

# Human Singularity and Symbolic Tree Structures

## The Demodularization Hypothesis

Stanislas Dehaene

### Abstract

Relative to other primates, humans exhibit a great variety of singular cognitive abilities for language, mathematics, music, tool use, theory of mind, and self-consciousness. What has brought about this singularity? This chapter examines the hypothesis that the human brain is unique in being endowed with a mental representation of nested, tree-like symbolic structures. Such syntactic structures are essential in the modern description of human languages, including natural languages as well as the artificial ones used in music or mathematics. Nonhuman animals may possess abstract representation of temporal sequences, but evidence suggests that those representations do not include the sort of nested tree structures typical of human grammars. Brain imaging, magnetoencephalography and intracranial recordings have begun to reveal the neural correlates of the nested structure of linguistic constituents, which involve Broca's area and the superior temporal sulcus of the left hemisphere. Importantly, the mental manipulation of musical and mathematical structures, which also involves nested trees, is not confined to such classical language areas. Instead, high-level mathematics involves bilateral intraparietal areas involved in elementary number sense and simple arithmetic as well as bilateral inferotemporal areas involved in processing Arabic numerals. This chapter proposes that several distinct circuits of the human brain have become attuned to nested tree structures for different domains, such as language, mathematics, or music. According to the demodularization hypothesis, during human brain evolution, primitive tree structures may have emerged within specialized neural circuits (e.g., those involved in spatial or geometrical computations) and were later exapted toward a more general role in language processing and conscious verbal report.

### Introduction: Hypotheses about Human Singularity

In many cognitive domains, humans are special among other primates. They are, for example, the only species that

- produce and understand language, a highly combinatorial communication system,
- create and use complex tools,
- formulate and test complex scientific theories, expressed in formal mathematical notations,
- possess a complex representation of other minds,
- exhibit a sophisticated representation of their own selves, and
- educate each other based on a representation of the gap between self versus other knowledge.

This list is not exhaustive. Can cognitive neuroscience shed some light on the origins of these remarkable human singularities?

Darwin's view, as expressed in *The Descent of Man*, was that the "difference in mind between man and the higher animals, great as it is, certainly is one of degree and not of kind" (Darwin 1888). In all of these domains, he argued, we can find nonhuman precursors in other animals. Nevertheless, Darwin would probably not have denied the importance of a research program searching for the evolutionary changes that made the acquisition of language, mathematics, theory of mind, or education radically easier for humans than for other animals. Indeed, Darwin defended the view that language ability reflects "an instinctive tendency to acquire an art," a view close to Noam Chomsky's notion of an innate "language acquisition device" or Peter Marler's idea of "learning by instinct."

In this chapter, I consider the possibility that human singularity relies on a novel type of mental representation, recursively nested trees, which makes humans capable of fast learning whenever the structure of the learned domain conforms to such nested trees (Bolhuis et al. 2014). What makes human thought complex, according to the nested tree hypothesis, is that words or symbols are not just strung together into a sequence, but are mentally represented as hierarchical trees, thus offering a much greater combinatorial diversity, as in distinguishing the concepts  $3x + 1$  from  $3(x + 1)$ , or *un-lockable* from *unlockable*. I speculate that the human brain is singular because, during evolution, it acquired the ability to represent complex internal tree structures, evaluate their adequacy as hypotheses for a given domain, and even impose them onto simple incoming sequences of stimuli.


## Five Types of Sequence Representations in the Human Brain

As early as the 1950s, the problem of serial order in behavior was identified by Karl Lashley (1951) as one of the pressing questions that behavioral and neural sciences should address. The problem can be stated succinctly: How does the brain encode temporal sequences of actions, such that this knowledge can be used to retrieve a sequence from memory, recognize it, anticipate

forthcoming items, and generalize this knowledge to novel sequences with a similar structure?

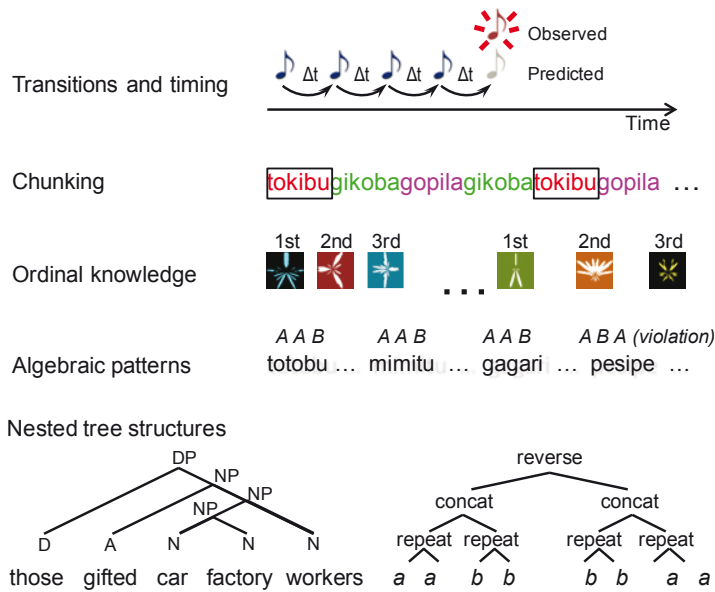
Lashley (1951) showed that a simple “associative chain,” based on the hypothesis that each element in a sequence served as a cue for the next, was insufficient to account for the complexity of motor sequences that were produced by various animal species. In a recent review of sequence learning research (Dehaene et al. 2015), my colleagues and I have proposed a typology of five different levels of sequence representations, each representing sequences at increasingly high levels of abstraction (Figure 15.1):

1. *Transition and timing knowledge*: knowledge of the specific transitions from one item to the next; that is, the identity and approximate timing of the next item relative to the preceding ones.
2. *Chunking*: the grouping of several contiguous items into a single “chunk” which can be manipulated as a whole at the next hierarchical level.
3. *Ordinal knowledge*: knowledge of which item comes first, which comes second, and so on, independently of their timing.
4. *Algebraic patterns*: abstract schemas that capture the sequential regularities underlying a sequence of items; for instance, noticing that the word “beriberi” comprises repeated syllables that conform to the ABAB pattern.
5. *Nested tree structures generated by symbolic rules*: at this level, characteristic of human languages, a sequence can be “parsed” according to abstract grammatical rules into a set of groupings, possibly embedded within each other, forming a nested structure of arbitrary depth, and involving the recursive use of the same types of element at multiple levels; an example is the parsing of the mathematical equation  $a + b \sin \omega t$  as a nested set of parentheses ( $a + (b(\sin(\omega t)))$ ) or, equivalently, a tree structure:



A variety of experiments have demonstrated that the first four levels in this typology are present in various nonhuman animals, particularly in primates. To give just a few examples:

1. Transition and timing knowledge is evident in a variety of experiments which show that monkeys develop precise temporal expectations of auditory or visual stimuli, and exhibit a neural “mismatch response” to violations of these expectations (Meyer and Olson 2011; Uhrig et al. 2014; Wilson et al. 2017).
2. Chunking knowledge was demonstrated, for instance, in tamarin monkeys using the same procedure as in human infants (Hauser et al. 2001). Both species can detect recurrent “words,” that is, chunks



**Figure 15.1** A typology of sequence representations. Five types of mental representations of sequences are postulated (for further details, see Dehaene et al. 2015). The first four are present in nonhuman animals, but a capacity to quickly acquire and manipulate nested tree structures may be unique to the human brain. From Dehaene et al. (2015), used with permission.

of syllables in a continuous stream such as *tokibugikobagopilatipolutokibugopilatipolutokibugikobagopila*.

3. Ordinal knowledge is part of a general endowment for number processing in various animal species (Nieder and Dehaene 2009). Macaques, for instance, know which image in a list came first, second, third, or last (Chen et al. 1997), and they can string together the first item of a list with the second of another, and so on.
4. Algebraic pattern knowledge has been observed in rats, macaques, and chimpanzees. All of these species can recognize when sequences of sounds or images obey a fixed pattern, such as *ABA* or *xxxY* (Murphy et al. 2008; Sonnweber et al. 2015; Wang et al. 2015).

Whether the fifth level (nested tree structures generated by symbolic rules) is present in nonhuman animals is a matter of great contention. The bulk of the evidence suggests that animals, even macaques or apes, do not easily learn natural or artificial languages whose structure is based on nested trees, or equivalently a context-free grammar (Penn et al. 2008; Fitch and Friederici 2012). Many primate species use vocal communication systems in the wild, but these do not seem to make use of a sophisticated syntax based on nested trees; at

most, they involve the sequential concatenation of vocal sounds and possibly the use of a single modifier or “suffix” (Ouattara et al. 2009). For apes who have been taught artificial languages, comprising tens or even hundreds of token words in visual or sign language, it is noteworthy that they were unable to combine them systematically according to a complex grammar (Terrace et al. 1979; Yang 2013). Recently, it has been claimed that several species of birds have mastered a simple center-embedded context-free grammar called  $A^nB^n$ , yet this claim has been heavily criticized and, in my opinion, fully refuted (Beckers et al. 2016). Parrots may exhibit superior performance (Pepperberg 2013), including the comprehension of questions comprising several words, such as “how many green keys.” However, such experiments are limited to a very small number of specific animals, with many years of training, and may still involve only a concatenation or intersection of abstract concepts, without genuine syntax.

### **Evidence for Nested Tree Structures in Human Language**

Comparatively, although still contested by a few linguists (Frank et al. 2012), there is considerable evidence that human language is “special” and requires covert, internal representations of abstract, symbolic syntactic structures (Haegeman 2005). Evidence for this arises from at least three distinct areas of research: linguistics, psycholinguistics, and cognitive neuroscience.

#### **Linguistic Descriptions of Language**

As early as the 1950s, Noam Chomsky convincingly refuted the Skinnerian view of language as a chain of conditioned responses. The current view is that language syntax involves a nested tree of phrases or “constituents,” which can be represented by parentheses: (the(big animal)(with (two horns))))). It is important to understand that this is a shorthand notation, as linguists today postulate that syntactic trees are not made of the raw input words, but of abstract covert objects, such as focus, topic, tense, complementizer, and trace. The case for nested constituent structures in linguistics rests on many observations:

- Cases of syntactic ambiguity, such as “looking at a man with binoculars,” demonstrate that the same sequence of words can have two distinct internal representations, depending solely on tree attachment. This makes it impossible for those representations to rely solely on temporal order. Such tree-based ambiguity exists even within a single word whose morphemes can be ambiguously attached, as in un-(lock-able) versus (un-lock)-able.
- Ellipsis or substitution, whereby some strings of words—specifically those forming a subtree of the entire sentence structure (also called a

“phrase” or “constituent”)—can be replaced by a single word: he (went (to (the store))) → he went to it, he went there, he did.

- Syntactic movement, whereby the same constituents can be moved to a distinct sentential location in order to form questions and relatives or to emphasize a specific topic: “*to the store*, that’s where he went” or “*where* did he go”).
- Long-distance dependencies, whereby the properties of two constituents (or rather, their top-level nodes or “heads”) must be matched. One example is agreement in number (singular or plural) between a subject and a verb: he *was* tall, but they *were* tall. Such agreement relationships require skipping over intermediate subconstituents (e.g., “the car which passed the trucks *is* red”) and therefore cannot be captured by linear structure alone. Many other long-distance phenomena, such as pronoun binding, also require considering a sentence’s tree structure.
- Evidence that languages can be distinguished based on minimal differences in the language-specific rules or “parameters” that govern the formation of such tree structures. The syntactic differences between English and Japanese, for instance, can be largely accounted for by a single parametric difference in the ordering of words when tree structures are linearized during speech production: the head-first versus head-last parameter (for a very accessible account, see Baker 2001).

## Behavioral Studies

Psycholinguistic studies (i.e., behavioral studies of language comprehension and production) have regularly observed that human verbal behavior reflects the underlying syntactic tree structures. To give but a few examples: During the comprehension of sentences with syntactic movements, human adults reactivate the antecedent of the trace of the displaced tree at precisely the moment when linguists postulate that such a trace should occur (e.g., Friedmann et al. 2008). During sentence production, they lengthen the duration of words and the intervals between sentences in direct, linear relation to the depth of the corresponding syntactic tree structure (Breen 2018). Finally, children’s acquisition of language involves a systematic and nontrivial generalization over the tree structures governing, for instance, the placements of auxiliaries and verbs relative to negation and adverbs (Déprez and Pierce 1993).

## The Search for Neural Correlates of Syntactic Trees

If the syntactic tree hypothesis is correct, there must be a set of specific brain circuits and neural codes, possibly unique to the human brain, that is engaged whenever humans process the syntactic structures of language. Functional MRI has provided strong converging evidence that a dedicated left hemispheric network is systematically associated with the formation and manipulation of

nested syntactic and semantic structures (see Table 15.1 and Figure 15.2). These regions are part of an amodal network, spread all along the left superior temporal sulcus and inferior frontal gyrus, which is activated whenever humans process spoken or written sentences. Indeed, they are also active when deaf people process their native sign language.

Amidst this network, a core set of regions formed by the left inferior frontal gyrus and the left posterior superior temporal sulcus, often associated with a node in the left basal ganglia (Moreno et al. 2018), appears to be specifically involved in syntax. These regions exhibit a level of activation that is monotonically related to the number of nested constituents in the stimuli: they do not activate strongly to lists of words but show increasing levels of activity as the words are combined into syntactically correct constituents of 2, 3, 4, or 6 words, all the way to a full sentence of 12 words. Furthermore, they continue to respond in this manner even when the stimuli are delexicalized and comprised only of the function words and grammatical morphemes needed to parse them syntactically (Table 15.1, Jabberwocky condition; Pallier et al. 2011).

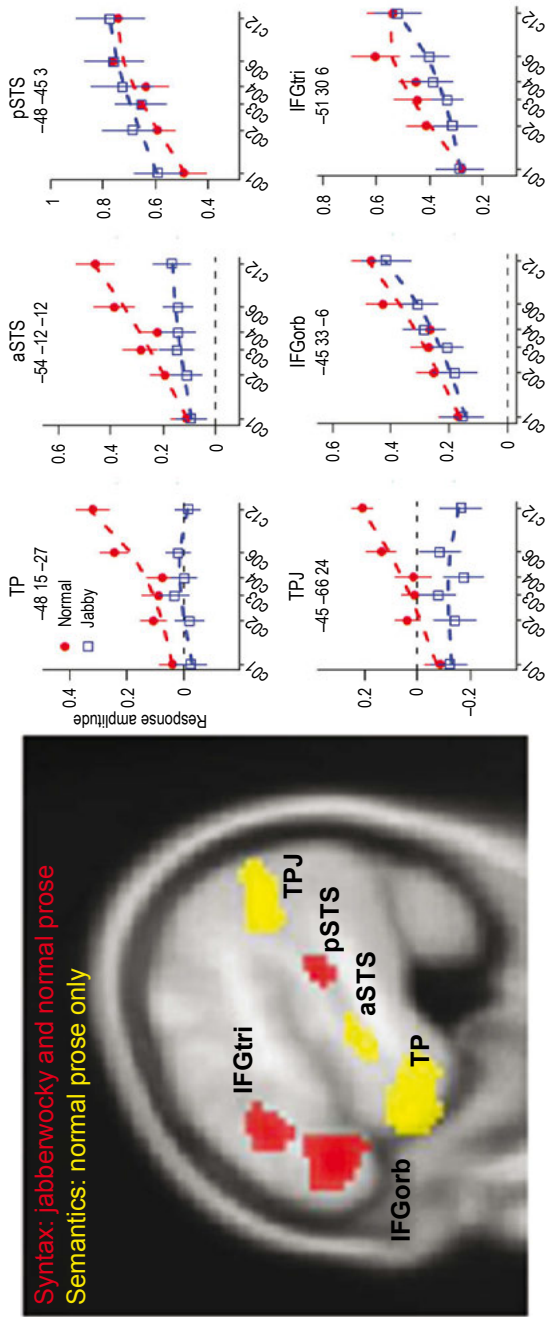
Many functional MRI studies indicate that these core regions are active whenever a subject represents or processes syntactic structures (even the simplest ones comprising a few words). Activation in those areas is proportional to syntactic complexity and to the amount of syntactic movement. Furthermore, they exhibit a clear dissociation from other more generic regions involved in cognitive effort and working memory (Fedorenko et al. 2011). They are also engaged when tree structures need to be manipulated internally to recover

**Table 15.1** During functional MRI, adult volunteers read 12-word sequences of different length; corresponding brain activity is shown in Figure 15.2.

| Constituents   | Examples (normal prose)  |
|----------------|--|
| 12 words (c12) | I believe that you should accept the proposal of your new associate  |
| 6 words (c06)  | the mouse that eats our cheese two clients examine this nice couch   |
| 4 words (c04)  | mayor of the city he hates this color they read their names          |
| 3 words (c03)  | solving a problem repair the ceiling he keeps reading will buy some  |
| 2 words (c02)  | looking ahead who dies important task his dog few holes they write   |
| 1 word (c01)   | thing very tree where of watching copy tensed they states heart plus |

| Constituents   | Examples (Jabberwocky)  |
|----------------|---|
| 12 words (c12) | I tosieve that you should becept the tropufal of your tew viroflate |
| 6 words (c06)  | the couse that rits our treeve fow plients afomine this kice bloch  |
| 4 words (c04)  | tuyor of the roty he futes this dator they gead their wames         |
| 3 words (c03)  | relging a grathem regair the fraping he meeps bouding will doy some |
| 2 words (c02)  | troking ahead who mies omirpant fran his gog few biles they grite   |
| 1 word (c01)   | thang very gree where of wurthing napy gused they flotes blart trus |





**Figure 15.2** Evidence for a core network for syntax in the human brain (from Pallier et al. 2011). The relevant language areas can be easily identified by having adult volunteers read 12-word sequences of varying length (Table 15.1) during functional MRI. Along the superior temporal sulcus and in the left inferior frontal region (Broca's area), activity increases monotonically with the complexity of the phrasal constituents. Red regions are more specifically involved in the syntax, because they continue to activate even when the phrases are formed of pseudo-words (a condition called "jabberwocky," in honor of the eponymous poem by Lewis Carroll).

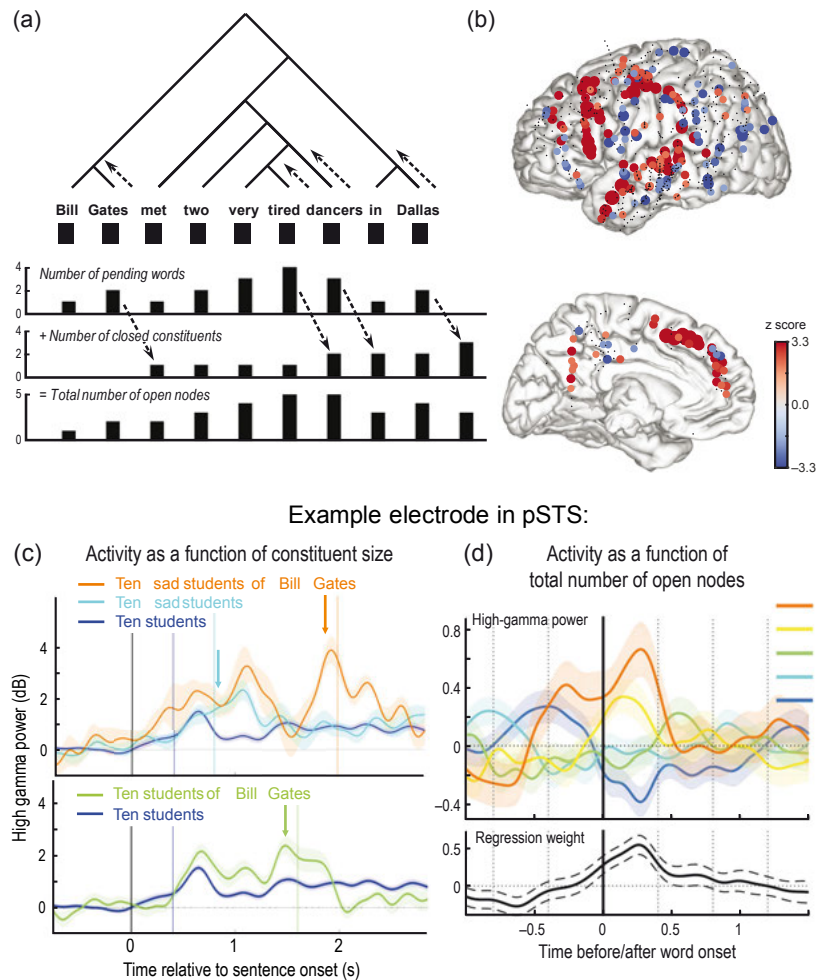


“who did what to whom” (Pattamadilok et al. 2015), or when syntactic ambiguities must be lifted. Finally, they appear to play a central role in agrammatism: patients with lesions to this core set of regions or the associated fiber tracts linking them are much more likely to develop agrammatic aphasia than other patients; even if these regions are not lesioned, they are strongly hypoactive in agrammatic patients (Tyler et al. 2011).

These studies pinpoint a set of highly specialized cortical and subcortical regions for sentence parsing and structural representation. Furthermore, they also largely validate some of the most important theoretical constructs postulated by linguists (e.g., the existence of nested constituents, syntactic movement, non-accusative verbs), since their presence or absence appears to be a key determinant of brain activity level in those regions. However, functional MRI has not, until now, revealed how these constructs are encoded by populations of neurons. Efforts are underway to clarify this point using patients with epilepsy whose brain signals can be recorded directly using epidural or intracortical electrodes.

In a recent study, we tracked the word-by-word changes in high gamma activity as patients processed each successive word in written sentences of controlled syntactic complexity. We observed that a subset of electrodes, largely confined to known cortical language areas, exhibited a systematic pattern reflecting the constituent structures of the stimulus sentences: their activity rose whenever a new word appeared, but it also decreased whenever several consecutive words or constituents could be merged into a large constituent structure. To take a simplified example: upon reading “two...sad...girls...often...cried,” the activation progressively increased, but it collapsed after “girls” because those three words could be combined into a single noun phrase, the subject of the subsequent verb phrase. Importantly, the activity at any given time was proportional to the number of open nodes; that is, the items (words or multi-word constituents) that had not yet been merged together. Remarkably, in a regression that accounted for high gamma activity, similar weights were given to individual words (e.g., sad) as well as to temporary constituents (e.g., “two sad girls,” “often cried”). The results were compatible with a bottom-up parsing system which applied rules of grammar in order to group words into nested constituent structures. Alternative models, for instance based solely on transition probabilities between individual words or between their grammatical categories, could be formally rejected, as they did not provide an equally good fit to the observed neurophysiological responses.

These findings provide strong evidence that constituents are the relevant units for syntactic structures. They also support the existence of the Merge operation, the most fundamental hypothesis of the recent “minimalist” approach to language. Merge is the basic tree-building operation hypothesized to take two words or constituents as input and form the binary tree whose leaves are those two constituents. Because this operation is recursive, it can represent an entire sentence as a tree with a nested set of embedded sub-trees (Figure 15.3).



**Figure 15.3** Neurophysiological evidence for the online formation of constituent structures during sentence reading (from Nelson et al. 2017): (a) The proposed model hypothesizes that brain activity builds up in language areas for each successive word in a sentence, but collapses at phrase boundaries, when several words can be compressed into a closed constituent. (b) This prediction was upheld in the intracranial electrodes shown in red. (c) In an example electrode in left posterior superior temporal sulcus (pSTS), high gamma activity rises and falls according to constituent boundaries and (d) strongly correlates with the number of remaining open nodes shortly after the presentation of each word.

This is analogous to the multiscale compression operation which is performed, for instance, by the JPEG image compression algorithm. The human brain possesses a capacity to compress a sentence by identifying nested groups of words that act as constituents or phrases and operate together as a unit.

Indeed, this compression hypothesis, grounded on the observed rise and fall of intracranial brain signals during constituent structure building (Nelson et al. 2017), can account for two basic observations about language processing:

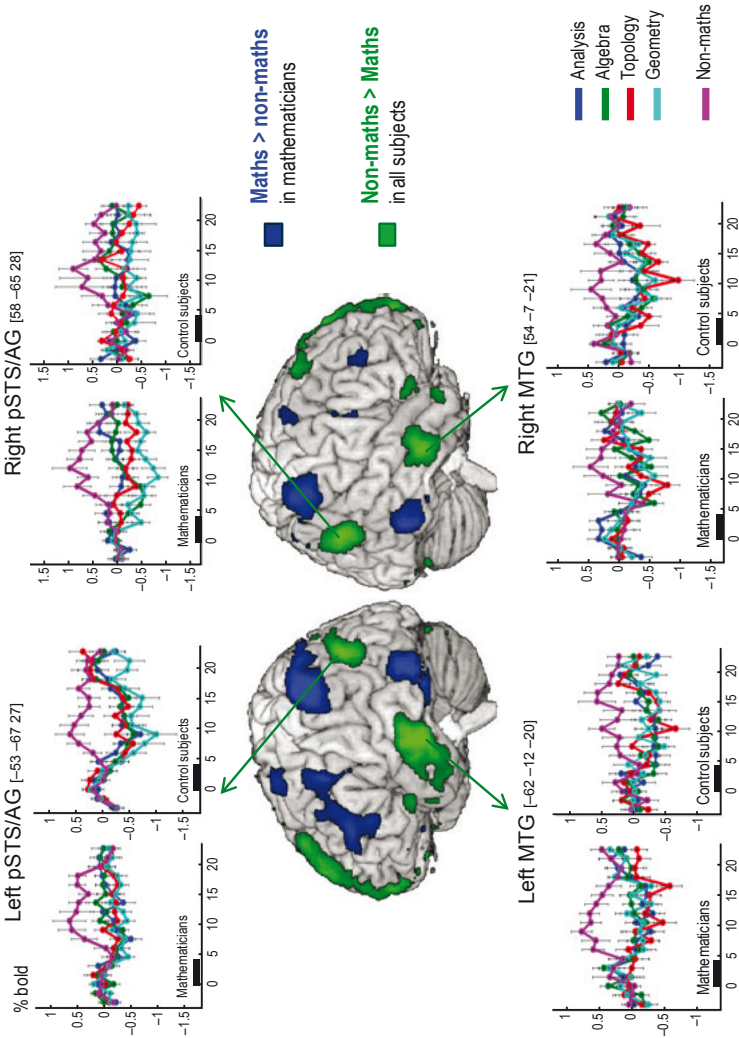
1. fMRI activity increases sublinearly, indeed close to logarithmically, with the number of words in a sentence (Pallier et al. 2011).
2. Memory for a sentence can be much better than a list of words of the same length and far exceed the typical working memory limit of ~7 items.

We do not yet know how syntactic trees are represented at the neural level, but it seems that the postulates of modern linguistics are largely vindicated.


### **The Language of Mathematics**

As noted earlier, outside of the human species, there is simply no behavioral or neural evidence so far to suggest that nonhuman animals are capable of a similar tree-based representation or compression operation. Indeed, the inferior frontal and superior/middle temporal regions involved in language processing are enormously expanded in the human brain, even relative to our closest cousins, the great apes (Smaers et al. 2017). Obviously, these regions were deeply transformed during evolution, and I speculate, as do others (Hauser et al. 2002; Fitch and Friederici 2012), that these regions may have acquired a novel representational tool: recursive tree structures. Is this property, however, unique to language? Did it evolve just once, in a single language-related brain circuit, thus placing language at the heart of human singularity? Is the emergence of language the sole factor responsible for all of our other talents for tool use, science, music, mathematics, or theory of mind?

To investigate this issue, my colleagues and I have performed a series of investigations of mathematical abilities and their brain mechanisms (Dehaene 2011). The results are very clear: the bulk of mathematics resides in a network of brain regions quite different from language areas. These regions encode nonlinguistic concepts of space, time, and number and are preempted or “re-cycled” during education to higher-level mathematics. Recently, for instance, we used functional MRI to investigate the brain networks for language and mathematics in professional mathematicians (Amalric and Dehaene 2016). During a brief period of reflection on mathematical statements, such as “the sine function is periodical: true or false?” mathematicians activated a bilateral dorsal network of parietal and frontal regions, as well as the bilateral lateral inferior temporal gyrus—regions that show no overlap with any language regions, as determined from a distinct localized area for written or spoken sentence processing (Figure 15.4). Indeed, the math-responsive network was also entirely different from the set of semantic regions involved in resolving similar general-knowledge verbal statements, such as “London buses are red: true or false?” (Figure 15.5).



**Figure 15.4** Mathematical reflection recruits a network of brain areas distinct from language areas (from Amalric and Dehaene 2016). When mathematicians listen to mathematical statements and evaluate their veracity, the activated network (blue) differs from the one recruited when evaluating non-mathematical propositions (green).



**Figure 15.5** Functional MRI demonstrates that the network active during mathematical operations (yellow) bypasses classical language areas (red).

Further experiments indicated that high-level mathematics activated the very same voxels involved in basic number recognition and arithmetic calculation, which are present in mathematicians and nonmathematicians alike, and can be traced to nonverbal representations of approximate number in infants. Thus, the results indicate a nonverbal origin for higher mathematical abilities and support the cortical recycling hypothesis, which postulates that they are grounded in prior circuits for processing number and space.

Our fMRI experiments with professional mathematicians did not explicitly study the format of representation for mathematical facts. However, there is general agreement that mathematic knowledge is organized as a language with nested tree structures. This point was already made by Galileo, who noted that “the book [of the universe] is written in the mathematical language, whose symbols are triangles, circles and other geometrical figures, without whose help it is impossible to comprehend a single word of it.” Several behavioral and brain-imaging experiments have confirmed that to understand how subjects manipulate even the most basic structures of mathematics, an appeal to language-like syntax is indispensable. For instance, when processing two-digit numbers, subjects quickly and automatically assign each of the two digits to a syntactic position as decade versus unit, and weigh their quantity accordingly. When processing algebraic equations such as  $3x^2 + y^3$ , they automatically parse it into a nested set of constituents, thus failing to recognize that “ $x^2 + y$ ” is, in fact, a subset of consecutive symbols in this string. In addition, when scanned with fMRI while manipulating such nested structures, they show a modulation of brain activity according to the number of nested constituents, not in language areas, but in bilateral intraparietal regions also involved in number sense (Maruyama et al. 2012).

In recent work, we investigated whether a “language of thought” comprising nested constituent structure is involved even in extremely basic tasks of

spatial working memory (Amalric et al. 2017). We asked subjects to view a sequence of spatial locations on an octagon, and to remember the order in which the eight locations were presented. This is a variant of the “Corsi blocks” task for spatial working memory, except that we systematically vary the amount of geometrical structure in the stimuli. At one extreme of simplicity, the successive locations could simply run around the octagon in a serial manner, clockwise or counterclockwise. However, we also presented more complex but still regular patterns: a zigzag, two nested squares, two rectangles...all the way to a completely irregular sequence without any noticeable regularity.

Our results showed that even in this simple spatial memory setting, subjects already deployed a mental “language of thought” for geometry. Indeed, several aspects of their behavior were predicted by the amount of regularity in the sequence: with regular sequences, their memory was better; they could anticipate items that had not yet even been presented; and their eyes automatically moved to the next location, in a manner that directly reflected the underlying tree structure. Similar results were obtained with preschoolers and uneducated adults from the Amazon (the Mundurucu); for related experiments on the development of the perception of fractals, see Martins et al. (2014).

We formalized the results by proposing that during the spatial memory task, subjects search for a minimal formula for the spatial sequence within a mental “language of geometry.” The idea is that subjects encode the sequence in memory according to its minimal description length (MDL), the shortest program capable of generating it. They use geometrical regularities such as repetitions and symmetries to compress the incoming sequence into a tree-like representation with a recursive structure (comprising, for instance, three nested levels of repetitions with variations). Indeed, MDL (also known as Kolmogorov complexity) was an excellent predictor of memory for spatial sequences.

We then used fMRI to identify which brain areas contribute to this “language of geometry.” While subjects simply moved their eyes to the next target, sequence complexity (MDL) predicted activation in a broad set of nonlinguistic brain areas, primarily in dorsal parietal and prefrontal cortex, extending into the dorsal part of Brodmann’s area 44. The latter area was the only region to survive once controls were introduced for eye movement length and for working memory per se. Thus, we believe that this dorsal part of the inferior frontal gyrus, distinct from other sectors of Broca’s area involved in natural language processing, was involved in representing a spatial sequence as a rule-based tree structure.

Overall, the data suggest that Darwin was wrong when he noted that “a complex train of thought can be no more carried out without the use of words, whether spoken or silent, than a long calculation without the use of figures or algebra” (Darwin 1888). Mathematical reasoning requires nested structures of symbols, but those symbols need not be the words of natural language. Tree structures are present outside of the language domain, as in mathematics, but they recruit distinct brain networks.



I therefore suggest that the ability to represent and to manipulate recursive tree structures is not the property of a single brain circuit, but of multiple parallel systems. A similar conclusion has been reached by Peter Hagoort and his colleagues even *within* the language system (Hagoort 2013): distinct levels of phonology, syntax, and semantics may involve parallel temporofrontal circuits. In each of these domains, different sectors of the left inferior frontal gyrus may play a similar role of “unifying” the linguistic elements provided by distinct posterior brain, yet at different timescales and with different units of computation (phonemes, words, meanings).

### **Conclusion: A Demodularization Hypothesis for the Emergence of Universal Tree Structures**

The hypothesis that multiple circuits of the human brain exhibit a capacity for recursion raises the issue of how those circuits, during evolution, acquired this novel property. Two possibilities may be envisaged. First, it is possible that a single genetic event affected multiple brain regions. A mutation in the genes that control radial neuronal migration or cortical layer formation, for instance, may have jointly affected a large set of brain areas, possibly conferring novel computational properties to many parallel circuits at once. According to this possibility, preexisting circuits for spatial memory, auditory memory, and so forth would all have acquired, at some point in the human lineage, the ability to represent tree structures.

An alternative possibility is that tree structures evolved first within a specific brain circuit, and that this property was later extended to other circuits. The dissociation between cortical circuits for mathematical knowledge and natural language processing makes it possible for those circuits to have evolved at a different time in hominization. Is there any archeological evidence for this speculation? A sophisticated form of language is usually thought to have emerged only with *Homo sapiens*, accounting for this species’ sudden cultural and geographical expansion. This recent origin for language (though speculative) contrasts sharply with the existence of very ancient proto-mathematical human-made artifacts. Tools with remarkable symmetry and geometrical regularity were already crafted as early as 2 million years ago and must therefore have been made not by *H. sapiens*, but by *H. ergaster* or archaic *H. erectus* (see Figure 15.6). Bifaces, for instance, are carefully crafted stone tools that present two orthogonal planes of symmetry and often a geometrically regular contour (e.g., an egg shape with a highly regular variation in curvature; or a pointed shape made of two lines and a circle). Around the same time, early humans also created polyhedral and spherical stone artifacts, sometimes close to a perfect sphere, suggesting that they could already conceive of regular mathematical objects before carefully sculpting them out of stone. While no other animal species creates such tools, it is still unclear whether these objects necessarily





**Figure 15.6** Early in evolution, humans produced many “proto-mathematical” objects with regular shapes based on symmetry and geometry, as evidenced by bifacial and spheroid artifacts (top panels) dating back ca. 1.6 million yr BCE. Complex combinations of these forms became evident as early as 70,000 years ago, as illustrated by the engraved patterns of an ochre artifact (bottom), reported by Henshilwood et al. (2002).

imply a “language of geometry” capable of combining multiple concepts according to nested tree structures. A combinatorial system of geometry is only clearly attested in symbolic drawings with parallel lines and equilateral triangles on ochres from the Blombos cave in South Africa dating back at least 70,000 years (Henshilwood et al. 2009). The rudiments of the “language of geometry” must have been present at that time and, depending as to whether one accepts bifaces as evidence, perhaps as much as 1.5 million years before the emergence of spoken language.

Do these objects imply the existence of a spoken language? Can their emergence in the human species be explained by a broader capacity for mental

representation based on a “language of thought?” I suggest that the capacity to form complex representation using nested tree structures first evolved in the domain of mathematical/scientific thought, using combinations of numbers and shapes, perhaps as early as 2 million years ago, and was later exapted as a broader, universal, nonmodular, generative ability for language.

There are several other possible origins for a nonverbal representation of nested tree structures prior to the emergence of language. Understanding human and animal bodies may require representing them as flexible tree structures; that is, as nested sets of “parts within parts within parts” (e.g., fingers within hands within limbs), in a recursive manner. Parsing visual scenes involves understanding the relative locations of objects, which may require nesting (e.g., a bird on the bush left of the rock). Similarly, navigating in space may require representing spaces as a nested structure of embedded places of increasing size (e.g., a spot in a room in a cave). Encoding of action plans may involve a nested tree structure of goals and subgoals. Finally, understanding social groups may involve representing the nested trees of family and dominance relationships.

For all these reasons, I find it plausible to propose a “demodularization hypothesis” for the emergence of universal tree structures in humans. The proposal is that a capacity for recursive tree-based representations emerged early on during primate evolution or hominization, within a specific module, before being extended to linguistic communication. Primitive tree structures would have emerged within specialized neural circuits (e.g., those involved in spatial or geometrical computations). Only much later would they have been exapted toward a more general role in language processing.

The peculiarity of language, indeed, is that its semantic structures span over and bring together a vast array of mental representations. Within a sentence, we can combine together any object, any action, any person, any idea, any logical connector or quantifier in countless manners. It seems that any conscious object of thought can be integrated in the language system. Indeed, this is why the criterion of verbal reportability, in humans at least, is considered by many as the primary evidence that some information is conscious (Weiskrantz 1997). The highest level of language processing is therefore nonmodular. As noted by Fodor (1983), it is not organized in a modular manner but as a “horizontal” system capable of interconnecting many, indeed virtually any, mental processors, and therefore participating in a nonmodular conscious “global neuronal workspace” (Dehaene 2014). This property, however, need not have been present in protohumans and may constitute an exaptation of a simpler, modular system.

What is exciting about this possibility is that, if true, we may reasonably hope to find precursors of the tree structures of language in nonhuman species. The neural code for linguistic structures lies beyond our reach, because obvious ethical reasons currently restrict our access to the massively parallel recordings of human neurons that would be needed to characterize it. However, by studying

the representation of space or body parts in nonhuman animals, we may perhaps identify a simpler but similarly organized tree-based neural system.