

The role of language in structure-dependent cognition

Department of Psychology, University of California Los Angeles, Los Angeles, CA USA.

Martin M. Monti*

September 18, 2013

*Martin M. Monti, PhD, UCLA Department of Psychology, Los Angeles CA, 90095, USA. Tel.
+1 (310) 825-8546 Fax. +1 (310) 206-5895 Email: monti@psych.ucla.edu.

In, Mody M. (ed), *Neural Mechanisms of Language*, Springer:New York NY.

1. Introduction

The ability to construct an indefinite number of ideas by combining a finite set of elements in a hierarchically structured sequence is a signal characteristic of human cognition. To illustrate, consider the sentence *The girl who kissed the boy closed the door*. It is immediately clear to any proficient English speaker that the state of affairs described by this sentence is that the girl is doing the closing. This specific interpretation is as effortless as automatic, and if anybody interpreted it any differently it might be sufficient grounds for doubting her proficiency of the English language. Nonetheless, one might wonder why, for example, we do not interpret the noun phrase *the boy* as being the subject of the verb phrase *closed the door*. After all, *the boy* is linearly much more proximal to the verb phrase than is *the girl*. Furthermore, the sentence actually even contains the well-formed fragment [...] *the boy closed the door*, in which, of course, it is the boy doing the closing. Yet, when we consider the full sentence, the relative linear proximity of its component elements does not appear to guide our interpretation. How is it then that we so effortlessly and automatically interpret the sentence above as describing a state of affairs in which a (certain) girl, who just so happens to have given a kiss to a (certain) boy, has closed the door? One explanation, which is perhaps the founding intuition of the modern study of language as a mental phenomenon, is that despite the fact that language is typically manifested as a temporally linear sequence of utterances, in our mind we spontaneously build a rich abstract hierarchical representation of how each discrete element within the sequence relates to every other element. It is the building of these abstract representations that allows us to assign meaning to strings of utterances. Although this ability is most prominently displayed in our use of natural language, it also characterizes several other aspects of human cognition such as logic reasoning, number and music cognition, action sequences and spatial relations, among others. As I will describe below, at least at an intuitive

level, these seemingly distant domains of human cognition all appear to be organized at an abstract level and might therefore share, hidden behind a linear surface structure, the hierarchical and recursive features that are most commonly described by the syntactic trees built by linguists (see Figure 1 for an example).

The objective of this chapter is to address the relationship (if any) between the mental computations that underlie the abstract structures we create when using natural language and those that underlie similar computations in other domains of human cognition. In what follows, I will first briefly trace the theoretical backdrop of this debate, and then present a dominant hypothesis concerning the role of language in human cognition, typically referred to as the supramodal hierarchical parser (SHP) hypothesis. According to this view a specific part of the human brain – traditionally considered to be a center for language processing – might in fact be involved in processing hierarchical structures across domains of human thought. The remainder of the chapter will review a number of functional neuroimaging experiments that have addressed this issue within the domains of logic reasoning and algebraic cognition. Finally the chapter will conclude by bringing together the different streams of evidence and evaluating the SHP hypothesis, as well as the overall debate concerning the role of language in human cognition.

1.1 Framing the debate: Theoretical background

Does language make us special? The extent to which the mechanisms of language contribute to shaping and organizing human cognition has been the focus of a longstanding debate. On the one hand, it is undeniable that language is one of the most characterizing aspects of the human mind. On the other hand, however, it is not clear whether the processes and properties of language are but one manifestation of the properties of our cognitive apparatus, or whether it is the emergence of language itself, in the human brain, that has endowed it with the ability to construct abstract

representations, a computational infrastructure that might then have been relied upon by other domains of human cognition.

Taking a more general perspective, the debate concerning the intertwining of language and thought in the human mind has a long-standing tradition of proponents that fall somewhere in-between two extreme positions. According to one view, as formulated by Wilhelm von Humboldt, “language is the formative organ of thought. [. . .] Thought and language are therefore one and inseparable from each other” (Losonsky, 1999, p. 99). As most frequently described, this view encompasses two complementary hypotheses. The first, often referred to as the *linguistic relativity* hypothesis, is a conjecture concerning the mechanism by which language exerts its influence on thought. As conceived by one of its most prominent proponents, Benjamin Lee Whorf, language shapes thought by providing the concepts around which perception of the world is organized:

We dissect nature along lines laid down by our native languages. The categories and types that we isolate from the world of phenomena we do not find there because they stare every observer in the face; on the contrary, the world is presented in a kaleidoscopic flux of impressions which has to be organized by our minds and this means largely by the linguistic systems in our minds. We cut nature up, organize it into concepts, and ascribe significances as we do, largely because we are parties to an agreement to organize it in this way (Whorf, 1940, p. 213).

Under this hypothesis, as noted by Edward Sapir, meanings are imposed upon us by “the tyrannical hold that linguistic form has on our orientation in the world” (Sapir, 1931, p.578), rather than discovered through experience. The second hypothesis, typically referred to as *linguistic determinism*, builds upon linguistic relativism and the observation of variability across languages, and concerns the effects of the influence of language on thought:

The fact of the matter is that the ‘real world’ is to a large extent unconsciously built upon the language habits of the group. No two languages are ever sufficiently similar to be considered as representing the same social reality. The worlds in which different societies live are distinct worlds, not merely the same world with different labels attached. (Sapir, 1929, p.209)

At the other end of the spectrum sits a view according to which “thought is mediated by language-independent symbolic systems, often called the language(s) of thought. [. . . W]hen humans learn a language, they learn to express in it concepts already present in their prelinguistic system(s)” (Gelman and Gallistel, 2004, p. 441). Under this hypothesis words are just symbols for mental experiences gained through experience with the world:

Spoken words are the symbols of mental experience and written words are the symbols of spoken words. Just as all men have not the same writing, so all men have not the same speech sounds, but the mental experiences, which these directly symbolize, are the same for all, as also are those things of which our experiences are the images (Aristotle, *De Interpretatione* Chapter I, 16a).

This view of language is also evident in John Locke’s work where words are conceived as signs of internal conceptions, and “stand as marks for the ideas within [the] mind, whereby they may be made known to others, and the thought of men’s minds be conveyed from one to another” (Locke, 1824, Book III, Chapter I). At its core, this view proposes that thought is prior to language (Pinker, 1984). As explained by Li and Gleitman:

Language has means for making reference to the objects, relations, properties, and events that populate our everyday world. It is possible to suppose that these linguistic categories and structures are more-or-less straightforward mappings from

a preexisting conceptual space, programmed into our biological nature. [. . .]

Humans invent words that label their concepts (Li and Gleitman, 2002, p.265).

Importantly, as explained by Gelman and Gallistel (2004), properties that are typically regarded as essential to language, such as compositionality, are in fact already present in our preexisting (i.e., prelinguistic) systems. The properties of language might thus just be one incarnation of the properties of thought (which might resemble those discussed by Fodor, 1975), no differently than the properties of other structure-dependent aspects of human cognition. In other words, we have the language we have in order to express the thoughts we have (Pinker and Jackendoff, 2005).

2 The relationship between language and thought

2.1 The algebraic mind: Structure dependence in human cognition

As mentioned in the introduction, one of the most fundamental hypothesis concerning how our minds process language is the idea that as we perceive serially ordered sequences of (linguistic) utterances, we spontaneously build non-linear (i.e., hierarchical) abstract representations which underlie our ability to assign meaning to a string of linguistic utterances. The psychological reality of these hierarchical constructs is well demonstrated by the first two lines of Lewis Carroll's famous poem Jabberwockie: *Twas brillig, and the slithy toves did gyre and gimble in the wabe*. Although semantically non-interpretable, the sentence feels structurally well-formed, while its reverse, for instance, despite featuring all the same words, does not: *Wabe the in gimble and gyre did toves slithy the and, brillig twas*. The relationship tying abstract structures and interpretation of linguistic statements is illustrated by Groucho Marx's famous statement *I once shot an elephant in my pajamas*. The hypothesis being, that the two possible interpretations of this sentence reflect two different abstract representations, each binding the elements of the sentence in a different way.

If, for instance, our mental structure directly connects the verb phrase *I shot* to the propositional phrase *in my pajamas*, that would represent an understanding that the shooter was wearing his night apparel as the events took place. On the other hand, if our mental representation of the sentence puts in direct relation the noun phrase *an elephant* with the propositional phrase *in my pajamas*, that would lead us to the (more puzzling) image that the shootee somehow managed to sneak into Groucho's nightwear before the fatal event.

Central to this view are a set of properties that lie at the heart of what we consider language (cf., Boeckx, 2010, p. 32). The first property stresses the *abstract* nature of linguistic representations and relates to the so-called "type-token" distinction, which is to say the ability to recognize the difference between classes of elements, such as verbs and nouns, and the specific elements within a class, such as the verb *to go* or the noun *boy*. This property is crucial to the idea of an "algebraic mind" because it allows defining combinatorial rule mappings that apply over classes of items (i.e., types) rather than individual tokens, rendering the rules abstract. The second property, *compositionality* (or *structure dependence*) refers to the fact that the meaning of a complex expression is derived from the meanings of its constituents as well as the specific relationships by which they are bound together. This implies that there is a level of organization of the elements within a structure that confers meaning and exists independently of the semantics of the individual symbols, the psychological reality of which is well captured by Chomsky's famous example *Sleepless green ideas sleep furiously*. It is this property of language that allows us to distinguish "between the boring news *dog bit man*, and the much more newsworthy *man bit dog*" (Boeckx, 2010, p. 32). Third, *quantification* (or *bracketing*) refers to the ability to properly assign the right set of brackets around certain groups of elements within a statement and therefore understand the hierarchy by which elements within a structure bind. Finally, *recursion* refers to the property by

which a rule can be applied to its own output or to the output of other rules, repeatedly, without any limit but for those imposed by the processing capacity of the individual (Corballis, 1992).

While these properties are most prominently displayed in our use of natural language, they are also central to other domains of human cognition (cf., Tettamanti and Weniger, 2006; Fadiga et al., 2009; Udden and Bahlmann, 2012), among others. As explained in Varley et al. (2005), for example, there exists an intuitive parallel between the interpretation of the two sentences “John kissed Jill” and “Jill kissed John” and the interpretation of the two algebraic statements “ $(5 - 3)$ ” and “ $(3 - 5)$.” Despite the fact that, within each pair, the same elements are used, their different combination leads to different interpretation (e.g., who is kissing who, and whether the result is positive or negative, respectively). A similar analogy can be drawn between the recursive application of rules in language, as in the sentence “The man saw the boy who kicked the ball,” and the recursive application of rules in algebra, as in “ $2 \times (5 - 3)$.” In both examples structures can be embedded within structures of the same kind, and understanding of the hierarchical ordering by which elements within the structure bind is key to correct interpretation. Finally, both domains are infinitely generative in that it is possible for a finite set of elements to combine into a potentially unbound set of well-formed expressions. It is indeed always possible to generate a new sentence by prefixing, for example, “Mary thinks [that ...]” to any well-formed sentence just as it is always possible to generate a novel and meaningful algebraic statement by adding “ $2 + [...]$ ” to any well-formed algebraic expression. Whether this analogy is superficial or substantial is the topic of wide discussion, and will be the central concern of this chapter.

2.2 The “supramodal hierarchical parser” hypothesis

Given the *prima facie* similarities that can be drawn between language and other aspects of human thought, one might wonder whether a common set of computations lie at the heart of all structure-

dependent cognition. This very intuition was already formulated in the work of Thomas Hobbes according to whom thinking amounted to performing arithmetic-like operations on internal structures (i.e., mental representations). While he did recognize that some forms of thought can exist outside of language, he believed that “mental processes where generality and orderly concatenation of thought are involved require the use of internal linguistic means” (Boeckx, 2010). Linguistic computations might therefore be seen as central to our modes of thought, and as providing the very fabric of structure-dependent cognition. Following this view, it has been recently proposed that the human brain encapsulates, in the left inferior frontal gyrus (LIFG; and specifically its *pars opercularis* and *pars triangularis*), a supramodal hierarchical parser (SHP; Tettamanti and Weniger, 2006; Fadiga et al., 2009; Udden and Bahlmann, 2012). The core of this proposal, which is depicted in Figure 1, is the hypothesis that the LIFG is involved in the computations necessary for processing and representing abstract, hierarchical, “syntax-like,” structured sequences across domains of human cognition. Following the explanation of Tettamanti and Weniger (2006), given three non-identical elements ‘X,Y,Z’, their arrangement in a hierarchical fashion, as pictured in Figure 1a, allows establishing univoque relations between the elements, thereby pinpointing one out of the 3! possible arrangement they can take, and thus assigning a specific meaning to the string. As shown in Figure 1b, the same hierarchical structure can be employed to describe the natural language statement *The girl runs*. Specifically, the first two elements, the determiner *the* and the noun *girl*, are first bound together (into a noun phrase) and then, as a unit, bound to the verb *runs*, thereby imparting a specific interpretation to the sentence. Importantly, any other ordering of the three elements, obtained by permuting their position, would either yield an ill-formed sentence, or a well-formed sentence with a different intension. This same hierarchical structure can also be employed to describe, as shown in Figure

1c, the algebraic statement “ $(X + Y) \times Z$ ” where the first elements are bound together by an addition operator, and the result of that operation is then bound to the third element by a multiplication operator. As for the linguistic statement, if the hierarchical structure binding the elements were different, the statement would have a different interpretation (i.e., result). A similar case can be made for the logic statement “ $(X \wedge Y) \rightarrow Z$ ” (which translates to the natural language statement “If X and Y then Z”), as depicted in Figure 1d, as well as several other domains of human thought including music cognition, spatial relations and action sequences (e.g., Tettamanti and Weniger, 2006; Fadiga et al., 2009).

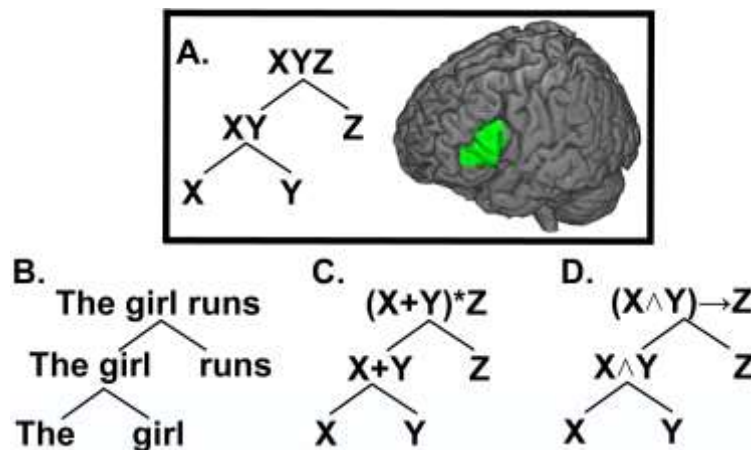


Figure 1: The LIFG supramodal hierarchical parser hypothesis (partially adapted from Tettamanti and Weniger, 2006, with permission from Elsevier).

Although the specific function of the LIFG is still a matter of debate (cf., Hagoort, 2005; Grodzinsky and Santi, 2008), several lines of evidence suggest that it is crucial to the processing of natural language (cf., Bookheimer, 2002). Particularly relevant to our discussion, this region of the brain is consistently found to be active for complex syntactic statements, such as object-embedded sentences (e.g., “The man that the girl is talking to is happy”), as compared to simpler subject-embedded sentences (e.g., “The girl that is talking to the man is happy”; cf., Just et al., 1996). Indeed, understanding a sentence’s argument hierarchy construction (i.e., who did what to

whom; Bornkessel et al., 2005) as well as whether thematic roles remain unchanged through syntactic transformations (e.g., active to passive; Monti et al., 2009, 2012) consistently recruits the LIFG (although not exclusively). Furthermore, the LIFG appears to be recruited for processing the long-distance dependencies and hierarchical structures (Friederici, 2004) that are at the heart of natural language (Chomsky, 1957). Although it is clear that the LIFG plays a crucial role in natural language (and its “structural” aspect in particular), there are several contrasting hypotheses concerning which computation(s) are specifically embedded within the neural circuitry of this region. While according to some this region might carry out computations specific to establishing long-range dependencies (Friederici, 2005) and syntactic movement (Grodzinsky and Santi, 2008), it has also been suggested that it might be involved in unifying lexical information (Hagoort, 2005) or, more generally, selecting among competing representations (Novick et al., 2010). I should stress, however, that the question of which role (if any) linguistic computations play in other domains of cognition is neutral with respect to this debate. For, inasmuch as it is agreed that whichever the specific computation carried out within the LIFG this is core to our processing of language, all that matters is evaluating whether this neural circuitry is also involved in the computations of other aspects of human cognition.

Of course, there is certain a definitional component to establishing whether a given domain of human thought is “linguistic.” For, how language is defined in the human brain might significantly affect which aspects of human cognition might be considered as resting on linguistic computations. As discussed in other chapters of this book, language encompasses, broadly conceived, a rich and wide set of cognitive processes. From a neural point of view, linguistic stimuli can thus elicit activations in several areas outside the LIFG as well as other “traditional” perisylvian language regions. Several fMRI and clinical studies, for example, have highlighted the role of motor cortices

in processing action-related words (Hauk et al., 2004), medial temporal lobe regions in semantic processing (Hoenig and Scheef, 2005), right hemispheric fronto-temporal areas in processing prosodic cues (Wildgruber et al., 2006), temporo-parietal and subcortical reward-related regions in processing humor (Bekinschtein et al., 2011), among many others. In fact, if the pragmatics of message selection (Grice, 1989) is counted as a core linguistic capacity, then virtually any neural area implicated in cognition could be considered a language structure (Monti et al., 2009). In the following discussion, however, I will focus on the set of processes underlying the construction of rule-governed relationships that allow generating the unbound range of possible expressions within a language (Chomsky, 1983).

3 The role of language in structure dependent cognition

3.1 Disentangling “language” and “thought” with fMRI

Before discussing the neuroimaging evidence concerning the role of language in structure-dependent cognition, it is worth considering some of the crucial features of functional magnetic resonance imaging (fMRI) as it is employed today to uncover the neural basis of human cognition. In particular, it is important to note that the fMRI signal (typically referred to as the blood oxygenation level dependent signal, BOLD; Ogawa et al., 1990) is difficult to interpret per se. Knowing that a mental activity elicits a BOLD signal of, say, 850 units in a given part of the brain is not very meaningful. More meaningful is the comparison of the BOLD signal between two different tasks. Hence, most task-based fMRI studies are based on the so-called subtraction principle whereby the metabolic response to a task of interest is compared to the metabolic response to a control (or “baseline”) task. If this latter task contains all the same cognitive processes as the task of interest, except for the one mental process of interest (often referred to as

the “pure insertion” hypothesis), subtracting the metabolic response to the baseline task from that of the main task should isolate the metabolic response specific to the cognitive process of interest. Evaluating an experiment’s baseline is thus critical to correctly interpreting a functional neuroimaging result. Imagine, for example, being interested in the neural basis of single word repetition. As a main (or target) task, one might present visually a set of words, one at a time, and ask participants to repeat them. As a baseline task, one might decide to employ periods of rest during which the participant is not performing any (overt) task. Subtracting the metabolic activity observed during the latter periods from that observed during the main task will likely isolate a wide set of neural foci including both the neural substrate of single words repetition and several other ancillary processes tied to processing visual stimuli as well as words. This baseline task is highly sensitive, because it captures all the neural structures that are elicited by the target task, but not very specific, because it captures many processes that are not specifically related to cognitive process of interest. In other words, of all the regions of the brain that might be uncovered by this subtraction, it is difficult to tell which are directly involved in word repetition and which are tied to the many other processes that factor into the target task. Given the same main task, adopting a baseline task in which participants are presented with strings of letters that do not form a meaningful word might be more effective in filtering out basic visual processes from the activations elicited by the main task. Nonetheless, the subtraction is still likely to uncover phonetic and semantic processes tied to reading meaningful words as well as the process of interest (i.e., single-word repetition). One might thus employ a baseline task in which subjects are presented with single words, as in the target task, and asked to read them. This baseline should be very effective in filtering out from the main task all the ancillary activations related to processing visual stimuli as well as the phonetics and semantics of reading single words. Ideally, the comparison

would thus pinpoint only regions involved in word repetition. However, what if participants spontaneously and automatically repeat the words subvocally as they read them? In this case, the baseline could elicit the same neural substrate as the target task filtering out, partially or entirely, the metabolic response related to the cognitive process of interest.

While the above examples might appear extreme and unlikely to exist in actual practice, as I will discuss in the next section, similar circumstances continuously arise in cognitive neuroscience research often resulting in substantial divergence of results across studies. Furthermore, this issue is particularly severe within the domain of higher cognitive functions where eliciting the process(es) of interest (e.g., reasoning) often requires eliciting several other ancillary processes.

3.2 Language and logic reasoning

Background. Deductive reasoning is the attempt to draw secure conclusions from prior beliefs, observations and suppositions (Monti and Osherson, 2012). This aspect of human cognition has been the focus of vigorous investigation within the fields of philosophy and psychology (e.g., Beall and van Fraassen, 2003; Osherson, 1975). It is typically regarded as a central feature of human intelligence (Rips, 1994), although some forms of deduction (e.g., transitive inference) have also been reported other species (e.g., Grosenick et al., 2007). With respect to the role of language in deductive reasoning, different views have been expressed.¹ On some accounts, language plays a central role in the deductive inference making process (Polk and Newell, 1995). According to others, reasoning is fundamentally based on processes other than the syntactic interpretation of sentences (Cheng & Holyoak, 1985; Osherson, 1975).² In considering the neurobiology of

¹ As described below, given that deductive reasoning is most often elicited by the means of verbal stimuli, it is trivial that linguistic processes are needed to apprehend the stimuli. What is under discussion is whether linguistic processes play a role in deductive reasoning beyond initial encoding of verbal materials.

² It might be worth clarifying that so-called Mental Rules theories of deduction (e.g., Osherson 1975), despite being sometimes portrayed as language-based (see Goel et al., 1998, 2000), might in fact be better understood as

deductive competence (as well as algebraic cognition – see next section), it is important to distinguish two potential roles for linguistic processing. At a minimum, the (verbal) stimuli typically employed to elicit deductive reasoning must be apprehended before deduction can take place. At a neural level, language and reading areas (Price, 2000; Bookheimer, 2002) would thus be expected to be involved in this stage. What is more controversial (and under discussion in this chapter) is whether language plays a part in the subsequent inferential process itself.

Neuroimaging studies of deductive reasoning. Overall, neuroimaging studies of reasoning have defended a variety of positions including the thesis that all deductive reasoning is left-hemispheric and language based (e.g., Goel et al., 1997, 1998; Reverberi et al., 2007), along with the contrary suggestion that none of it is (e.g., Goel and Dolan, 2001; Parsons and Osherson, 2001; Knauff et al., 2003). Different yet, other studies have been interpreted as supporting a “dual process” view of deduction according to which, depending on whether the reasoner has prior beliefs over, or familiarity with, the contents of the argument she is reasoning about, language resources may or may not be recruited (e.g., Goel and Dolan, 2003). This dramatic variance of results highlights the complexities of disentangling “thought” from linguistic processes using correlational methods such as fMRI, and is, to a significant extent, tied to the subtraction problem discussed in the previous paragraph. Knauff et al. (2003), for example, recorded the metabolic response of healthy volunteers while they judged whether each of a number of arguments featuring two premises and one conclusion, were deductively valid (for example: “The dog is cleaner than the cat.” “The ape is dirtier than the cat.” Does it follow: “The dog is cleaner than the ape?”). Comparison of the metabolic response during the target task to that observed during rest periods uncovered activations in some left hemispheric language regions, among others. As discussed above, due to the non-

describing deductive inference as a “syntax-like,” algebraic, computation, rather than a linguistic one (cf., Monti et al., 2007).

specific nature of the baseline task, it difficult to assess whether the involvement of posterior perisylvian language regions reflects the engagement of linguistic resources during the deductive inference process or during the initial processing of verbal stimuli. In a set of pioneering neuroimaging studies, Goel et al. (1997) and Goel et al. (1998) employed a baseline task in which subjects were asked to determine how many of the three sentences in a given argument had people as their subject. While this baseline does, to some extent, filter out ancillary processes related to encoding visual and verbal stimuli, the minimal amount of linguistic processing required is likely to be less than that required to read the same sentences in view of inferential reasoning. Thus, again, it is difficult to tell whether the observed activations in linguistic regions reflect simple reading or the involvement of linguistic mechanisms in deductive reasoning. Other experimental design factors, such as the timing of the task of interest (as well as the baseline one), can also substantially effect the interpretation of neuroimaging findings. Goel et al. (2000), for example, employ a baseline task that is isomorphic to the target task but included a conclusion unrelated to the premises. To illustrate, consider the two deductive trials presented in Table 1 (each consisting of two premises and one conclusion).

Table 1: Sample deductive stimuli from Goel et al. (2000). (‘P_{1,2}’ indicate Premise 1 and Premise 2; ‘C’ indicates the Conclusion of the argument.)

| | Argument #1 | Argument #2 |
|----------------------|-------------------------------------|-----------------------------------|
| P₁ | All poodles are pets. | All poodles are pets. |
| P₂ | All pets have names. | All pets have names. |
| C | [Therefore] All poodles have names. | [Therefore] No napkins are white. |

The idea of comparing the metabolic activity in response to the two arguments is very clever because the status of a trial (with respect to being a target or baseline trial) depends on whether the conclusion is related to the premises, as in Argument #1, allowing deduction to take place, or not, as in Argument #2. The participant, however, is unaware of such distinction and performs all trials under the same set of instructions, namely to assess whether the conclusion follows from the premises. This experimental set-up, however, has two very problematic and unwanted consequences. First, the presence of extraneous materials early in the conclusion statement (i.e., napkins) is sufficient for the participant to recognize, with very little reading, the invalidity of the trial (and therefore that it is a baseline trial). Thus, as for the two previous experiments discussed above, the baseline task might not sufficiently filter out linguistic processes tied to sentence reading during deductive trials. Second, the slow sequential presentation of each statement, serially added to the display at 3 seconds intervals, allows deductive processes to start taking place as soon as the second premise is displayed and, crucially, before the conclusion is presented. Thus, until the conclusion is presented, target and baseline trials might elicit comparable amounts of deductive reasoning. As a result, this baseline task may subtract essential elements of deductive reasoning from deduction trials, while not adequately filtering reading activations (cf., Monti and Osherson, 2012). It is not surprising, then, that the authors report engagement of linguistic regions in the LIFG for the target task (compared to the baseline).

A language-independent network for deductive reasoning As the above discussion illustrates, characterizing the neural substrate of deductive reasoning presents several complexities which have prevented the field from reaching a consensus on what role (if any) language plays in this aspect of human thought. In a recent series of experiments, however, the case has been made for deductive reasoning recruiting a language-independent distributed network of brain regions (see

Monti and Osherson, 2012, for a review). In an attempt to avoid many of the experimental pitfalls described above, Monti et al. (2007) adopted a “cognitive load” design in which participants were instructed to assess whether each of a number of logic arguments were deductively valid. Half the arguments were simple to assess (e.g., “If the block is either round or large then it is not blue.” “The block is round.” “The block is not blue.”), whereas the other half were more complex (e.g., “If the block is either red or square then it is not large.” “The block is large.” “The block is not red.”). From a cognitive perspective, complex and simple deductions can be expected to recruit the same kind of mental operations, but in different number, repetition or intensity. If linguistic structures are involved in the inferential process, complex deductions should recruit them significantly more than simple ones. On the other hand, if the role of language is confined to initial encoding of stimuli, simple inferences can be expected to require similar levels of reading compared to their complex counterparts. This expectation is reinforced by the fact that the statements included in the simple and complex arguments are matched for linguistic complexity (as the two sample inferences above demonstrate) and can be expected to prompt for similar amounts of initial language processing. Thus, should any language-related activation be apparent, it cannot be considered to reflect differences in initial reading or comprehension. Subtraction of the metabolic response observed during simple trials from that observed during complex trials was thus expected to adequately filter-out the initial reading of verbally presented materials while revealing areas of the brain correlating with increased deductive reasoning. The authors reported two main findings. First, the complex minus simple subtraction did not reveal any activation in the LIFG supramodal hierarchical parser (as well as in posterior temporal regions), indicating that although the region was active at the beginning of each trial, in correspondence with initial reading, it was equally recruited by simple and complex inferences. Second, the subtraction uncovered a

distributed network of areas spanning regions that are believed to perform cognitive operations that sit at the “core” of deductive reasoning (in frontopolar and fronto-medial cortices) as well as several other “cognitive support” frontal and parietal regions known to be related to working memory and spatial attention process. Using a different approach, Monti et al. (2009) compared logic inferences based on sentential connectives (i.e., “if . . . then . . .”, “and”, “or”, “not”) to inferences based on the syntax and semantics of ditransitive verbs (e.g., “give”, “say”, “take”). To illustrate, a valid linguistic inference might include the premise “X gave Y to Z.” and the conclusion “Z was given Y by X.” Similarly, a valid logic inference might feature the premise “If both X and Y then not Z.” and the conclusion “If Z then either not X or not Y.” In this design, logic and linguistic arguments were each compared to a matched baseline in which the very same sentences evaluated for inferential validity were also evaluated for grammatical well-formedness. Occasionally, in order to ensure that participants fully encoded the sentences during the baseline trials, they were presented with “catch trials” in which statements contained grammatical anomalies (e.g., “X gave to Y to Z.”). If logical inference is based on mechanisms of natural language that go beyond mere reading for meaning, then the comparison of each type of inference to its matched baseline should uncover common activations in regions known to underlie linguistic processing, and particularly within the LIFG SHP. Subtraction of the metabolic response during grammatical judgments from linguistic inferences uncovered extensive activation within the LIFG, confirming its role in evaluating semantic equivalence of distinct sentences (Dapretto and Bookheimer, 1999), morphological processing (Sahin et al., 2006), detecting semantic roles (Bornkessel et al., 2005), transforming sentence syntax (Ben-Shachar et al., 2003). Conversely, when the same comparison was performed on logic arguments, no activity was detected in the LIFG (or posterior temporal perisylvian language regions), indicating that logic reasoning does

not recruit language resources beyond what is necessary for simple reading. When the logic and linguistic inference tasks were compared directly (over the full brain, as well as on a region-by-region basis), only the latter were shown to engage the LIFG (as well as posterior temporal perisylvian regions). Furthermore, the activations detected during the logic inference trials replicated exactly those seen, with a different task, in Monti et al. (2007) (see Monti and Osherson, 2012, Figure 1).

Overall, several studies have failed to uncover activation within the LIFG during deductive reasoning tasks (e.g., Noveck et al., 2004; Canessa et al., 2005; Fangmeier et al., 2006; Prado and Noveck, 2007; Rodriguez-Moreno and Hirsch, 2009; Prado et al., 2010b). These findings can be interpreted as implying that the role of language is confined to initial encoding of verbal statements into mental representations suitable for the inferential calculus. The representations themselves, as well as the deductive operations, however, are not linguistic in nature.

Neuropsychological evidence. Despite the above data, the view just formulated is still very debated inasmuch as a number of fMRI studies have reached the opposite conclusion (e.g., Reverberi et al., 2007, 2010; Prado et al., 2010a). As discussed in Monti and Osherson (2012), several factors relating to experimental design can explain this fracture within the literature. To address this point, however, it might be worth considering evidence from the neuropsychological literature, which has the distinct advantage, over neuroimaging methods, to uncover causal relationships between cognitive processes and neural circuitry. In particular, Reverberi et al. (2009) assessed inferential abilities in patients with left lateral frontal damage, right lateral frontal damage or medial frontal damage. Interestingly, patients with damage in right frontal cortex exhibited no apparent difficulty in assessing the validity and judging the difficulty of deductive inferences. Patients with damage in left frontal cortex, instead, appeared impaired in assessing

deductive problems, but only inasmuch as their working memory was affected (i.e., patients with intact working memory, as tested with standard neuropsychological measures, were able to correctly perform the inferential task). Finally, patients with medial frontal damage were unable to solve deductive inferences (a finding recently replicated by Kosciak and Tranel, 2012). Two aspects of these results are crucial. First, no patient had damage to the LIFG, nor appeared to have language deficits, a fact difficult to reconcile with the claim that language processes are at the heart of this domain of cognition. Second, the cortical damage that impaired deductive reasoning falls in the same prefrontal areas that have been previously characterized as “core” to deductive inference (Monti et al., 2007, 2009; Rodriguez-Moreno and Hirsch, 2009). Overall, then, although the question is still very debated, neuropsychological evidence suggests that the neural mechanisms within the LIFG are not sufficient for supporting this kind of structure dependent cognition, while the fMRI evidence, together with the patient data, suggest that medial and polar frontal cortex might be necessary and sufficient (perhaps among other regions of the brain) for processing the hierarchical dependencies imposed by logic structure.

3.3 Language and arithmetic cognition

Background. The relationship between natural language and mental arithmetic has also long been debated (e.g., Spelke and Tsivkin, 2001; Gelman and Gallistel, 2004). In particular, as discussed above (and more extensively in Varley et al., 2005), there is a certain analogy between the operations of natural language and those of mental arithmetic. Indeed, it has been argued that both language and number rely on a recursive computation that exploits the same neural mechanism operating over linguistic structures (Hauser et al., 2002). Recursion might thus have evolved over time from a process that was highly domain specific (to natural language) to a domain general process that gave the human mind the unique ability to use recursion to solve non-linguistic

problems such as numerical manipulation. A similar view is also embraced by Spelke and Tsivkin (2001) who stated that natural language is the “most striking combinatorial system” of the human mind and that formal mathematics might be one of this system’s “richest and most dramatic outcomes” (p. 84). The view that arithmetic reasoning might have co-opted the recursive machinery of language is also explicit in Chomsky (1998) where he argues that the human faculty for arithmetical reasoning can be thought of as being abstracted from language and that it operates by “preserving the mechanisms of discrete infinity and eliminating the other special features of language” (p. 169). Similarly, Fitch et al. (2005) state that the only clear demonstrations that recursion operates in human cognitive domains come from mathematical formulas and computer programming, which clearly employ the same reasoning processes that language does. Overall, as for other domains of human cognition, the debate is whether “the generative power of grammar might provide a general cognitive template and a specific constitutive mechanism for ‘syntactic’ mathematical operations involving recursiveness and structure dependency” (Varley et al., 2005, p. 3519).

Neuroimaging studies of number cognition. In a landmark study by Dehaene et al. (1999), the relationship tying language and mathematical knowledge was addressed with a joint behavioral and neuroimaging approach. In the behavioral study, bilingual speakers learned and trained arithmetic facts, including exact and approximate calculations, in one of two languages. When tested on trained and novel exact calculations, participants exhibited a language-switching penalty, which is to say, when the language in which calculations were trained mismatched the language in which they were later tested, participants exhibited slower reaction times. Under the same circumstances, however, approximate calculations did not exhibit a comparable language switching cost, suggesting that only exact arithmetic knowledge hinges on linguistic

knowledge/mechanisms. When the two tasks were compared using fMRI, approximate calculations recruited mainly parietal regions, while exact calculations recruited regions associated with some aspects of linguistic processing, in the angular gyrus as well as the inferior frontal regions (although in an area that appears more frontal than the LIFG SHP). In a subsequent study, Stanescu-Cosson et al. (2000) compared neural activation for exact and approximate calculation using both small and large numbers. Approximate calculations elicited activity in the bilateral intraparietal sulci (among several others), confirming the view that linguistic resources are not recruited by this aspect of number cognition. Exact calculation, on the other hand, mainly recruited (among others) the angular gyrus as well as the anterior section of the inferior frontal gyrus. However, as for the results presented in Dehaene et al. (1999), the region of the inferior frontal gyrus that was found active is clearly anterior to that observed during syntactic processing of natural language and may thus be related to control of verbal retrieval processes implemented in more posterior cortico-subcortical verbal networks including (as explained in Stanescu-Cosson et al., 2000, p. 2252). The authors do point out that the LIFG proper was found to be specifically recruited by exact calculations with large numbers, however they speculate that its involvement is mainly a function of the increased effortfulness associated with lesser known facts involved in operations with large numbers.

Overall, the results presented above, together with several others, have coalesced in a coherent view according to which language may well play a role in exact calculation, but in connection with the verbal coding of rote exact arithmetic facts (Dehaene et al., 2003, 2004).

A language independent network for processing algebraic structure. Although most studies in this domain are primarily concerned with the neural basis of the representation of quantity, numbers, and calculation, two concurrent studies, in 2012, have begun investigating the role of

language in processing and manipulating the syntax-like structures of algebra (Maruyama et al., 2012; Monti et al., 2012, respectively). Maruyama et al. (2012), for example, looked at the cortical representation of simple algebraic expressions such as “ $1 + (4 (2 + 3))$ ” to assess whether the neural structures responsible for parsing these nested structures are indeed shared with the neural mechanisms responsible for computing the syntax of linguistic expressions. In this experiment, participants were shown strings of algebraic symbols variously arranged so to form expressions with high levels of nesting (e.g., “ $4 + (1 (3 + 2))$ ”) versus strings with little or no nesting (e.g., “ $(3 + 2) + 4(1)$ ”) or algebraically meaningless strings (e.g., “ $4 + 3)(+2(1)$ ”). Importantly, participants were not asked to resolve the equations, but rather just to encode the expressions sufficiently for a probe matching task (i.e., to decide whether the string matches a probe presented at a short delay). Contrary to the LIFG supramodal parser hypothesis, all the regions that were increasingly recruited by greater nesting fell outside the traditional left perisylvian language regions, and mainly included occipital, temporal and inferior parietal regions, indicating that processing syntactically complex algebraic expressions does not rely on the LIFG supramodal parser or other traditional language regions. In that same year, Monti et al. (2012) addressed the question of whether manipulating the syntax-like structure of algebraic expressions relies on the same neural structures required to manipulate the syntax of natural language statements. In an design similar to that used in Monti et al. (2009), discussed above, participants were presented with pairs of natural language statements (e.g., “X gave Y to Z.” and “Z was given Y by X.”) and algebraic statements (e.g., “X minus Y is greater than Z.” and “Z plus Y is smaller than X.”). Participants were instructed to evaluate whether the statements within each pair were equivalent. Although the two tasks are psychologically very similar, judging equivalence in the former kind of pairs depends on whether the principal verb assigns the same semantic roles (i.e., who did what to whom) to X, Y, and Z across a syntactic

transformation. Conversely, judging the equivalence in the algebraic pairs depends on the properties of elementary algebraic operations (i.e., addition, subtraction) and relations (i.e., equality, inequality). Therefore, if left IFG truly acts as a supramodal parser of hierarchical structure, this region should be equally involved in processing and manipulating the hierarchical dependencies of linguistic and algebraic expressions. As a baseline, participants were shown the same statements but asked to assess whether they were grammatically correct (as in Monti et al., 2009). When the linguistic trials were analyzed, subtraction of the grammar judgment task from the equivalence task revealed extensive activation in the LIFG as well as other perisylvian linguistic areas, as expected. However, when the same comparison was performed on the algebraic pairs, no activation was detected in the supramodal parser (or any other perisylvian language region), indicating that beyond initial reading and comprehension of stimuli, the neural substrate of language does not intervene in algebraic reasoning (consistent with the findings of Monti et al., 2009; Maruyama et al., 2012). Conversely, extensive activation was detected in the infra-parietal sulci, consistent with the number cognition literature (cf., Dehaene et al., 2003).

Neuropsychological evidence. In the domain of mental arithmetic there is a relatively rich neuropsychological and developmental disorder evidence that supports the neuroimaging findings. Indeed, some patients with semantic dementia and global aphasia have been shown to retain mathematical competence (e.g., Cappelletti et al., 2001; Delazer et al., 1999, respectively), while individuals with, for example, developmental dyscalculia and William's Syndrome, have been shown to have severe impairment in the number domain while retaining normal language skills (Butterworth, 2005; Bellugi et al., 1993, respectively). One case in particular demonstrated that the ability to process the structured hierarchy of algebraic expression can be retained in patients suffering from agrammatic aphasia, which is to say patients who appear to be unable to process

the structured hierarchy of natural language (Varley et al., 2005). Indeed, when three patients with extensive left hemispheric damage were presented with the reversible sentences such as “The lion killed the man.” they were unable to match it to the appropriate figure (when having to choose between the figure of a man killing a lion and that of a lion killing a man). However, when presented with reversible algebraic expressions, such as “ $(3 - 5)$ ”, patients had no trouble judging whether the result was positive or negative. This dissociation indicates that while the patients appeared to have lost structure sensitivity in the domain of language, they did retain it in the domain of algebra. Similarly, patients also appeared to have lost the ability to perform bracketing in language, as assessed by their inability to judge whether a sentence is grammatical or not. Nonetheless, they did retain the ability to correctly resolve algebraic expressions observing the hierarchy expressed by parenthetical structures. Overall, this brief overview of some of the cardinal results in the neuropsychological literature suggest, in accord with the neuroimaging evidence, that the structured hierarchy of algebra is not processed, in the human brain, by the neural mechanisms of natural language.

4 Discussion

Overall, at least with respect to the domains of deductive reasoning and mental algebra, the neuroimaging evidence fails to support the “supramodal hierarchical parser” hypothesis, and thus that the view that the LIFG encapsulates neural circuitry tuned to detect and represent complex hierarchical dependencies regardless of the specific domain of cognition (Tettamanti and Weniger, 2006; Fadiga et al., 2009). Furthermore, the above evidence also suggests that other left hemispheric perisylvian regions, in posterior temporal cortices, typically considered to be at the heart of language processing, also do not contribute to the processing and manipulation of deductive and algebraic hierarchical structures. The neuropsychologic and developmental disorder

literature confirms the view that language mechanisms are not sufficient for deductive reasoning and number cognition. Indeed, as discussed above, the patients with fronto-medial damage described by Reverberi et al. (2009) and Kosciak and Tranel (2012), as well the dyscalculic patients discussed in Butterworth (2005), are unable to process the structured hierarchies of deductive inference and number cognition despite showing normal language skills. Furthermore, the patients described by Varley et al. (2005), who retain the ability to understand the structured-hierarchy of algebra while being at chance in language comprehension, suggest that language is also not necessary to access the representations and computations of algebra. (To date, there is no parallel finding in the domain of deductive inference, although Varley and Siegal, 2000, report the case of an agrammatic aphasic patient who, despite profound language impairment, could successfully perform causal reasoning.) Taken together, the presented data suggest that the involvement of language in structure-dependent cognition might be limited to what Polk and Newell (1995) termed a “transduction” role. That is, the mechanisms of language might be involved in decoding verbally presented information into representations suitable for algebraic and deductive computations (and encoding their output into language, if needed). The representations and “syntax-like” operations themselves, however, are not in linguistic format.

While the above data establishes that structure-dependent cognition is not parasitic on language *in the mature cognitive architecture*, it is still possible that the generative power of language plays an enabling role through ontogeny (e.g., Bloom, 1994; Spelke, 2003), or has played such a role through phylogenetic history (e.g., Corballis, 1992). The case for homology, however, faces the complication of having to explain how the neural mechanisms of language extended into other domains, re-implemented their circuitry in distant brain regions (fronto-polar and fronto-medial cortices for deductive inference, Monti et al., 2009; Rodriguez-Moreno and Hirsch, 2009; and the

infra-parietal sulci for algebra, Monti et al., 2012; Maruyama et al., 2012), to then entirely disconnect from them.

In conclusion, while there is a certain beauty and efficiency in the hypothesis that language, the most characterizing and striking aspect of the human mind, provided the computations enabling generative cognition across domains of cognition, the available data does not support this view. In the words of Albert Einstein, “[w]ords and language do not seem to play any part in my thought processes. The physical entities which seem to serve as elements in thought are signs and images which can be reproduced and combined at will” (Hadamard, 1954).

Acknowledgments

This work was funded by the James S. McDonnell Foundation Scholar Award. Parts of this work have appeared in Monti et al. (2009), Monti et al. (2012), and Monti and Osherson (2012).

References

- Beall, J. C. and van Fraassen, B. C. (2003). *Possibilities and Paradox: An Introduction to Modal and Many-Valued Logic*. Oxford University Press, Oxford, England.
- Bekinschtein, T. A., Davis, M. H., Rodd, J. M., and Owen, A. M. (2011). Why clowns taste funny: the relationship between humor and semantic ambiguity. *J Neurosci*, 31(26):9665–9671.
- Bellugi, U., Marks, A. B., Bihrl, A., and Sabo, H. (1993). Dissociation between language and cognitive functions in Williams syndrome, pages 177–189. Psychology Press.
- Ben-Shachar, M., Hendler, T., Kahn, I., Ben-Bashat, D., and Grodzinsky, Y. (2003). The neural reality of syntactic transformations: evidence from functional magnetic resonance imaging. *Psychol Sci*, 14(5):433–440.

- Bloom, P. (1994). Generativity within language and other cognitive domains. *Cognition*, 51:177-189.
- Boeckx, C. (2010). *Language in Cognition: Uncovering mental structures and the rules behind them*. Wiley-Blackwell, Oxford, UK.
- Bookheimer, S. (2002). Functional mri of language: new approaches to understanding the cortical organization of semantic processing. *Annu Rev Neurosci*, 25:151–188.
- Bornkessel, I., Zysset, S., Friederici, A. D., von Cramon, D. Y., and Schlesewsky, M. (2005). Who did what to whom? the neural basis of argument hierarchies during language comprehension. *Neuroimage*, 26(1):221–233.
- Butterworth, B. (2005). Developmental dyscalculia. In Campbell, J. I., editor, *Handbook of mathematical cognition*, pages 455–467. Psychology Press, New York, NY.
- Canessa, N., Gorini, A., Cappa, S. F., Piattelli-Palmarini, M., Danna, M., Fazio, F., and Perani, D. (2005). The effect of social content on deductive reasoning: an fmri study. *Hum Brain Mapp*, 26(1):30–43.
- Cappelletti, M., Butterworth, B., and Kopelman, M. (2001). Spared numerical abilities in a case of semantic dementia. *Neuropsychologia*, 39(11):1224–1239.
- Cheng, P. W. and Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive psychology*, 17(4):391–416.
- Chomsky, N. (1957). *Syntactic structures*. Mouton, The Hague, Berlin.
- Chomsky, N. (1983). Noam Chomsky's views on the psychology of language and thought. In Rieber, R. W. (ed.), *Dialogues on the psychology of language and thought: conversations with Noam Chomsky, Charles Osgood, Jean Piaget, Ulric Neisser, and Marcel Kinsbourne*. Plenum Press.

- Chomsky, N. (1986). *Knowledge of language: its origin, nature and use*. New York: Praeger.
- Chomsky, N. (1998). *Language and Problems of Knowledge*. The MIT Press, Cambridge, MA.
- Corballis, M. C. (1992). On the evolution of language and generativity. *Cognition*, 44:197–226.
- Dapretto, M. and Bookheimer, S. Y. (1999). Form and content: dissociating syntax and semantics in sentence comprehension. *Neuron*, 24(2):427–432.
- Dehaene, S., Molko, N., Cohen, L., and Wilson, A. J. (2004). Arithmetic and the brain. *Curr Opin Neurobiol*, 14(2):218–224.
- Dehaene, S., Piazza, M., Pinel, P., and Cohen, L. (2003). Three parietal circuits for number processing. *Cogn Neuropsychol*, 20(3):487–506.
- Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., and Tsivkin, S. (1999). Sources of mathematical thinking: behavioral and brain-imaging evidence. *Science*, 284(5416):970–974.
- Delazer, M., Girelli, L., Semenza, C., and Denes, G. (1999). Numerical skills and aphasia. *J Int Neuropsychol Soc*, 5(3):213–221.
- Fadiga, L., Craighero, L., and D’Ausilio, A. (2009). Broca’s area in language, action, and music. *Ann N Y Acad Sci*, 1169:448–458.
- Fangmeier, T., Knauff, M., Ruff, C. C., and Sloutsky, V. (2006). fMRI evidence for a three-stage model of deductive reasoning. *J Cogn Neurosci*, 18(3):320–334.
- Fitch, W. T., Hauser, M. D., and Chomsky, N. (2005). The evolution of the language faculty: clarifications and implications. *Cognition*, 97(2):179–210; discussion 211–25.
- Fodor, J. A. (1975). *The language and thought*. Harvard University Press.
- Friederici, A. D. (2004). Processing local transitions versus long-distance syntactic hierarchies. *Trends Cogn Sci*, 8(6):245–247.

Gelman, R. and Gallistel, C. R. (2004). Language and the origin of numerical concepts. *Science*, 306(5695):441–443.

Goel, V., Buchel, C., Frith, C., and Dolan, R. J. (2000). Dissociation of mechanisms underlying syllogistic reasoning. *NeuroImage*, 12(5):504–514. PMID: 11034858.

Goel, V. and Dolan, R. J. (2001). Functional neuroanatomy of three-term relational reasoning. *Neuropsychologia*, 39(9):901–909.

Goel, V. and Dolan, R. J. (2003). Explaining modulation of reasoning by belief. *Cognition*, 87(1):B11–B22.

Goel, V., Gold, B., Kapur, S., and Houle, S. (1997). The seats of reason? an imaging study of deductive and inductive reasoning. *Neuroreport*, 8:1305–1310.

Goel, V., Gold, B., Kapur, S., and Houle, S. (1998). Neuroanatomical correlates of human reasoning. *Journal of Cognitive Neuroscience*, 10:293–302.

Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.

Grodzinsky, Y. and Santi, A. (2008). The battle for broca’s region. *Trends Cogn Sci*, 12(12):474–480.

Grosenick, L., Clement, T. S., and Fernald, R. D. (2007). Fish can infer social rank by observation alone. *Nature*, 445(7126):429–432.

Hadamard, J. (1954). *An Essay on: The Psychology of Invention in the Mathematical Field*, volume 107. Princeton University Press, Princeton, NJ.

Hagoort, P. (2005). On broca, brain, and binding: a new framework. *Trends Cogn Sci*, 9:416-23.

Hauk, O., Johnsrude, I., and Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron*, 41(2):301–307.

Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). The faculty of language: what is it, who has it, and how did it evolve? *Science*, 298(5598):1569–1579.

Hoenig, K. and Scheef, L. (2005). Mediotemporal contributions to semantic processing: fmri evidence from ambiguity processing during semantic context verification. *Hippocampus*, 15(5):597–609.

Just, M. A., Carpenter, P. A., Keller, T. A., Eddy, W. F., and Thulborn, K. R. (1996). Brain activation modulated by sentence comprehension. *Science*, 274(5284):114–116.

Knauff, M., Fangmeier, T., Ruff, C. C., and Johnson-Laird, P. N. (2003). Reasoning, models, and images: behavioral measures and cortical activity. *J Cogn Neurosci*, 15(4):559–573.

Koscik, T. R. and Tranel, D. (2012). The human ventromedial prefrontal cortex is critical for transitive inference. *J Cogn Neurosci*, 24(5):1191–1204.

Li, P. and Gleitman, L. (2002). Turning the tables: language and spatial reasoning. *Cognition*, 83(3):265–294.

Locke, J. (1824). *The works of John Locke in nine volumes; Vol 1*. Rivington, London, UK.

Losonsky, M. (1999). *Humboldt: 'On Language': On the Diversity of Human Language Construction and Its Influence on the Mental Development of the Human Species*, CUP.

Maruyama, M., Pallier, C., Jobert, A., Sigman, M., and Dehaene, S. (2012). The cortical representation of simple mathematical expressions. *Neuroimage*, 61(4):1444–1460.

Monti, M. M. and Osherson, D. N. (2012). Logic, language and the brain. *Brain Res*, 1428:33–42.

Monti, M. M., Osherson, D. N., Martinez, M. J., and Parsons, L. M. (2007). Functional neuroanatomy of deductive inference: a language-independent distributed network. *Neuroimage*, 37(3):1005–1016.

- Monti, M. M., Parsons, L. M., and Osherson, D. N. (2009). The boundaries of language and thought in deductive inference. *Proc Natl Acad Sci U S A*, 106(30):12554–12559.
- Monti, M. M., Parsons, L. M., and Osherson, D. N. (2012). Thought beyond language: neural dissociation of algebra and natural language. *Psychol Sci*, 23(8):914–922.
- Noveck, I. A., Goel, V., and Smith, K. W. (2004). The neural basis of conditional reasoning with arbitrary content. *Cortex*, 40(4-5):613–622.
- Novick, J. M., Trueswell, J. C., and Thompson-Schill, S. L. (2010). Brocas area and language processing: Evidence for the cognitive control connection. *Language and Linguistics Compass*, 4(10):906–924.
- Ogawa, S., Lee, T. M., Kay, A. R., and Tank, D. W. (1990). Functional brain mapping by blood oxygenation level dependent contrast magnetic resonance imaging: A comparison of signal characters with a biophysical model. *Proc Natl Acad Sci U S A*, 87:9868 – 9872.
- Osherson, D. (1975). Logic and models of logical thinking. *Reasoning: Representation and process in children and adults*, pp 81–91. Lawrence Earlbaum Associates Inc.
- Parsons, L. M. and Osherson, D. (2001). New evidence for distinct right and left brain systems for deductive versus probabilistic reasoning. *Cereb Cortex*, 11(10):954–965.
- Pinker, S. (1984). *Language learnability and language development*.
- Pinker, S. and Jackendoff, R. (2005). The faculty of language: what’s special about it? *Cognition*, 95(2):201–236.
- Polk, T. A. and Newell, A. (1995). Deduction as verbal reasoning. *Psychological Review*, 102(3):533.

- Prado, J., Henst, J.-B. V. D., and Noveck, I. A. (2010a). Recomposing a fragmented literature: how conditional and relational arguments engage different neural systems for deductive reasoning. *Neuroimage*, 51(3):1213–1221.
- Prado, J. and Noveck, I. A. (2007). Overcoming perceptual features in logical reasoning: a parametric functional magnetic resonance imaging study. *J Cogn Neurosci*, 19(4):642–657.
- Prado, J., Noveck, I. A., and Henst, J.-B. V. D. (2010b). Overlapping and distinct neural representations of numbers and verbal transitive series. *Cereb Cortex*, 20(3):720–729.
- Price, C. J. (2000). The anatomy of language: contributions from functional neuroimaging. *J Anat*, 197 Pt 3:335–359.
- Reverberi, C., Cherubini, P., Frackowiak, R. S. J., Caltagirone, C., Paulesu, E., and Macaluso, E. (2010). Conditional and syllogistic deductive tasks dissociate functionally during premise integration. *Hum Brain Mapp*, 31(9):1430–1445.
- Reverberi, C., Cherubini, P., Rapisarda, A., Rigamonti, E., Caltagirone, C., Frackowiak, R. S. J., Macaluso, E., and Paulesu, E. (2007). Neural basis of generation of conclusions in elementary deduction. *Neuroimage*, 38:752–762.
- Reverberi, C., Shallice, T., D’Agostini, S., Skrap, M., and Bonatti, L. L. (2009). Cortical bases of elementary deductive reasoning: inference, memory, and metaduction. *Neuropsychologia*, 47:1107–1116.
- Rips, L. (1994). *The Psychology of Proof*. MIT Press, Cambridge, MA.
- Rodriguez-Moreno, D. and Hirsch, J. (2009). The dynamics of deductive reasoning: an fmri investigation. *Neuropsychologia*, 47(4):949–961.
- Sahin, N. T., Pinker, S., and Halgren, E. (2006). Abstract grammatical processing of nouns and verbs in broca’s area: evidence from fMRI. *Cortex*, 42(4):540–562.

Sapir, E. (1929). The status of linguistics as a science. *Language*, pages 207–214. Sapir, E. (1931). Conceptual categories in primitive languages. *Science*, 74(1927):578.

Spelke, E. S. (2003). What makes us smart? Core knowledge and natural language, In *Language in mind: Advances in the study of language and thought*, p 277–311. MIT Press Cambridge, MA.

Spelke, E. S. and Tsivkin, S. (2001). Language and number: a bilingual training study. *Cognition*, 78(1):45–88.

Stanescu-Cosson, R., Pinel, P., van De Moortele, P. F., Bihan, D. L., Cohen, L., and Dehaene, S. (2000). Understanding dissociations in dyscalculia: a brain imaging study of the impact of number size on the cerebral networks for exact and approximate calculation. *Brain*, 123:2240-55.

Tettamanti, M. and Weniger, D. (2006). Broca’s area: a supramodal hierarchical processor? *Cortex*, 42(4):491–494.

Udden, J. and Bahlmann, J. (2012). A rostro-caudal gradient of structured sequence processing in the left inferior frontal gyrus. *Philos Trans R Soc Lond B Biol Sci*, 367(1598):2023–2032.

Varley, R. and Siegal, M. (2000). Evidence for cognition without grammar from causal reasoning and ‘theory of mind’ in an agrammatic aphasic patient. *Curr Biol*, 10(12):723-726.

Varley, R. A., Klessinger, N. J. C., Romanowski, C. A. J., and Siegal, M. (2005). Agrammatic but numerate. *Proc Natl Acad Sci U S A*, 102(9):3519–3524.

Whorf, B. L. (1940). Science and linguistics. *Technology Review*, 42(6):229–31,247–8.

Wildgruber, D., Ackermann, H., Kreifelts, B., and Ethofer, T. (2006). Cerebral processing of linguistic and emotional prosody: fMRI studies. *Prog Brain Res*, 156:249–268.