Title:

Composition: The neurobiology of syntactic and semantic structure building

Liina Pylkkänen, Departments of Linguistics and Psychology, New York University; NYU
Abu Dhabi Institute, NYU Abu Dhabi

Jonathan R. Brennan, Department of Linguistics, University of Michigan

Abbreviated title: Composition

Corresponding author:

Liina Pylkkänen
Departments of Linguistics and Psychology
New York University
New York, NY 10003, USA
Tel: 212-992-8764
liina.pylkkanen@nyu.edu

total number of words: 6303 (inc. title pg)

total number of figures: 1

Acknowledgements:

Support for the writing of this chapter and research summarized within it was provided by the National Science Foundation Grants BCS-1221723 (LP), IIS-1607251 (JRB), and grant G1001 from the NYUAD Institute, New York University Abu Dhabi (LP).

Keywords: syntax; semantics; composition; language comprehension; language production; computational models; left anterior temporal lobe

ABSTRACT

Language is a combinatorial system able to create an infinite array of complex meanings from memory representations in our mental dictionary. How is this integrative process neurally implemented? Studies on the processing of structured vs. unstructured language stimuli have identified a largely left-lateral “combinatorial network,” with the anterior temporal lobe as its most consistent integrative node, likely contributing the first stage of a multi-stage combinatorial process. This chapter summarizes our current understanding of the neurobiology of composition, with a focus on three paradigms from extant literature that most directly address the basic process of combining words into phrases and sentences: the so-called sentence vs. list paradigm, the two-word phrase paradigm and approaches using model comparison with natural narratives as stimuli.

1. WHAT IS COMPOSITION? MANY CORRELATED COMPUTATIONS EXECUTED ACCORDING TO OUR GRAMMATICAL KNOWLEDGE

While the retrieval of stored representations is a process shared by many domains of cognition – this is how you distinguish a dog barking from a baby crying or recognize your favorite hat – only in language do the memory representations of elementary building blocks compose into infinite meaningful configurations along an intricate rule system tacitly living in our brains, the grammar. The composition of structured meanings is the essence of language, the source of its expressive power. What do we know about the brain basis of this remarkable ability?

When approaching the neuroscience of composition, one is immediately faced with a principle challenge: the composition of words into complex messages is achieved by a cascade of tightly correlated and possibly simultaneous computations. Thus, understanding this process requires ways to unpack the constituent processes. For example, the comprehension of *cat* in the context *someone fed the...*, will elicit activity reflecting the syntactic combination of *cat* with the article *the* to form a noun phrase, the processes by which this noun phrase completes both the verb phrase *fed the cat* and the processes by which the verb phrase completes the whole sentence. But of course, the sole purpose of building these syntactic structures is to determine how the words semantically combine with each other. Thus, each step of syntactic structure building is also paired with a potentially multifaceted process of meaning composition. Further, these processes are inter-dependent: the word *fed* both contributes to building a verb phrase and also allows one to predict semantically compatible words to complete the sentence, namely things that can eat. Clearly then, careful experimental design resting on a solid theoretical foundation is paramount for the neuroscience of composition. For each possible neural correlate of

composition, we must ask: which of the tightly correlated computations can be ruled out as a functional hypothesis for this particular neural activity? This type of workflow is illustrated in Figure 1, which first shows a broader network implicated for some type of composition-relevant operations, dubbed the “combinatory network,” and then depicts the extent to which the various network nodes are engaged for more specific computations. In this spirit, our chapter will review the extent to which the brain basis of composition is understood within current neurobiology of language, with a focus on the ability of extant results to rule out specific functional hypotheses about activity associated with composition.

In addition to multiple correlated computations, the neural signals elicited by composition also reflect online structure building computations that conform to the *grammar of some particular language*. Thus, to correctly model this process, we would need not only the right description of the online computational algorithm – i.e., exactly the steps by which structure gets built – but also the right description of the grammar – i.e., a correct model of the representations that get built. Since both of these are in a sense the “end questions” of entire fields within linguistics, we can never assume that our answers to them are correct; our answers will always be work in progress. In practice, our understanding has progressed along two fronts: (i) by examining the effects of contrasts that differ in composition under any theory and (ii) by comparing the ways in which brain data correlate with different combinations of grammars and algorithms. These two approaches are discussed in sections 2 and 3 below, respectively.

2. COMPOSITION IN CONTROLLED EXPERIMENTS

2.1 Sentence vs. List: Extracting the combinatorial network

When stimuli that engage composition in language are contrasted with maximally similar ones that do not, what brain activity is affected? A large literature has addressed this question, yielding descriptions of what could be called the “combinatorial network.” The most common experimental paradigm within this literature contrasts well-formed sentences with random unstructured lists of words. Starting with the classic findings of Mazoyer et al. (1993) and Stowe et al. (1998), the most replicated result from this contrast is that sentences elicit higher activation than lists in the temporal poles, often bilaterally (Humphries et al., 2006; Rogalsky and Hickok, 2009). Other than this finding, the results of these studies have been variable. This is not surprising, given that different studies have used different sentence materials, and thus, to the extent that different sentences may require different interpretive processes, the only consistency we can expect across all studies is activity reflecting processes that are shared across most sentences. However, broadly speaking, the sentence vs. list literature has also identified the posterior temporal lobe (e.g. Friederici et al., 2000; Pallier et al., 2011; Snijders et al., 2009; Vandenberghe et al., 2002) and the inferior frontal gyrus (Pallier et al., 2011; Snijders et al., 2009), as relatively common loci of increased activity for sentence stimuli. Finally, some evidence also exists for the middle temporal gyrus (Brennan & Pytkänen, 2012; Pallier et al., 2011), the temporo-parietal junction (Pallier et al., 2011) and medial parts of ventral prefrontal-cortex (Brennan & Pytkänen, 2012) as part of this combinatorial network.

How do all these regions contribute to composing words into complex structures? At the outset, numerous hypotheses are capable of explaining each of the activations: they could reflect any aspect of syntax, semantics or referential processing. Our understanding of the specific contributions of most of these nodes is still nascent, but in what follows we summarize the extent to which the hypothesis space has been narrowed down. In particular, we will focus on two questions: (i) which of these activations survive if we simplify the stimulus to just one instance of composition and (ii) to what extent can syntactic or semantic computations be ruled out for any of these activities?

2.2 Two word phrases: Identifying the minimal composition network for comprehension and production

One way to rule out large classes of hypotheses for the sentence vs. list results is to get rid of the sentences and ask what activity is elicited by the composition of the smallest possible structures. Results for such minimal combinations would not be reflective of any processes pertaining to the sentence level – such as the formation of long distance dependencies (as in, for example, relative clauses in which the object of a verb is expressed outside its canonical object position: *the ball that the dog ate..*), agreement, resolution of co-reference, and so forth – but rather, would be promising correlates of building a single phrase. From such a result, one could scale up to assess what additional activity is recruited for the construction of larger structures.

This approach was taken in Bemis and Pykkänen (2011), who created a type of two-word version of the sentence vs. list paradigm to test what neural activity would be sensitive to the presence of just a single step of composition. In this paradigm, magnetoencephalography (MEG)

activity is recorded while subjects comprehend pairs of words that either form a phrase or not, as well as single word controls. The words are followed by pictures that either match or mismatch the verbal description; indicating the match or mismatch served as a comprehension task. The initial studies used just color-object combinations (*white lamp*, *red boat*, etc.), which made it possible to run parallel production experiments on the same expressions (Pykkänen et al., 2014). In the production versions, the same pictures were used as a production prompt, with subjects naming the pictures as “green cup” and so forth. MEG activity was analyzed in response to the second word (e.g., *cup* in *blue cup*) in the comprehension versions and in response to the picture in the production versions, where the temporal resolution of MEG allows one to capture the planning of speech production prior to the onset of motion artifacts. Across a series of studies, a stable pattern was observed in the left anterior temporal lobe (LATL) and the ventromedial prefrontal cortex (vmPFC), showing increased activity for phrases relative to for non-combinatory controls both in comprehension and production. Unsurprisingly, the time courses of the effects were different between comprehension and production. In comprehension, the LATL increase elicited by composition was early (200-250ms after the second word) and the vmPFC increase late (~400ms) whereas in production, both effects of phrasal structure onset early and relatively in parallel, starting around 200ms.

Crosslinguistic generality for this pattern has been demonstrated for Arabic (Westerlund et al., 2015) and American Sign Language (Blanco-Elorrieta et al., 2018), and the LATL effects generalize to predicate-argument configurations as well, such as *eats meat* (Westerlund et al., 2015). This generality is compatible with the hypothesis that the LATL reflects syntactic aspects of composition, but as we discuss in the next section, subsequent research has compellingly ruled out this account.

2.3. A thorny question: Syntax vs. semantics

Manipulations of the combinatorial structure of sentences are most often described as manipulations of syntax. But in almost all cases, changing structure also changes the semantic combinatorics. What do we know about the processes by which syntactic structures get built, versus the combinatorial steps that build complex meanings?

2.3.1. Isolating syntax with “semantics-free” stimuli: Jabberwocky and artificial grammar studies

Since semantic confounds make the pure study of syntax so difficult with natural language, much research on this question has given up on natural language as stimuli. Instead, many groups have chosen so called Jabberwocky sentences as “semantics-free” expressions (inspired by Lewis Carroll’s Jabberwocky poem). In these stimuli, the grammatical elements of sentences are preserved, but the conceptual/lexical items are replaced with pseudowords, as in *the solims on a sonting grilloted a yome and a sovir* (Humphries et al., 2006). Although these expressions are standardly labeled as “minus SEMANTICS”, most researchers would acknowledge that these stimuli are not void of meaning – they have a rich relational structure including tense, argument structure, anaphoric relationships and so forth even while lacking actual conceptual labels. Further, exactly what subjects do when comprehending such expressions is unclear: we do not have theoretical or psycholinguistic models that speak to this. Nevertheless, comparing such stimuli to “jabberwocky lists” – as in *rooned the sif inot lilf and the foig aurene to* (Humphries et al., 2006) – has been a common approach to the study of syntactic composition. However, given

that what is missing from Jabberwocky sentences is specifically the conceptual labels, as opposed to *all* semantics, the hypothesis-killing power of these stimuli is in fact rather narrow: increases for Jabberwocky sentences must not reflect the composition of complex conceptual content (since the conceptual labels are missing). But, subjects' strategies for dealing with these unnatural stimuli could substantially complicate things.

Indeed, no consistent result emerges from this literature: While many classic studies on Jabberwocky sentences found them to elicit increased activation in anterior temporal cortex (Mazoyer et al., 1993; Friederici et al., 2000; Humphries et al., 2006), more recent studies have found left ATL increases only in the presence of natural semantics (Pallier et al., 2011; Matchin et al., 2017). We are not aware of aims to reconcile these findings—understanding this pattern would require a within-subjects elicitation of the difference to begin with, and then careful hypothesizing about possible contrasts in the stimuli and tasks.

But, the Jabberwocky literature has given rise to another candidate for purely structural processing: the pars opercularis (BA 44). For example, increased BA 44 activity was found both for jabberwocky determiner-phrases (such as “Diese Flirk” in German, Zaccarella & Friederici, 2015) and for natural ones (*this ship* vs. *ship*, Schell et al., 2017), as compared to control stimuli. However, under the hypothesis that these regions build syntactic structure, their insensitivity to the sentence vs. list contrast and minimal composition manipulations in numerous studies remains a puzzle. As a possible solution, Zaccarella et al. (2017) propose that the list conditions of many sentence vs. list studies may in fact have elicited combinatory processing due to the mixing of content and function words in the lists (e.g., *her eyes during close the she ceremony*). In a meta-analysis, they show that in studies employing only content words or only function words in the lists, BA 44 increases for sentences are systematically elicited. This highlights the

importance of experimental design and the challenges in choosing appropriate non-combinatory control conditions.

Posterior superior temporal regions have also emerged as possibly relevant in Jabberwocky studies; for example, both Pallier et al. (2011) and Matchin et al. (2017) found increased pSTS activity for sentences over lists but while in the Pallier study this effect was shared by Jabberwocky stimuli, in the Matchin one it was not. One account links activity in this region with thematic role assignment (Frankland & Greene, 2015), but numerous specific functions are compatible with existing results. All this suggests that Jabberwocky manipulations may not be a fruitful way to approach the syntax vs. semantics question.

Another literature that has moved away from natural language in order to isolate syntactic processing has used artificial grammars, containing no natural language words at all. This literature also points towards an involvement of BAs 44 and 45 in the processing of artificial grammars (Bahlmann et al., 2008; Friederici et al., 2006; Petersson & Hagoort, 2012; Uddén et al., 2008). Ultimately, viable hypotheses for any activity sensitive to the combinatorial properties of language – and language-like artificial stimuli – must explain not only the positive results in the literature but also the negative ones. Given this, the hypothesis that regions within Broca's area perform hierarchical structure building is too general: these regions are often silent when syntactic structures are built (e.g., Mazoyer et al., 1993; Stowe et al. 1998; Humphries et al., 2006; Rogalsky and Hickok, 2009; Bemis & Pyllkanen, 2011, 2012, 2013; Pyllkanen et al., 2014; Blanco-Elorrieta et al., 2018). If neural activity reflects structure building, it in principle should engage every time words are combined into larger structures. Currently, this behavior has not been demonstrated for either the posterior temporal or inferior frontal areas discussed above.

2.3.2. Possible oscillatory reflexes of syntactic phrase building

One promising candidate that does show the type of generality required for a correlate of syntactic composition has been recently identified in the frequency domain: Ding et al. (2016) presented subjects with words, phrases and sentences at predictable rates, with no acoustic cues to structure, and demonstrated that despite the monotonicity of the stimulus, cortical activity as measured by MEG peaked exactly at the frequencies at which words, phrases and sentences occurred in their design, specifically at 4, 2 and 1Hz, respectively. Unpacking this initial result will require elaborate follow-up work that tests how these effects can be eliminated and whether they are semantically sensitive, but as things currently stand, the extant data are compatible with the hypothesis that the 2Hz peak in some sense reflects syntactic structure building (though the result is also equally compatible with the hypothesis that it reflects semantic composition).

2.3.3. Evidence for a non-syntactic role for the LATL and vmPFC

Given that composition is many correlated computations, each experiment on the brain basis of composition should have the potential to narrow down the hypothesis space in some informative way. The original minimal composition findings of Bemis & Pylkkänen (2011) revealed an early (~200ms) increase in the LATL and a later one in the vmPFC (~400ms) in response to a word that completes a two-word phrase. That study was informed by prior work that had showed the vmPFC to be sensitive to the semantic complexity of expressions even when their surface syntax was kept constant (reviewed in Pylkkänen et al., 2011), suggesting that the vmPFC activity elicited by the small phrases most likely reflects semantic aspects of composition.

Since purely semantic manipulations of composition had not engaged the LATL, Bemis & Pylkkänen (2011) conjectured that the LATL activity elicited by adjective-noun combinations might reflect the syntactic side of composition. This hypothesis was compatible with sentence vs. list findings and also with deficit/lesion data available at the time (Dronkers et al. 2004). However, subsequent research has compellingly ruled out any purely structural interpretation of the early LATL activity: it does not show the generality required for syntactic composition. For example, although number phrases such as “two boats” require the application of syntactic composition, number phrases do not appear to engage the LATL (e.g., Blanco-Elorrieta & Pylkkänen, 2016). This finding eliminates any straightforward account of the LATL in terms of syntactic composition, and instead points towards explanations that perhaps depend more on the conceptual content of the composing elements (cf., Baron & Osherson, 2011): while the modifier *blue* adds a feature to *boats*, the number term *two* enumerates the number of tokens in the boat-set, but adds no (obvious) conceptual content.

This general idea connects with prior literature on semantic dementia, which had shown that left ATL atrophy affects the specificity of an individual’s conceptual space, leading them to lose specific conceptual labels such as *poodle* and to resort to more general ones, such as *dog* or *thing*. ATL atrophy has not been linked with deficits in phrase-structure processing (Wilson et al. 2014). This invites the hypothesis that perhaps LATL effects of composition in fact reflect not composition but rather the increase in conceptual specificity created by the addition of an adjectival modifier. Subsequent studies have, however, shown that LATL amplitudes do not linearly increase as a function of conceptual specificity, blind to the single word vs. phrase distinction. Instead, when lexical frequency is controlled, single word specificity tends to not affect the LATL reliably and composition is required to elicit an increase (Westerlund &

Pylkkänen, 2014; Zhang & Pylkkänen, 2015). However, the size of the composition effect is modulated by conceptual specificity, with conceptually richer modifiers increasing the amplitudes of their head nouns more. Thus LATL activity at ~200-250ms appears to act as a specificity driven conceptual combiner. Additional findings have also shown that its function may be restricted to relatively simple, perhaps mostly intersective conceptual combinations (Ziegler & Pylkkänen, 2016; Poortman & Pylkkänen, 2016). This makes sense given its early timing: the activity may reflect an early concept combiner that applies whenever the input meanings have been sufficiently accessed, the idea being that more complex input meanings may not have been sufficiently accessed by 200ms (Pylkkänen, 2016). Thus, for the LATL, we take a purely syntactic explanation to be off the table, leaving on the table a much more nuanced conceptually based explanation (surely to be refined in the future).

2.4. Summary

Simple manipulations of sentence and phrase structure have yielded a description of the so-called “combinatory network,” covering much of the left temporal cortex (LATL, MTG, pSTS) and potentially extending to areas of prefrontal and temporoparietal cortex (vmPFC, LIFG, TPJ). However, our functional understanding of this network is still rudimentary. A sequence of studies on the LATL shows that activity sensitive to a contrast such as sentence vs. list may, in fact, reflect specificity modulated conceptual composition. The bigger picture is that a more detailed understanding of the functional contributions of the various nodes requires subtler contrasts in the stimulus materials, and systematic efforts to modulate the effects observed for

gross contrasts. Only this way can we understand the computational limits of each type of neural activity.

3. MODELING COMPOSITION FOR NATURALISTIC STIMULI

Simple linguistic comparisons allow researchers to target neural composition operations that must be present under any theory. One question we might ask is whether conclusions about localization and timing from such studies generalize to richer linguistic contexts. A second question is whether these neural signals offer insight into the representations and algorithms that are implemented within these neural circuits beyond the coarse-grain labels of syntax versus semantics. Both questions may be answered by leveraging computational models to characterize neural signals recorded from participants as they process natural, ecologically rich, linguistic stimuli.

3.1. Computational models of composition

A study by Brennan and colleagues (2012) illustrates the basic idea. They model composition word-by-word by counting the number of new phrases that have been completed after each word. For example, the sentence “Eleanor's sister pet the dog” includes the phrases, “Eleanor's sister,” “the dog,” and “pet the dog” as well as the trivial phrases made by each word and the entire sentence. Counting phrases that are completely processed at each word yields: *<eleanor's, 1>* *<sister, 2>* *<pet, 1>* *<the, 1>* *<dog, 4>*. This is a simple extension of Pallier et al., (2011) and

Bemis & Pylkkänen (2011), mentioned above, in that each phrase evokes composition operations.

Brennan et al. apply a broad-coverage account of phrase-structure (Marcus et al. 1993) to annotate every sentence of a selection of real-word language: a chapter from *Alice's Adventures in Wonderland*. Participants listened to this chapter while undergoing fMRI scanning.

Convolving word-by-word composition steps with the hemodynamic response function yields an estimator for a neural time-series that would be recorded from a brain region that was sensitive to the number of phrases completed word-by-word. A linear regression between this estimator and fMRI data revealed LATL activity that were positively correlated with the number of phrases being composed. Such a correlation was not seen in other regions of the combinatorial network, like the LIFG.

The same method can also be applied with electrophysiology (Frank et al., 2015; Brennan & Pylkkänen, 2017; Nelson et al. 2017). For example, Brennan & Pylkkänen (2017) compare the number of phrases posited incrementally with MEG data that were recorded during story-reading. Cross-correlating the number of completed phrases with the evoked signal revealed a significant effect in the LATL at 350 ms after word onset. Altogether, there is a promising alignment in terms of both localization and, to some extent, timing between results using rich natural stimuli and results from cases of minimal composition, reviewed above. Interestingly though, while LATL effects emerge at around ~200ms for simple phrases, they occur about 100ms later within a narrative (time-locked to word onset). Given the larger amount of top-down information present in sentences, the opposite could easily be true, i.e., effects could be faster in sentence contexts. Alternatively, the later timing could be due to the more complex meanings of

sentences as compared to small phrases. Careful study is needed to clarify the factors that may affect the timing of composition operations.

One assumption of the previous studies is that hemodynamic signals vary in proportion to the number of phrases composed. This proportionality is a linking hypothesis that connects the cognitive states that are computationally modeled to the neural signals that are experimentally measured. Other linking hypotheses between composition and brain signals can also be explored. For example, Henderson et al. (2016) use the linking hypothesis of surprisal. This quantity, which comes from information theory, reflects the (un)expectedness of a word given the preceding linguistic context (Hale, 2001). Using naturalistic story-reading and the same account of phrase-structure mentioned above, Henderson et al. report that LATL activity increases for unexpected words. Different linking hypotheses like surprisal and phrase-counting tap into different cognitive operations. Ongoing work aims to more clearly specify this mapping, for instance by probing how expectedness affects structure-building operations (Hale, 2014). Pairing computational models with naturalistic data is, in general, a highly flexible approach: by specifying different kinds of linking hypotheses, or different kinds of linguistic information in the models, multiple aspects of composition can be studied simultaneously (Wehbe et al., 2014).

3.2. Comparing alternative accounts of composition

Computational models of naturalistic processing also furnish a way to rigorously compare alternative hypotheses of composition operations. The basic logic is to specify a family of models that share the same linking hypothesis but differ on some parameter of interest, say the

grammar rules used to define phrases. Such models can be ranked in terms of their fit to a target neural signal.

Nelson et al. (2017) apply this model comparison logic to probe how predictively the brain composes sentences. A non-predictive “bottom-up” strategy holds that phrases are only postulated once all of the words that belong to that phrase have already been encountered. A more predictive strategy is also plausible. For example, upon encountering the word “the” at the beginning of a sentence, the composition system might reasonably predict that the next word will be a noun. Subsequently encountering a noun, like “cat”, verifies this prediction and also licenses a new prediction that a verb phrase will come next. This second approach is the “left-corner” strategy. The most eager “top-down” strategy composes syntactic structure prior to encountering any of the words that belong to a particular phrase. Under this strategy, for example, the composition system would postulate a noun phrase and verb phrase prior to encountering any words in a new sentence.

Nelson et al. (2017) compare these three different strategies by correlating each with electrocorticography (ECoG) signals recorded from patients undergoing monitoring for severe epilepsy. Patients read sentences with a variety of grammatical structures and answered comprehension questions. The goodness-of-fit between the three models shows that signals recorded from left anterior temporal, left inferior frontal, and midline frontal recording sites, are better fit when phrases are incrementally postulated according to either a left-corner or a bottom-up strategy but do not fit well with a top-down strategy.

The strategy for composing words together is just one of many parameters needed for a fully explicit computational account. Another parameter is the grammar that guides licit composition. Brennan et al. (2016) probe two different grammars using fMRI. One is a simple phrase-structure

grammar that captures major constituent structures in English (Marcus et al., 1993). A second grammar, adapted from a generative syntax textbook (Sportiche et al., 2013), incorporates more abstract rules to account for linguistic regularities and dependencies within English and also in other languages. Model comparison reveals better fits to LATL and left posterior temporal lobe fMRI signals using the more abstract grammar. A key take-away is that composition operations unfold differently under alternative syntactic theories of what is being composed and neural signals that reflect composition are sensitive to such differences.

3.3. Summary

Computational models of incremental sentence comprehension offer a tool to tease out neural signal related to composition when participants process rich natural stimuli. Studies have applied these models to hemodynamic and electrophysiological data collected under a variety of tasks. Thus far, results show general agreement with those from minimal linguistic comparisons: the LATL and also the posterior temporal lobes stand out as sensitive to composition (Brennan et al., 2012; Brennan et al., 2016; Henderson et al., 2016; Nelson et al., 2017). One question that remains open is whether a continuous predictor of conceptual combination (created based on findings from the two-word phrase literature) would explain LATL activity better than the grammatical predictors used in extant computational modeling work.

Because computational models are explicit, the fit between any particular model and neural signals can be quantified. Comparing different models in terms of their fit offers a novel way to test claims about the functions that are instantiated across the composition network, such as the nature of the syntactic representations that guide composition (Brennan et al., 2016; Nelson et

al., 2017; Frank et al., 2015). But, the explicitness required by this approach is also a limitation. To offer quantitative predictions, one must commit to a number of specific assumptions about composition, including the syntactic and semantic grammar, the composition strategy, how multiple possibilities are navigated when the input is ambiguous, and other parameters. The literature has just begun to systematically explore the hypothesis space defined by these many parameters.

4. CONCLUSION

Our understanding of the neural basis of sentence processing must rely on a solid characterization of the basic processes by which the brain composes complex structure and meaning from elementary building blocks. Current literature teaches us that the left anterior temporal cortex is the most consistent locus of combinatory effects across many methodologies, likely contributing an early process of conceptual combination. How more grammatically based composition proceeds is still a mystery, though studies using artificial stimuli that perhaps enhance the subject's awareness of grammatical processing point towards a role for the left inferior frontal cortex (e.g, Zaccarella & Friederici, 2015; Petersson & Hagoort, 2012). Results from cortical stimulation mapping have also shown that disrupting processing in the left inferior frontal gyrus causes errors in grammatical encoding (Chang, Kurteff & Wilson, 2017), suggesting that the LIFG does in some way interface with grammatical knowledge. But simple studies varying composition and model comparisons using naturalistic stimuli have failed to support a role for left inferior frontal cortex in structure building (Bemis & Pylkkanen, 2011, 2012; Brennan et al., 2012, 2016), creating a tension in the literature that is yet to be resolved. A

second area of tension concerns the potential role of posterior temporal regions in composition, as they show mixed sensitivity to simple manipulations of composition. Overall, research using model comparisons on narrative processing has shown that activity within the combinatorial network is generally better explained by parsing strategies that take lexical content into account, at least to some extent, and by grammars that are more abstract as opposed to simpler. Going forward, the naturalistic methodology should include systematic testing of hypotheses arising from more tightly controlled experiments, both to test the extent to which those findings “scale up,” and to help us better connect the two bodies of literature.

Figure and legend:

	Combinatory network shows more activity for sentences compared to word lists	Simple composition network shows more activity for simple phrases compared to words	Sensitive to semantics in syntactically parallel expressions	Activity better fit by hierarchical models compared to sequence-based models
LATL 200-300 ms	☒ ☒ ☒	☒ ☒ ☒	☒ ☒ ☒	☒
LIFG 300-500 ms	☒ ☒ ☐	☒ ☐ ☐	☐ ☐	☐
LPTL/AG 200-400 ms	☒ ☒ ☐	☒ ☐ ☐	☐ ☐	☒
vmPFC 400-500 ms	☒ ☐ ☐	☒ ☐	☒ ☒	(no data)

An informal depiction of our current understanding of brain regions supporting composition and the extent to which the functional roles of individual network nodes are understood. Lack of understanding can result either from a lack studies or from a lack of generalizations across studies. Here, the number of boxes in each cell represents the general quantity of studies addressing the role of the region and the checks inside the boxes the amount of positive evidence for the generalization in the first row. Timing estimates primarily reflect results from MEG studies comparing sentence vs. list activation (e.g., Brennan & Pykkänen, 2012) or phrase vs. word activation (e.g., Bemis & Pykkänen, 2011). The table does not separate results according to method; thus for example, positive results for the LIFG come primarily from fMRI (and are thus ambiguous as regards timing) and ones for the vmPFC from MEG. Connecting separate findings from different methods is a major goal for future research. In all, the only network node showing a high degree of consistency across the literature is the LATL.

REFERENCES

- Bahlmann J, Schubotz RI, Friederici AD (2008) Hierarchical artificial grammar processing engages Broca's area. *Neuroimage* 42:525–534.
- Bemis, D. K. & Pylkkänen, L. (2011). Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience*, 31(8), 2801–2814.
- Blanco-Elorrieta, E., Kastner, I., Emmorey, K. & Pylkkänen, L. (2018). Shared neural correlates for building phrases in signed and spoken language. *Scientific Reports*. DOI:10.1038/s41598-018-23915-0.
- Blanco-Elorrieta, E., & Pylkkänen, L. (2016). Composition of complex numbers: delineating the computational role of the left anterior temporal lobe. *NeuroImage*, 124, 194-203.
- Brennan, J. R., Nir, Y., Hasson, U., Malach, R., Heeger, D. J., & Pylkkänen, L. (2012). Syntactic structure building in the anterior temporal lobe during natural story listening. *Brain and Language*, 120, 163–173.
- Brennan, J. R. & Pylkkänen, L. (2017). MEG evidence for incremental sentence composition in the anterior temporal lobe. *Cognitive Science*, 41(S6), 1515–1531.

Brennan, J. R., Stabler, E. P., Van Wagenen, S. E., Luh, W.-M., & Hale, J. T. (2016). Abstract linguistic structure correlates with temporal activity during naturalistic comprehension. *Brain and Language*, 157-158, 81–94.

Chang, E. F., Kurteff, G., & Wilson, S. M. (2018). Selective interference with syntactic encoding during sentence production by direct electrocortical stimulation of the inferior frontal gyrus. *Journal of cognitive neuroscience*, 30(3), 411-420.

Ding, N., Melloni, L., Zhang, H., Tian, X., & Poeppel, D. (2016). Cortical tracking of hierarchical linguistic structures in connected speech. *Nature neuroscience*, 19(1), 158.

Dronkers, N. F., Wilkins, D. P., Van Valin, R. D., Redfern, B. B., & Jaeger, J. J. (2004). Lesion analysis of the brain areas involved in language comprehension: Towards a new functional anatomy of language. *Cognition*, 92(1-2), 145–177.

Frank, S. L., Otten, L. J., Galli, G., & Vigliocco, G. (2015). The ERP response to the amount of information conveyed by words in sentences. *Brain and Language*, 140(0), 1–11.

Frankland, S. M. & Greene, J. D. (2015). An architecture for encoding sentence meaning in left mid-superior temporal cortex. *Proceedings of the National Academy of Sciences U S A*, 112(37), 11732–7.

Friederici AD, Bahlmann J, Heim S, Schubotz RI, Anwender A (2006) The brain differentiates

human and non-human grammars: Functional localization and structural connectivity. *Proc Natl Acad Sci USA* 103:2458–2463.

Friederici, A. D., Meyer, M., & von Cramon, D. Y. (2000). Auditory language comprehension: an event-related fMRI study on the processing of syntactic and lexical information. *Brain and language*, 74(2), 289-300.

Hale, J. T. (2001). A probabilistic Earley parser as a psycholinguistic model. In *North American Chapter of The Association For Computational Linguistics* (pp. 1–8).: Association for Computational Linguistics Morristown, NJ, USA.

Hale, J. T. (2014). *Automaton Theories of Human Sentence Comprehension*. CSLI Publications.

Henderson, J. M., Choi, W., Lowder, M. W., & Ferreira, F. (2016). Language structure in the brain: A fixation-related fMRI study of syntactic surprisal in reading. *NeuroImage*, 132, 293–300.

Marcus, M., Marcinkiewicz, M., & Santorini, B. (1993). Building a large annotated corpus of English: The Penn treebank. *Computational linguistics*, 19(2), 313–330.

Matchin, W., Hammerly, C., & Lau, E. F. (2017). The role of the IFG and pSTS in syntactic prediction: Evidence from a parametric study of hierarchical structure in fMRI. *Cortex*, 88, 106–123.

Nelson, M. J., El Karoui, I., Giber, K., Yang, X., Cohen, L., Koopman, H., Cash, S. S., Naccache, L., Hale, J. T., Pallier, C., & Dehaene, S. (2017). Neurophysiological dynamics of phrase- structure building during sentence processing. *Proc Natl Acad Sci U S A*, 114(18), E3669–E3678.

Pallier, C., Devauchelle, A.-D., & Dehaene, S. (2011). Cortical representation of the constituent structure of sentences. *Proc Natl Acad Sci U S A*, 108(6), 2522–2527.

Petersson, K. M., Folia, V., & Hagoort, P. (2012). What artificial grammar learning reveals about the neurobiology of syntax. *Brain and language*, 120(2), 83-95.

Petersson, K. M., & Hagoort, P. (2012). The neurobiology of syntax: beyond string sets. *Phil. Trans. R. Soc. B*, 367(1598), 1971-1983.

Poortman, E. B., & Pylkkänen, L. (2016). Adjective conjunction as a window into the LATL's contribution to conceptual combination. *Brain and language*, 160, 50-60.

Pylkkänen, L. (2015). Composition of complex meaning: Interdisciplinary perspectives on the left anterior temporal lobe. In *Neurobiology of language* (pp. 621-631).

Pylkkänen, L., Bemis, D. K., & Blanco Elorrieta, E. (2014). Building phrases in language production: an meg study of simple composition. *Cognition*, 133(2), 371–84.

Pylkkänen, L., Brennan, J., & Bemis, D. K. (2011). Grounding the cognitive neuroscience of semantics in linguistic theory. *Language and Cognitive Processes*, 26(9), 1317-1337.

Sportiche, D., Koopman, H., & Stabler, E. (2013). *An Introduction to Syntactic Analysis and Theory*. Wiley-Blackwell, West Sussex.

Uddén, J., Folia, V., Forkstam, C., Ingvar, M., Fernández, G., Overeem, S., ... & Petersson, K. M. (2008). The inferior frontal cortex in artificial syntax processing: An rTMS study. *Brain research*, 1224, 69-78.

Wehbe, L., Murphy, B., Talukdar, P., Fyshe, A., Ramdas, A., & Mitchell, T. (2014). Simultaneously uncovering the patterns of brain regions involved in different story reading subprocesses. *PLoS One*, 9(11), e112575.

Westerlund, M., Kastner, I., Al Kaabi, M., & Pylkkänen, L. (2015). The LATL as locus of composition: MEG evidence from English and Arabic. *Brain and language*, 141, 124-134.

Westerlund, M., & Pylkkänen, L. (2014). The role of the left anterior temporal lobe in semantic composition vs. semantic memory. *Neuropsychologia*, 57, 59-70.

Wilson, S. M., DeMarco, A. T., Henry, M. L., Gesierich, B., Babiak, M., Mandelli, M. L., Miller, B. L., & Gorno-Tempini, M. L. (2014). What role does the anterior temporal lobe play in

sentence-level processing? Neural correlates of syntactic processing in semantic variant primary progressive aphasia. *Journal of Cognitive Neuroscience*, 26(5), 970–85.

Zaccarella, E., & Friederici, A. D. (2015). Merge in the human brain: A sub-region based functional investigation in the left pars opercularis. *Frontiers in psychology*, 6, 1818.

Zaccarella, E., Schell, M., & Friederici, A. D. (2017). Reviewing the functional basis of the syntactic Merge mechanism for language: A coordinate-based activation likelihood estimation meta-analysis. *Neuroscience & Biobehavioral Reviews*, 80, 646-656.

Zhang, L., & Pylkkänen, L. (2015). The interplay of composition and concept specificity in the left anterior temporal lobe: An MEG study. *NeuroImage*, 111, 228-240.

Ziegler, J., & Pylkkänen, L. (2016). Scalar adjectives and the temporal unfolding of semantic composition: An MEG investigation. *Neuropsychologia*, 89, 161-171.