima that would do for the task at hand, but the brain appears to be structured in an optimal way that is closer or indeed at the global maximum, asymptotically close to being the best of all possible brains given the constraints at hand and the initial conditions (Cherniak et al. 2002).

---

[12] Others include Chomsky's (2013a) famous example of auxiliary inversion ("Can eagles that fly swim?" / "Instinctively, eagles that fly swim."). As he argues, the relationship between the two clauses satisfy the condition of minimal structural distance and not the much simpler computational operation of minimal linear distance. Cf. also Ross (1986), which outlines a set of constraints that limit syntactic computations.

The optimal structure of biological neural systems was first confirmed by studying the neural system of the millimetre-long roundworm, *C. elegans*, which is one of the only animals for which we have a complete neuroanatomical map, down to the synapse level (Cherniak 1995; 2005). A dozen computers running for an entire week exhaustively searched all possible ganglion layouts of the roundworm, yielding the result that, as Cherniak puts it, the actual is the ideal and optimal: the roundworm's actual ganglion layout appears to be the most optimal of the millions of possible layouts. Subsequent studies have observed even finer wiring optimisation in the layout of the cerebral cortex of rats, cats and macaque monkeys (Cherniak et al. 2004). Since the way in which the human brain works is optimised in the above sense, and since the language faculty is in the brain, there is no reason to expect that the language faculty would not also respect the principle of efficient computation. Incidentally, such conclusions of optimality, developmental constraints and the like, are not limited to neuroscience nor to the language faculty (Boeckx & Piattelli-Palmarini 2005). Rather, they can be found in work across many species, from the inception of modern biology (Thompson 1942; Turing 1952; Leiber 2001) to current thinking (Maynard Smith 1985; Kauffman 1993; Stewart 1998; Gould 2002; Fox Keller 2002).

The asymmetry in favour of computational efficiency is dramatically illustrated by the presence in natural language of structural ambiguity and garden path sentences. These clearly cause problems for communication and so one might ask why natural language has them at all. Phillips (1996) shows that many well-known ambiguities (such as *John said Bill left yesterday*) can in fact be explained by the same computational principle (Branch Right) that forces certain structural biases in the parsing of expressions (cf. also Phillips & Lewis 2013).[13] This is a default structural choice that the parser, all things being equal, opts for – and this choice has nothing to do with communicative efficiency. As Phillips puts it, experimental results show "that given the choice between a local attachment which is structurally more complex, not supported by discourse and not required by syntax or semantics, and a non-local attachment which is structurally 'simpler' and involves an obligatory syntactic constituent, the parser opts for the more local attachment" (Phillips 1996: 20).[14] The normal operation of the computational system has all sorts of unwanted effects, such as ambiguities and garden path sentences, but only when communication is involved. This suggests that externalising language with intent to communicate is a peripheral aspect of language.

If the primary function of language were communication then one would expect that the underlying mechanisms of language will be structured in a way that favours successful communication. But we find that not only is this not the case, but that the underlying mechanisms of language are in fact structured in a way to maximise computational efficiency, which ends up causing communicative problems. There is the computational system and there is the parser, the latter of which is part of the systems that externalise language. The closer these two are to each other, the more that is gained in terms of overall efficiency. The further away they are from each other, however, the more questions that arise as to why? I have suggested that the two are further away than is commonly thought. In order to externalise language, the parser must respect the computational and structural features of the syntax-LF bundle. The latter, however, causes problems for the externalisation of language in communication (ambiguity, garden path sentences, etc.). The problems for communication arise when the computational system, which operates

---

[13] Explanations for the presence of ambiguity in natural language in terms of communicative function have of course also been proposed (e.g., Piantadosi et al. 2012), but space does not permit to review them here.

[14] Notice that evidence of this kind also gives weight to the working assumption that the parser, which is part of the sensorimotor systems used for the externalisation of language, is separate from the computational system of language.

along an independent path, is used in the act of communicating. In other words, impediments to successful and smooth communication are the result of the computational operations of the underlying mechanisms of language being asked to perform a function, externalising language with intent to communicate, that is not their primary function.

To recap, there is a conflict between computational efficiency and communicative efficiency. If the function of language were primarily for communication, one would predict that the language system would opt for an operation that aids in the communication of propositional thoughts, or at the very least one that does not hinder parsing or interpretation. But a look at the evidence from linguistic, comparative, neuropathological, developmental, and neuroscientific research shows that this is not the case. This suggests that the language system is composed of computational operations that are optimised for computation, not for communicative efficiency.

5 Concluding remarks

It follows from the above that the nature of language, when taking into account its design features and its internal structure, is not as it is widely assumed to be. That is, language is meaning with externalisation (in sound, sign, etc.); it is not, as the orthodox conception would have it, sound and meaning combinations structured for communicative efficiency. Speech, sign, or any other kind of externalisation are secondary properties of language. The fundamental property of language is the internal construction of indefinitely many expressions by a generative procedure that yields a uniquely human perspective (in the form of a conceptual structure) on the world. It is in this sense that language (not a particular natural language but rather its underlying computational mechanisms) is an instrument of thought; it provides us with a unique way of structuring the world around us, which we use for various purposes such as thinking and talking about the world.

Acknowledgements

Competing interests

The author has no competing interests to declare.

References

Akmajian, Adrian, Richard A. Demers & Robert M. Harnish. 1980. Overcoming inadequacies in the "Message-Model" of linguistic communication. *Communication and Cognition* 13. 317–336. Reprinted in Asa Kasher (ed.). 1998. *Pragmatics: Critical concepts. 6 vols.* New York, NY: Routledge.

Asoulin, Eran. 2013. The creative aspect of language use and the implications for linguistic science. *Biolinguistics* 7. 228–248.

Bach, Kent & Richard M. Harnish. 1979. *Linguistic communication and speech acts.* Cambridge, MA: MIT Press.

Barker, Chris & Geoffrey K. Pullum. 1990. A theory of command relations. *Linguistics and Philosophy* 13. 1–34. DOI: http://dx.doi.org/10.1007/BF00630515

Baronchelli, Andrea, Nick Chater, Romualdo Pasto-Satorras & Morten H. Christiansen. 2012. The biological origins of linguistic diversity. *PLoS ONE* 7(10). e48029. DOI: http://dx.doi.org/10.1371/journal.pone.0048029

Berwick, Robert C., Kazuo Okanoya, Gabriel J.L. Beckers & Johan J. Bolhuis. 2011. Songs to syntax: The linguistics of birdsong. *Trends in Cognitive Sciences* 15(3). 113–121. DOI: http://dx.doi.org/10.1016/j.tics.2011.01.002

Boeckx, Cedric. 2011. Biolinguistics: A brief guide for the perplexed. *Linguistic Sciences* 10(5). 449–463.

Boeckx, Cedric & Massimo Piattelli-Palmarini. 2005. Language as a natural object: Linguistics as a natural science. *The Linguistic Review* 22. 447–466. DOI: http://dx.doi.org/10.1515/tlir.2005.22.2–4.447

Brattico, Pauli & Lassi Liikkanen. 2009. Rethinking the Cartesian theory of linguistic productivity. *Philosophical Psychology* 22(3). 251–279. DOI: http://dx.doi.org/10.1080/09515080902957357

Burton-Roberts, Noel. 2011. On the grounding of syntax and the role of phonology in human cognition. *Lingua* 121. 2089–2102. DOI: http://dx.doi.org/10.1016/j.lingua.2011.08.001

Caplan, David, Gloria Waters, Gayle DeDe, Jennifer Michaud & Amanda Reddy. 2007. A study of syntactic processing in aphasia I: Behavioral (psycholinguistic) aspects. *Brain and Language* 101. 103–150. DOI: http://dx.doi.org/10.1016/j.bandl.2006.06.225

Carruthers, Peter. 2002. The cognitive functions of language. *Behavioral and Brain Sciences* 25. 657–726.

Cherniak, Christopher. 1994. Philosophy and computational neuroanatomy. *Philosophical Studies* 73. 89–107. DOI: http://dx.doi.org/10.1007/BF01207659

Cherniak, Christopher. 1995. Neural component placement. *Trends in Neuroscience* 18. 522–527. DOI: http://dx.doi.org/10.1016/0166–2236(95)98373–7

Cherniak, Christopher. 2005. Innateness and brain-wiring optimization: Non-genomic nativism. In António Zilhão (ed.), *Evolution, rationality and cognition: A cognitive science for the twenty-first century*, 103–112. New York, NY: Routledge.

Cherniak, Christopher, Zekeria Mokhtarzada, Raul Rodriguez-Esteban & Kelly Changizi. 2004. Global optimization of cerebral cortex layout. *Proceedings National Academy of Sciences* 101. 1081–1086. DOI: http://dx.doi.org/10.1073/pnas.0305212101

Cherniak, Christopher, Zekeria Mokhtarzada & Uri Nodelman. 2002. Optimal-wiring models of neuroanatomy. In Giorgio A. Ascoli (ed.), *Computational neuroanatomy: Principles and methods*, 71–82. Totowa, New Jersey: Humana Press. DOI: http://dx.doi.org/10.1385/1-59259-275-9:71

Chklovskii, Dmitri B., Thomas Schikorski & Charles F. Stevens. 2002. Wiring optimization in cortical circuits. *Neuron* 34. 341–347. DOI: http://dx.doi.org/10.1016/S0896–6273(02)00679–7

Chomsky, Noam. 1995. *The minimalist program.* Cambridge, MA: MIT Press.

Chomsky, Noam. 2000. Minimalist inquiries: The framework. In Roger Martin, David Michaels, & Juan Uriagereka (eds.), *Step by step: Essays on minimalist syntax in honor of Howard Lasnik*, 89–155. Cambridge, MA: MIT Press.

Chomsky, Noam. 2007. Biolinguistic explorations: Design, development, evolution. *International Journal of Philosophical Studies* 15(1). 1–21. DOI: http://dx.doi.org/10.1080/09672550601143078

Chomsky, Noam. 2013. What kind of creatures are we? Lecture 1: What is language? *The Journal of Philosophy* 90(12). 645–700. DOI: http://dx.doi.org/10.5840/jphil2013110121

Chomsky, Noam. 2013a. Problems of projection. *Lingua* 130. 33–49. DOI: http://dx.doi.org/10.1016/j.lingua.2012.12.003

Crain, Stephen. 2012. *The emergence of meaning.* Cambridge: Cambridge University Press. DOI: http://dx.doi.org/10.1017/CBO9780511842863

Cummins, Robert. 1975. Functional analysis. *The Journal of Philosophy* 72(20). 741–765. DOI: http://dx.doi.org/10.2307/2024640

Cummins, Robert. 1996. Systematicity. *The Journal of Philosophy* 93(12). 591–614. DOI: http://dx.doi.org/10.2307/2941118

Cummins, Robert, James Blackmon, David Byrd, Pierre Poirier, Martin Roth & Georg Schwarz. 2001. Systematicity and the cognition of structured domains. *The Journal of Philosophy* 98(4). 167–185. DOI: http://dx.doi.org/10.2307/2678476

Cummins, Robert & Martin Roth. 2009. Traits have not evolved to function the way they do because of a past advantage. In Francisco J. Ayala & Robert Arp (eds.), *Contemporary debates in philosophy of biology*, 72–85. Oxford: Wiley-Blackwell. DOI: http://dx.doi.org/10.1002/9781444314922.ch4

Davidson, Donald. 1975. Thought and talk. In Samuel D. Guttenplan (ed.), *Mind and language: Wolfson College lectures 1974*, 7–23. Oxford: Clarendon Press.

Davidson, Donald. 1982. Rational animals. *Dialectica* 36(4). 317–328. DOI: http://dx.doi.org/10.1111/j.1746–8361.1982.tb01546.x

Deacon, Terrence W. 1997. *The symbolic species: The co-evolution of language and the brain.* New York, NY: W. W. Norton & Company.

de Waal, Frans. 2001. *The ape and the sushi master: Reflections of a primatologist.* New York, NY: Basic Books.

Ding, Nai, Lucia Melloni, Hang Zhang, Xing Tian & David Poeppel. 2016. Cortical tracking of hierarchical linguistic structures in connected speech. *Nature Neuroscience* 19(1). 158–164. DOI: http://dx.doi.org/10.1038/nn.4186

Di Sciullo, Anna Maria & Cedric Boeckx (eds.). 2011. *The biolinguistic enterprise: New perspectives on the evolution and nature of the human language faculty*. Oxford: Oxford University Press.

Dummett, Michael. 1981. *The interpretation of Frege's philosophy.* London: Duckworth.

Dummett, Michael. 1989. Language and communication. In Alexander George (ed.), *Reflections on Chomsky,* 192–212. Oxford: Blackwell.

Dyer, Fred & Jeffrey Dickinson. 1994. Development of sun compensation by honeybees: How partially experienced bees estimate the sun's course. *Proceedings of the National Academy of Sciences, USA* 91. 4471–4474. DOI: http://dx.doi.org/10.1073/pnas.91.10.4471

Ellis, Donald G. 1999. *From language to communication.* London: Lawrence Erlbaum Associates.

Fitch, Tecumseh W., Marc D. Hauser & Noam Chomsky. 2005. The evolution of the language faculty: Clarifications and implications. *Cognition* 97. 179–210. DOI: http://dx.doi.org/10.1016/j.cognition.2005.02.005

Fodor, Jerry A. 1975. *The language of thought.* New York: Crowell.

Fodor, Jerry A. 1998. *Concepts: Where cognitive science went wrong.* Oxford: Clarendon Press. DOI: http://dx.doi.org/10.1093/0198236360.001.0001

Fodor, Jerry A. 2001. Language, thought, and compositionality. *Mind & Language* 16. 1–15. DOI: http://dx.doi.org/10.1111/1468–0017.00153

Fodor, Jerry A. 2008. *LOT 2: The language of thought revisited.* Oxford: Oxford University Press. DOI: http://dx.doi.org/10.1093/acprof:oso/9780199548774.001.0001

Fox Keller, Evelyn. 2002. *Making sense of life.* Cambridge, MA: Harvard University Press.

Gallistel, Randy Charles. 1990. *The organization of learning.* Cambridge, MA: MIT Press.

Gallistel, Randy Charles (ed.). 1991. *Animal cognition.* Cambridge, MA: MIT Press.

Gallistel, Randy Charles. 2009. The foundational abstractions. In Massimo Piattelli-Palmarini, Juan Uriagereka & Pello Salaburu (eds.), *Of minds and language,* 58–73. Oxford: Oxford University Press.

Gallistel, Randy Charles. 2011. Prelinguistic thought. *Language learning and development 7(4)*. 253–262. DOI: http://dx.doi.org/10.1080/15475441.2011.578548

Gallistel, Randy Charles & Adam Philip King. 2010. *Memory and the computational brain: Why cognitive science will transform neuroscience*. Oxford: Wiley-Blackwell.

Garey, Michael & David Johnson. 1979. *Computers and intractability: A guide to the theory of NP-completeness*. New York, NY: W. H. Freeman.

Gauker, Christopher. 2002. No conceptual thought without language. *Behavioral and Brain Sciences* 25. 687. DOI: http://dx.doi.org/10.1017/S0140525X02360126

Gould, Stephen Jay. 2002. *The structure of evolutionary theory*. Cambridge, MA: Harvard University Press.

Grice, H. Paul. 1957. Meaning. *Philosophical Review* 66. 377–388. DOI: http://dx.doi.org/10.2307/2182440

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Syntax and semantics 3: Pragmatics*, 41–58. New York, NY: Academic Press.

Hauser, Marc D. 2000. *Wild minds.* New York, NY: Henry Holt and Company.

Hauser, Marc D., Charles Yang, Robert C. Berwick, Ian Tattersall, Michael J. Ryan, Jeffrey Watumull, Noam Chomsky & Richard C. Lewontin. 2014. The mystery of language evolution. *Frontiers in Psychology* 5. 1–12. DOI: http://dx.doi.org/10.3389/fpsyg.2014.00401

Hauser, Marc D., Noam Chomsky & Tecumseh W. Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298. 1569–1579. DOI: http://dx.doi.org/10.1126/science.298.5598.1569

Hinzen, Wolfram. 2006. *Mind design and minimal syntax*. Oxford: Oxford University Press. DOI: http://dx.doi.org/10.1093/acprof:oso/9780199289257.001.0001

Hinzen, Wolfram. 2013. Narrow syntax and the language of thought. *Philosophical Psychology* 26(1). 1–23. DOI: http://dx.doi.org/10.1080/09515089.2011.627537

Hurford, James R. 2007. *The origins of meaning: Language in the light of evolution*. Oxford: Oxford University Press.

Jackendoff, Ray. 1972. *Semantics in generative grammar*. Cambridge, MA: MIT Press.

Jackendoff, Ray. 1999. Possible stages in the evolution of the language capacity. *Trends in Cognitive Sciences* 3(7). 272–279. DOI: http://dx.doi.org/10.1016/S1364–6613(99)01333–9

Jackendoff, Ray. 2002. *Foundations of language: Brain, meaning, grammar, evolution*. Oxford: Oxford University Press.

Jackendoff, Ray & Steven Pinker. 2005. The nature of the language faculty and its implications for evolution of language (Reply to Fitch, Hauser, and Chomsky). *Cognition* 97. 211–225. DOI: http://dx.doi.org/10.1016/j.cognition.2005.04.006

Karlsson, Fred. 2010. Syntactic recursion and iteration. In Harry van der Hulst (ed.), *Recursion and human language,* 43–67. Berlin: Mouton de Gruyter. DOI: http://dx.doi.org/10.1515/9783110219258.43

Kauffman, Stuart. 1993. *The origins of order*. Oxford: Oxford University Press.

Langacker, Ronald W. 1969. On pronominalization and the chain of command. In David A. Reibel & Sanford A. Schane (eds.), *Modern Studies in English*, 160–186. Englewood Cliffs, NJ: Prentice-Hall.

Lasnik, Howard. 1976. Remarks on coreference. *Linguistic Analysis* 2. 1–22.

Leiber, Justin. 2001. Turing and the fragility and insubstantiality of evolutionary explanations: A puzzle about the unity of Alan Turing's work with some

larger implications. *Philosophical Psychology* 14(1). 83–94. DOI: http://dx.doi.org/10.1080/09515080120033553

Lenneberg, Eric. 1962. Understanding language without ability to speak: A case report. *Journal of Abnormal and Social Psychology* 65(6). 419–425. DOI: http://dx.doi.org/10.1037/h0041906

Mangum, Wyatt A. 2010. Animal communication: The "language" of honey bees. In Susan J. Behrens & Judith A. Parker (eds.), *Language in the real world: An introduction to linguistics*, 255–273. New York, NY: Routledge.

Maynard Smith, John, Richard Burian, Stuart A. Kauffman, Pere Alberch, John H. Campbell, Brian Goodwin, Russell Lande, David M. Raup, & Lewis Wolpert. 1985. Developmental constraints and evolution. *The Quarterly Review of Biology* 60(3). 265–287. DOI: http://dx.doi.org/10.1086/414425

McDowell, John. 1994. *Mind and world.* Cambridge, MA: MIT Press.

McLaughlin, Brian P. 2009. Systematicity redux. *Synthese* (170). 251–274. DOI: http://dx.doi.org/10.1007/s11229–009–9582–0

Millikan, Ruth Garrett. 2005. *Language: A biological model.* Oxford: Clarendon Press. DOI: http://dx.doi.org/10.1093/0199284768.001.0001

Moravcsik, Julius M. 1981. Frege and Chomsky on thought and language. *Midwest Studies In Philosophy* 6(1). 105–124. DOI: http://dx.doi.org/10.1111/j.1475–4975.1981.tb00431.x

Moro, Andrea. 2013. *The equilibrium of human syntax: Symmetries in the brain.* New York, NY: Routledge.

Moro, Andrea. 2014. On the similarity between syntax and actions. *Trends in Cognitive Sciences* 18(3). 109–110. DOI: http://dx.doi.org/10.1016/j.tics.2013.11.006

Origgi, Gloria & Dan Sperber. 1998. Evolution, communication, and the proper function of language: A discussion of Millikan in the light of pragmatics and of the psychology of mindreading. In Peter Carruthers and Andrew Chamberlain (eds.), *Evolution and the human mind: Language, modularity and social cognition*, 140–169. Cambridge: Cambridge University Press.

Parker, Anna R. 2006. Evolving the narrow language faculty: Was recursion the pivotal step? In *Proceedings of the 6th International Conference on the Evolution of Language.* 239–246.

Petitto, Laura-Ann. 2005. How the brain begets language: On the neural tissue underlying human language acquisition. In James McGilvray (ed.), *The Cambridge companion to Chomsky*, 84–101. Cambridge: Cambridge University Press.

Petitto, Laura-Ann, Robert Zatorre, Kristine Gauna, E. J. Nikelski, Deanna Dostie & Alan Evans. 2000. Speech-like cerebral activity in profoundly deaf people while processing signed languages: Implications for the neural basis of human language. *Proceedings of the National Academy of Sciences* 97(25), 13961–13966.

Phillips, Colin. 1996. *Order and structure* (MIT Working Papers in Linguistics). Cambridge, MA: Massachusetts Institute of Technology dissertation.

Phillips, Colin & Matthew Wagers. 2007. Relating structure and time in linguistics and psycholinguistics. In Gareth Gaskell (ed.), *Oxford handbook of psycholinguistics*, 739–756. Oxford: Oxford University Press. DOI: http://dx.doi.org/10.1093/oxfordhb/9780198568971.013.0045

Phillips, Colin & Shevaun Lewis. 2013. Derivational order in syntax: Evidence and architectural consequences. *Studies in Linguistics* 6. 11–47.

Piantadosi, Steven T., Harry Tily, & Edward Gibson. 2012. The communicative function of ambiguity in language. *Cognition* 122(3). 280–291. DOI: http://dx.doi.org/10.1016/j.cognition.2011.10.004

Piattelli-Palmarini, Massimo, Juan Uriagereka, & Pello Salaburu. 2009. *Of minds and language: A dialogue with Noam Chomsky in the Basque Country.* Oxford: Oxford University Press.

Pinker, Steven & Paul Bloom. 1990. Natural language and natural selection. *Behavioral and Brain Science* 13. 707–784. DOI: http://dx.doi.org/10.1017/S0140525X00081061

Pinker, Steven & Ray Jackendoff. 2005. The faculty of language: What's special about it? *Cognition* 95. 201–236. DOI: http://dx.doi.org/10.1016/j.cognition.2004.08.004

Povinelli, Daniel J. 2000. *Folk physics for apes.* Oxford: Oxford University Press.

Reddy, Michael. 1979. The conduit metaphor: A case of frame conflict in our language about language. In Andrew Ortony (ed.), *Metaphor and thought,* 284–310. Cambridge: Cambridge University Press.

Reinhart, Tanya. 1983. *Anaphora and semantic interpretation.* Beckenham: Croom Helm.

Reinhart, Tanya. 2006. *Interface strategies: Optimal and costly computations.* Cambridge, MA: MIT Press.

Roeper, Tom & Margaret Speas (eds.). 2014. *Recursion: Complexity in cognition.* Heidelberg: Springer.

Ross, John Robert. 1986. *Infinite syntax!* Norwood, NJ: Ablex.

Ryle, Gilbert. 1968. A puzzling element in the notion of thinking. In Peter F. Strawson (ed.), *Studies in the philosophy of thought and action,* 7–23. Oxford: Oxford University Press.

Sauerland, Uli & Hans-Martin Gärtner. 2007. *Interfaces + recursion = Language? Chomsky's Minimalism and the view from syntax-semantics.* Berlin: Mouton de Gruyter.

Searle, John. 1974. Chomsky's revolution in linguistics. In Gilbert Harman (ed.), *On Noam Chomsky: Critical essays,* 2–33. New York, NY: Anchor.

Semino, Elena. 2006. A corpus-based study of metaphors for speech activity in British English. In Anatol Stefanowitsch & Stefan T. Gries (eds.), *Corpus-based approaches to metaphor and metonymy,* 36–62. Berlin: Mouton de Gruyter.

Shettleworth, Sara J. 2010. *Cognition, evolution, and behaviour.* Oxford: Oxford University Press.

Sigurðsson, Halldór Ármann. 2003. The silence principle. In Lars-Olof Delsing, Cecilia Falk, Gunlög Josefsson, & Halldór Á. Sigurðsson (eds.), *Grammar in focus, Festschrift for Christer Platzack 18 November 2003,* Vol. 2, 325–334. Lund: Wallin & Dalholm.

Sigurðsson, Halldór Ármann. 2004. Meaningful silence, meaningless sounds. *Linguistic Variation Handbook* 4(1). 235–259.

Simon, Herbert A. 1956. Rational choice and the structure of the environment. *Psychological Review* 63(2). 129–138. DOI: http://dx.doi.org/10.1037/h0042769

Slezak, Peter. 2002. Thinking about thinking: Language, thought and introspection. *Language & Communication* 22. 353–373. DOI: http://dx.doi.org/10.1016/S0271–5309(02)00012–5

Smith, Neil & Ianthi-Maria Tsimpli. 1995. *The mind of a savant.* Oxford: Blackwell.

Sperber, Dan & Deirdre Wilson. 1986. *Relevance: Communication and cognition.* Cambridge, MA: Harvard University Press.

Sprouse, Jon & Norbert Hornstein (eds.). 2013. *Experimental syntax and island effects.* Cambridge: Cambridge University Press.

Stewart, Ian. 1998. *Life's other secret: The new mathematics of the living world.* New York, NY: Wiley.

Stromswold, Karin. 1999. The cognitive neuroscience of language acquisition. In Michael Gazzaniga (ed.), *The new cognitive neurosciences,* 909–932. Cambridge, MA: MIT Press.

Thompson, D'Arcy. 1942 [1917]. *On growth and form.* Cambridge: Cambridge University Press.

Tomalin, Marcus. 2007. Reconsidering recursion in syntactic theory. *Lingua* 117. 1784–1800. DOI: http://dx.doi.org/10.1016/j.lingua.2006.11.001

Tomasello, Michael. 2006. Acquiring linguistic constructions. In Deanna Kuhn & Robert Siegler (eds.), *Handbook of child psychology (Sixth Edition)*, 255–298. Hoboken, New Jersey: John Wiley & Sons, Inc.

Turing, Alan. 1952. On the chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London,* Series B, 237. 37–72. DOI: http://dx.doi.org/10.1098/rstb.1952.0012

Vanparys, Johan. 1995. A survey of metalinguistic metaphors. In Louis Goossens, Paul Pauwels, Brygida Rudzka-Ostyn, Anne-Marie Simon-Vandenbergen & Johan Vanparys (eds.), *By word of mouth,* 1–34. Amsterdam: Benjamins. DOI: http://dx.doi.org/10.1075/pbns.33.02van

von Humboldt, Wilhelm. 2000 [1836]. *On language.* Cambridge: Cambridge University Press.

Walsh, Denis M. 2008. Function. In Stathis Psillos & Martin Curd (eds.), *The Routledge Companion to Philosophy of Science*, 349–357. Abingdon, UK: Routledge.

Williams, Joanna Radwanska. 1993. Expression and communication as basic linguistic functions. *Intercultural Communication Studies 3*. 91–100.

Yamada, Jenny Ellen. 1990. *Laura: A case for the modularity of language.* Cambridge, MA: MIT Press.

Zwart, Jan-Wouter. 2011. Recursion in language: A layered-derivation approach. *Biolinguistics* 5(1–2). 43–56.