

FEATURES AND FUNCTIONS: DECOMPOSING THE NEURAL AND COGNITIVE BASES  
OF SEMANTIC COMPOSITION

Christine Boylan

A DISSERTATION

in

Psychology

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2015

Supervisor of Dissertation

Co-supervisor of Dissertation

---

Sharon L. Thompson-Schill  
Christopher H. Browne Distinguished Professor of Psychology

---

John C. Trueswell  
Professor of Psychology

Graduate Group Chairperson

---

Sara Jaffee  
Professor of Psychology

Dissertation Committee:

Russell A. Epstein, Professor of Psychology

Florian Schwarz, Assistant Professor of Linguistics

Sharon L. Thompson-Schill, Christopher H. Browne Distinguished Professor of Psychology

John C. Trueswell, Professor of Psychology

FEATURES AND FUNCTIONS: DECOMPOSING THE NEURAL AND COGNITIVE BASES  
OF SEMANTIC COMPOSITION COPYRIGHT

2015

Christine Emily Boylan

This work is licensed under the  
Creative Commons Attribution-  
NonCommercial-ShareAlike 3.0  
License

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/2.0/>

## ABSTRACT

### FEATURES AND FUNCTIONS: DECOMPOSING THE NEURAL AND COGNITIVE BASES OF SEMANTIC COMPOSITION

Christine Boylan

Sharon L. Thompson-Schill

John C. Trueswell

In this dissertation, I present a suite of studies investigating the neural and cognitive bases of semantic composition. First, I motivate why a theory of semantic combinatorics is a fundamental desideratum of the cognitive neuroscience of language. I then introduce a possible typology of semantic composition: one which involves contrasting feature-based composition with function-based composition. Having outlined several different ways we might operationalize such a distinction, I proceed to detail two studies using univariate and multivariate fMRI measures, each examining different dichotomies along which the feature-vs.-function distinction might cleave. I demonstrate evidence that activity in the angular gyrus indexes certain kinds of function-/relation-based semantic operations and may be involved in processing event semantics. These results provide the first targeted comparison of feature- and function-based semantic composition, particularly in the brain, and delineate what proves to be a productive typology of semantic combinatorial operations. The final study investigates a different question regarding semantic composition: namely, how automatic is the interpretation of plural events, and what information does the processor use when committing to either a

distributive plural event (comprising separate events) or a collective plural event (consisting of a single joint event).

## TABLE OF CONTENTS

<b>ABSTRACT .....</b>	<b>III</b>
<b>LIST OF TABLES .....</b>	<b>IX</b>
<b>LIST OF ILLUSTRATIONS .....</b>	<b>X</b>
<b>I. INTRODUCTION.....</b>	<b>1</b>
<b>WHY CONCERN OURSELVES WITH COMPOSITIONALITY? .....</b>	<b>1</b>
<b>ATTRIBUTIVE VS. RELATIONAL COMBINATION .....</b>	<b>2</b>
<b>ADJUNCTS VS. ARGUMENTS, AND WHETHER VERB SEMANTICS IS PRIVILEGED IN THE ANGULAR     GYRUS.....</b>	<b>3</b>
<b>COLLECTIVE VS. DISTRIBUTIVE INTERPRETATIONS OF PLURAL SETS.....</b>	<b>5</b>
<b>AIMS OF THE DISSERTATION .....</b>	<b>8</b>
<b>II. RELATIONAL VS. ATTRIBUTIVE INTERPRETATION OF NOMINAL COMPOUNDS DIFFERENTIALLY ENGAGES ANGULAR GYRUS AND ANTERIOR TEMPORAL LOBE .....</b>	<b>10</b>
<b>INTRODUCTION.....</b>	<b>10</b>
<i>A tale of two hubs: the angular gyrus and anterior temporal lobe.....</i>	<i>13</i>
<i>Relational vs. attributive interpretation of nominal compounds.....</i>	<i>16</i>
<b>MATERIAL AND METHODS .....</b>	<b>22</b>
<i>Experimental Task and Design.....</i>	<i>25</i>
<i>Image acquisition.....</i>	<i>25</i>
<i>Analysis .....</i>	<i>26</i>

<b>RESULTS .....</b>	<b>28</b>
<i>Task-responsive voxels in anatomical ROIs .....</i>	<i>28</i>
<i>Combinatorial bias predicts activity in left and right AG .....</i>	<i>29</i>
<i>BOLD timecourse of attributive interpretation differs by ROI.....</i>	<i>31</i>
<b>DISCUSSION.....</b>	<b>34</b>
<i>Angular gyrus and thematic relations .....</i>	<i>35</i>
<i>Anterior temporal lobe and semantic composition.....</i>	<i>43</i>
<i>Limitations, future directions, and conclusions .....</i>	<i>45</i>
 <b>III. COMPOSITIONALITY AND THE ANGULAR GYRUS: A MULTI-VOXEL SIMILARITY</b>	
<b>ANALYSIS OF THE SEMANTIC COMPOSITION OF NOUNS AND VERBS .....</b>	<b>47</b>
<b>INTRODUCTION.....</b>	<b>47</b>
<i>Roles of left anterior temporal lobe and left angular gyrus in semantic composition .....</i>	<i>47</i>
<i>A feature vs. function dichotomy?.....</i>	<i>52</i>
<i>Multi-voxel pattern similarity as a window to various dimensions of compositionality .....</i>	<i>57</i>
<b>MATERIAL AND METHODS.....</b>	<b>59</b>
<i>Experimental Task and Design .....</i>	<i>67</i>
<i>Image acquisition and pre-processing.....</i>	<i>68</i>
<i>Analysis .....</i>	<i>69</i>
<b>RESULTS .....</b>	<b>73</b>
<i>Continuous similarity analyses between subjects' pairwise similarity rankings and neural</i> <i>similarity scores for pairs of two-word phrases.....</i>	<i>75</i>
<b>DISCUSSION.....</b>	<b>77</b>
<i>Angular gyrus and thematic relations .....</i>	<i>78</i>

<i>Semantic representations or semantic control?</i> .....	81
<i>Limitations, future directions, and conclusions</i> .....	83
<b>IV. A MATTER OF AMBIGUITY? USING EYE MOVEMENTS TO EXAMINE COLLECTIVE VS. DISTRIBUTIVE INTERPRETATIONS OF PLURAL SETS</b> .....	<b>90</b>
<b>INTRODUCTION</b> .....	90
<b>EXPERIMENT 1</b> .....	96
<i>Method</i> .....	96
<i>Results</i> .....	103
<i>Summary of Experiment 1</i> .....	109
<b>EXPERIMENT 2</b> .....	110
<i>Method</i> .....	110
<i>Results</i> .....	113
<i>Summary of Experiment 2</i> .....	118
<b>DISCUSSION</b> .....	119
<i>A matter of ambiguity</i> .....	119
<i>Conjoined NPs and plural determiners differ in selecting between collectivity and distributivity</i> .....	121
<b>V. GENERAL DISCUSSION</b> .....	<b>128</b>
<b>REVIEWING THE AIMS OF THE DISSERTATION</b> .....	128
<b>THE NEURAL BASES OF THE FEATURE-FUNCTION DIVISION</b> .....	129
<b>THE LATERALIZATION PUZZLE</b> .....	131
<b>THE AG AND VERB/EVENT SEMANTICS</b> .....	131
<b>CONCLUSION</b> .....	132

<b>APPENDIX.....</b>	<b>134</b>
<b>BIBLIOGRAPHY.....</b>	<b>141</b>



## LIST OF TABLES

Table 1. An example set of two-word compositional and non-compositional phrases sharing a given noun or verb.....	72
Table 2. Conjoined NP sentences.....	109
Table 3. Plural NP sentences.....	122

(see Appendix for additional tables)

## LIST OF ILLUSTRATIONS

Figure 1. Plot of relational-attributive combinatorial bias .....	24
Figure 2. Task-responsive voxel activity in left ATL and left AG .....	29
Figure 3. Beta coefficients from subject-wise parametric analysis of voxel activity in each ROI.....	31
Figure 4. Peristimulus FIR curves of attributive vs. relational activation in each ROI...33	
Figure 5. ROI boundaries of left ATL and left AG.....	71
Figure 6. Diagram of similarity analysis.....	72
Figure 7. Pairwise correlations between relevant pairs of word phrases in left AG.....	74
Figure 8. Two typical scene pairs for collective and distributive events.....	100
Figure 9. Plots of collective advantage for conjoined NP subjects.....	105
Figure 10. Plots of collective advantage for conjoined NP sentences in four time windows.....	106
Figure 11. Timecourse of sentences with distributive-biased verbs.....	109
Figure 12. Plots of collective advantage for plural NP subjects.....	114
Figure 13. Plots of collective advantage for plural NP sentences in four time windows.....	116

## I. INTRODUCTION

### **Why concern ourselves with compositionality?**

Language owes its infinite expressive capacity to our ability to take simple building blocks, such as words or concepts, and combine them into complex representations. How such conceptual combination might be realized in the brain, and whether formal accounts of syntactic and semantic composition are useful in characterizing the neural system underpinning conceptual combination, are still highly debated questions, and the emergence of compositional meaning from units such as morphemes, words, or concepts, is largely a mystery. However, understanding the engine of compositionality in the brain is a fundamental desideratum of any cognitive neuroscientific model of “concepts.”

This dissertation examines how the compositional system accomplishes conceptual and grammatical semantic combinatorics along a number of possible dimensions. We investigate several ways in which a natural dichotomy may exist between feature-/property-based composition, which concerns attributive modification operations *on* concepts, and relation-based composition, which concerns grammatical and thematic relations *between* concepts. Chapters 2 and 3 examine two different ways of characterizing this dichotomy that engage two regions of the brain we describe as combinatorial hubs: the angular gyrus (AG) and the anterior temporal lobe (ATL). Chapter 3 details the possibility that event semantics are particularly privileged in the putative “relation-based” semantic space supported by (left) AG. Chapter 4 investigates a different question regarding semantic composition: namely, how automatic is the

interpretation of plural events, and what information does the processor use when committing to either a distributive plural event (comprising separate events) or a collective plural event (consisting of a single joint event).

### **Attributive vs. relational combination**

For one intuitive explanation of the distinction between what we call feature-based and function-based compositional operations, we first look to the literature on conceptual combination. Here, the distinction is often cast in terms of “properties” instead of “features,” and “(thematic) relations” instead of “functions,” as in Wisniewski and Love’s (1998) discussion of two different interpretations of a noun-noun combination like *robin hawk*:

One kind of interpretation involved a *thematic relation* between the referents of the modifier and head concepts. For example, a *robin hawk* could mean “a hawk that preys on robins.” In *property* interpretations, people asserted that one or more properties of the modifier concept apply in some way to the head concept, as in “hawk with a red breast,” for *robin hawk*. Sometimes these interpretations refer to an entity which shares many properties of both constituents (e.g., a robin hawk could refer to a bird that is a cross between a robin and a hawk).

(Wisniewski & Love, 1998, p.178)

A thematic interpretation of *robin hawk* entails that each entity play a different functional role (or thematic role): the modifier refers to a robin as the object, or “patient,” of the action “to prey on,” while the head noun *hawk* refers the agent of “to prey on.” When *robin hawk* is interpreted as a property combination, however, the modifier *robin* refers to a property of the robin – e.g. its red breast – rather than the robin itself. The property-

combining interpretation involves no specification of functional roles or argument structure as does the relational interpretation.

This fundamental distinction between property- or feature-based combination on the one hand and functional, thematic relation-based composition on the other is ubiquitous across linguistic and psychological theory, albeit under various different guises. While chapter 2 uses functional magnetic resonance imaging (fMRI) to test how property-based, or “attributive,” and relational nominal compounds like *robin hawk* differentially engage ATL and AG, chapter 3 tests two more ways of operationalizing the “feature-function” distinction, while still examining ATL and AG as neural substrates.

### **Adjuncts vs. arguments, and whether verb semantics is privileged in the angular gyrus**

Chapter 3 uses fMRI multi-voxel pattern analysis to test two different hypotheses about the role of AG in compositional semantics. Our verb-centric hypothesis states that if AG is preferentially sensitive to event-denoting verbs and their thematic relations, then phrases that share a given verb (like *eats meat* and *eats quickly*) will evoke similar patterns of activation in AG. Phrase pairs that do not share verb semantics will not engage similar patterns of activation in AG. We find evidence that the verb is indeed privileged in AG semantic space.

Another hypothesis asks whether AG may be sensitive to relational information independent of the verb. The dichotomy between “feature” and “function” has a good deal of traction in more formal theories of semantics, particularly in the distinction between what linguists call “adjuncts” and “arguments,” respectively. The difference

between arguments and adjuncts, however, is more subtle than that between verbs and non-verbs, and so we unpack the terms a bit below.

Intuitively speaking, the verb is the central predicate of a sentence, and predicates, such as *eat*, must be able to take arguments (like *meat*) in order to participate in well-formed sentences ( Heim & Kratzer, 1998; Pylkkänen, Brennan, & Bemis, 2011). While the case of verbs and their direct objects is perhaps the most canonical example of so-called “argument saturation”, there are other types of function-argument relations: for instance, the composition of prepositional phrases (e.g. *with meat*, where the preposition *with* takes the argument *meat*).

Many well-formed linguistic expressions can be derived merely by iterative application of arg-type composition; however, there are other sorts of linguistic expressions it cannot derive. Take, for instance, the following sentence:

(2) John ate a sandwich quickly in the kitchen.

The semantics of the verb “*ate*” can be represented as the function  $ate(x,y)$ , where  $x$  will be the object argument (here “a sandwich”) and  $y$  the subject (here, “John”). Leaving aside the subject argument, we can see that applying argument-type composition will get us “ate a sandwich,” since it is relatively intuitive that “a sandwich” is an argument of “ate.” However, it is not so clear that we can derive the meaning of “ate a sandwich quickly in the kitchen” in terms of argument-type composition alone. In order for argument-type composition to accomplish this, all three underlined constituents in (2) above must be arguments of “ate.” However, the case in (3) below shows us that “a sandwich” does not behave like the other constituents do.

(3) John ate a sandwich quickly in the kitchen, and

- (a) Bill did (so), too.
- (b) \*Bill did (so) a salad.
- (c) Bill did (so) slowly.
- (d) Bill did (so) in the backyard.

It would seem that there is something about the status of the constituent “a sandwich” that cannot be left out or replaced (see 3b), while the other constituents “quickly” and “in the kitchen” can (3c, 3d). This is because the *only* argument of “ate” in 2-3 is “a sandwich,” whereas the other underlined constituents are modifier, or “adjunct,” phrases that modify the predicate “ate.” How do we derive the compositional meaning of these modifying phrases?

Adjunct-type composition allows us to combine such modifiers with their “heads,” where “ate” is the head in 2-3, and “in the kitchen” and “at midnight” are the modifiers. These optional elements serve to further specify features or properties of the head they modify (while “ate” is an event-denoting element, “quickly” and “in the kitchen” are *when-* and *where-* type properties of that event).

In Chapter 3, we discuss evidence that specifically event-denoting *verb* argument structure may selectively engage the AG, while it remains to be seen whether AG and/or ATL differentiate arguments and adjuncts more generally.

### **Collective vs. distributive interpretations of plural sets**

In Chapter 4, we turn to a different set of questions regarding events and semantic composition, this time examining the degree to which the decision to interpret plural events as either distributive (occurring separately) or collective (occurring jointly) might

be automatic and immediate, and also whether this decision is biased to be collective or distributive at certain points along the timecourse of interpretation. Examining sentences like “John and Bill carried a box,” where this sentence could entail John and Bill each carrying a separate box (a distributive plural) or John and Bill together carrying the same box (a collective plural), we ask whether the collective or distributive quality of the predicate is something the processor can leave underdetermined, or whether the grammar requires a commitment to one interpretation or the other early on in the sentence.

Using eye-tracking and the visual world paradigm, we tested a hypothesis (Frazier et al.'s (1999) Minimal Semantic Commitment hypothesis) that distinguishes between two types of mental representations the processor might entertain upon encountering an underdetermined semantic constituent (like a plural event that could be collective or distributive): if the representation is ambiguous, the processor will commit to one of its interpretations and later revise it if necessary, but if the representation is vague, the processor refrains from committing to an interpretation, leaving some features underdetermined until further information is made available. The crucial difference between these two proposed representation types is that an ambiguous representation necessitates a decision about an interpretation, while a vague representation tolerates unspecified features. This is somewhat similar to the notion that adjuncts have an “optional” status in the grammar, while arguments are elements of the grammar that are necessary in order for a parse to be successful, as described in the section above.

The following two sentences illustrate a case when a linguistic item can be said to tolerate “underspecification,” or vagueness:

(4) a. John ate.



b. John ate quickly.

The difference in meaning between (4a,b) resides in the adverb “quickly.” In sentence (4a), the manner in which John ate is left unspecified, and yet a reader will find sentence (4a) perfectly interpretable without knowing the manner in which John ate. On the other hand, if the sentence lacked information specifying other propositional content, including, for instance, the number (singular) or tense (past) of the verb, this sentence would be grammatically uninterpretable (forcing a case of ambiguity when the processor would have to commit to a number or tense even if uncertain). Thus, while some information about the verb must be determined in order to parse the sentence, other features ostensibly need not be, as in the case of the adverbial adjunct “quickly” (1b). Adjunct information, such as the manner in which John ate, is therefore characterized as vague, rather than ambiguous.

Given the prediction that an ambiguous item will prompt the processor to converge on one particular interpretation even in the absence of disambiguating information, chapter 4 seeks to test whether, and when, sentences underdetermined for collective/distributive plurality will nonetheless immediately converge on one interpretation. We recorded participants’ eye movements as they interrogated two scenes, one collective and one distributive, while listening to corresponding sentences either with or without the early disambiguating adverbs “together” (explicitly indexing a collective event) and “each” (indexing a distributive event). Experiment 1 queries the timecourse of interpreting sentences beginning with conjoined noun phrases (conjoined NPs), such as “John and Bill (each/together) are carrying a box,” while Experiment 2 uses the same scenes but with sentences starting with simple plural NPs (“The boys (each/together) are

carrying a box”). We find evidence that the collectivity/distributivity of NPs like “John and Bill” and “The boys” is a matter of ambiguity (a grammatically forced, necessary choice). However, whereas early fixations are biased towards collective scenes in conjoined NP *John and Bill* sentences, they are biased towards distributive scenes in *The boys* sentences. We discuss possible reasons for this difference in the final section of chapter 4.

### **Aims of the dissertation**

In this dissertation, I seek to demonstrate the following:

- i. Semantic compositionality can be characterized along a “feature-vs.-function” dichotomy in several possible dimensions (chapters 2 and 3)
- ii. The AG supports various aspects of function-/relation-based composition, among them verb-specific argument structure around events (chapter 3) and more general thematic relation-based composition (chapter 2).
- iii. Semantic decisions on the collective or distributive representation of plural events occur early on in sentence processing – at the verb phrase – even in the absence of disambiguating evidence. However, the bias of this decision is largely determined by the affordances of the subject noun phrase, and not just the verb phrase (chapter 4).

While we find some evidence for the distinction between feature-based composition and functional relation-based composition within left ATL itself (see Chapter 2), we find stronger evidence for at least a single dissociation of relation-based composition in AG. In chapter 5, we review the several ways we have operationalized such “relation-based”

composition, and we consolidate evidence that AG subserves certain of these relation-based operations. We close by reviewing evidence that AG may also be involved in the representation of plural sets, and we speculate that further study of AG may not only shed light on relation-based composition, but also on the distinction between distributive and collective plural sets.

## II. RELATIONAL VS. ATTRIBUTIVE INTERPRETATION OF NOMINAL COMPOUNDS DIFFERENTIALLY ENGAGES ANGULAR GYRUS AND ANTERIOR TEMPORAL LOBE

### **Introduction**

Language's infinite generative capacity allows us to produce utterances ranging from the prosaic, as in "Close the door," to the ridiculous, as in "Hold the newsreader's nose squarely, waiter, or friendly milk will countermand my trousers" (Stephen Fry, *A Bit of Fry and Laurie*). Less ridiculous, but no less novel, sentences are uttered every day, and the ability of a reader or listener to understand such novel sentences, the propositional meanings of which cannot be retrieved from memory, requires a compositional algorithm that takes word meanings and combines them in such a way as to produce a more complex meaning. The neural substrate of this compositional algorithm remains elusive. Earlier work sometimes considered this engine of composition a more or less undifferentiated mechanism working to combine elements at all levels of language, whether that be syntax, semantics, or phonology (Hagoort, 2005; *inter alia*). While it is entirely plausible that the brain co-opted a basic, domain-general combinatorial mechanism for language, and implements it across multiple brain regions, such a monolithic theory of linguistic composition has its limitations.

For instance, many approaches to the study of composition benefit from a clear distinction between syntactic and semantic composition. Numerous psycho- and neurolinguistic studies investigating the syntax-semantics interface have included so-called Jabberwocky phrases, in which nonsense words replace content words while

function words remain in place, as semantically vacuous (or at least impoverished) syntactic controls for linguistic phrases (e.g. “the mouse that eats our cheese” vs. “the mouse that eats our cheese” (example from Pallier, Devauchelle, & Dehaene, 2011)). These studies demonstrate that subjects parse Jabberwocky phrases into hierarchical constituents similar to their natural language counterparts, even without knowing what the phrase means. Studies of complement coercion also suggest that syntactic and semantic argument structures are not isomorphic: evidence from behavioral, eye-tracking, and electrophysiological measures demonstrate a processing cost where semantic material unexpressed in the syntax must be inserted in order to coerce a coherent argument structure; e.g. “The man began the book” is interpreted as “The man began [reading/writing] the book” via implicit insertion of some event information (Baggio, Choma, van Lambalgen, & Hagoort, 2010; Kuperberg, Choi, Cohn, Paczynski, & Jackendoff, 2009; Kuperberg, Sitnikova, & Lakshmanan, 2008; McElree, Pyllkkänen, Pickering, & Traxler, 2006; Pyllkkänen & McElree, 2006, 2007; Traxler, McElree, Williams, & Pickering, 2005).

In this study, we proceed one step further, and suggest that within the domain of semantic composition, there is evidence for a distinction between two basic combinatorial operations, even when syntax is held constant. Specifically, we investigate the case of noun-noun compounds, in which the syntax is always a modifier noun followed by a head noun (e.g. *mountain lake*, where the syntax dictates this is a lake (in the mountains), not a mountain (in a lake)). Noun-noun compounds are a particularly tractable case of minimal composition (Bemis & Pyllkkänen, 2011), isolating the instance of combination rather

than investigating compositionality in the context of multi-word phrases or sentence stimuli, where multiple types of combinatorial operations occur simultaneously or in quick succession. Noun-noun compounds are particularly interesting because the first noun – the modifier noun – can be either predicating/“attributive” (as in *zebra clam*, where *zebra* denotes the attribute “striped”) or non-predicating/“relational” (as in *mountain lake*, where “mountain” is not an attribute but an object bearing a spatial relation with “lake”). Predicating combinations can be paraphrased as “a [noun] that is [adjective],” such as *red ball* (“a ball that is red”). Non-predicating combinations cannot be paraphrased this way: e.g. *tennis ball* is not “a ball that is tennis,” but rather is “a ball for playing tennis” (Downing, 1977; Gagné & Shoben, 1997; Levi, 1978). Attributive noun-noun compounds are predicating in that they can be paraphrased as “a [head noun] that is [modifier noun]-like”, as in *zebra clam* – “a clam that is zebra-like” (“a clam that is striped”). Relational noun-noun compounds are more complex in that they are non-predicating, and derive their meaning from some extrinsic predicating relation (e.g. “a ball *for playing* tennis”) (Levi, 1978; Murphy, 1990).

We find evidence that relational and attributive interpretations of noun compounds differentially engage two regions of the brain otherwise broadly implicated in semantic composition: the angular gyrus (AG) and the anterior temporal lobe (ATL). Below, we discuss how the distinction between attributive and relational combination may shed light on the functional differences between these two putative neural “hubs” of semantic composition.

*A tale of two hubs: the angular gyrus and anterior temporal lobe*

Mounting evidence suggests ATL and AG are involved in semantic processing, generally, and in semantic composition, specifically; however, only recently has there been effort to characterize their division of labor. Both have been characterized as "semantic hubs," owing to functional and anatomical patterns that are consistent with multimodal convergence (Binder & Desai, 2011; Lambon Ralph, 2014; Patterson et al., 2007; Seghier, 2012). The ATL is uniquely situated at the end of a caudal-to-rostral stream of information processing feeding from primary sensory and motor areas and association cortex (Binder et al., 2009; Binder & Desai, 2011; Binney, Parker, & Lambon Ralph, 2012; Felleman & Van Essen, 1991). Moving anteriorly along the temporal lobe, one finds a caudal-to-rostral hierarchy emerge as neuronal responses are more tuned to complex stimuli and more invariant to low-level sensory variation; such a hierarchy has been established along both visual (Felleman & Van Essen, 1991) and auditory (Rauschecker & Scott, 2009) streams. This "graded convergence" may provide a mechanism both for attributive feature combination and, in the limit, for maximally invariant amodal, abstract conceptual representations. The culmination of this graded convergence up the temporal lobe (Rauschecker & Scott, 2009; Stringer & Rolls, 2002) is a basal rostral region of ATL shown to have very limited extra-temporal connectivity and high intra-temporal connectivity (Binney et al., 2012). Such neuroanatomical sequestration is a necessary condition for a region to be able to represent abstract, modality-invariant semantics. Thus, ATL is a prime candidate for attributive semantic composition.

In one of the first studies investigating the neural correlates of minimal two-word composition, Baron and colleagues (2010) found evidence from fMRI pattern analyses that the left ATL subserved the combination of concepts such that the superimposition of individual patterns of the simplex concepts YOUNG and MAN (as represented by various face stimuli) reliably predicted the activation pattern for the complex concept YOUNG MAN. Consistent with this finding, a magnetoencephalography (MEG) study of visually presented two-word phrases comparing nouns in minimal compositional contexts (*red boat*) with nouns in non-compositional contexts (in which a non-word letter string was concatenated with a real word, e.g. *xkq boat*) found increased composition-related activity in left ATL (Bemis & Pylkkänen, 2011). There is a growing body of functional and tractographic studies to suggest that the representational unit of property-based composition in left ATL may be multimodal sensorimotor features, particularly *visual* concrete properties of object-concepts in more ventromedial regions of ATL, and possibly more abstract auditory-visual properties in more dorsolateral regions of ATL (Coutanche & Thompson-Schill, 2014; Hoffman, Binney, & Lambon Ralph, 2015), corroborating the notion of the left ATL as hub of the so-called ventral “what” pathway.

In addition to the ATL, researchers have also ascribed the label “semantic hub” to the AG, as it lies at the junction between temporal, parietal, and occipital lobes and thus receives a confluence of auditory, somatosensory, spatial, and visual inputs. Conceptual combination studies of the sort described above (Bemis & Pylkkänen, 2012) have demonstrated involvement of both left AG and left ATL, and several studies implicate bilateral AG in the contrast between well-formed sentences on the one hand and word lists, pseudowords, or scrambled sentences on the other (Bavelier et al., 1997; Bottini et



al., 1994; Humphries, Binder, Medler, & Liebenthal, 2007; Humphries, Binder, Medler, Liebenthal, & others, 2006). Left AG also shows greater activity for semantic violations vs. congruent well-formed sentences (Friederici, Rüschemeyer, Hahne, & Fiebach, 2003; Kang, Constable, Gore, & Avrutin, 1999; Kuperberg et al., 2000; Luke, Liu, Wai, Wan, & Tan, 2002; Ni et al., 2000) and for connected discourse vs. unrelated sentences (Fletcher et al., 1995; Homae, Yahata, & Sakai, 2003; Xu, Kemeny, Park, Frattali, & Braun, 2005). This broad profile of effects has led some to suggest that the AG may play a potentially domain-general role in semantic information integration structured around events.

Not all studies investigating conceptual combination find activation in both left ATL and bilateral AG. Of those stimuli that elicit differential activity in AG but not in left ATL, one finds that the type of composition may more often be based on thematic relations rather than attributive combination. Graves et al. (2010) compared familiar meaningful noun-noun compounds, such as *lake house*, with reversed phrases, such as *house lake*, the meanings of which were not obvious; they found that right AG, along with other right-lateralized temporoparietal areas, showed greater activation for processing the more obviously combinatorial phrases. Interestingly, the authors noted that most of their noun-noun stimuli were interpreted as denoting thematic relations between head and modifier nouns (see below for further explanation); that is, most compounds consisted of nouns participating in some spatial relation (as in “a house *on a lake*”) or event-based relation rather than picking out an attribute of the modifier noun. It is likely that these stimuli were probing semantic thematic relations in particular rather than combinatorial semantics in general.

### *Relational vs. attributive interpretation of nominal compounds*

In order to further distinguish between property-based associations and relation-based associations between concepts, consider the following nominal compound: *robin hawk*. Wisniewski (1996) found that people's interpretations of a novel compound of this sort could be characterized in one of two ways. Some individuals applied a property of the concept "robin," such as a red breast, to the head noun "hawk," to arrive at an interpretation like "a red-breasted hawk." Others found a thematic relation between the two birds, noting that a hawk might hunt a robin, and interpreted "robin hawk" as "a hawk that preys on robins." In the first type of interpretation, "robin" indicated some attribute or feature commensurate with the head noun "hawk," while in the second type of interpretation, the modifier noun "robin" was not broken down into features, but rather participated in a thematic relation with the head noun "hawk."

It is worth noting here that the terminology "predicating and non-predicating" is perhaps more precise than the terms "relational" and "attributive," which bear the misfortune of being both very common and denoting very different qualities depending on the theoretical framework. (For instance, there are certain non-deverbal "relational" nouns, like *sister*, *boss*, *edge*, *height*, etc., that seem to take implicit semantic arguments (Partee & Borschev, 2003), but we do not consider such relational nouns here.) However, "predicating and non-predicating" are terms that perhaps imply an overly syntactic typology for a phenomenon we argue arises from conceptual-semantic affordances; that is, *robin hawk* is interpretable as "a hawk that hunts robins" only because of the association between *hawk* and *hunting* (*robin chicken* would be unlikely to

invite such a predator-prey relation). Therefore, we will refer to these nominal compound types as “relational” and “attributive.”

Earlier literature assumed that the thematic relation was the primary means of interpreting nominal compounds: the consensus was that only after failing to find a plausible thematic relation binding the modifier and head nouns did people derive a property-based interpretation (Downing, 1977; Gagné & Shoben, 1997; Shoben & Gagné, 1997; Wisniewski & Gentner, 1991). Later proposals recognized attributive interpretations as somewhat distinct, but still considered property-based combinations too infrequent to be considered a different process; while Wisniewski & Love (1998) reported that attributive interpretations accounted for 29% of their nominal compounds, other samples reported attributives occurring as little as 1% of the time in corpora (Downing, 1977; Gagné, 2000; Warren, 1978). Parsimony dictated that attributive interpretations were simply another kind of relation, namely a resemblance relation (where *zebra clam* is merely a clam that resembles a zebra), and a single-process model prevailed (Costello & Keane, 2000; Gagné, 2000). One prominent formulation of single-process conceptual combination is the Competition Among Relations in Nominals (CARIN) theory (Gagné & Shoben, 1997). Under this account, the modifier noun (*zebra* in *zebra clam*, or *mountain* in *mountain lake*) is not incorporated into the head noun’s representation, but rather a relation (e.g. noun RESEMBLE modifier, or noun LOCATED modifier, respectively) is inserted that links the two concepts. Moreover, under CARIN, some relations will be considered prior to others, depending on the lexical items being combined (e.g. LOCATED is a more apt relation than ABOUT when construing the

compound *mountain lake*). This model stipulates that the RESEMBLE relation (that is, the attributive interpretation) is largely dispreferred.

However, further study found evidence for a categorical distinction between a relation-linking process like that described under CARIN, and another process, by which a property or attribute of the modifier is “transferred” to the head noun (attributive combination). Several studies found that the interpretation of ambiguous nominal compounds (such as *robin hawk*) could be manipulated based on priming the ambiguous item with relational or attributive compounds. Wisniewski & Love (1998) found that ambiguous targets were more likely to be interpreted attributively when following an attributive-biased compound, but more likely to be interpreted as relational when preceded by a relational prime. While this suggested that attributive and relational processes were actually distinct, dual processes, Gagné (2000) failed to replicate this effect using the same stimuli and procedure. Estes (2003) followed this work with another priming study, testing whether relational interpretations occurred serially prior to attributive interpretations, or whether these processes occurred in parallel. He found that both comprehension and reaction times were facilitated when target combinations matched prime combinations in attribution or relation. Moreover, Estes (2003) tested the CARIN model’s serial relation prediction that there should be an interaction between prime type and target type: if relation precedes attribution, then an attributive prime should interfere with interpretation of a relational target while a relational target should not hinder comprehension of an attributive target. That is, under the serial CARIN account, an attributive prime would involve additional (attributive) processing not

otherwise induced during a relational prime, and this extra processing would interfere with comprehension of a relational target. Estes (2003) did not find such an interaction, and interpreted these findings as inconsistent with a serial model like CARIN. However, this null result is hardly damning to a serial, single-process account, and evidence arbitrating between single and dual process models of relational and attributive compound interpretation remains equivocal.

### *Relational and attributive semantics in the brain*

The distinction between property- and relation-based semantic processes is not unique to work on conceptual combination, but also appears in the neuropsychological semantic memory literature. Here, one abiding question has been: Do relational and attributive conceptual combinations arise from neuroanatomically separable components of the semantics, or are they subsumed by the same combinatorial operation? If the latter, are these operations hierarchically disposed in some way: that is, do attempts at relational interpretation precede attributive interpretation, or vice-versa?

In the semantic memory literature, the distinction between so-called taxonomic and thematic associations serve as a parallel to what we have described as attributive and relational associations, respectively. Contrary to the conceptual combination literature, where models like CARIN suggest relational associations may precede attributive associations, some connectionist accounts appear to suggest that taxonomic (attributive) semantic knowledge may be logically prior to thematic (relation) knowledge. By some accounts, taxonomically defined concepts are the constituents of thematic relations: for instance, the thematic relation between *dog* and *bone* is supervenient on knowledge of

these concepts' properties (Lewis, Poeppel, & Murphy, 2015). Also, under many connectionist frameworks, thematic relations are reified into features, such that the concept *dog* might be linked not only to nodes for “furry,” “warm-blooded,” and “loyal,” but also nodes for explicitly relation-based facts, such as “bears live young,” “is led on a leash,” and “chews/buries bones” (Rogers & McClelland, 2004). This approach contrasts with a model whereby thematic relations constitute a qualitatively different level of representation from taxonomic features or properties, where the unit of representation is the event rather than the feature.

Consistent with an account whereby property-based associations take precedence over relation-based associations (contrary to CARIN), a recent MEG priming study found that activity in left ATL was only sensitive to property-based taxonomic associations, while both taxonomic and thematic associations predicted activity in the left temporoparietal junction (TPJ), inclusive of AG (Lewis et al., 2015). Note that this study supports a theory of *logical* precedence for taxonomic associations over thematic associations; it does not, however, provide evidence for a temporal precedence of one type of association over the other (as CARIN does for relational/thematic operations over attributive/taxonomic operations).

Another study, however, found evidence of a double dissociation between taxonomic and thematic semantic errors in left ATL and left AG, respectively (Schwartz et al., 2011). Speakers' semantic errors can be divided into either taxonomic category errors (that is, uttering an incorrect word, but one which has commensurate features, such as when “apple” is named as “pear”) or thematic relation errors (that is, uttering “dog”

when “bone” was intended, reflecting the thematic relation between “dog” and “bone”). Schwartz et al. (2011) examined the taxonomic and thematic errors produced by 86 post-stroke aphasics in a picture-naming task and conducted voxel-based lesion-symptom mapping (VLSM) on each error type separately (with shared variance between error types regressed out). Taxonomic errors were mapped to left ATL lesions, while thematic errors were localized to left AG. This double dissociation between ATL and AG supports the view that the ATL and AG support distinct semantic computations, corresponding to property-based and relation-based operations, respectively.

The current study examines the neural dissociation between property- and relation-based conceptual combination in order to discern (1) whether these processes are indeed functionally distinct and (2) whether they might allow us to better characterize the roles of AG and ATL in semantic combination. While the double dissociation of the sort reported in Schwartz et al. (2011), would be indicative of entirely dissociable systems, it is also possible that a common underlying semantic process derives both sorts of combination such that the distinction is moot. Midway between these two hypotheses is the possibility that these two types of combination are both functionally and neurally distinct, but recruit overlapping brain networks.

We find evidence for (1) a single dissociation in bilateral AG showing more task-responsive activity for relational compounds than attributive compounds, and (2) a timing difference in ATL, specifically an earlier ATL response to attributive compounds than relational compounds.. This serial temporal order is directly contrary to that put forth in

the CARIN model, but compatible with a different serial model: one whereby attributive/taxonomic operations are logically prior to thematic operations.

## **Material and methods**

### Participants

Eighteen subjects (eleven female) participated in this study. Subjects ranged in age from 18 to 42 years, and all were right-handed native speakers of English with normal or corrected-to-normal vision and no reported history of neurologic problems. Subjects gave written informed consent and were provided monetary compensation (\$20/hour) for their time. The human subjects review board at the University of Pennsylvania approved all experimental procedures.

### *Stimuli*

#### *Stimuli Design*

We drew our nominal compound stimuli from two studies investigating the effects of attributive- and relational-biased compounds (Estes, 2003; E. J. Wisniewski & Love, 1998). Of the stimuli used in these studies, we chose the 64 most attributive-biased and 64 most relational biased items according to a norming study we conducted via Amazon Mechanical Turk (Buhrmester, Kwang, & Gosling, 2011). Subjects (n=17) were asked to write their interpretations of each noun compound and indicate their familiarity with the noun compound on a 1-7 Likert scale, and three independent coders designated these interpretations as either attributive or relational. Coder agreement was over 95%, and where coders' designations diverged, the primary author's designation was used (Boylan). Criteria for definitions of relational and attributive compounds were taken from



Estes, 2003, and Wisniewski & Love, 1998. Relational and attributive items were matched on unigram frequency ( $t(254) = 1.09, p = 0.28$ ), compound length ( $t(126) = 0.85, p = 0.40$ ), and compound familiarity ( $t(126) = 1.67, p = 0.10$ ) (Brysbaert & New, 2009).

#### *Norming for relational-attributive bias*

In order to measure the variability in the interpretations of our noun-noun compound stimuli, and thus the degree of bias toward attributive and relational combination, we combined responses from the Mechanical Turk stimulus norming survey with responses taken from a survey of our fMRI subjects after they left the scanner. Both surveys asked subjects to describe what they thought each noun compound meant, with the added instruction to the fMRI participants that they write down the interpretations they had entertained while viewing the stimuli inside the scanner. Responses from a total of 35 subjects per item (17 from Mechanical Turk norming, 18 from fMRI subjects) were coded as either attributive or relational.

Figure 1 shows the distribution of attributive and relational bias by item, order-ranked from unanimously attributive interpretations to unanimously relational interpretations. The average “relational bias” for an item categorically labeled as relational in the Estes and Wisniewski & Love studies was 89.6% (SD = 0.14); likewise, the average “attributive bias” for an attributive item was 94.5%, (SD = 0.10). Despite having identified 128 items from the Mechanical Turk survey as relatively biased towards either relational or attributive meanings, additional responses from fMRI subjects demonstrated that some items were much more ambiguous than others. For instance,

while the item “cow parsnip” was originally designated as relational-biased (example interpretation: “a parsnip fed to a cow”), as was “pine mushroom” (e.g. “mushroom that grows on pine trees”), additional responses indicated these items were equally likely to have attributive readings, where interpretations such as “a parsnip shaped like a cow,” and “mushroom that looks like a pine cone” were offered for “cow parsnip” and “pine mushroom,” respectively.




Figure 1. Plot of relational-attributive combinatorial bias (normed on 35 participants, inclusive of 18 fMRI subjects). Figure 1 shows the distribution of attributive and relational combinatorial bias by item, order-ranked from unanimously attributive (relational bias = 0) interpretations to unanimously relational (relational bias = 1) interpretations. Blue indicates those items marked as relational in Estes (2003) and

Wisniewski & Love (1998) studies and red indicates those marked as attributive in those studies.

### *Experimental Task and Design*

The subject's task on each trial was to read two simultaneously centrally presented words constituting a nominal compound and indicate by button press (1) when they had decided on a coherent meaning for the compound and (2) whether a subsequent "probe" matched the meaning they had in mind. The nominal compound was presented for 2 seconds, and was immediately followed by a fixation cross, on screen for 6 or 8 seconds, during which time the subject need only passively view the screen. This fixation period was followed by a probe phrase, which either matched the modal interpretation for a given item (as determined in the Mechanical Turk survey; e.g. "a prickly carpet" for "cactus carpet") or was a dispreferred interpretation (e.g. "a carpet on which a cactus stands" for "cactus carpet"), where one out of eight trials had a dispreferred probe. This was followed by a 6-8-second fixation-cross ITI. The entire experiment consisted of 8 runs of 16 trials each. After leaving the scanner, subjects were given a questionnaire which asked them to write down what they had thought each nominal compound meant when they had viewed them in the scanner.

### *Image acquisition*

FMRI data were collected at the Hospital of the University of Pennsylvania on a 3T Siemens Trio System using a 32-channel multiple-array head coil. Four types of image sequences were collected for each participant: (1) a standard low-resolution

anatomic localizer; (2) a high-resolution, T1-weighted sequence for localization of fMRI activity in standard stereotactic space; (3) T2\*-weighted images from 9 experimental runs; (4) a B<sub>0</sub> field map sequence for subsequent geometric unwarping of T2\*-weighted images.

After acquiring T1-weighted anatomical images (TR=1630 ms, TE=3.11 ms, TI = 1100 ms, voxel size = 0.9 mm x 0.9 mm x 1.0 mm, flip angle 15°), we collected T2\*-weighted images using a gradient-echo echoplanar pulse sequence (TR=2000 ms, TE=30 ms, voxel size=2 mm x 2 mm x 2 mm, flip angle = 60°, BW = 1578 Hz/Px, 60 slices, with a multi-band acceleration factor of 3).

### *Analysis*

#### *Image analysis and ROIs*

fMRI data were pre-processed offline using the AFNI (Cox & Jesmanowicz, 1999) software package. The first four volumes of each functional run were removed so as to allow the signal to reach steady-state magnetization. Functional images were slice-time corrected, and a motion correction algorithm employed in AFNI registered all volumes to a mean functional volume. Images were then unwrapped via B<sub>0</sub> field maps (using FSL software; <http://www.fmrib.ox.ac.uk/fsl>) to reduce non-linear magnetic field distortions. We applied a high-pass filter of 0.01 Hz on each run to remove low frequency trends. Functional data were registered to the individual subject's anatomical MRI. Transient spikes in the signal were removed using AFNI's 3dDespike.

Our *a priori* ROIs were left and right anterior temporal lobes and left and right angular gyri, which we delimited using AFNI's CA\_ML\_18\_MNIA atlas. Our anterior temporal ROIs spanned labels “left/right temporal pole” and “left/right medial temporal pole”, while our angular gyrus ROIs circumscribed only the atlas's “left/right angular gyrus” ROI (see Figure 2).

Using AFNI's TENT function, we modeled the hemodynamic response function (HRF) as a finite impulse response (FIR) basis set fit to each condition, with bin-width equal to the 2-second TR, and 9 knots (TRs) modeled for a given trial. We used the full individually fitted 9-knot FIR HRFs to assess differences in the shapes and timecourses of the BOLD responses to attributive and relational compounds (see below); however, for our initial voxel selection and our analysis of relational-attributive combinatorial bias, we collapsed the FIR output: While the FIR model outputs 9 TENT functions and thus 9 beta estimates per condition per voxel, we selected the beta estimate of the largest magnitude (positive or negative) within a given TENT series such that our design matrix had one beta estimate per condition per voxel.

To identify task-activated voxels for inclusion in further analysis, we first conducted a GLM with FIR regressors for task and fixation ITI. The task TENT series was time-locked to the onset of the nominal compound, and the ITI TENT began at the onset of the fixation ITI (10-12 seconds after the onset of the nominal compound). To investigate effects of relational-attributive combinatorial bias in each ROI, we used a model with covariates for task and jittered probe event, where task TENTs were again time-locked to the nominal compound presentation and probe TENTs synced to the onset

of the probe question (8-10 seconds post compound onset), along with a continuous covariate for the relational-attributive bias (see Figure 1) of the noun compound in each trial. Head movement and global signal were included as covariates of no interest in both models.

In addition to investigating effects of relational-attributive bias, we also utilized subjects' post-scanning surveys to label each item/trial as relational or attributive based on individual subjects' responses. We then extracted the peristimulus BOLD signal timecourse starting at the onset of each trial, where TR0 was the onset of presentation of the noun compound, to TR8 post-stimulus onset (total of 9 TRs). TR0 was subtracted from each condition so that the starting point of the BOLD time series was aligned across conditions (Staresina, Fell, Do Lam, Axmacher, & Henson, 2012). This is analogous to the procedure of "baseline-correcting" in EEG analysis. Thus, only TRs 1–8 (2–16 seconds post stimulus onset) entered statistical timecourse analysis. We averaged FIR parameter estimates across voxels in each ROI in the participant's native space, and the resulting values entered into subsequent BOLD timecourse analyses.

## **Results**

### *Task-responsive voxels in anatomical ROIs*

In a group-level contrast targeting bilateral AG and bilateral ATL, we found several clusters of voxels with a reliable ( $p < 0.01$ , uncorrected) activation difference between task and ITI fixation baseline, where the task condition collapsed attributive and relational trials together. Two distinct clusters of activity were revealed in left AG, where one

cluster was positively activated for task relative to baseline (cluster centroid: [-45 -54 38] Talairach coordinates) and another more posterior, inferior cluster that was more active during baseline relative to task (cluster centroid: [-39 -67 35] Talairach coordinates). This motivated us to treat positively and negatively activated task-responsive voxels as distinct functional subregions within the anatomical left AG ROI. Other clusters in right AG and left ATL were largely positively task-responsive, and so did not prompt any functional division between above- and below-fixation task activation (see Figure 2). No significant task-responsive clusters survived even a liberal threshold in right ATL, and so this anatomical region was not analyzed further.




Figure 2: Task-responsive voxel activity (task vs. ITI fixation baseline) in left ATL (red), left AG (yellow), and right AG (blue) at Talairach coordinates [-45 -57 38].

### *Combinatorial bias predicts activity in left and right AG*

For each subject, we identified the 50 most positive task-responsive voxels in each of our three ROIs: right AG, left AG, and left ATL. We also identified the 50 most negative task-responsive voxels in left AG, for a total of four functional ROIs. The location and distribution of these top 50 task-responsive voxels varied greatly across subjects for each

ROI, precluding clear anatomical delineation of positive vs. negative task-responsive regions within left AG for a given subject, even though group-level clusters suggest an anatomical divide along PGa/PGp (Noonan, Jefferies, Visser, & Lambon Ralph, 2013; see Discussion). Within each of four 50-voxel ROIs in each subject, we examined the subject-wise effect of relational-attributive bias (referred to henceforth as simply “combinatorial bias,” see Fig. 1) on BOLD signal amplitude.

Within the 50 most task-responsive voxels in right AG, we found a significant main effect of combinatorial bias ( $t(17) = 2.44, p=0.01$ ), and this effect obtained for a wide range of ROI sizes within right AG (see Fig. 3a). No such effect was observed in task-responsive voxels in left ATL ( $t(17) = 0.84, p=0.20$ ; see Fig. 3b) or in positive task-responsive voxels in left AG ( $t(17) = 0.38, p=0.35$ ; see Fig 3c). Activity in those 50 voxels that responded most negatively to task (relative to baseline) in left AG was marginally predicted by combinatorial bias ( $t(17) = 1.23, p=0.11$ ). A significant main effect of combinatorial bias emerges when the size of the negatively task-responsive left AG ROI is increased to 90 voxels ( $t(17) = 2.73, p=0.005$ ), and this effect is also reliable for ROI sizes larger than 90 voxels ( $p<0.01$ , see Fig. 3d).

In order to compare combinatorial bias effects across ROIs, we conducted an ANOVA between ROI (right AG, left AG positive, left AG negative, and left ATL) and the degree to which combinatorial bias predicted BOLD signal amplitude. Combinatorial bias coefficients differed significantly across ROIs ( $F(3, 34) = 3.79, p=0.02$ ). Post hoc t-tests between pairs of ROIs revealed that combinatorial bias coefficients in right AG were significantly different from coefficients in the negative task-responsive left AG ROI



( $t(34) = 2.54, p=0.02$ ), while coefficient differences between other ROI pairs were not significant. When disregarding the direction of task-responsiveness, there are no significant differences in the magnitude of these effects.

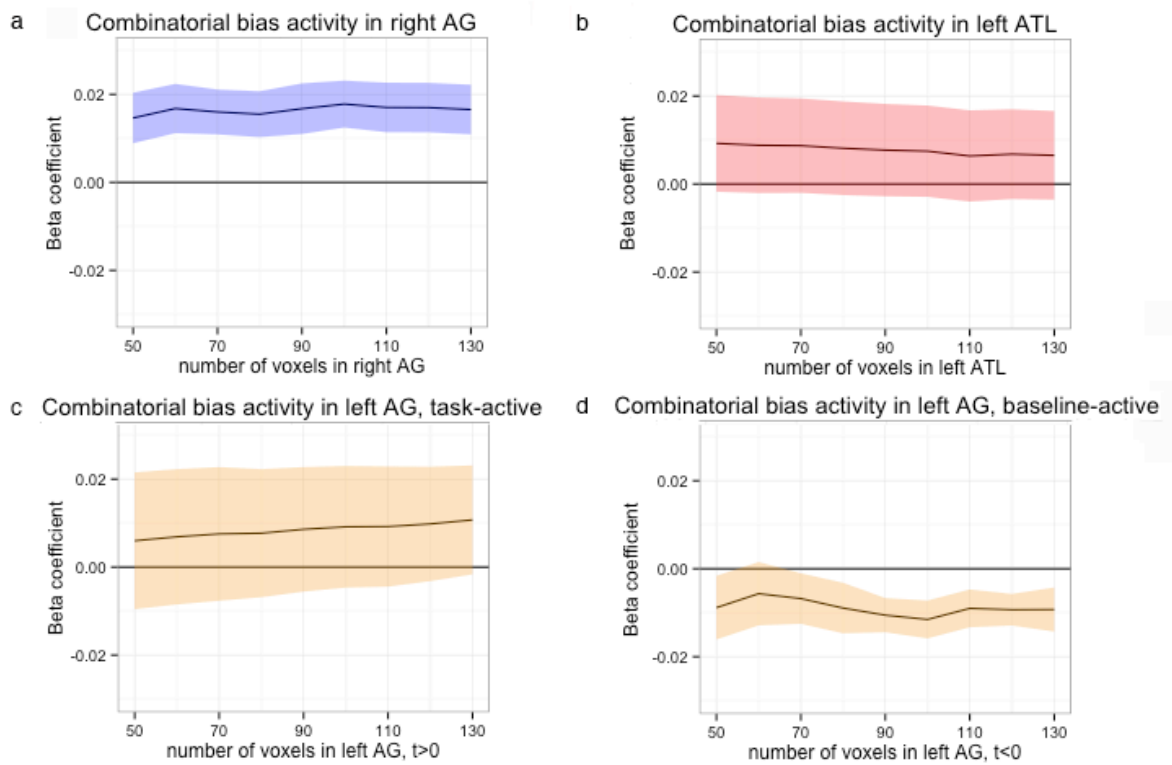


Figure 3. Beta coefficients from subject-wise parametric analysis of voxel activity by combinatorial bias (see Fig. 1) across a range of voxels within (a) right AG, (b) left ATL, (c) left AG, positively task-responsive voxels only, and (d) left AG, negatively task-responsive voxels only. Voxels chosen by most positive (or, in (d), negative)  $t$  statistics for the task-vs.-baseline contrast (see Fig. 2) in each anatomical ROI. Error ribbon indicates  $\pm 1$  SEM.

*BOLD timecourse of attributive interpretation differs by ROI*

In light of the evidence that combinatorial bias predicts the magnitude of the response in both left and right AG, we further investigated whether relational and attributive combination effects might also show distinct temporal BOLD profiles across ROIs. Using the same task-responsive voxel selection criteria as above, we compared the timecourse of BOLD activity for relational and attributive combinations in the 50 most task-responsive voxels in right AG, left AG (negatively task-responsive voxels only), and left ATL. In this analysis, we treated relational and attributive combination categorically, coding each trial condition based on individual subjects' responses in a post-scan survey.

Activity for relational combination significantly differed from attributive combination between 6 and 12 seconds post stimulus onset in both right AG and left AG ( $p < 0.05$ ; see Figure 4). Interestingly, activity associated with attributive combination was greater than relational activity in left ATL at a markedly early 4 seconds post-stimulus onset ( $p < 0.05$ ).

Examining the latencies of each subject's effect peak (using a nonparametric Wilcoxon signed rank test), we found that the response to attributive trials peaked significantly earlier than the relational effect in left ATL ( $p = 0.01$ ). This is not due to differences in subjects' attributive and relational response times, as we observed no significant difference in RT between conditions ( $t(17) = 0.51$ ,  $p = 0.61$ ;  $M_{\text{attributive}} = 2.65$  s ( $SD = 0.28$ );  $M_{\text{relational}} = 2.70$  s ( $SD = 0.26$ )). Time-to-peak analysis also revealed that both conditions peak earlier in left ATL than in right AG ( $p < 0.01$ ). While inferring temporal properties of neural activity from BOLD timecourses has its limitations, this

provides intriguing evidence that left ATL may also encode a distinction between attributive and relational combination in addition to bilateral AG.

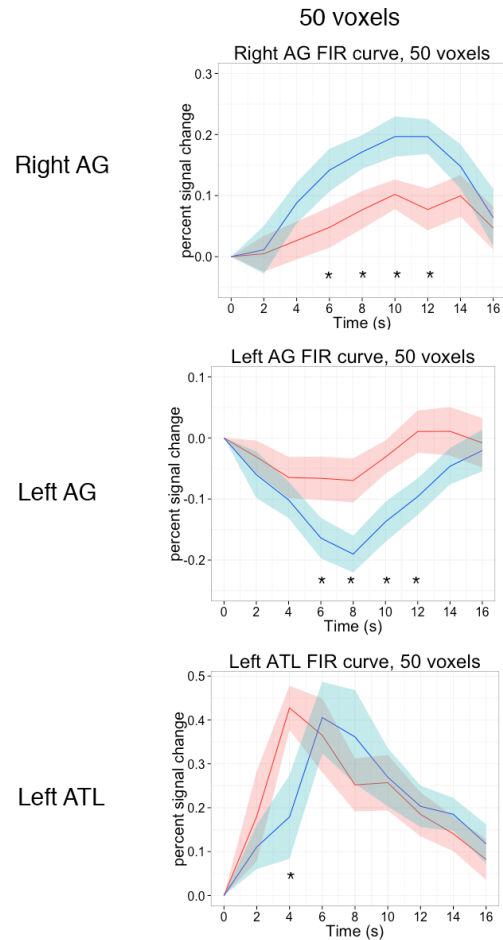


Figure 4. Peristimulus FIR curves of attributive (red) vs. relational (blue) activation in each ROI. FIR curves shown are from the 50 most positively task-responsive voxels for the task-vs.-fixation contrast in each anatomical ROI, except for left AG, where only the 50 most negatively task-responsive voxels are shown. Attributive and relational interpretations were coded based on a given subject's interpretations taken from post-test

surveys. Error bars indicate +/-1 SEM. Asterisks indicate significant effect of condition at given time point ( $p < 0.05$ ).

Upper panel: Right AG; Middle panel: Left AG; Lower panel: Left ATL.

## **Discussion**

This study sought to determine whether, and how, relational and attributive interpretations of nominal compounds differentially engaged putative “semantic hubs,” the left ATL and bilateral AG. We found evidence that both relational and attributive processes engaged ATL and AG, but that each brain region responded very differently to the relational-attributive dichotomy. Both right and left AG showed differential responses to relational and attributive compounds, with relational compounds diverging more from the baseline period than attributive compounds. However, while right AG responded more to both compound types than to baseline, the direction of this activation was reversed in left AG, such that left AG responded more at baseline than to compound interpretation. This profile of activation in left AG was consistent with its role in the so-called default network. Left ATL did not show a combinatorial bias effect *per se*, but the time course of individual subjects’ BOLD response curves indicated that attributive interpretations induced an earlier peak response than relational interpretations. Thus, while the magnitude of response in left and right AG was greater to relational combination than attributive combination, the timing, but not the magnitude, of left ATL response varied across the two combination types .

These combined findings support an account whereby relational and attributive operations are not dissociable by a coarse neuroanatomical divide, but rather are encoded

differently in different regions. The multiple, potentially redundant, instantiations of a combinatorial relational-attributive code across the brain provide compelling evidence that the relational-attributive distinction is a productive one. While these data suggest that bilateral AG is more engaged in computing relational combination, comparison with the left ATL profile does not constitute a true double dissociation between relational and attributive processes. While we do not find evidence to support a full dual-route, parallel process model of relational and attributive processing, we find that these combinatorial operations may be instantiated in overlapping networks across ATL and AG. The time course of attributive and relational BOLD response in the left ATL is also consistent with a serial, potentially single-process model. Interestingly, the serial process implicated here is the opposite of the CARIN model (see Introduction): while the latter predicts that attributive “relations” are the interpretations of last resort – i.e. attributive associations are only analyzed after other relations are considered – we instead find evidence that attributive processing precedes relational processing. This finding is more consistent with accounts holding that attributive feature extraction is necessarily prior to computing functional relations between concepts.

#### *Angular gyrus and thematic relations*

##### *AG sensitivity to thematic relations as verbs*

Bilateral AG, and more prominently left AG, have been implicated in a wide range of linguistic and non-linguistic semantic processes (see Seghier, 2012, for review), but our study pursues an emerging hypothesis that AG specifically subserves the semantics of thematic relations. There is increasing evidence that AG may be selectively

activated by thematic role information carried on verbs in particular. For instance, in one group of studies, experimenters looking at 1-, 2-, and 3-argument verbs (that is, intransitive, transitive, and ditransitive verbs, respectively) found that activation in bilateral angular and supramarginal gyrus (BA 39 and 40) correlated parametrically with the number of thematic roles that can attach to a given verb, even when the verb was presented in isolation (Meltzer-Asscher, Schuchard, den Ouden, & Thompson, 2013; Thompson et al., 2007; Thompson, Bonakdarpour, & Fix, 2010). Boylan, Trueswell, & Thompson-Schill (2015) also found that multi-voxel patterns in left AG tracked information relating to the presence of a shared verb in pairs of two-word phrases, demonstrating that AG represents information specific to verbs, perhaps event structure or thematic relations mediated by verbs.

While left AG has been implicated in the detection of syntactic errors (Embick, Marantz, Miyashita, O'Neil, & Sakai, 2000), it is also involved in the detection of semantic incongruities (Friederici et al., 2003; Newman, Pancheva, Ozawa, Neville, & Ullman, 2001; Ni et al., 2000) as well as the processing of connected discourse as opposed to unrelated sentences (Fletcher et al., 1995; Homae et al., 2003; Xu et al., 2005). This suggests that the sensitivity of AG to thematic roles and verb structure is not limited to the syntactic composition alone, but also to the semantic content (Pallier et al., 2011). It may be that left AG acts as an interface between semantic memory and syntactic structure, mapping semantic-thematic relations onto structural constraints surrounding verbs and their arguments. Indeed, electrophysiological and neuroimaging studies support an overlap between (morpho-)syntactic and semantic-thematic verb

violations. Kuperberg et al. (2008) compared three different types of verb violations: (1) semantic–thematically violated verbs (e.g. “at breakfast the eggs would eat”) (2) morphosyntactically violated verbs (e.g. “at breakfast the boys would eats”) and (3) real-world violations (e.g. “at breakfast the boys would plant”). They found that, unlike real-world violations, both semantic-thematic and morpho-syntactic violations elicited activity in a frontal/inferior parietal/basal ganglia network, in accord with previous electrophysiological findings that semantic-thematic and syntactic violations evoked P600 event-related potentials highly similar in latency and scalp distribution (Hoeks, Stowe, & Doedens, 2004; Kuperberg, 2007). The authors concluded that this frontal/AG/basal ganglia activity reflected attempts to integrate structural constraints of the verb with semantic properties of the Agent NP argument (Buccino et al., 2001; Chao & Martin, 2000; Damasio et al., 2001; Fogassi et al., 2005).

#### *Lateralized effects of thematic relations in AG*

If thematic role knowledge, particularly on the verb, is privileged content of the semantic space of AG, then how do we account for the AG activation profile of stimuli like nominal compounds, which do not contain any verb? Likewise, taxonomic and thematic errors that localize to lesions in ATL and AG comprise errors on nouns, not verbs. We argue that the verb functions as a “spell-out” of a thematic relation, and stimuli like relational nominal compounds require positing implicit verbs and events (as in “a hawk that *hunts* robins” for *robin hawk*). It is also interesting to note that word pairs in a thematic error, such as “dog” and “bone,” can be described as related via such an implicit verb/event; in the “dog-bone” case, “chews” or “buries,” etc. That is, thematic

knowledge is precisely knowledge of verbs and their arguments. We speculate that the verb may be the minimal linguistic expression of the fundamental thematic and event-based concepts that AG subserves.

It may also be that the degree to which a thematic relation is grammaticalized – for instance, whether it is spelled out in an explicit verb – accounts for the subtly different profiles of activity between right and left AG, and even within left AG itself. Graves and colleagues' (2010) study of nominal compounds vs. their non-attested reversals found BOLD activity in right AG, but not left, increased for attested nominal compounds like *lake house* (as compared with *house lake*). The authors offer a connectionist account of how noun-noun compounds might engage right AG but not left AG. They suggest that left and right AG can be modeled as attractor networks, where such a network settles into an attractor basin when it optimizes the error space in the mapping between inputs and outputs. Whereas left AG is suggested to have relatively narrow attractor basins, reflecting highly specific and constrained mappings between words and meanings, right AG may contain wider, shallower basins, representing more extensive overlap in meanings. This would accommodate “looser” meanings, effectively “filling in” the extrinsic relation necessary for the interpretation of compounds like *dog bone* that lack the explicit (morpho)syntactic information (i.e. “a bone that a dog chews on”) that would spell out the relation between the two nouns in the phrase. This version of Beeman and colleagues' (1994) “coarse semantic coding hypothesis” would account for why verb-based minimal composition, where the thematic relation is grammaticalized and explicit, would be derived in the narrow attractor basins of left AG, while the more



fluid meanings afforded by nominal compounds are derived in the wider semantic net of right AG.

*Functional heterogeneity of left AG: semantic representation and semantic control*

The effect of combinatorial bias we found in right AG accords with the Graves et al. (2010) nominal compound effect in that region, and the distinction between relational and attributive interpretations adds another dimension to the characterization of right AG as a combinatorial hub. However, our nominal compound effect in negatively task-activated voxels in left AG is rather more novel.

It has been noted that left AG is functionally heterogeneous, and recent work has begun to map this heterogeneity to subregions within left AG and surrounding areas. In a meta-analysis of studies comparing semantic tasks with high-vs.-low demands on executive control, Noonan et al., 2013, found a functional divergence between dorsal AG (bilateral, including dorsal/anterior AG and boundaries with superior marginal gyrus (SMG) and inferior parietal sulcus (IPS)) and left mid AG (somewhat closer to PGp than PGa), with respect to executive and representational roles in semantic processing. Dorsal AG showed reliably greater activation in high > low semantic conditions, and was characterized as allocating attention to semantic representations in a task-dependent and goal-driven manner. This characterization is not in itself inconsistent with a model of AG as a site of conceptual combination, as such compositional operations require selective attention to certain properties of events in order to construct higher-order derived concepts. However, dorsal AG's role in semantic control was contrasted with mid AG, the activity of which was modulated by the semantic representational content of stimuli

even when matched on task demands. Crucially, mid AG was associated with the “default mode network” (Raichle et al., 2001), showing more positive activation in the absence of a task. Seghier et al. (2010), found that left mid AG was a region within the default network that responded more negatively to perceptual decisions than semantic decisions, though both stimuli elicited a negative divergence from baseline (as did both types of nominal compound stimuli relative to baseline in our study), and other studies found this region showed more negative activity to abstract items than to concrete items (similar to our finding that relational compounds elicited more negative activity relative to attributive items in left mid AG) (Binder, Westbury, McKiernan, Possing, & Medler, 2005; Noonan et al., 2013; Wang, Conder, Blitzer, & Shinkareva, 2010).

Noonan et al. (2013) note that centers of activation in the putatively functionally distinct regions of dorsal AG and mid AG lie at Talairach coordinates [-41-55 45] and [-39 -65 30], respectively. These align well with the two clusters of activation we observed for the task-vs.-baseline contrast in left AG: a more dorsal cluster of positive activation at [-45 -54 38] and a more posterior, inferior cluster of more negative activation at [-39 -67 35]. We find that the latter cluster aligns both anatomically and functionally with the left mid AG region of the default network. Unlike the more dorsal cluster, the left mid AG cluster showed a combinatorial bias effect, suggesting that, like the concrete > abstract item effect reported in this region, there is also an attributive > relational effect in mid AG (where abstract items and relational items, respectively, induce more negative activation).

*A stimulus-specific negative BOLD response?*

While we characterize the left mid AG and right AG effects as similarly indexing a greater magnitude of activation for relational interpretations, this is based on the absolute value of activation relative to baseline. It is entirely possible that the left mid AG effect should rather be characterized as an “attributive > relational” effect when considered as a local vascular and/or neural suppression. Proposed hemodynamic mechanisms for negative BOLD responses include local “blood stealing” or “blood sharing,” whereby blood is diverted to (local or remote) active regions and away from the site of the negative BOLD response. This vascular effect could occur without necessarily reflecting a change in neural activity in the negative BOLD region (Bressler, Spotswood, & Whitney, 2007; Shmuel et al., 2002; Smith, Williams, & Singh, 2004). However, these purely vascular explanations are not sufficient to predict the negative BOLD response profiles found in several visual studies. For instance, Smith et al. (2004) found that a visual stimulus exciting primary cortex in one hemisphere caused a sustained negative BOLD response in the opposite hemisphere visual cortex. “Blood stealing” is not likely to explain this interhemispheric effect, as the blood supplies of the two hemispheres are largely independent of one another. Shmuel et al. (2006), while showing that negative BOLD activity is correlated with decreased neuronal activity measured via simultaneous electrical recordings, also demonstrated that the local decrease in neuronal activity predicted the spatiotemporal properties of the negative BOLD pattern better than surrounding positive BOLD activity did. This suggests that local neural suppression drives the negative BOLD response.

Moreover, recent studies suggest that such neural suppression may be more stimulus-specific than most attentional accounts predict. It is well known that attention directed to a given location corresponds with an increase in neural activity at that region (even in the absence of visual/auditory stimulation) and a decrease in neural activity at unattended regions. This predicts that attentional neural suppression will occur over relatively broad swathes of (unattended) sensory space, however (Bressler et al., 2007). Recent studies show that negative BOLD responses to stimuli are spatially (retino- and tonotopically) constrained and highly tuned to certain properties of the stimulus (Bressler et al., 2007; Linke, Vicente-Grabovetsky, & Cusack, 2011). Thus, it could be that neural suppression itself carries stimulus-specific information. On the other hand, for a reduction in signal to carry such stimulus-specific information, it may be that the firing trace of center-surround responses in a sub-voxel population of neurons is dominated by the surround suppression, even though the most narrowly tuned neurons are being excited (Bressler et al., 2007; Linke et al., 2011; Müller & Kleinschmidt, 2004).

While the mechanism for stimulus-tuned negative BOLD activity requires further study, negative BOLD responses across the brain are as potentially equally informative as positive BOLD responses. In the current study, the mechanism by which information relevant to relational compounds might be more “suppressed” in left AG is unclear: the pattern of suppression may itself be part of the representation of the relation between the two nouns, or it could be that more inhibition of unrelated information is required to resolve the meaning of a relational compound compared with an attributive compound. Further study is required to disentangle these options.

### *Anterior temporal lobe and semantic composition*

While bilateral AG appears to index the relational-attributive distinction by way of different overall activation for the duration of the task, left ATL may derive both relational and attributive combinations, but with a temporal offset in activation peak for the relational combination as compared with the attributive. This presents an intriguing alternative to a strict dual-process hypothesis for attributive and relational combination. Rather, this temporal delay in computing relational combinations in left ATL is consistent with a theory whereby property-based, taxonomic concepts are logically prior to relation-based, thematic concepts. It is also consistent with a model where “properties” are predicates, and thematic relations can be reified into features of object concepts such that “walks on a leash” is as much a property of *dog* as simple adjectives like “furry” or “loyal.” Since language is flexible in this regard, able to express the verb-/event-based relational concept “bears live young” in the adjective “viviparous,” and to convert between verbs (relations between nouns) and participles (properties of nouns), etc., it is perhaps unsurprising that there be a means of converting relational and attributive concepts along a single dimension. While relational and attributive combinations may be qualitatively distinguished in AG, it might be that these differences are collapsed in left ATL.

Indeed, an emerging view of left ATL is that its primary function is binding distributed multimodal features (where these features might be attributive or relational) such that an object-concept can be specified at various levels of categorization. For instance, patients with semantic dementia (SD) tend to show greater impairment of

specific names of objects than of more general names (Hodges, Graham, & Patterson, 1995; Rogers et al., 2004). These patients are slower to categorize objects at the basic level than at a more general level; e.g., they find it easier to categorize a swallow as an animal than as a bird (Hodges et al., 1995; Timothy T. Rogers et al., 2006). Conversely, healthy controls tend to prefer the basic, ‘bird’ level of classification (Murphy & Brownell, 1985; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Furthermore, features that are specific to a given concept are more damaged in SD cases than general features shared by several concepts (Papagno & Capitani, 2001; Patterson et al., 2007; Rogers et al., 2004, 2006), a fact which may underlie the attested categorization difficulties in SD.

This pattern has also been found in healthy subjects, in which the left ATL has shown increased activity for categorization of specific objects (e.g. ‘sparrow’) compared with categorization at the basic (e.g. ‘bird’) level or domain general (e.g. ‘animal’) level (Grabowski et al., 2001; Rogers et al., 2006; Tyler et al., 2004). In a study directly comparing how conceptual combination and object-concept specificity engage left ATL, Westerlund & Pylkkänen (2014) concluded that combination and specificity effects in left ATL likely arise from a single feature-binding operation.

Note that relational nominal compounds are exactly like attributive nominal compounds in that the first (modifier) noun is indicating what kind of thing the second (head) noun is. That is, regardless of whether one interprets *robin hawk* attributively or relationally, it is still a *hawk* object-concept of the *robin* type. Thus attributive and relational modification might be of a piece when the operation is to determine the level of

specificity *robin hawk* has relative to *hawk*. The crucial difference is that relational combination requires the integration of some extrinsic relation, while attributive combination rather selects which feature of the head noun is to be emphasized or substituted with the salient attribute of the modifier noun. In the attributive case, the modifier noun selects from among features already denoted in the set  $\{robin, hawk\}$  to yield “a hawk that is robin-like”, whereas a relational compound requires retrieval of some “feature” outside the set  $\{robin, hawk\}$ : namely, the “prey” or “hunt” relation (“a hawk *that hunts* robins”). This integration of an extrinsic relation/feature may account for the delay in left ATL BOLD response for relational combination relative to attributive combination.

#### *Limitations, future directions, and conclusions*

In this study, we find evidence that both left and right AG treat relational combination as distinct from attributive combination. Given that activity in left AG has been found to track certain properties of verb argument structure (Boylan et al., 2015) we propose that left AG might subserves more *explicit* thematic relations, particularly when expressed as verbs. This profile of left AG contrasts with that of right AG, which has previously been shown to track the combinatorial strength of nominal compounds (Graves et al., 2010) which lack a verb or explicit thematic relation, but which nonetheless might accommodate integration of *implicit* thematic relations. The function of left and right AG in composing explicit and implicit thematic relations aligns with a version of the “coarse semantic coding hypothesis” whereby left AG subserves narrower attractors such that connectivity between (thematic) associations might constitute a tighter, denser semantic

networks, while right AG supports more flexible (thematic) associations between concepts in wider, shallower attractor basins. While this model of attractor dynamics is an intriguing one, it remains to be directly tested in right and left AG. These results invite further study to pinpoint how right and left AG differ with regard to semantic combination.

We also found evidence that relational and attributive combination are temporally differentiated in left ATL. Inferring temporal signatures of neural activity from BOLD response curves has its caveats, however, given the potential nonlinearities between neural structures and the vasculature (Henson, Shallice, Josephs, & Dolan, 2002). Thus the left ATL timecourse would benefit from further study using methods, such as EEG or MEG, with higher temporal resolution than fMRI.

Also, while we find evidence that AG is involved in computing argument structure in the context of an event, there is also evidence that ATL indexes the thematic roles of nouns (e.g. agents and patients). Frankland & Greene (2015) were not only able to decode agenthood vs patienthood in ATL, but found that the patterns for these conceptual variables localized to neuroanatomically separable regions of ATL (the upper bank of the superior temporal sulcus and the lateral superior temporal gyrus, respectively). The degree to which relational compound interpretation deposes thematic roles to the constituent nouns, such as agent (*hawk*), theme/patient (*hunted robin*), instrument, location, etc., may also be read out in ATL or AG. While the division of labor between AG and ATL in semantic composition has yet to be made clear, our data offer a compelling reason why future study of the neural bases of combinatorial language would benefit from a distinction between attributive and relational operations.



### III. COMPOSITIONALITY AND THE ANGULAR GYRUS: A MULTI-VOXEL SIMILARITY ANALYSIS OF THE SEMANTIC COMPOSITION OF NOUNS AND VERBS

#### **Introduction**

Language owes its infinite expressive capacity to our ability to take simple building blocks, such as words or concepts, and combine them into complex representations. In linguistics, such “semantic composition” refers expressly to the combination of words into complex linguistic expressions, the meanings of which are a function of both the constituent building blocks (words) and the “rules” used to combine them (the grammar). Whether, and how, such grammatical operations might be realized in the brain are still highly debated questions, and the emergence of compositional meaning from units such as morphemes, words, or concepts, is largely a mystery. However, understanding the engine of compositionality in the brain is a fundamental desideratum to any cognitive neuroscientific model of semantics.

#### *Roles of left anterior temporal lobe and left angular gyrus in semantic composition*

The psycholinguistically motivated neuroanatomical models of semantic processing that have emerged in the past few years involve several brain areas which roughly cluster into four main regions: left inferior frontal, left anterior temporal, left posterior temporal, and left temporo-parietal (Ben Shalom & Poeppel, 2007; Binder & Desai, 2011; Binder, Desai, Graves, & Conant, 2009; Lau, Phillips, & Poeppel, 2008; Pallier, Devauchelle, & Dehaene, 2011; Patterson, Nestor, & Rogers, 2007; Price, 2010). Of these regions, there are two – left anterior temporal lobe (ATL) and left angular gyrus

(AG) – that are prime candidates to support composition, because both show greater activation for well-formed sentences than for non-compositional lists of words (Pallier et al., 2011, *inter alia*). In addition, both have been characterized as “semantic hubs”, owing to functional and anatomical patterns that are consistent with multimodal convergence (Binder & Desai, 2011; Lambon Ralph, 2014; Patterson et al., 2007; Seghier, 2012). While the mechanism by which different modalities converge on a single given conceptual representation is still unclear, it is likely that the mechanism that can encode the binding of modality-specific features into a given concept also accomplishes the binding of words into higher-level linguistic constructs (Westerlund & Pylkkänen, 2014). We begin with a brief review of findings relating to composition involving ATL in order to motivate contrasting ideas we will consider in the current study regarding composition in AG.

The ATL is uniquely situated at the end of a caudal-to-rostral stream of information processing feeding from primary sensory and motor areas and association cortex (Binder et al., 2009; J. R. Binder & Desai, 2011; Binney, Parker, & Lambon Ralph, 2012; Felleman & Van Essen, 1991). It is thus located at a prime “convergence zone” for inputs from many different modalities. Moving anteriorly along the temporal lobe, one finds a caudal-to-rostral hierarchy emerge as neuronal responses are more tuned to complex stimuli and more invariant to low-level sensory variation; such a hierarchy has been established along both visual (Felleman & Van Essen, 1991) and auditory (Rauschecker & Scott, 2009) streams. This “graded convergence” may provide a mechanism both for “feature combination” and, in the limit, for maximally invariant

amodal, and thus abstract, conceptual representations. The culmination of this graded convergence up the temporal lobe (Rauschecker & Scott, 2009; Stringer & Rolls, 2002) is a basal rostral region of ATL shown to have very limited extra-temporal connectivity and high intra-temporal connectivity (Binney et al., 2012). Such neuroanatomical sequestration is arguably a *sine qua non* for a region able to represent abstract, modality-invariant semantics. Thus, ATL is a prime candidate for semantic composition.

In one of the first studies investigating the neural correlates of minimal two-word composition, Baron and colleagues (Baron et al., 2010) found evidence from fMRI pattern analyses that the left ATL subserved the combination of concepts such that the superimposition of individual patterns of the simplex concepts *young* and *man* (as represented by various face stimuli) reliably predicted the activation pattern for the complex concept *young man*. Consistent with this finding, a magnetoencephalography (MEG) study of visually presented two-word phrases comparing nouns in minimal compositional contexts (*red boat*) with nouns in non-compositional contexts (in which a non-word letter string was concatenated with a real word, e.g. *xkq boat*) found increased composition-related activity in left ATL (Bemis & Pylkkänen, 2011). Thus, there is a growing body of functional and tractographic studies to suggest that the left ATL substrate of composition may be multimodal sensorimotor features, and particularly *visual* features of object-concepts (Coutanche & Thompson-Schill, 2014), corroborating the notion of the left ATL as hub of the so-called ventral “what” pathway.

While the left ATL has recently received much attention as a potential semantic hub, it is not the only region to invite this label. Researchers have also ascribed the role of a semantic hub to the AG, as it lies at the junction between temporal, parietal, and

occipital lobes and thus receives a confluence of auditory, somatosensory, spatial, and visual inputs. Left AG has been implicated along with left ATL in conceptual combination studies of the sort described above (Bemis & Pylkkänen, 2012), and several studies demonstrate bilateral AG sensitivity to manipulations whereby well-formed sentences are contrasted with word lists, pseudowords, or scrambled sentences (Bavelier et al., 1997; Bottini et al., 1994; Humphries et al., 2007, 2006). Left AG also shows greater activity for connected discourse vs. unrelated sentences (Fletcher et al., 1995; Homae, Yahata, & Sakai, 2003; Xu, Kemeny, Park, Frattali, & Braun, 2005). This broad profile of effects has led some to suggest that the AG may play a potentially domain-general role in semantic information integration structured around events (Binder & Desai, 2011; Binder et al., 2009; Lau et al., 2008; but cf. Noonan, Jefferies, Visser, & Lambon Ralph, 2013, and Discussion below for evidence that certain sites within AG are involved in semantic control processes, not representations).

Not all studies investigating conceptual combination find activation in both left ATL and bilateral AG. Upon closer inspection of those stimuli that elicit differential activity in AG but not in left ATL, one finds that the type of composition involved is invariably based on thematic relations rather than feature combination *per se*. For instance, Graves et al. (2010) compared familiar meaningful noun-noun pairs, such as *lake house*, with reversed phrases, such as *house lake*, the meanings of which were not obvious; they found that AG, along with other temporoparietal areas (mostly right-lateralized), showed greater activation for processing the more obviously combinatorial phrases. In characterizing the compositional operation employed in interpreting their particular noun-noun stimuli, the authors noted that most of their noun-noun stimuli were

interpreted as denoting thematic relations between head and modifier nouns; that is, most compounds consisted of nouns participating in some spatial relation (as in “a house *on a lake*”) or event-based relation rather than sharing some common feature (as in, for instance, a nominal compound like *cactus carpet*, which is more likely to be interpreted as “a carpet *that is prickly like a cactus*” than as some sort of relational compound, like “a carpet *with a cactus placed on it*”) (Estes, 2003b; E. J. Wisniewski & Love, 1998). This raises the question as to whether these stimuli were probing combinatorial semantics in general, or semantic thematic relations in particular.

In another group of studies, experimenters looking at 1-, 2-, and 3-argument verbs (that is, intransitive, transitive, and ditransitive verbs, respectively) found that activation in bilateral angular and supramarginal gyrus (BA 39 and 40) correlated parametrically with the number of thematic roles that can attach to a given verb, even when the verb was presented in isolation (Meltzer-Asscher et al., 2013; Thompson et al., 2007, 2010). Whereas Graves et al. (2010) indicates AG involvement in processing spatial and event-based relations, broadly construed, the work on verb adicity suggests a more selective sensitivity to verbs’ thematic relations and/or event complexity carried on the verb. While AG has been found to be sensitive to both linguistic event structure and non-verbal events depicted in scenes and mini-movies (Sitnikova, Holcomb, Kiyonaga, & Kuperberg, 2008a; Sitnikova, Holcomb, & Kuperberg, 2008b), it could be that the verb is the minimal linguistic expression of fundamental thematic relation-based or event-based concepts that AG subserves. This would predict that verb semantics would be particularly privileged in AG semantic space.

### *A feature vs. function dichotomy?*

Given that both ATL and AG are implicated in semantic composition, we might start with the hypothesis that *any* kind of semantic similarity between two concepts might influence the similarity of neural (in our case, voxel) patterns evoked by the concepts in these two regions. For instance, regions that encode the meaning of a two-word phrase (such as “eats meat”) ought to elicit a similar neural response to other two-word phrases that share either of these two words as compared to a phrase that shares none of the words. In this study, we go one step further and explore possible restrictions on this prediction. We suggest that whereas the left ATL may be involved in structuring semantic knowledge around commensurate features of (object-) concepts, the AG builds semantic knowledge based on functional/thematic relations between concepts. Of course, this distinction could be operationalized in a number of different ways. In this study, we test two possible dimensions along which the left ATL and AG might cleave “feature-based” and “function-based” composition: one dimension respects whether two concepts share an event (and since verbs can denote events, we operationalize this as two phrases that share a verb), while another dimension concerns whether two concepts share an argument, which we will explain in greater detail below.

### *Angular gyrus and event-denoting verbs*

It is widely agreed that verbs and nouns constitute meaningfully distinct linguistic forms, but it is less clear whether their processing engages different brain areas. In a meta-analysis using hierarchical clustering to identify regions associated with nouns and verbs, Crepaldi et al. (2013) identified several clusters associated with noun or verb

processing across all tasks, the left AG among them; however, the left AG was associated with nouns, not verbs (though this might be due to the authors' inclusion of activation peaks from studies examining nouns not directly contrasted with verbs). The authors conclude from the distribution of their clusters that the neural circuits of noun and verb processing are highly contiguous across a wide network of frontal, parietal, and temporal regions, and that there is little evidence suggesting that verb processing relies primarily on embodied motor representations.

These conclusions contrast with a prominent theory stating that, while not divided by the grammatical class of the noun or verb *per se*, neural areas are divided by the semantic primitives of objects and actions (Bird, Howard, & Franklin, 2000; Vigliocco, Vinson, Druks, Barber, & Cappa, 2011). Under this account, only prototypical nouns (object nouns) and prototypical verbs (action verbs) dissociate neural areas: specifically, action/verb processing recruits more of the fronto-parietal network while object/noun processing recruits more temporal regions (Cappa, Sandrini, Rossini, Sosta, & Miniussi, 2002; Cappa & Perani, 2003; Damasio & Tranel, 1993; Shapiro, Moo, & Caramazza, 2006). However, several studies fail to support such specific roles for frontoparietal areas and temporal regions in action/verb and object/noun processing, respectively (Crepaldi, Berlingeri, Paulesu, & Luzzatti, 2011; Liljeström et al., 2008; Tranel, Martin, Damasio, Grabowski, & Hichwa, 2005; Tyler, Russell, Fadili, & Moss, 2001). Also, a study by Bedny and colleagues (2008) found that while a certain region of left AG responded more to verbs than nouns, activity in this region did not distinguish between high motion and low motions words, whether nouns and verbs were included together or

queried separately. This dissociation between grammatical class and action content, along with evidence for a dissociation between grammatical class and imageability (Bedny & Thompson-Schill, 2006), motivates treating the distinction between most nouns and verbs as a dimension of meaning separate from just action content, particularly when characterizing left AG.

Instead, we propose investigating the combinatory properties of *events* as a means to probe the semantics of AG. Our verb-centric hypothesis states that if AG represents events, then phrases that refer to similar events should evoke similar patterns of neural activity in AG. Even more specifically, phrases that share a given event-denoting verb (such as *eats*) should evoke similar patterns of activity in AG; events involving eating should have more similar neural patterns than they do to other kinds of events. Crucially, however, phrase pairs that differ along this dimension, in that one phrase has a verb and the other phrase does not (e.g. *eats meat* and *tasty meat*), will not evoke similar patterns of activation in AG, even though these latter two-word phrases also share a content word (*meat*).

Note that events are necessarily compositional constructs, as they represent not only information about event participants (thematic relations), but also temporal information vis-a-vis tense, aspect, etc. As we have mentioned, AG is more active in conditions which involve tracking narrative and discourse structure as compared with disconnected sentences (Fletcher et al., 1995; Homae, Yahata, & Sakai, 2003; Xu, Kemeny, Park, Frattali, & Braun, 2005). The connective tissue of such narrative and discourse structure is temporal order and temporally mediated causal relations. This



suggests that temporal information, particularly that carried on a verb, is critical to engaging AG.

While there are “event” words that are not verbs, such as deverbal nouns and event-denoting nouns like *party* or *hurricane*, these lexical items do not often have temporal information. Inflected verbs necessarily carry information on tense and aspect, while nouns generally do not (rare exceptions being nouns that arguably denote a tense, like *ex-wife*, *former champion*, *husband-to-be*, *president-elect*). There is good reason to expect that event-denoting verbs dissociate from event-denoting nouns in both semantic and neural space: a study crossing grammatical class (nouns vs. verbs) with event denotation (events vs. objects) found that while left posterior middle temporal gyrus responded to both event nouns and verbs over objects, a region at the left temporoparietal junction (inclusive of AG) responded more to verbs than to any nouns, including event-denoting nouns (Bedny, Dravida, & Saxe, 2013).

Therefore, because we focus on events as spatiotemporal denotations, the natural way to examine this is to analyze verb similarity patterns. In this study, we do not directly test which component of event composition – thematic relations (theta roles) or temporal event information – might drive verb sensitivity in AG, but previous evidence implicates both.

#### *Angular gyrus and function-argument composition*

Under an alternative hypothesis, AG is sensitive to relational information independent of the verb. The dichotomy between “feature” and “function” has a good deal of traction in more formal theories of semantics, particularly in the distinction

between adjuncts and arguments, respectively. The verb is the central predicate of a sentence, and predicates, such as *eat*, take arguments (like *meat*) (Heim & Kratzer, 1998; Pylkkänen, Brennan, & Bemis, 2011). While the case of verbs and their direct objects is perhaps the most canonical example of what we will call “argument-type composition,” there are other types of function-argument relations: for instance, the composition of prepositional phrases (e.g. *with meat*, where the preposition *with* takes the argument *meat*). We are particularly interested in the status of function-argument relations because, up until now, studies on two-word minimal composition have focused almost exclusively on another type of composition: adjunct-type composition (e.g. *red boat*, *old man*). However, two recent magnetoencephalography (MEG) studies suggest ATL may subserve both adjunct-type and argument-type composition. Linzen, Marantz, & Pylkkänen (2013) found that left ATL is sensitive to verb-argument structure, specifically the subcategorization frames of verbs. Westerlund, Kastner, Kaabi, & Pylkkänen (2015) examined several different instances of both argument- and adjunct-type composition: namely verb-argument (*eats meat*), preposition-argument (*in Italy*) and determiner-argument (*Tarzan’s vine*) composition for function-argument composition and adjective-noun (*black sweater*), adverb-verb (*never jogged*) and adverb-adjective (*very soft*) composition for adjunct-type composition. The authors found that function-argument composition (inclusive of verb phrases, prepositional phrases, and possessives) drove increased activation in left ATL. In our study, we extend this paradigm to examine whether AG responds specifically to argument-type composition.

Rather than rely only on univariate measures of activation for adjunct- vs. argument-based composition across a given brain region, we treat composition as it applies to the multivariate patterns of particular base words (in this case, nouns and verbs). We constructed sets of two-word phrases such that particular pairs in each set would allow us to test our hypotheses. Consider the phrases *eats meat*, *eats quickly*, *with meat*, and *tasty meat*. Using the logic of Baron et al. (2010), who found that additive superimposition of the voxel patterns underlying simplex concepts like *young* and *man* could predict the complex pattern of *young man*, we ask whether the complex concepts *eats meat* and *with meat* might be acting on the base word *meat* in the same way, as (1) both phrases are instances of function-argument composition and (2) *meat* is an argument in both (while it is not an argument in, say, *tasty meat*). On the other hand, if AG represents information carried on the verb (that is, the event, *ex hypothesi*), then the neural pattern evoked by the complex concept *eats meat* is expected to be more similar to that of *eats quickly* than to any other phrase.

#### *Multi-voxel pattern similarity as a window to various dimensions of compositionality*

Until recently, the prevailing approach in the neuroscientific study of concepts was to employ univariate tests of fMRI data, using a brain region's average metabolic response to discriminate stimulus conditions based on locations of peak activation, potentially at the expense of voxel-level signal variation distributed across a given region. However, increasing use of multivariate methods to harness this voxel-level neural variability has revolutionized the study of object concepts. Current multi-voxel pattern analysis (MVPA) methods are predicated on the idea that information is instantiated in a

spatially distributed pattern of neural activity. While some MVPA methods use various classification techniques over voxel patterns to *discriminate* between stimulus conditions (Boylan, Trueswell, & Thompson-Schill, 2014; Coutanche & Thompson-Schill, 2014; Polyn, Natu, Cohen, & Norman, 2005; *inter alia*), other methods analyze voxel patterns with respect to the strength of *similarities* between stimuli within given dimensions (e.g. shape, color, animacy, etc.) (Clarke & Tyler, 2014; Connolly et al., 2012; Fair et al., 2009; Fairhall & Caramazza, 2013; Haxby, 2001; Kriegeskorte, 2008; Weber, Thompson-Schill, Osherson, Haxby, & Parsons, 2009). With the latter approach, neural pattern variation can extend across a more continuous space than is sought in nominal/dichotomous classification techniques. Such MVPA techniques can be remarkably powerful tools, and have been used to query neural patterns using only a few TRs per stimulus event, and with each stimulus event modeled as a single unique regressor in a GLM (the beta values of which enter a correlation matrix or other similarity analysis) (Musz & Thompson-Schill, 2014).

The current study employs MVPA pattern similarity measures to query the relatively high-level semantic similarity space of two-word compositional phrases. We compare fMRI multi-voxel patterns associated with pairs of two-word compositional phrases - e.g. *eats meat* compared with *eats quickly*, *tasty meat*, or *with meat* – in which we hold constant a single word (here, either *meat* or *eats*), but manipulate (1) whether the word shared is a noun or a verb and (2) whether the two compositional phrases share a composition type (both argument-type or adjunct-type composition) or differ in their composition type. As we explain in greater detail below, our strategy is to observe the

extent to which neural patterns evoked by two-word phrases are altered by relative isolated manipulations in their content, allowing us a means of inferring the principles that govern neural coding in different cortical regions. This is analogous to how a vision scientist can observe the tuning properties of a neuron by varying a stimulus dimension such as wavelength. Here, the dimensions we manipulate are (1) shared verb (thus shared event) and (2) shared composition type, which allows us to examine (1) how information carried on the verb may be critical to certain regions involved in semantic composition but not to others, and (2) whether the putative grammatical rules that distinguish various two-word phrases are differentially instantiated in regions of the brain implicated in basic semantic composition. We also investigate a corollary to the first hypothesis: if verbs are somehow privileged in certain brain regions (namely AG), then we might expect to see that the neural similarity between voxel patterns associated with phrases sharing a given verb might be predicted by subjects' ratings on how similar the meanings of these verb-sharing phrases are to one another. We also test whether nouns might likewise drive pattern similarity of noun-sharing phrases in left ATL.

## **Material and Methods**

### *Participants*

Twenty-one subjects participated in this study. Two participants' data were excluded due to excessive motion, and one subject was found to have an anatomical anomaly. Data from the remaining eighteen subjects are reported here. Subjects ranged in age from 18 to 28 years, and all were right-handed native speakers of English with

normal or corrected-to-normal vision and no reported history of neurologic problems. Subjects gave written informed consent and were provided monetary compensation for their time. The human subjects review board at the University of Pennsylvania approved all experimental procedures.

### *Stimuli*

#### *Stimuli Design*

Crossing type of composition (argument-type vs. adjunct-type) with presence/absence of verb, we chose compositional word phrases that conformed to four different types:

- 1) +verb\_arg: a word phrase that composed via argument-type composition and included a verb, e.g. *eats meat*
- 2) -verb\_arg: a phrase that composed via argument-type composition, the head of which was a preposition instead of a verb; e.g. *with meat*
- 3) +verb\_adj: a phrase that composed via adjunct-type composition and included a verb, e.g. *eats quickly* (note that adjective-noun phrases are not the only type of adjunct-type composition)
- 4) -verb\_adj: a phrase that composed via adjunct-type composition and did not include a verb; e.g. *tasty meat*

where +verb\_arg, -verb\_arg, and -verb\_adj always had the same noun, and +verb\_arg and +verb\_adj always had the same verb. These four types of compositional phrases are further illustrated in Table 1.

Stimuli consisted of 36 sets of four compositional phrases and two non-compositional items. We implemented the “minimal composition paradigm,” where composition is isolated to two-word phrases and contrasted with one-word non-compositional items consisting of an unpronounceable letter string and a real word (Bemis & Pylkkänen, 2011a; Westerlund et al., 2015). Each non-compositional item was presented in one of two possible word orders, for a total possible four one-word items. The format for the non-compositional one-word items ([noun/verb] + [non-pronounceable letter string]) was counter-balanced for the real word being in phrase-initial or phrase-final position.

Table 1: Stimuli Design

	<i>Argument (argument-type)</i>	<i>Adjunct (adjunct-type)</i>
	<i>composition</i>	<i>composition</i>
<b>+verb</b>	<i>eats meat</i>	<i>eats quickly</i>
<b>verb control</b>	<i>eats fghjl / fghjl eats</i>	<i>eats fghjl / fghjl eats</i>
<b>-verb</b>	<i>with meat</i>	<i>tasty meat</i>
<b>-verb (noun)</b>	<i>meat fghjl / fghjl meat</i>	<i>meat fghjl / fghjl meat</i>
<b>control</b>		

Table 1: An example set of two-word compositional and one-word non-compositional items sharing a given noun or verb.

Our hypothesis concerns verbs that denote events, so we selected verbs that were eventive rather than stative. Of our 36 verbs, we have one traditionally stative verb – *love* – but this verb is sometimes used in the continuous aspect (*I am loving, I was loving*), so it is not as strongly stative as *have* or *own*. The other 35 verbs are strongly eventive (e.g. *I am eating, I am kicking, I am buying, etc.*)

A given verb in each compositional set had to be able to compose with a direct object, as in “eats meat,” but, conversely, could not be so strongly transitive as to *require* a direct object. Therefore, all verbs were chosen to be optionally transitive. This optional transitivity allowed for compositional phrases of the +verb\_adj type, as in *eats quickly*, where there is no direct object. Moreover, all verbs had present tense inflection in order



to ensure they were interpreted in the active voice. A verb in a +verb\_arg phrase like *chews gum*, when presented in the past tense – *chewed gum* – might be read as a passive participle – as in “*gum that is chewed.*” In such a case, an adjectival participle would compose via adjunct-type rather than argument-type composition.

Similar constraints were placed on the noun stimuli. A given noun had to be able to compose in a variety of contexts (with an adjective, a preposition, and a verb) without requiring extra plural suffixes or determiners (*throw stones, throw the stone*), the addition of which might involve another type of compositional operation. Therefore, only mass (non-count) nouns were included. Also, since one of our hypotheses concerns the status of events as denoted by verbs, we avoided event-denoting nouns, though marginally event-denoting nouns in our stimuli included the nouns *traffic, crime, opera, praise, and pardon.*

Note that, due to the constraints of our particular two-word phrase sets (see Table 1), we could not compare two adjunct-type phrases. The nature of our similarity analysis required that a given two-word phrase be compared only to another two-word phrase that shared either a noun or a verb. This is because we are investigating the changes in the voxel pattern of a given base word – e.g. *eats* or *meat* – when composed with a function head, an argument, or an adjunct. While it is possible that the operations of argument composition and adjunct composition each have their own stable and distinctive patterns in the brain regardless of what words are composed, it is more likely that the instantiation of the operation is highly dependent on the words composed, and thus highly distributed spatially in the brain. Under this view, a voxel pattern for function-argument

composition of *eats meat* cannot be expected to be the same voxel pattern instantiating the composition of *plays guitar*. For this reason, we need to make all similarity comparisons relative to an “anchor” word – either a shared noun or a shared verb in this case. Our two types of adjunct-type compositions did not share a word: an event-denoting adjunct-type phrase like *eats quickly* cannot be directly compared with the non-event-denoting adjunct-type phrase like *tasty meat*, even though these two phrases are in the same set and can each be compared with an event-denoting argument-type phrase like *eats meat*. Therefore, we had no means of assessing adjunct-type similarity profiles. However, because a number of previous studies have shown that adjective-noun compositional phrases like *tasty meat* activate left ATL, we were less interested in replicating such a result than querying other types of compositionality, specifically argument-type pattern similarity (*eats meat* compared with *with meat*) and verb-based pattern similarity (*eats meat* compared with *eats quickly*). This allowed us to test two different hypotheses about the role of AG specifically: On the one hand, if AG subserves argument-type composition in general (that is, application of arguments to any type of function head; in this case, either a verb or a preposition), then we predict argument-type pairs will elicit highly similar patterns of activation in AG. On the other hand, if AG is preferentially sensitive to event-denoting verbs, then specifically verb-sharing phrases will elicit similar patterns in AG.

Similarly, if left ATL is specialized for feature-based composition, either noun-sharing or verb-sharing phrases might elicit shared activation patterns. Moreover, if left ATL is sensitive to composition-type, we predict conserved patterns of activation for

phrases that share either adjunct- or argument-type composition (where only argument-type shared composition is tested here.)

### *Stimuli norming*

All words and word phrases were matched for length except for adverbs (in +verb\_adj, e.g. “eats *quickly*”) compared with nouns (in +verb\_arg, e.g. “eats *meat*”), where adverbs were significantly longer than nouns ( $M_{adv}=7.6$ , s.d. = 1.96;  $M_{noun}= 6.2$ , s.d = 1.97). Only adverbs and nouns were matched on frequency. We also collected imageability ratings on a 1-7 Likert scale (1 being lowest) on one- and two-word items. There was no significant difference between our noun-based stimuli and our verb-based stimuli (not including *eats meat* stimuli that included both a verb and a noun):  $M_{noun\_items} = 4.74$ , s.d. = 0.83;  $M_{verb\_items} = 4.39$ , s.d. = 0.58.

As noted above, all verbs appeared in the present tense in this study. When presented in a non-compositional one-word stimulus (e.g. *bvref picks*), some of these present tense verbs might be ambiguous between a verb and noun interpretation. However, in all such cases of possible ambiguity, the dominant form of the base word was a verb (assessed using Google Books Ngram Corpus, American English; Lin et al., 2012).

We also normed our two-word phrase pairs’ dissimilarity using a survey posted on Amazon Mechanical Turk (Buhrmester et al., 2011). Instructions were as follows:

“In this survey, you will be asked to indicate how alike two instances of a word in two different contexts are. You will do this using

a slider bar. 0 means “exactly the same” and 7 means “substantially changed. For example, a minimal change might be something like “stare at the cash register” and “ask about the cash register.” A moderate change would be the difference between “buys rice” and “grows rice.” In “buys rice”, the rice is likely packaged in a container, while in “grows rice” the rice is in a field on a plant.”

In this way, we were able to extract pairwise dissimilarity scores for phrases that shared a given noun or a given verb, and also average these scores to yield a measure of how much, on average, a given noun or verb changed depending on the word it was composed with. That is, pairwise dissimilarity scores for the phrase pairs (*meat, eats meat*), (*meat, tasty meat*), (*meat, with meat*), (*eats meat, with meat*), (*tasty meat, with meat*), and (*tasty meat, eats meat*) were averaged together to yield a mean dissimilarity score for “*meat*”; likewise, the pairwise dissimilarity scores for (*eats, eats meat*), (*eats quickly, eats meat*), and (*eats, eats quickly*) were averaged together to yield a mean dissimilarity score for “*eats*.” We could then compare the pairwise dissimilarity norms with “neural similarity” scores (see below for discussion), where “neural similarity” was the measure of how much a multi-voxel evoked pattern for a given noun/verb changed depending on the word it was composed with (where less change in the patterns indicates greater similarity). We could also analyze the mean dissimilarity norms to look for coarse similarity differences between noun-containing phrases and verb-containing phrases.

After filtering responses for English as a first language, completeness, time to response, fluency, and neurological disorders, between 16 and 40 subjects per each of 9

lists (one list for each phrase pair in a given set) remained. Responses from 16 randomly chosen subjects per list were then analyzed and used to calculate the pairwise and mean similarity norm scores. Interestingly, phrase pairs sharing a noun were rated as significantly ( $p < 0.05$ ) more similar on average ( $M = 3.16$ ,  $s.d. = 0.25$ ; where mean similarity norm on a 0-7 Likert scale, 0 most similar) than phrase pairs sharing a verb ( $M = 3.44$ ,  $s.d. = 0.46$ ).

### *Experimental Task and Design*

The subject's task on each trial was to read two simultaneously centrally presented words constituting either a compositional or one-word item. This phrase was presented for 3 seconds. The critical phrase was immediately followed by either a 9-second fixation cross, during which the subject need only passively view the screen, or a two-word phrase probe presented for 3 seconds and followed by a 6-second fixation-cross ISI. This probe was presented in capital letters and terminated with a question mark so as to distinguish it from the preceding critical phrase. If the initial critical phrase was compositional (e.g. *asks nicely*), then the probe was also compositional (e.g. *INQUIRES POLITELY?*); otherwise, the probe consisted of a noun/verb and a non-word letter string in the same order as the non-compositional one-word item it followed (e.g. *asks xblrdc* followed by *INQUIRES PCXFDL?*; or *xblrdc asks* followed by *PCXFDL INQUIRES?*). The subject was instructed to indicate by button press (yes/no) whether the probe phrase (or, in the case of a non-compositional trial, the probe word) was synonymous with the preceding phrase or word. If a trial had no probe, no response was required from the participant. 10% of trials had a probe phrase or word, and 30% of these catch trials had

probes that were not synonymous. Probes from catch trials were excluded from analysis. The entire experiment consisted of 9 runs of 24 trials each.

### *Image acquisition and pre-processing*

FMRI data were collected at the Hospital of the University of Pennsylvania on a 3T Siemens Trio System using a 32-channel multiple-array head coil. Four types of image sequences were collected for each participant: (1) a standard low-resolution anatomic localizer; (2) a high-resolution, T1-weighted sequence for localization of fMRI activity in standard stereotactic space; (3) T2\*-weighted images from 9 experimental runs; (4) a B<sub>0</sub> field map sequence for subsequent geometric unwarping of T2\*-weighted images.

After acquiring T1-weighted anatomical images (TR=1630 ms, TE=3.11 ms, TI = 1100 ms, voxel size = 0.9 mm x 0.9 mm x 1.0 mm, flip angle 15°), we collected T2\*-weighted images using a gradient-echo echoplanar pulse sequence (TR=3000 ms, TE=25 ms, voxel size=2 mm x 2 mm x 2 mm, flip angle = 90°, 41 axial slices). Slices were collected at 20° counter-clockwise to the anterior commissure to posterior commissure (AC-PC) plane. This slice orientation was chosen so as to maximize the volume of anterior temporal as well as temporo-parietal cortex within the acquisition, since the former region is particularly prone to signal loss from proximity to sinuses (known as “susceptibility artifact”, Patterson et al., 2007).

FMRI data were pre-processed offline using the AFNI (Cox & Jesmanowicz, 1999) software package. The first four volumes of each functional run were removed so

as to allow the signal to reach steady-state magnetization. Functional images were slice-time corrected, and a motion correction algorithm employed in AFNI registered all volumes to a mean functional volume. Images were then unwarped via  $B_0$  field maps (using FSL software; <http://www.fmrib.ox.ac.uk/fsl>) to reduce non-linear magnetic field distortions. We applied a high-pass filter of 0.01 Hz on each run to remove low frequency trends. Images were transformed to Talairach standardized space (Talairach & Tournoux, 1988) and voxels were resampled in the process to 3.5 mm x 3.5 mm x 3.5 mm.

## Analysis

### *ROIs and image analysis*

Using AFNI (Cox, 1996), functional data were registered to the individual subject's anatomical MRI. Transient spikes in the signal were removed using AFNI's 3dDespike. Our *a priori* ROIs were left anterior temporal pole and left angular gyrus, which we delimited using AFNI's CA\_ML\_18\_MNIA atlas. Our anterior temporal ROI spanned labels "left temporal pole" and "left medial temporal pole", while our angular gyrus ROI circumscribed only the atlas's "left angular gyrus" ROI (see Fig. 5).

For those voxels within a given ROI, multiple regression was used to generate parameter estimates ( $\beta$ ) representing each voxel's activity in each stimulus item condition within subject. Voxels'  $\beta$ s were calculated by convolving all variables with a gamma-variate hemodynamic response function and entering them into a general linear model (GLM) (AFNI; Cox & Jesmanowicz, 1999). Motion estimates were included as

regressors of no interest. After implementing our voxel selection criteria (see below), the per-voxel  $\beta$  values were entered into the similarity analysis.

### *Voxel selection and similarity analysis*

In order to query the similarity space of the various composition conditions in each of our ROIs, we first had to identify which voxels to include in subsequent similarity analyses. For each subject, we selected those voxels which varied the most with respect to the contrast between compositional phrases (e.g. *eats meat*) and one-word items (*fghjl eats*), using a GLM at each voxel within bilateral ATL and bilateral AG. Because even a liberal t-threshold on the compositionality contrast revealed no differential activity in right ATL, we did not further analyze this region. The 100 voxels with the highest unsigned (positive and negative; see Discussion for motivation for including both) t-values from the compositionality localizer for each subject for each remaining ROI made up the pattern template for the similarity analysis. Having chosen which 100 voxels would constitute our per-subject, per-ROI vectors, we then modeled each stimulus event as a unique regressor in a GLM, and entered the stimulus item GLM  $\beta$ s for those previously chosen 100 best voxels into vectors of 100 values per condition per ROI per subject (see Fig. 6). We then conducted a correlation analysis over these pattern vectors using Pearson's  $r$ . Our initial regions of interest (ROIs) included bilateral ATL and bilateral AG. Because right AG demonstrated no significant pattern similarity results in any voxel group, only results for left-lateralized ATL and AG are reported below.

While we report results for the 100-voxel set below, we also used three other



voxel sets (50, 200, and 500) to confirm that our results were not idiosyncratic to an arbitrary feature selection criterion. These were entered into the similarity analysis as described for the 100 best voxels, and we report results for these voxel sets in Appendix Tables 1-6.




Figure 5. ROI boundaries of left ATL and left AG.




Figure 6. Diagram of similarity analysis. A vector of the N (50, 100, 200, and 500) best voxels'  $\beta$  values for a given condition (e.g. +verb\_arg, or *eats meat*) was correlated against another vector of  $\beta$  values from the same voxel array for a different condition in the same composition set (e.g. +verb\_adj, or *eats quickly*).

#### *Neural similarity scores for individual nouns and verbs*

In order to have a neural measure to relate to the Amazon Mechanical Turk similarity norms we had calculated for each noun and verb (see Section 2.2.2), we calculated the correlations from the 100 voxels previously chosen for analysis between evoked patterns of every noun-containing phrase and every verb-containing phrase. That is, we calculated multi-voxel pattern correlations for the phrase pairs (meat, eats meat), (meat, tasty meat), (meat, with meat), (eats meat, with meat), (tasty meat, with meat), (tasty meat, eats meat), (eats, eats meat), (eats quickly, eats meat), and (eats, eats quickly). We then correlated these noun- and verb-specific neural similarity scores with the respective pairwise similarity norm scores from the Amazon Mechanical Turk survey. We predicted that the neural similarity score, which was higher the more consistent a given noun or verb pattern was when being composed with other words, would be negatively correlated with the dissimilarity norm score, which was higher for nouns or

verbs the meanings of which differed more depending on what words the nouns/verbs were composed with.

## Results

### *Categorical similarity analyses of shared verb and shared composition type*

We evaluated the similarity of the multi-voxel patterns evoked by each item across the set of voxels that differentially responded to compositional and non-compositional conditions in a given ROI. That is, we chose the 100 voxels per person with the highest unsigned (positive and negative) *t* values from the composition-vs-non-composition contrast in each ROI, estimated the beta value for a given item at each voxel, calculated the correlation across the 100 voxels between pairs of items that shared a common concept, and averaged those correlations across the 36 items. Specifically, we contrasted two hypotheses of the role of left AG in two-word composition: (1) that the left AG is specialized for combinations involving argument-type composition, and/or (2) that the left AG is specialized for event/verb semantics.

We compared correlations between pairs like (a) *eats meat* and *with meat*, where both a noun is shared and putative composition type (argument saturation) is shared with pairs like (b) “*eats meat*” and “*tasty meat*,” where only a noun is shared, and pairs like (c) “*eats meat*” and “*eats quickly*,” where the verb is shared. We found a main effect of *Condition* ((a)shared noun and composition; (b) shared noun only; (c) shared verb only) in left AG ( $F(2,48) = 6.23, p=0.004$ ). As shown in Figure 7, the shared verb correlations ( $r = 0.17$ ) are significantly greater than correlations between noun-sharing phrases ( $r = -0.03$ ; Welch’s  $t(25.1) = 8.03, p<0.001$ ) or noun+composition-sharing phrases ( $r = 0.04$ ;

$t(24.6) = 5.29, p < 0.001$ ). The only correlation significantly different from chance in left AG is that between verb-sharing phrases. This pattern was robust across several other voxel selection criteria (see Tables 1 and 2 in Appendix). These findings favor the hypothesis that the left AG is specialized for verb semantics.




Figure 7. Pairwise correlations between relevant pairs of word phrases in 100 best voxels in left AG. Error bars are 95% confidence intervals.

The same set of comparisons between pairs of two-word phrases in left ATL yields a different pattern from that seen in left AG (see Tables 3 and 4 in Appendix for

left ATL comparisons). The noun-sharing phrases and the noun+composition-sharing phrases were significantly greater than both chance (Appendix Table 3) and the verb-sharing condition (Appendix Table 4), though this was not robust across other voxel selection criteria. *Prima facie*, this might suggest that left ATL is tuned to information carried on the noun and to shared composition type, consistent with Westerlund et al., 2015. However, the overall pattern of ranked correlations was not robust across voxel sizes in left ATL as it was in the case of left AG: that is, the *eats meat ~ with meat* was not consistently the highest correlation across voxel groups in left ATL, and there was much greater variability in the ordering and magnitude of correlations in left ATL. More importantly, because the current stimuli were not well suited to exploring the full similarity profile of compositionality in left ATL, particularly because we did not have a pair of phrases that shared both adjunct-type composition and either a noun or a verb (see intro), we cannot tell from these data alone whether left ATL is sensitive to composition type, shared noun, or some other dimension of conceptual similarity.

In order to compare the overall similarity structure in left AG with that in left ATL, we performed a 3x2 ANOVA over the Fisher's z-transformed subject means of Pearson's correlation values in the 100-voxel group, with factors *Condition* (shared noun and composition; shared noun only; shared verb only) and *ROI* (left ATL; left AG). We found a significant interaction between *Condition* and *ROI* ( $F(2,96) = 5.02, p=0.008$ ).

*Continuous similarity analyses between subjects' pairwise similarity rankings and neural similarity scores for pairs of two-word phrases*

Implicit so far in our treatment of the categorical dimensions “+/- shared verb” and “+/- shared composition” is the assumption that the only change between pairs of word phrases in a given set is the words constituting those phrases. However, it might be that the meanings of our two-word phrases differ based on factors related to a syncategorematic “context” of the word phrases. On the one hand, meanings could vary idiosyncratically; for instance, the meaning of “*meat*” in *eats meat* vs. *with meat* might not differ much between these two contexts, but “*rice*” in *grows rice* vs. *on rice* might differ much more from one context to the next (where *on rice* calls to mind rice in an edible state, while *grows rice* is more evocative of farming the plant). On the other hand, nouns and verbs might differ systematically in their “changeability” across phrasal contexts; that is, noun-sharing phrases might be more or less variable around a noun than verb-sharing phrases are around a verb.

In order to check for this potential source of similarity structure in our stimuli, we first looked at the pairwise norming scores taken from the Amazon Mechanical Turk survey. We found that there was a significant difference ( $p < 0.05$ ) between noun-sharing pairs and verb sharing pairs ( $M = 3.16$  and  $M = 3.44$ , respectively, on a Likert scale of 0-7, 7 being maximal difference between a pair of word phrases), indicating that, on average, noun-sharing phrases were rated as more similar than verb-sharing phrases. This normed similarity measure was not predictive of neural similarity in either left ATL or left AG at any feature level (50, 100, 200, or 500) when taking verb- and noun-based correlations together, nor was there a main effect of noun- vs. verb-sharing on neural similarity scores across ROIs. However, we found a significant correlation between AMT similarity norms

of verb-containing phrases and neural similarity scores in the 100 best voxels in left AG ( $r = -0.12, p < 0.05$ ; see Tables 5 and 6 of Appendix for AG and ATL correlations across voxel selection criteria).

## **Discussion**

While functional neuroimaging studies have made great strides in mapping brain areas involved in language processing, a model of the neural bases of semantic processing is still in its nascence. This may be in part due to the fact that the cognitive neuroscience of semantics does not always utilize linguistic theory. Indeed, it sometimes does not need to. After all, the legacy of model-theoretic semantics has concerned itself primarily with formalizing “the metaphysics of truth in natural language” rather than the various constraints on language processing or the representations of concepts (Seuren, 2009). Cognitive neuroscientists are often more explicitly interested in semantics as it deals with binding sensorimotor features of object-concepts, or how different categories of objects are represented with regard to action-oriented events, e.g., function vs. manipulability (Yee, Drucker, & Thompson-Schill, 2010), etc. For instance, early attempts to define “category-specific” regions of cortex using lesion studies provided evidence that damage to ATL was associated with deficits specific to the knowledge of living things (Gainotti, 1996), while damage to left temporo-parietal junction affected knowledge of man-made artifacts (e.g. *wrench*, *hammer*, etc. which, interestingly, are often also verbs) (Tranel, Damasio, & Damasio, 1997). However, such emphasis on accounting for sensorimotor and action-based properties of language may neglect the more abstract significations our language is capable of expressing, and thus miss

potential means of generalizing certain embodied aspects of cognition. In this study, we expanded the purview of conceptual semantics from the domain of object concepts and action semantics to more abstract dimensions – here, composition type, argument structure, and event semantics.

We were particularly interested in two regions of the brain – the left ATL and left AG – implicated neuroanatomically as “convergence zones,” and also as “semantic hubs” for their involvement in processing compositional language. We found that the left AG displayed a markedly different pattern-similarity profile from that of left ATL. The only dimension of stimulus similarity that produced a detectable effect on neural similarity in AG was shared verb, and by extension, shared event. Left AG appears to be invariant to composition type, and therefore the level at which AG tracks argument structure may not be as general as that described by “argument-type” composition as denoted above, but rather may explicitly subserve *verb* argument structure, namely thematic relations. This is an important distinction, as there are many more types of argument structure in language than verbs and their arguments, and these data now behoove us to examine AG’s selective involvement in composing verbs and their arguments.

#### *Angular gyrus and thematic relations*

It is still unclear exactly what information carried on the verb might be engaging AG. Evidence that bilateral AG activity is parametrically modulated by the valency of a verb – that is, the number of arguments a verb can take (Meltzer-Asscher et al., 2013; Thompson et al., 2007, 2010) – suggests that the AG may read out the syntactic



complexity of a verb constituent, rather than, or in addition to, the semantic content of the verb itself. Left AG has also been implicated in the detection of syntactic errors (Embick et al., 2000). However, AG is also involved in the processing of connected discourse as opposed to unrelated sentences (Fletcher et al., 1995; Homae et al., 2003; Xu et al., 2005), suggesting that AG participates in the construction or analysis of event semantics. Thus it may be that AG acts as an interface between semantic memory and syntactic structure, mapping semantic-thematic relations onto structural constraints surrounding verbs and their arguments. Indeed, electrophysiological and neuroimaging studies support an overlap between (morpho-)syntactic and semantic-thematic verb violations. Kuperberg et al. (2008) compared three different types of verb violations: (1) semantic-thematically violated verbs (e.g. “at breakfast the eggs would eat”) (2) morphosyntactically violated verbs (e.g. “at breakfast the boys would eats”) and (3) real-world violations (e.g. “at breakfast the boys would plant”). They found that, unlike real-world violations, both semantic-thematic and morpho-syntactic violations elicited activity in a frontal/inferior parietal/basal ganglia network, in accord with previous electrophysiological findings that semantic-thematic and syntactic violations evoked P600 event-related potentials highly similar in latency and scalp distribution (Hoeks et al., 2004; Kuperberg, 2007). The authors concluded that this frontal/AG/basal ganglia activity reflected attempts to integrate structural constraints of the verb with semantic properties of the Agent NP argument (Buccino et al., 2001; Chao & Martin, 2000; Damasio et al., 2001; Fogassi et al., 2005).

Evidence from lesion analyses also suggests that such thematic role knowledge is

privileged in bilateral AG. The literature on semantic knowledge has long distinguished between so-called taxonomic semantic knowledge, or knowledge of shared/commensurate features, and thematic semantic knowledge, or knowledge of the relations between object-concepts (crucially from different taxonomic categories) that play complementary roles in events. Speakers' semantic errors tend to reflect either taxonomic fidelity (that is, uttering an incorrect word, but one which has commensurate features, such as when "apple" is named as "pear") or co-occurrence fidelity (that is, uttering "dog" when "bone" was intended, reflecting the thematic relation between "dog" and "bone") (Schwartz et al., 2011). Schwartz and colleagues (2011) analyzed the error typologies of 86 individuals with post-stroke aphasia and conducted voxel-based lesion-symptom mapping (VLSM) on each error type separately (with shared variance between error types regressed out). Taxonomic errors were mapped to left ATL lesions, while thematic errors were localized to left AG. This double dissociation between ATL and AG supports the view that the ATL and AG support distinct semantic computations, corresponding roughly to feature-based and relation-based operations, respectively (but cf. Lewis, Poeppel, & Murphy (2015) for evidence that both taxonomic and thematic associations engage AG, while ATL subserves taxonomic associations specifically).

It is interesting to note that word pairs in a thematic error, such as "dog" and "bone," can be described as related by virtue of some implicit verb/event; in this case "chews" or "buries," etc. That is, thematic knowledge is precisely knowledge of verbs and their arguments. Our current study provides evidence that verbs in particular, not nouns, and not just any argument-type composition, may indeed be the representational

substrate of semantic knowledge in AG.

*Semantic representations or semantic control?*

There is some debate, however, as to whether AG is a hub for mapping syntactic and semantic representations, or if it is rather part of an extended regulatory “semantic control” network. Indeed, the functional heterogeneity of bilateral AG apparently defies neat description. While AG activity is most consistently and robustly elicited by tasks involving semantic processing, both in auditory and visual modalities (see Seghier, 2012, for review), AG is also implicated in the default network, where AG is deactivated during goal-oriented tasks (Shehzad et al., 2009); number processing (Dehaene, Piazza, Pinel, & Cohen, 2003); attention and spatial cognition, where AG may play a role in shifting attention toward particular stimuli having greater salience in terms of motion, value, emotion, and meaning (Gottlieb, 2007); and verbal working memory retrieval and episodic memory retrieval (Vilberg & Rugg, 2008). Generally, AG activation increases with the amount of semantic information that can be retrieved from a given input, whether exogenously generated or self-generated during mentation (Binder & Desai, 2011; Seghier, 2012).

The ostensible functional heterogeneity of AG in the literature may arise more from ROI definitions that (unintentionally) obscure neuroanatomical divisions within AG. In a meta-analysis of studies comparing semantic tasks with high-vs.-low demands on executive control, Noonan et al., 2013, found a functional divergence between dorsal AG (bilateral, including dorsal/anterior AG and boundaries with superior marginal gyrus

(SMG) and inferior parietal sulcus (IPS)) and left mid AG (somewhat closer to PGp than PGa), with respect to executive and representational roles in semantic processing. Dorsal AG showed reliably greater activation in high >low semantic conditions, and was characterized as allocating attention to semantic representations in a task-dependent and goal-driven manner. This characterization is not in itself inconsistent with a model of AG as a site of conceptual combination, as such compositional operations require selective attention to certain properties of events in order to construct higher-order concepts. However, dorsal AG's role in semantic control was contrasted with mid AG, the activity of which was modulated by the semantic representational content of stimuli even when matched on task demands. While mid AG is associated with the "default mode network" (Raichle et al., 2001), and thus shows more positive activation in the absence of a task, it shows more negative activation for abstract as compared with concrete concepts (Binder, Westbury, McKiernan, Possing, & Medler, 2005; Wang, Conder, Blitzer, & Shinkareva, 2010), and more positive deflection from baseline for semantic as compared with phonological decisions matched on executive demands (Binder et al., 1999, 2009).

The left and right AG ROIs drawn in our study encompassed both dorsal and mid AG regions, but voxel features across subjects were highly dispersed across the ROI in standard space such that we could not determine a difference in pattern similarity profiles between dorsal and middle aspects of AG. It is thus possible that the patterns we report here captured a combination of executive demand and semantic-thematic representation similarities.

However, if there were differences across our dimensions of interest in executive

demand, or syntactic/semantic complexity broadly construed, we might expect to have seen evidence for this in univariate contrasts. Yet we found no evidence of significant univariate differences across our dimensions of interest: there were no significant clusters in our ROIs for the noun-based vs verb-based phrase contrast (where *eats quickly* and *eats* were verb-based phrases, and *tasty meat*, *with meat*, and *meat* were noun-based phrases), and neither did the argument-vs-adjunct-type contrast reveal any significant differences in either ROI. Nevertheless, we cannot rule out the possibility that the voxel pattern correlations in our AG ROI also reflect semantic control processes.

#### *Limitations, future directions, and conclusions*

Our claim is primarily in regard to composition, rather than lexical effects per se, because we are investigating the changes in the voxel pattern “template” of a given base word – e.g. *eats* or *meat* – when composed with a function head, an argument, or an adjunct. The voxel selection criteria we used specifically targeted composition-responsive voxels; that is, voxels the activity of which changes maximally when adding another word to a given base word (e.g. adding the argument *meat* to *eats*.) Given that our dependent measure is change (or similarity) in voxel *patterns*, rather than univariate changes in activity across a cluster or ROI, we chose to include voxels that responded both maximally positively and maximally negatively to an instance of composition. Sampling the ends of both tails allows us to capture a greater range of possible variance in voxel patterns, a range that would be limited if we only looked at positive voxel changes. In addition to facilitating the pattern analysis, the inclusion of composition-negative voxels was motivated by emerging evidence that a region of left AG is part of

the so-called “default network,” as discussed above. Seghier et al. (2010), found that the left mid AG was a region within the default network that responded more negatively to perceptual decisions than semantic decisions, though both stimuli elicited a negative divergence from baseline. Mid AG has also been found to be less active for more “difficult” semantic stimuli: e.g. more negative activation for items with longer decision or processing times (Binder et al., 2009; Noonan et al., 2013). Since compositional items might be understood as more semantically “rich,” or more “difficult,” the mid AG region might be expected to index compositionality, but in the composition-negative direction. We did not want to exclude this region when we cast our net over AG voxels, and so we included both composition-positive and composition-negative voxels. It should be noted that analysis of similarity patterns derived from composition-positive voxels alone yield highly similar profiles in both left ATL and left AG, though the shared verb correlation was slightly weaker in left AG.

As mentioned above, the voxels we selected for similarity analysis were highly spatially distributed across subjects, and we were not able to define a particular region of AG (dorsal vs. med) driving similarity patterns. Further study into the functional differences (1) between dorsal and mid AG and (2) between composition-positive and composition-negative voxels across the brain, may clarify whether voxels responding negatively to composition reflect attention-based or representation-based information about verb composition.

It is also important that we discern lexical effects from composition effects when characterizing the role of left AG in verb semantics. When examining voxel pattern

similarity across composition-positive and composition-negative voxels, we do not find a correlation between the non-compositional one-word item *eats* and the compositional phrase *eats meat*. However, as the voxels were selected to maximize the differences between exactly this contrast, a lack of correlation is not only unsurprising, but expected. Instead of choosing voxels most sensitive to the composition vs. non-composition contrast, we collapsed the two-word compositional and one-word non-compositional conditions together and contrasted this combined “word condition” with the ITI fixation period. We then selected the 100 voxels and the 500 voxels within left AG with the highest positive t statistic for combined word task over fixation baseline. This more agnostic selection criterion allowed us to assess the *eats* ~ *eats meat* correlation and also compare it with the *eats meat* ~ *eats quickly* correlation. Using the word-vs.-fixation selection criterion, these correlations were neither significantly different from chance nor from one another. This indicates that the word task>baseline contrast is not optimal for testing the substrate of our verb-based effect, and that this effect is indeed driven by composition-sensitive voxels

While this study provides evidence that left AG contains patterns representing information specific to verbs, regardless of whether these verbs are composed with adverbial adjuncts or noun arguments, we cannot entirely rule out the possibility that AG is also involved with argument-type composition in general. In addition to “eats meat”-type verb phrases, the other argument-type compositional phrase included here was the prepositional phrase (e.g. “with meat”). Prepositions have several unique properties. High-frequency, semantically vacuous/impoverished prepositions might have a very

different combinatorial effect than adjectives or verbs when composing with nouns (“tasty meat,” “eats meat,” respectively). Indeed, the preposition is little more than a function word, and lacks the semantic content carried on adjectives and verbs. Not only do the prepositions in our stimuli set have the highest average item frequency, but prepositions as a class may also combine with many more surface forms than either nouns or adjectives. This “compositional diversity” may render prepositions, and prepositional phrases, qualitatively different from the other parts of speech used here, and this diversity may make extraction of stable patterns from the prepositional phrase items less likely. Nevertheless, further study is needed to examine whether preposition function heads engage AG in the same way we found verbs do. While this study provides evidence that a shared *argument* (*meat* in *eats meat* and *with meat*) is not sufficient to drive pattern similarity in AG, it does not query whether a shared *preposition* (*with* in *with x* and *with y*) is possibly sufficient to drive similarity in the same way a shared verb is (*eats* in *eats meat* and *eats quickly*).

The current study only investigates cases of minimal composition: that of two words isolated from a sentence or discourse. However, it is unlikely that AG is only tracking this level of composition. There is abundant evidence that AG may engage in domain-general event processing in event structures as broad as discourse and in non-verbal depictions of events. Indeed, both ATL and AG are best described as “hubs” at a domain-general level. While the current study did not directly test the manner in which left ATL might subserve feature-based combination, a large body of literature suggests as much. In contrast to ATL, we find increasing evidence that AG is engaged in semantic



integration of relation-based event structure, and we must now consider whether this distinction between ATL and AG is ultimately reducible to the well-attested difference between the ventral “what” pathway (the combinatorial hub of which is the ATL) and the dorsal “how/where” pathway (the integrational hub of which is the AG) (Binder & Desai, 2011). The AG is surrounded by the dorsal spatial attention networks, the posterior temporal regions involved in motion perception, and the anterior parietal regions involved in representing action (Kravitz, Saleem, Baker, & Mishkin, 2011). While AG may have originated as a dorsal “where/how” convergence zone of spatial, goal-oriented, and action information, it may have been co-opted by language to represent increasingly abstract relational information. These relations might be learned merely by tracking co-occurrence statistics (“dog” often co-occurs with “bone”; “eggs” often co-occur with “breakfast), or, more likely, these thematic relations are learned part and parcel of hierarchical structures arising in natural language syntax. Thus, the emergence of event and argument structure in thought and language may have been an extension of the already extant dorsal pathways underpinning action and goal understanding.

It is interesting to note, however, that while the fronto-temporal language network may have evolved to be strongly left-lateralized, it is less clear the degree to which right and left AG diverged with regard to processing events and representing thematic relations. Graves et al. (2010) offer a connectionist account of how noun-noun compounds, such as “lake house,” when compared with their less compositional reversals (“house lake”), show differential activity in right AG but not left AG. They suggest that left and right AG can be modeled as attractor networks, where such a network is said to

settle into an attractor basin when it optimizes the error space in the mapping between inputs and outputs. Whereas left AG is suggested to have relatively narrow attractor basins, reflecting highly specific and constrained mappings between words and meanings, right AG may contain wider, shallower basins, representing more extensive overlap in meanings. This would accommodate “looser” meanings, and thus enable the interpretation of compounds like “dog bone” that lack the explicit (morpho)syntactic information (i.e. “a bone that a dog chews on”) that would otherwise aid in resolving the relation between the two nouns in the phrase. This attractor network account of the difference between left and right AG accords with Beeman and colleagues' (1994) “coarse semantic coding hypothesis” of the right hemisphere. In this study, “summation primes,” three words weakly related to a target word, were found to better prime a target when the triplet was presented to the left visual hemifield (right hemisphere, RH) than the right visual hemifield (left hemisphere, LH), while the converse was true for “direct primes,” where there was one strongly associated prime flanked by two unrelated primes. This was taken as evidence that RH contains larger semantic fields weakly activating concepts more distantly related to an input word, whereas LH contains smaller semantic fields that conservatively activate concepts highly related to an input word. This distinction might account for why we found evidence of verb-specific pattern conservation in left AG but not right AG. It may be that left AG subserves specifically strong thematic relations, while right AG weakly activates to a wider variety of compositional items. Further study is needed to examine whether such strong vs. weak relations between words in compositional phrases might differentially engage right and left AG.

We suggest that multi-voxel pattern similarity analysis is uniquely suited to address such questions. Our study has demonstrated the sensitivity of this technique to compositional operations even at the level of minimal two-word phrases. We are only just beginning to characterize the AG with respect to its involvement in semantic composition, and this study suggests that the “feature-function” dichotomy may be a fruitful distinction in beginning to operationalize the compositional processes occurring in both ATL and AG.

#### IV. A MATTER OF AMBIGUITY? USING EYE MOVEMENTS TO EXAMINE COLLECTIVE VS. DISTRIBUTIVE INTERPRETATIONS OF PLURAL SETS

##### **Introduction**

Human language processing is remarkably fast, and there is a growing consensus that this is because language comprehension is not only a function of receiving input, but also of anticipating and hypothesizing structures to be checked against that input. The facilitation afforded to the language processor when a prediction is correct, however, is balanced against the risk that the predicted linguistic form will be incorrect and require re-analysis. Given this trade-off, it may sometimes favor the processor to abstain from a prediction, for instance in cases when the preceding linguistic information has high Shannon entropy or is otherwise under-predictive. However, even in high-uncertainty contexts, the processor may often be induced to commit to a representation/interpretation in order to proceed with an incremental parse. As Frazier et al.'s Minimal Semantic Commitment (MSC) formalizes it, the processor will commit to a representation in the absence of specific evidence for that representation “when faced with alternative decisions that are grammatically incompatible with each other or when the failure to make a decision would violate a grammatical principle” (p. 88). These forced predictions differ from cases when the processor is merely opportunistic, committing to an interpretation where pragmatic factors, rather than purely grammatical inducements, provide evidence in favor of one interpretation over the other. Whereas the grammatical constraint on the processor prompts a “necessary decision,” the pragmatic factors produce an “invited decision.” Both these types of “decisions” operate over representations that are termed *ambiguous*, since the possible interpretations of that representation are finite

and limited. However, when such possible interpretations are not grammatically or pragmatically constrained – e.g. when the grammar does not require a particular representation or feature specification– the representation is termed *vague*. Vague representations do not prompt the processor to commit to a representation, and a representation can remain underspecified at no cost to the processor.

The following two sentences together illustrate a case in which certain linguistic material constitutes neither a “necessary” nor “invited” decision of ambiguity:

5 a. John ate.

b. John ate quickly.

The difference in meaning between (5a,b) resides in the adverb “quickly.” In sentence (5a), the *manner* in which John ate is left unspecified, and yet a reader will find sentence (5a) perfectly interpretable without knowing the manner in which John ate. On the other hand, if the sentence lacked information specifying other propositional content, including, for instance, the number (singular) or tense (past) of the verb, this sentence would be grammatically uninterpretable. Thus, while some information about the verb must be determined in order to parse the sentence, other features ostensibly need not be, as in the case of the adverbial adjunct “quickly” (5b). Adjunct information, such as the manner in which John ate, is therefore characterized as *vague*, rather than *ambiguous*.

While ambiguous representations require the processor to commit to one representation, which is then abandoned and re-analyzed if later found incorrect, vague representations incur no such processing costs. Given these predictions, Frazier et al. (1999) used eye-tracking measures of reading times to compare processing loads of sentences that were explicitly distributive (6a), explicitly collective (6b), and locally

indeterminate at the predicate (6c,d). 6a and 6b indicate early on in the sentence whether the “cake-eating” event is distributive (i.e. John is eating a piece of cake and Bill is eating a (separate) piece of cake) or collective (John and Bill are participating in the same cake-eating event), while (6c,d) specify this information only at the end of the sentence.

Frazier et al. state that if the “decision” as to whether an event is collective or distributive is a matter of vagueness, then (6c,d) should not incur any extra processing cost relative to (6a,b), since the presence of “each/together” would not prompt any revision on an earlier collective/distributive interpretation. If, on the other hand, the distinction is a matter of ambiguity, then a decision as to whether the event was collective or distributive is prompted earlier, even in the absence of the disambiguating “each/together” adverb. In the latter case, the sentence-final adverb in either 6c or 6d might incur a processing cost reflecting the revision of an earlier commitment to a collective/distributive reading.

- 6 a. John and Bill each ate a piece of cake.
- b. John and Bill together ate a piece of cake.
- c. John and Bill ate a piece of cake each.
- d. John and Bill ate a piece of cake together.

Finding increased processing load associated with sentences like (6c), Frazier et al. claimed that the distinction between the distributive and collective interpretations was one of ambiguity and not vagueness. They note a prevailing theory of distributivity (Heim et al., 1991) in which the distributive is the marked reading and involves stipulating a distributive operator *D* (the spell-out of which is “each”). Given the distributive representation is more complex, an early ambiguity should favor a commitment to the collective reading rather than the more complex distributive reading:

- 7 a. [John and Bill] ate a piece of cake. (default: collective)
- b. [John and Bill *D*] ate a piece of cake. (D: distributive operator)

However, given the account that the D operator adds structural complexity to the underlying representation, an increased reading time for the distributive form might be expected regardless of whether the decision itself is vague or ambiguous. Interestingly, Frazier et al. found an increased processing load for explicit as well as locally indeterminate distributives, suggesting that the distributive operator is sufficient to increase processing load, with or without an occasion for revision. They claim, however, that the increased processing load found for locally indeterminate distributives could not be accounted for by the presence of the distributive alone, and suggested that this processing cost was due to the processor having committed to the collective reading and then revising that commitment.

One major limitation of the Frazier et al. study was that it tested a representational hypothesis using processing load measures. In the present study, we improved upon Frazier et al.'s basic design by employing the visual world paradigm. With this method, we monitored the eye movements of participants as they listened to explicit or locally indeterminate collective and distributive sentences and considered collective and distributive scenes on a computer screen. Rather than relying on processing times to infer representational commitments, we tracked which representations were considered along the time course of the sentence. Additionally, instead of finally disambiguating the locally indeterminate sentences with an adverb, we disambiguate all sentences on a sentence-terminal object or objects (for collective and distributive actions, respectively). We also gathered normed ratings for every verb in our critical sentence stimuli such that

we had measures of the verbs' lexical-pragmatic biases for collective vs. distributive events. This allowed us to query separately those verbs that might be lexically more biased towards distributive readings (e.g. *wearing, eating*), potentially interacting with a putatively a priori preference for collective interpretations.

Our first study targeted sentences beginning with conjoined noun phrases (NPs), such as “John and Bill...” (see Table 2 below), since there is now a rich literature investigating collective and distributive readings on these types of NP. In the case of a distributive reading, the “atomization” of a conjoined NP like *John and Bill* into its parts – *John* and *Bill* – is particularly intuitive. While the collective interpretation is formalized as applying a predicate to a “sum” of individuals *John and Bill* (Clifton & Frazier, 2012; Link, 1983; Moxey, Sanford, Wood, & Ginter, 2011), the distributive interpretation applies the predicate separately to the elements *John* and *Bill* (Heim, Lasnik, & May, 1991).

Our second study sought to compare the conjoined NP sentences from Experiment 1 with sentences using simple plural NP subjects, as in *The boys ate a piece of cake* (cf. 6a-d above). There has been some periodic debate as to whether *the<sub>plural</sub>* as in *The boys*, can sensibly accommodate a distributive reading, or whether it is unambiguously collective. Scha (1984; inter alia) argues that plural *the* can only be read as collective, while others treat it as ambiguously collective or distributive (Bennett, 1974; Hausser, 1974). The status of the definite plural amid other NPs is interesting. For instance, consider 4 below (examples from Roberts (1987)):

- 8 a. Four women brought a salad to the potluck.
- b. Jane and Mary brought a salad to the potluck.



- c. Each woman brought a salad to the potluck.
- d. The women brought a salad to the potluck.

While 8a and b are compatible with either a distributive or a collective interpretation, 8c can only be distributive. The plural NP subject in 8d is strongly collective, but it is not clear from this sentence alone that the collective is the only possible reading. In fact, given the case in (9) below (example also from Roberts (1987)), the assumption that *the<sub>plural</sub>* must be collective is defeasible:

9 Every woman brought a dish to the potluck.

The hostess asked those from Acton to bring a casserole.

*The women from Boxborough brought a salad*, and those from Littleton a dessert.

The italicized content in sentence (9) clearly conveys a distributive reading, even though such a reading for sentence 8d might be less accessible. Therefore, it would seem that *The boys* may be subject to the same ambiguity/vagueness between collective and distributive as the conjoined NP *John and Bill* from Experiment 1.

In the experiments reported below, we recorded the eye movements of listeners while they chose between a depicted distributive or collective scene when hearing sentences like those in Table 2 below (*John and Bill [each/together/ Ø] are carrying a ball/box*). If the decision to assign a collective or distributive reading to an event is a matter of vagueness, then looks to the collective and distributive scenes when hearing sentences like *John and Bill are carrying a ball/box* should not diverge until the disambiguation at the end of the sentence (in our case, the object – *ball* or *box*).

However, an earlier divergence in gaze would indicate either a “necessary” or “invited” decision, in which the processor encounters an ambiguity that it must resolve even if it

does not have enough information to choose between the collective and distributive meanings. We find evidence that the collective vs. distributive distinction is a matter of ambiguity, regardless of the type of subject NP. However, while conjoined NP subjects (*John and Bill*) prompt early looks to the collective, even when the verb is relatively more biased towards a distributive reading (e.g. wearing, eating), the plural determiner phrase subjects (*The boys*) produce looks to the distributive. We discuss reasons for this difference in gaze preference below.

## **Experiment 1**

### *Method*

#### *Participants*

Twenty-seven undergraduate students (15 female) participated in this study for course credit. All participants were native English speakers and undergraduates at the University of Pennsylvania. Two subjects were omitted due to incomplete data recording. Additionally, it was noted during the recording process that one subject was fixating on the center of the screen during the task, and since this was a marked departure from the instructions and potentially problematic for our dependent measure, this subject's data were not included in the analysis. Of the remaining 24 participants, 13 were female. This study was approved by the University of Pennsylvania Institutional Review Board, and subjects received course credit for their participation.

#### *Apparatus*

The images were presented on a 17" Samsung screen with 1680x1050 resolution, and the sound was played on Altec Lansing FX2020 ASIO speakers. The sentence onset was synchronized to 500ms after the appearance of the scene pair. The subjects'

responses were recorded using VPIxx ResponcePixx tabletop button box. Right eye gaze was recorded using an Eyelink 1000 eye tracker on a desktop mount at a sampling rate of 1kHz (re-sampled offline to 100Hz).

### *Procedure*

The experimental session began with a calibration procedure, which usually lasted no more than three minutes. Participants were then instructed to fixate on a small dot in the center of the screen, which also corresponded to a point equidistant from two scenes that would appear during the trial period. Participants were told to fixate on this point between trials but to consider the trial scenes freely and at their own pace.

### *Materials and Design*

In this study, we employed a visual world paradigm in which two scenarios, one distributive (two characters engaging in two distinct actions) and one collective (the same two characters engaging in the same activity together) were presented on a screen. While subjects contemplated a scene pair, they listened to a sentence describing a distributive or collective scenario. Subjects were instructed to indicate with a button press which scenario corresponded to the sentence they heard. Sentences were ultimately disambiguated by the object that underwent the action, though sentence items with an adverb (“each” or “together”) were effectively disambiguated earlier.

Target items consisted of 24 sets of six prerecorded sentences. The words and phrases in the sentences were recorded and coded for onset and offset times using PRAAT and FAVE (Boersma & Weenink, 2015; Rosenfelder, Fruehwald, Evanini, & Yuan, 2011). The six items in each of the 24 sets were derived from a base sentence that included a plural subject (e.g. “John and Bill”), a predicate that was ambiguous between a

collective and distributive reading (e.g. “are carrying”), and a direct object that could bias the sentence towards a collective or distributive reading. To extend the ambiguity period for the ambiguous sentences, all objects were modified by a neutral (non-disambiguating) adjective. The six sentences in each of 24 critical item sets were as follows:

Table 2

Conjoined NP	Disambiguator	Ambiguous region 1 (verb)	Ambiguous region 2 (adjective)	Object
1. John and Bill	∅	are carrying	a bright red	ball.
2. John and Bill	each (distributive)	are carrying	a bright red	ball.
3. John and Bill	together (collective)	are carrying	a bright red	ball.
4. John and Bill	∅	are carrying	a bright red	box.
5. John and Bill	each (distributive)	are carrying	a bright red	box.
6. John and Bill	together (collective)	are carrying	a bright red	box.

Table 2: Experiment 1: Conjoined NP sentences. Subjects listened to one of these 6 sentence types while viewing a scene diptych like that shown in Figure 8. Sentences began with conjoined NP subjects and were either not disambiguated until the object word (1,4) or were disambiguated early by distributive “each” (2, 5) or collective “together” (3,6).

Figure 8




Figure 8: Two typical scene pairs, A and B, only one of which would appear for a given item in a given list. See (1)-(6) in Table 2 for list of sentences matched to each scene pair. Scenes were clipart images edited in Gimp and Paintbrush software.

To avoid repetition and priming effects, we presented each given subject with only one of the six variants of each sentence set. In addition to the 24 test items, there were 24 fillers. These fillers were sentences of the same basic form as the test items, and each filler was presented along with two scenes (one target scene and one distracter scene). Eight of the fillers were disambiguated at the subject, eight at the predicate and eight at the object. Thus the targets and fillers were distributed across six lists in a Latin Square design, with targets counterbalanced on disambiguating term, object bias, and the side of the screen the target scenario appeared on. Items were pseudorandomized and item order was matched across lists.

Note that though the object (*ball/box*) was not uttered until the end of the sentence, the relative “collective bias” of a scenario was represented in the visual world independent of the sentence. If subjects find collective or distributive scenes more salient or attractive in some way, regardless of the ambiguity/vagueness of an accompanying sentence, then we should see preference for one or another type of event as depicted in the visual world alone. We conducted an online survey asking subjects (n=12) to indicate which of the two scenes was more “plausible.” The two scenes were paired as they were in the eye-tracking study: one collective, one distributive, and both depicting the same base event (*carrying, eating, etc.*) There were no significant differences in plausibility preference for collective vs. distributive scenes. We were thus able to proceed under the assumption that the collective and distributive scenes themselves did not bias looks independent of the sentence uttered.

Our hypothesis states that a late, adverb-prompted divergence in looks between collective and distributive scenes would indicate the decision was a matter of vagueness, while an earlier divergence in gaze (say, at the verb) would suggest a case of ambiguity. However, we have no explicit cut-off for what is “early” enough to be a forced grammatical decision or “late” enough to be a more opportunistic invited decision. In the case of pre-adverb gaze divergence, examination of other factors around the verb allows us to see how robust an early preference for a collective or distributive scene might be. Therefore, we conducted another survey to assess the verb’s “collective/distributive affordance;” that is, a normed measure of the likelihood a verb denotes a collective or distributive event.

To collect such a measure, we used a methodology introduced by Berent, Pinker, Tzelgov, Bibi, & Goldfarb (2005) and extended by Patson & Warren (2010) to examine how singular indefinite noun phrases in distributed predicates (as in *a ball/box*) can be interpreted as conceptually plural. When shown sentences like those in 2 and 3 in Table 2 above, subjects in the Patson & Warren study indicated how many balls were involved in the event: one or many. Instead of using a real-world object as the indefinite NP in our study, we replaced *a ball* with a nonce word, such as *a wug* (Albright, 2009; Berko, 1958). This allowed us to query subjects' intuitions about the verb independent of the affordances of real-world objects like balls, boxes, pieces of cake, etc. Assessing verb-level biases gives us a more local measure of participants' preferences at the verb, before the disambiguating object is uttered. We conducted a survey on Amazon Mechanical Turk (Buhrmester et al., 2011) asking subjects (n=114) to indicate their confidence on a 1-4 Likert scale that there was/were "definitely one wug" (1), "probably one wug" (2), "probably more than one wug" (3), and "definitely more than one wug" (4). We later used these ratings to identify the twelve most distributive-biased verbs (second quartile of 24 verbs; M=3.11, s.d. 0.34) and examine how their time course compared with that of all 24 verbs.

### *Analysis*

Areas of interest were drawn around the distributive and collective scenes, and using these boundaries, fixations were coded as "collective," "distributive," or "other." Eye-tracking data were first scrutinized for accuracy and track loss: any missing data points and trials which led to incorrect responses were removed from the data before further analysis. A sample was coded as "track loss" if the participant's eyes were



closed or otherwise occluded, and a trial was dropped from analysis if track loss accounted for more than 25% of the frames (this made up less than 5% of the data).

Timecourses of fixations to collective and distributive scenes were plotted for  $\emptyset$ -, each-, and together-disambiguated sentences for each subject. The dependent measure is termed the “collective advantage,” as it is the number of fixations on distributive scenes subtracted from the number of fixations for each sample in each disambiguation condition. Plotted in Figure 9 below are the timecourses time-locked to the onset of the verb (e.g. *carrying*), with a delay of 150 ms to account for saccade planning and execution. We also aggregated collective advantage measures across time windows defined by the onsets and offsets (delayed by 150ms) of the following phrases: *John and Bill | each/together/ $\emptyset$  | are carrying | a bright red | ball/box*. This time window analysis allowed us to time-lock to multiple time points in the sentence, rather than just to the predicate onset, in order to analyze gaze preferences at each phrase. Statistical analysis was conducted over an empirical logit (e-logit) transform of the collective advantage measure for each subject. The e-logit is a quasi-logit transformation optimized for handling cases for which the standard logit is too large or small (when the probabilities approach 0 or 1) (Barr, 2008).

## *Results*

### *Predicate-locked time-courses*

By the time the predicate is uttered, the subject has received disambiguating information for the *together* and *each* conditions, and so has enough information to resolve on a scene even while the predicate is heard. However, subjects hearing the non-disambiguated sentences have no prior disambiguating information. Figure 9 shows plots

of the collective advantage time-locked to the onset of the predicate. Since the predicate is hypothesized to carry the distributive operator, it is reasonable to treat the predicate as the effective disambiguator for the  $\emptyset$  condition. These plots provide a visual confirmation that once the disambiguating term (*each/together*) is uttered (~500ms prior to predicate onset), looks to the collective and distributive scenes diverge rapidly.  $\emptyset$  sentences appear to pattern initially with the *together*, prompting increased looks to the collective scene. Interestingly, though the *each* sentences initially prompt looks to the distributive scene, there is an increase in looks to the collective during the adjectival phrase (~100ms post-predicate). This reconsideration of the other scene is not found in the *together* trajectory, in which looks to the collective scene persist or increase monotonically throughout the evolution of the utterance post-adverb. The non-disambiguated sentences (purple and yellow in Fig 9) prompt an initial increase in looks to the collective during the utterance of the verb, returning to chance at the adjective phrase.

Figure 9

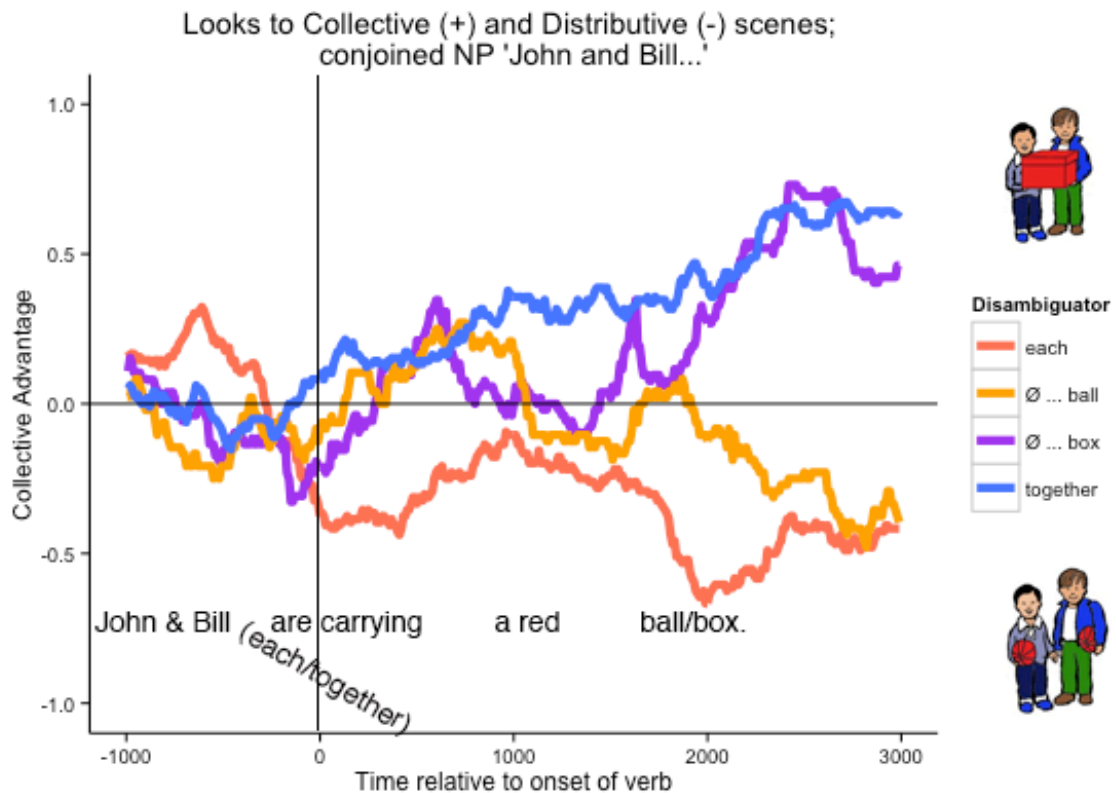


Figure 9: Plots of collective advantage to collective (positive) and distributive (negative) scenes over time for conjoined NP subjects, time-locked to the onset of the verb. Timecourses of each-disambiguated sentences are in red and timecourses of together-disambiguated sentences are in blue. Note that the  $\emptyset$  condition is split into those sentences eventually disambiguated by the object *ball/box* to a collective (purple) or distributive (yellow) scene.

#### *Time window analyses*

We also ran a by-subject 3x3 ANOVA on e-logit collective advantage measures, with factors Disambiguator (*each/together/Ø*) and Time window (*“John and Bill”/“are carrying”/“a bright red”*). Since the disambiguator (*each/together*) time window was

non-existent in the  $\emptyset$  disambiguator condition, this time window was not included in the ANOVA, though we do include it among the planned pairwise t-tests of disambiguator effects in each time window (see Figure 10). There was a main effect of both Disambiguator ( $F(2,46) = 18.33, p < 0.001$ ) and Time window ( $F(2,46) = 4.59, p = 0.02$ ), as well as a significant interaction between Disambiguator and Time window ( $F(4,92) = 10.2, p < 0.001$ ).

Figure 10

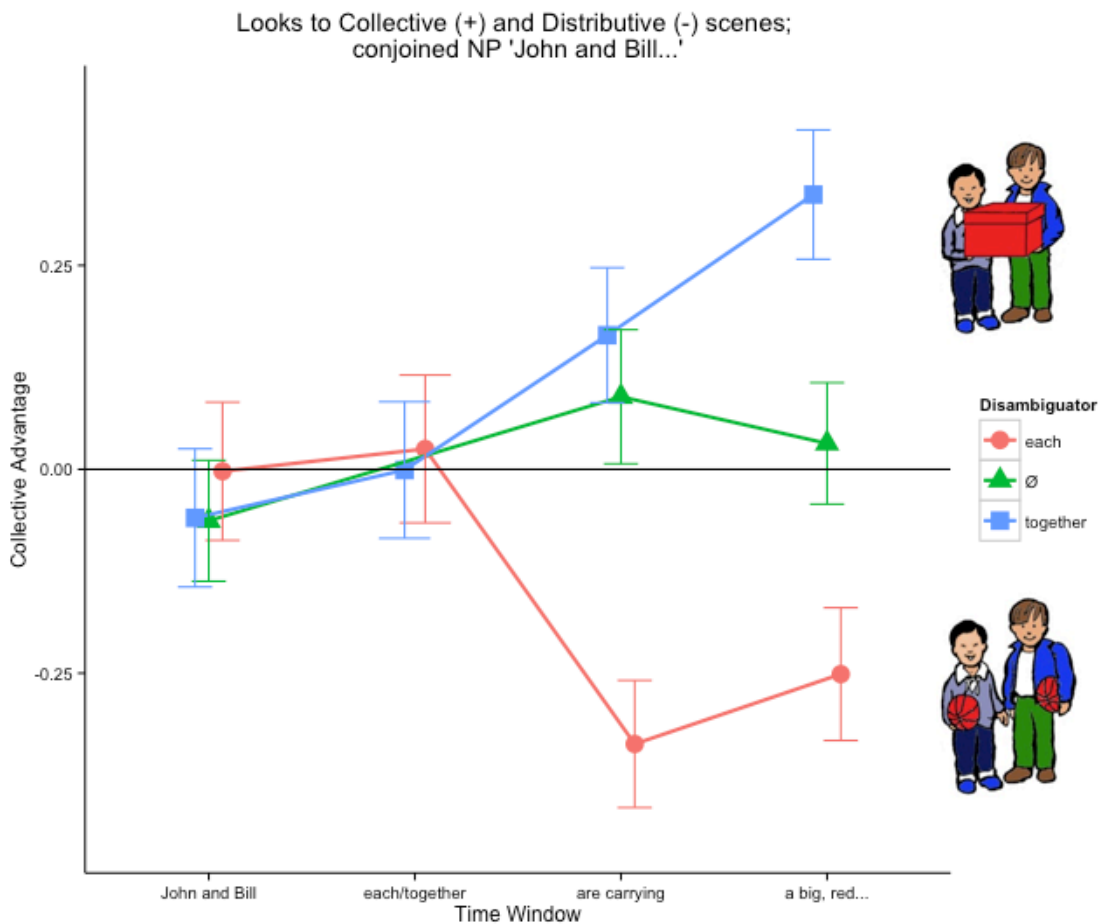


Figure 10: Plots of collective advantage for conjoined NP sentences in four time windows. Average fixations across subjects for each-disambiguated sentences are plotted

in red, together-disambiguated sentences in blue, and non-disambiguated (until object) sentences in green. The final convergence on ball or box is not shown here, and the non-disambiguated sentences that eventually converge on *ball* or *box* are collapsed here (purple and yellow timecourses from Figure 9 are collapsed and indicated in green here.)

We also ran pairwise two-sided t-tests on *each* vs. *together*, *each* vs.  $\emptyset$ , and *together* vs.  $\emptyset$ , using the Holm (1979) method of correction for multiple comparisons. Within the conjoined NP (*John and Bill*) window, we found no pairwise differences in collective advantage between *each*, *together*, and  $\emptyset$ , and none of the disambiguator conditions differed significantly from chance ( $\mu=0$ ). In the *each* vs. *together* time window, we still found no significant divergence in looks, and neither condition was different from chance. In the verb (*are carrying*) time window, however, we see that the *each* condition is significantly more likely to prompt looks to the distributive than the  $\emptyset$  condition ( $p<0.001$ ), *together* condition ( $p<0.001$ ), and chance ( $p<0.001$ ). The *together* and  $\emptyset$  conditions are not significantly different, and the *together* condition directs looks to the collective scene above chance ( $p=0.03$ ) while the  $\emptyset$  condition is marginally above chance ( $p=0.06$ ). Finally, in the adjective time window (*a bright red*), immediately preceding the object (*ball/box*), we find all three disambiguator conditions diverging from one another: the *each* condition is significantly more likely to direct looks to the distributive scene than chance,  $\emptyset$ , and *together* conditions (all  $p<0.001$ ), while *together* is also prompting looks to the collective scene above both chance and  $\emptyset$  ( $p<0.001$ ). The  $\emptyset$  condition is equally likely to direct looks to the collective and distributive scenes in this time window.

*Distributive-biased verbs*

The timecourse and time window data indicate that the *together* and  $\emptyset$  conditions pattern together early on, both favoring the collective scenario at the verb. While this collective preference is trivial for the *together* conditions, as the sentence has already explicitly indexed the collective meaning via *together*, it is interesting that the  $\emptyset$  condition timecourse initially evolves as if there were an implicit “together.” Does this collective preference also hold for the  $\emptyset$  condition of the subset of verbs in our stimuli for which we have evidence of a lexical/pragmatic *distributive* bias?

Figure 11

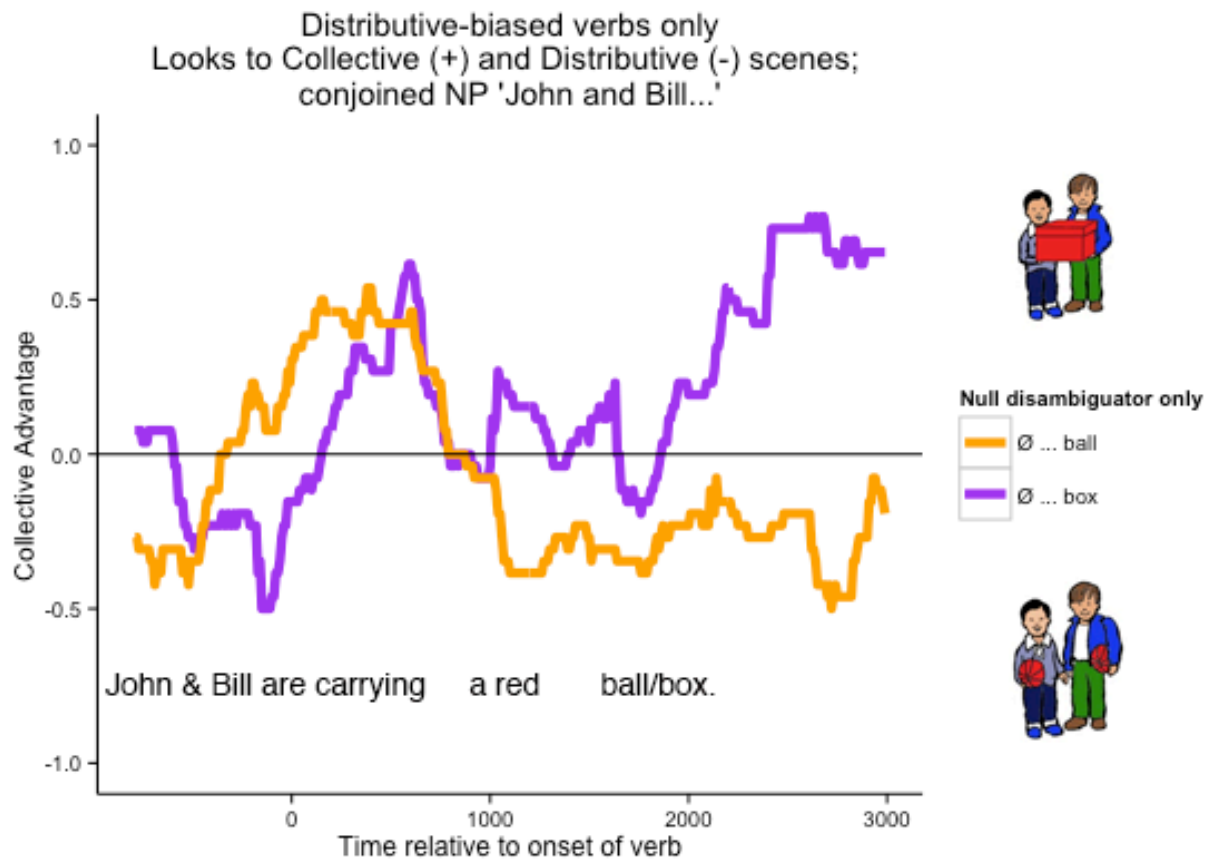


Figure 11: Timecourse of sentences with distributive-biased verbs. Note that *carrying* is not a distributive-biased verb, but is included here for ready comparison with previous figures.

Looking only at the twelve most distributive-biased verbs according to our Mechanical Turk survey scale (Figure 11), we find that even if a verb has a subcategorization preference for distributive events or plural objects, there is still a preference for the collective meaning at the predicate ( $p < 0.1$  against chance).

### *Summary of Experiment 1*

This profile of disambiguator effects across time windows demonstrates that the *together* and  $\emptyset$  conditions pattern together and diverge from the *each* condition as early as the onset of the verb. Thus, even though the ambiguous  $\emptyset$  sentences do not have information discerning between a possible collective or distributive interpretation, they nevertheless prompt looks to the collective, the time course for which is reliably different from the time course of distributive-directed *each* sentences. This provides evidence that the subject has committed to the collective interpretation in the absence of disambiguating information.

In this analysis of the dynamics of looks to collective and distributive scenes, we find evidence that the processor does not passively await information before committing to a collective representation. Rather, the processor appears to default to an underlying collective reading even in the absence of disambiguating information. Moreover, in the absence of an adverbial distributivity operator, the processor still pre-emptively commits to a collective meaning even when we consider those verbs more likely to denote a distributive event (Figure 11).

We now ask whether sentences with conjoined NP subjects like *John and Bill* differ from sentences with another type of subject NP. If the collective reading is indeed the default interpretation, we should also find evidence for preferential gaze to the collective when sentences begin with a plural determiner phrase such as *The boys*.

## **Experiment 2**

### *Method*



### *Participants*

Twenty-seven undergraduate students (18 female) participated in this study for course credit. None of these subjects had participated in Experiment 1. All participants were native English speakers and undergraduates at the University of Pennsylvania. This study was approved by the University of Pennsylvania Institutional Review Board, and subjects received course credit for their participation.

### *Apparatus*

The images were presented on a 17" Samsung screen with 1680x1050 resolution, and the sound was played on Altec Lansing FX2020 ASIO speakers. The sentence onset was synchronized to 500ms after the appearance of the scene pair. The subjects' responses were recorded using VPIxx ResponcePixx tabletop button box. Right eye gaze was recorded using an Eyelink 1000 eye tracker on a desktop mount at a sampling rate of 1kHz (re-sampled offline to 100Hz).

### *Procedure*

The experimental session began with a calibration procedure, which usually lasted no more than three minutes. Participants were then instructed to fixate on a small dot in the center of the screen, which also corresponded to a point equidistant from two scenes that would appear during the trial period. Participants were told to fixate on this point between trials but to consider the trial scenes freely and at their own pace.

### *Materials and Design*

This study employed the same visual world paradigm as in the first study. The exact same scene diptychs were presented to subjects: one distributive scene and one collective scene. This study differed from Experiment 1 only insofar as the sentences

played during presentation of the scenes included a plural determiner phrase, such as *The boys, the girls, or the friends* instead of conjoined NPs like *John and Bill*.

Target items consisted of 24 sets of six prerecorded sentences. The six sentences in each of 24 critical item sets were as follows:

Table 3

Conjoined NP	Disambiguator	Ambiguous region 1 (verb)	Ambiguous region 2 (adjective)	Object
1. The boys	∅	are carrying	a bright red	ball.
2. The boys	each (distributive)	are carrying	a bright red	ball.
3. The boys	together (collective)	are carrying	a bright red	ball.
4. The boys	∅	are carrying	a bright red	box.
5. The boys	each (distributive)	are carrying	a bright red	box.
6. The boys	together (collective)	are carrying	a bright red	box.

Table 3: Experiment 2: Plural NP sentences. Subjects listened to one of these 6 sentence types while viewing a scene diptych like that shown in Figure 8. Sentences began with plural NP subjects and were either not disambiguated until the object word (1,4) or were disambiguated early by distributive “each” (2, 5) or collective “together” (3,6).

## *Analysis*

Areas of interest were drawn around the distributive and collective scenes, and using these boundaries, fixations were coded as “collective,” “distributive,” or “other.” Eye-tracking data were first scrutinized for accuracy and track loss: any missing data points and trials which led to incorrect responses were removed from the data before further analysis. A sample was coded as “track loss” if the participant’s eyes were closed or otherwise occluded, and a trial was dropped from analysis if track loss accounted for more than 25% of the frames (this made up less than 5% of the data).

Timecourses of fixations to collective and distributive scenes were plotted for Ø-, each-, and together-disambiguated sentences for each subject. The dependent measure is termed the “collective advantage,” as it is the number of fixations on distributive scenes subtracted from the number of fixations for each sample in each disambiguation condition. Plotted in Figure 12 below are the timecourses time-locked to the onset of the verb (e.g. *carrying*), with a delay of 150 ms to account for saccade planning and execution. We also aggregated collective advantage measures across time windows defined by the onsets and offsets (delayed by 150ms) of the following phrases: *The boys | each/together/Ø | are carrying | a bright red | ball/box*

## *Results*

### *Predicate-locked time-courses*

Timecourses of conditional looks to collective and distributive scenes prompted by plural determiner (e.g. *The boys*) sentences show a marked departure from the pattern we observed in Experiment 1. Figure 12 shows plots of the collective advantage time-locked to the onset of the predicate. As expected, we see that *each* and *together* sentences

diverge rapidly, with  $\emptyset$  sentences appearing to pattern initially more with the *each* condition in prompting increased looks to the distributive scene.

Figure 12

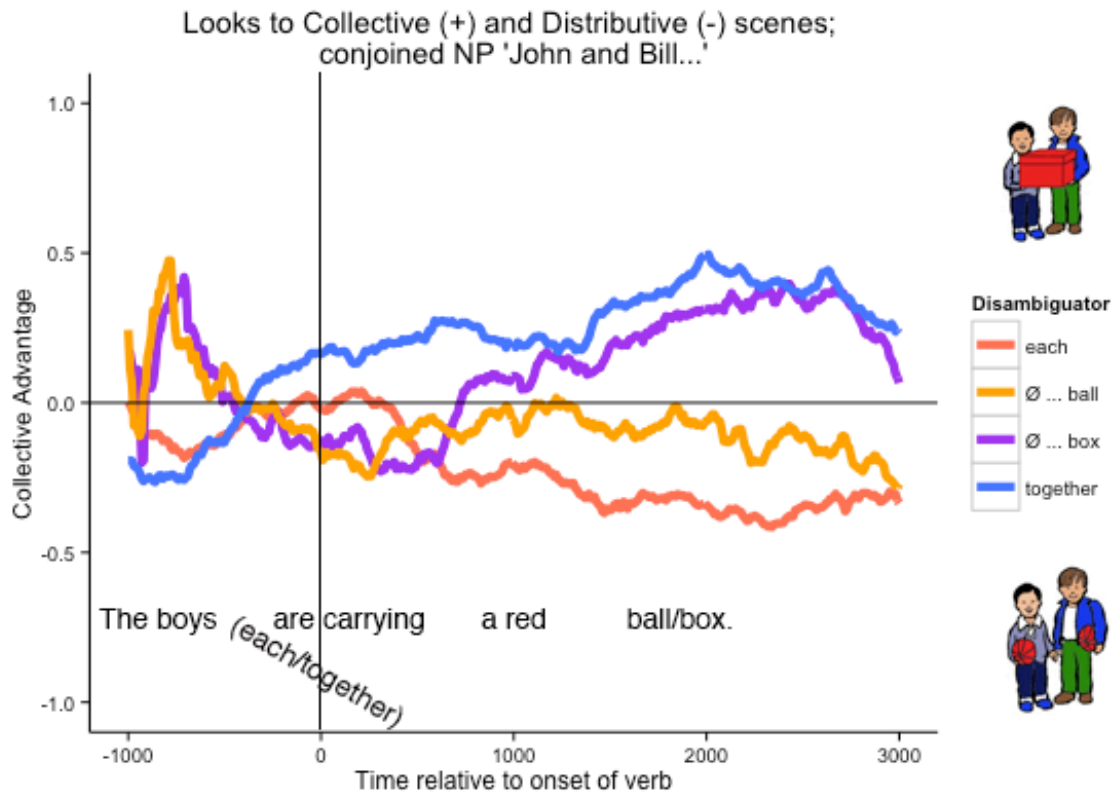


Figure 12: Plots of collective advantage to collective (positive) and distributive (negative) scenes over time for plural NP subjects, time-locked to the onset of the verb. Timecourses of each-disambiguated sentences are in red and timecourses of together-disambiguated sentences are in blue. Note that the  $\emptyset$  condition is split into those sentences eventually disambiguated by the object *ball/box* to a collective (purple) or distributive (yellow) scene.

### *Time window analyses*

We ran a by-subject 3x3 ANOVA on e-logit collective advantage measures, with factors Disambiguator (*each/together/∅*) and Time window (“*The boys*”/“*are carrying*”/“*a bright red*”). Since the disambiguator (*each/together*) time window was non-existent in the  $\emptyset$  disambiguator condition, this time window was not included in the ANOVA, though we do include it among the planned pairwise t-tests of disambiguator effects in each time window (see Figure 13). We found a main effect of Disambiguator ( $F(2,52) = 16.07, p < 0.001$ ) and a significant interaction between Disambiguator and Time window ( $F(4,104) = 12.14, p < 0.001$ ).

Figure 13




Figure 13: Plots of collective advantage for plural determiner NP sentences in four time windows. Average fixations across subjects for each-disambiguated sentences are plotted in red, together-disambiguated sentences in blue, and non-disambiguated (until object) sentences in green. The final convergence on ball or box is not shown here, and the non-disambiguated sentences that eventually converge on *ball* or *box* are collapsed here (purple and yellow timecourses from Figure 12 are collapsed and indicated in green here.)

We also ran pairwise two-sided t-tests on *each* vs. *together*, *each* vs.  $\emptyset$ , and *together* vs.  $\emptyset$ , using the Holm (1979) method of correction for multiple comparisons. As with the conjoined NP (*John and Bill*) window, we found no pairwise differences in collective advantage between *each*, *together*, and  $\emptyset$  in the *The boys* time window, and none of the disambiguator conditions differed significantly from chance ( $\mu=0$ ). In the *each* vs. *together* time window, we still found no significant divergence in looks, though the *together* condition was, interestingly, significantly more likely to direct looks to the distributive scene compared with chance. In the verb (*are carrying*) time window, we found that the *together* condition is marked, it being significantly more likely to prompt looks to the collective than the  $\emptyset$  condition ( $p<0.001$ ), *each* condition ( $p<0.001$ ), and chance ( $p<0.001$ ). The *each* and  $\emptyset$  conditions are not significantly different, and the  $\emptyset$  condition directs looks to the distributive scene above chance ( $p=0.006$ ) while the *each* condition does not differ significantly from chance ( $p>0.1$ ). Finally, in the adjective time window (*a bright red*), immediately preceding the object (*ball/box*), we find all three disambiguator conditions diverging from one another: the *each* condition is significantly more likely to direct looks to the distributive scene than chance ( $p<0.001$ ),  $\emptyset$  ( $p=0.001$ ), and *together* conditions ( $p<0.001$ ), while *together* is also prompting looks to the collective scene above both chance ( $p=0.004$ ) and  $\emptyset$  ( $p=0.03$ ). The  $\emptyset$  condition is equally likely to direct looks to the collective and distributive scenes in this time window.

#### *Distributive-biased verbs*

The timecourse and time window data for plural NP sentences indicate that the *each* and  $\emptyset$  conditions pattern together early on, both favoring the distributive scenario at the verb. While this distributive preference is expected for the *each* condition, it is rather

surprising, especially given the  $\emptyset$  condition's collective preference in Experiment 1, that the  $\emptyset$  condition directs looks to the distributive scene. We examined whether the subset of verbs showing a relative collective bias (twelve verbs with highest collective score on the Mechanical Turk survey) also favored the distributive scenes at the predicate time window. We found that collective-biased verbs were at chance for collective advantage. Thus, the more distributive-biased verbs in our stimuli appear to drive the overall distributive gaze preference at the verb in sentences beginning with plural NPs like *The boys*.

### *Summary of Experiment 2*

As in Experiment 1, we find evidence that verbs taking plural NP subject arguments demonstrate preference for a particular reading even in the absence of information determining a collective or distributive meaning. However, unlike previous studies and Experiment 1, we find that plural determiner NPs direct looks to the *distributive* reading, not the collective. Targeting the verb time window (*are carrying*), we indeed find an interaction between the type of subject NP (conjoined NP like *John and Bill* vs. plural determiner phrase like *The boys*) and disambiguator type (*each/together/  $\emptyset$* ) ( $F(2,98) = 8.68, p < 0.001$ ), reflecting the fact that  $\emptyset$  and *together* pattern together in conjoined NP sentences while  $\emptyset$  and *each* pattern together in plural determiner NP sentences.

This may appear to be inconsistent with theories stating that because an extra operator must be posited to “distribute” a predicate over individuals in a group, the introduction of such a distributive operator is dispreferred in the absence of an explicit (adverbial) quantifier (Clifton & Frazier, 2012; Roberts, 1987; inter alia). Below, we



discuss several possible reasons why sentences beginning with *The boys* might differ from sentences beginning with *John and Bill*.

## **Discussion**

### *A matter of ambiguity*

In both Experiments 1 and 2, we find fixation patterns consistent with what Frazier et al. (1999) call “necessary decision”: when the processor reaches the verb, it forces a commitment to one interpretation (the collective or the distributive). In Experiment 1, we found that even the time course of sentences with pragmatically distributive-biased predicates favor the collective reading until the object re-directs attention to the distributive scene (Figure 11), which suggests this is a grammatical forced choice rather than a pragmatic “invited decision.” In the case of an “invited decision,” a lexical or pragmatic bias might induce an immediate preference for one meaning or another based on a particularly strong collective or distributive “cloze” (where we intend this to mean a preference for a collective/distributive event rather than any one particular lexical item). However, since even distributive-biased predicates prompt looks initially to the collective scene, it is unlikely that the ambiguity is a lexical/pragmatic “invited decision.” Similarly, if there were a preference for the images depicting collective scenes over those depicting distributive scenes, we would have seen this in our image survey. Therefore, the immediacy of the preference for the collective scene in Experiment 1 is highly suggestive of a grammatical forced choice, or a “necessary decision” in the parlance of Frazier et al. (1999).

Both experiments support a theory that treats the collective/distributive distinction as ambiguous rather than vague. However, a limitation of these studies and others is the

treatment of the “vagueness” hypothesis as the null hypothesis. Whereas Frazier et al. do not suggest any positive test for vagueness (only negative results would “confirm” vagueness), our study is able to test one prediction of the vagueness hypothesis: if the collective/distributive distinction were a matter of vagueness only, we would expect increased looks to be a function of new information only. That is, we would expect increased looks to the collective scene when *together* was presented to be just as likely as increased looks to the distributive scene when *each* was presented. However, *each* was the only disambiguator for which there was a significant divergence from chance in both the *each/together* time window and the *are carrying* time window in the conjoined NP sentences (Experiment 1). Likewise, *together* was the only disambiguator for which there was a divergence from chance in both the *each/together* time window and the *are carrying* time window in the plural NP sentences. This suggests that *each* is somehow more informative than *together* in conjoined NP sentences, while *together* is more informative in plural NP sentences.

While the distinction between “vagueness” and “ambiguity” remains weak, and may ultimately be better characterized non-dichotomously, there are still other hallmarks of ambiguity to investigate with the current paradigm. For instance, the literature on so-called “digging-in effects” details evidence from grammaticality judgments and RTs that, as a period of ambiguity (say, in a garden-path sentence) is extended, the commitment to the initial interpretation strengthens with time and is more difficult to re-analyze if wrong (Ferreira & Henderson, 1991; Tabor & Hutchins, 2004). Interestingly, we fail to see such strengthening of an initial interpretation in our experiments: neither the initial commitment to the collective in Experiment 1 nor the initial commitment to the

distributive in Experiment 2 persist into the adjectival phrase (*a big red...*) This could indicate a fixation profile less consistent with ambiguity than with vagueness. We did not have enough variability in the length of the ambiguous adjectival phrases across our sentences to test whether ambiguity duration might affect re-analysis difficulty. It would be interesting to examine possible digging-in effects by extending the ambiguous region in our sentence – for instance, by adding more adjectives or disfluencies like “um, uh.” In such a replication of Experiment 1, we would expect that an increase in the length of the ambiguous period would not only allow for more looks to the collective scene to “accumulate” prior to disambiguation (given the increased time in which to pursue a collective bias), but, more importantly, a delayed or decreased “recovery” to the distributive scene upon disambiguation. A comparison of distributive recovery profiles in sentences with shorter ambiguities vs sentences with longer ambiguities would require a measure of proportion of *shifts* in fixation rather than a proportion of fixations to the distributive: given a fixation to the collective at time  $t$  immediately prior to distributive disambiguation, we would expect the probability of shifting to the distributive scene at time  $t+1$  (arbitrary units, post-disambiguation) would decrease with longer ambiguities. This “digging-in” dependent measure could be highly useful in further characterizing the nature and degree of the collective bias.

*Conjoined NPs and plural determiners differ in selecting between collectivity and distributivity*

There are a number of reasons why the collective interpretation should be preferred, when, in the absence of a disambiguator, the decision arises between a

collective and distributive meaning. For instance, under Referential Theory (Crain & Steedman, 1985), the distributive reading might be dispreferred not only because it must stipulate an operator, but also because it entails the existence of several distinct events. Thus both the grammatical and conceptual semantics of distributivity are characterized as more complex than the collective. Moreover, in both our experiments, the scenes depicting distributive events are also superficially more complex, in that there are two objects rather than one.

While several prevalent theories of distributivity hold that the distributive D operator applies over the VP (Crain & Steedman, 1985; Heim et al., 1991), we find evidence that the affordances on the subject NP may determine whether a D operator is stipulated at the predicate. Regarding the collective affordances of conjoined NPs, Moxey et al. (2011) found that readers were more likely to use a plural pronoun (such as *they*) when completing sentence fragments involving conjoined NPs (like *John and Bill arrived at the restaurant...*) than fragments introducing the same entities in separate phrases (*John arrived at the restaurant with Bill...*), even though both sentences denote collective events. That is, even when compared with unequivocally collective events like *John arrived at the restaurant with Bill*, ambiguous/vague sentences beginning with a conjoined NP like *John and Bill arrived at the restaurant* are more likely to prompt a joint *they* anaphor. This suggests that a collective reading is (even more) preferred when the subject NP is a conjoined NP, compared with when the collective event sentence involves a *John... with Bill* construction.

Consistent with this, other studies found that readers exhibit comprehension facilitation when members of a conjoined NP like *John and Bill* are referred to as *they*

compared with when just one member is referred to in isolation (as in *John* alone) (Albrecht & Clifton, 1998; Garrod & Sanford, 1988; Gordon, Hendrick, Ledoux, & Yang, 1999). This provides further evidence that conjoined NPs are particularly conducive to a collective reading, even in the absence of disambiguating information. Thus, even though we see in the current study that collective/distributive biases emerge at the onset of the VP, it may be the affordances of the subject NP that actually determine whether the predicate is distributed across multiple events or considered collectively.

While Experiment 1 is congruent with the expectation that the collective scene be preferred, Experiment 2 poses an interesting contrast. There are several possible explanations for why the plural determiner NP *The boys* should differ from *John and Bill* and prompt looks to the distributive scene at the onset of the verb. One possibility is that *The boys* is also highly selective for a collective interpretation, so much so that the presence of a distributive scene as an option is more salient and surprising than in the *John and Bill* case. Though we have reason to believe *John and Bill* should be biased towards a collective reading, the conjoined NP may yet be more accommodating of a distributive reading than is *The boys*. An eye-tracking study by Patson & Ferreira (2009) showed that participants' parsing strategy for sentences with anaphors and reciprocal verbs (e.g. *wrestle*) depended on whether the plural NP in the preceding sentence was a conjoined NP (*The trainer and the vet*) or a definite plural NP (*The trainers*); e.g. *The trainers / The trainer and the vet were near the swamp. While they wrestled the alligator watched them closely.* Reading times (RTs) at the disambiguating region *watched* indicated whether participants were garden-pathed. They found RTs were shorter at the disambiguating region when the preceding sentence had a conjoined NP compared with

sentences having a plural definite. Patson and Ferreira interpreted this as evidence that conjoined NPs, more so than definite plural NPs, are Complex Reference Objects (CROs), consisting of a representation of both (1) a single sum or group entity (conducive to a collective reading) and (2) sets of individuals within that entity (conducive to a distributive reading) (Barker, 1992; Moxey, Sanford, Sturt, & Morrow, 2004; Moxey et al., 2011).

It must be noted that the distributive option, however “surprising,” is not anomalous or uninterpretable in either our conjoined NP or plural NP experiments: participants continued to click on the “correct” scene in Experiment 2 as well as in Experiment 1 (that is, participants clicked on the scene that matched the object(s) uttered in the sentence, whether collective or distributive). As we saw in example 9 above (see Introduction), though the distributive reading of sentences like *The boys are carrying a box* is the marked reading, a context which defeases the collective reading of *The boys are carrying a box* is certainly possible. Therefore, it seems an unsatisfactory explanation that it is the “surprise” of seeing a distributive scene that drives the early distributive preference in Experiment 2

Given the immediate and automatic preference for the collective we see in the case of *John and Bill*, it is interesting that such a preference should be reliably overridden in *The boys* by what is an ostensive dispreferred distributive reading. What could override this semantic collective bias? The early distributive fixations in *The boys* sentences may instead be due to morphosyntactic properties of the definite plural. The definite plural NP has an explicit plural morpheme (as in the *-s* in *The boys*) where the conjoined NP does not. As such, the definite plural may place greater focus on the

plurality of the subject NP, if not also the event. Both intuition and the Patson & Ferreira, 2009, findings described above might tell us that it is instead the conjoined NP, by separating and naming *John and Bill* explicitly, that emphasizes the individuals in the plurality; however, we have also reviewed evidence that collective anaphors, like *they*, facilitate comprehension of conjoined NPs more than do distributed references to the individuals in the plural set (Albrecht & Clifton, 1998; Garrod & Sanford, 1988; Gordon, Hendrick, Ledoux, & Yang, 1999). Thus, though conjoined NPs accommodate distributive readings, they are still biased towards collective readings. We suggest that the collective bias of the conjoined NP *John and Bill* arises at the level of the semantics (as does the intuition that *The boys are carrying a ball* is also collective-biased), whereas the initial distributive bias of *The boys* arises at the level of the morphosyntax.

Under this interpretation of the results of Experiment 2, early fixation to the distributive is not necessarily a “commitment” to that scene’s meaning, though in the case of Experiment 1, early looks to the collective scene may be such. If the initial distributive fixation reflects an automatic but *transient* initial parse prompted by the plural -s morpheme, then this is potentially compatible with the notion of vagueness, whereby a subject refrains from committing to an interpretation until disambiguating information is available. (This should not be subject to a “digging in effect,” since the morphosyntactic plurality should arise and decay independent of the duration of the (semantic) ambiguity. Furthermore, in Experiment 2, the fact that the *together* condition directed looks to the *distributive* scene above chance at the *each/together* time window, and not to the collective scene, suggests that the initial consideration of the distributive is automatic and is induced even in the face of contrary evidence (*together* being an explicit lexical-

semantic marker of the collective reading). In this case, we suggest that morphosyntactic evidence for a plural distributive reading temporarily trumps, or is processed prior to, semantic/pragmatic biases for a collective reading.

This interpretation is consistent with so-called “syntax-first” models of language processing: various electrophysiological studies of language find that certain temporal signatures of syntactic processing occur prior to lexical-semantic processes. For instance, event-related potentials (ERPs) to syntactic violations are seen at ~125ms post stimulus onset (Early Left Anterior Negativity (ELAN); Friederici, Pfeifer, & Hahne, 1993; Neville, Nicol, Barss, Forster, & Garrett, 1991), while neural correlates of lexical semantic violations are observed later at ~400ms (N400; (Marta Kutas & Hillyard, 1980; M. Kutas & Federmeier, 2011; Lau, Phillips, & Poeppel, 2008). Some researchers have put forth strictly serial models, whereby syntactic and semantic processing are separately encapsulated (Friederici, 2002), while more recent models allow that these processes are highly interactive and parallel, but that syntactic forms facilitate processing at levels of representation that can be accessed more rapidly at word presentation than lexical semantic information (Boylan et al., 2014; Dikker & Pylkkanen, 2011; Dikker, Rabagliati, Farmer, & Pylkkanen, 2010).

### *Conclusions and Future Directions*

Taken together, these studies provide evidence of a productive distinction between two different types of NP with regard to the interpretation of plurality. Though we find prima facie evidence that plural determiner NPs like *The boys* prompt consideration of the distributive meaning before the collective, while conjoined NPs like *John and Bill* favor



the collective reading, we have reason to believe the mechanisms prompting early fixations might be different. We propose that the early distributive preference following *The boys* arises from the morphosyntactic plural –s marker, and that this preference dissipates over the time course of the sentence, giving way later to the semantically collective bias that is intuitive in the reading of both *John and Bill are carrying a ball* and *The boys are carrying a ball*. Further study examining other types of definite plural subject NPs may shed more light on the differences between definite plural and conjoined NPs. For instance, this line of inquiry may bridge a gap between the well-established literature on formal semantic theories of plurality and a growing body of work on children's acquisition of plurality and number words: there is mounting evidence that children appear to prefer distributive readings of sentences like *Two boys are pushing a car*, where adults prefer the collective interpretation (Syrett & Musolino, 2013). This will require further study of collectivity and distributivity across various different NPs and events.

## V. GENERAL DISCUSSION

### **Reviewing the aims of the dissertation**

In this dissertation, I sought to demonstrate the following:

- iv. Semantic compositionality can be characterized along a “feature-vs.-function” dichotomy in several possible dimensions (chapters 2 and 3)
- v. The AG supports various aspects of function-/relation-based composition, among them verb-specific argument structure around events (chapter 3) and more general thematic relation-based composition (chapter 2).
- vi. Semantic decisions on the collective or distributive representation of plural events occur early on in sentence processing – at the verb phrase – even in the absence of disambiguating evidence. However, the bias of this decision is largely determined by the affordances of the subject noun phrase, and not just the verb phrase (chapter 4).

While we refer to the typology of semantic composition in (i) as a division of “features” vs. “functions,” this shorthand belies a longstanding dichotomy appearing all the way from semiotics to developmental psychology. We discuss and test several possible instantiations of the division, including taxonomic vs. thematic associations, attributive vs. relational associations (chapter 2), object vs. event (verb) concepts (chapter 3), and adjuncts vs. arguments (chapters 3). However, other similar dichotomies we have not touched on include the following:

- i. Paradigmatic vs. syntagmatic distinction, where the syntagm can be said to be a structure-sensitive combination of “this-and-this-and-this” (as are the words

in the sentence *the man cried*), while the paradigm is the selection of “this-or-this-or-this” from a set of items with similar features (replacing “the man” with “the woman” to form the sentence *the woman cried*). (Chandler, 2007; de Saussure, 1916)

- ii. “vertical” vs. “horizontal” similarity distinction in language acquisition, where horizontal similarity describes the sort of within-sentence co-occurrence statistics toddlers could use to learn novel words (as in *man cried*) and also learn thematic relationships about nouns co-occurring in the same event. Vertical, or positional, similarity refers to words that can be used in the same position across sentences (*man* and *woman* in *the man/woman cried*), and by virtue of this are often in a similar category, sharing properties and features (Wojcik & Saffran, 2015)
- iii. predicate modification vs. function application, which are Merge operations over adjuncts and arguments, respectively (Heim & Kratzer, 1998)

This is by no means an exhaustive list, but we include it here to illustrate that the fundamental “feature-function” division has been a productive one long before our current line of research found potential neural substrates for the division. We now review the research detailed in the above chapters on the feature-function dissociation in the ATL and AG.

### **The neural bases of the feature-function division**

The studies in chapters 2 and 3 focused on characterizing the role of AG in semantic composition, contrasting this function with the manner in which left ATL might

subserve feature-based combination (Bemis & Pykkänen, 2011b; Westerlund et al., 2015). In contrast to ATL, we find increasing evidence that AG is engaged in semantic integration of relation-based event structure. While the current studies only investigate cases of minimal composition – that of two words isolated from a sentence or discourse – it is unlikely that AG and ATL are only tracking this level of composition. There is abundant evidence that AG, for one, engages in domain-general event processing in event structures as broad as discourse and in non-verbal depictions of events.

Indeed, both ATL and AG are best described as “hubs” at a domain-general level, and the typology of semantic composition described in (ii) above may cleave along the dorsal-ventral streams. As described in earlier chapters, the left ATL is a “convergence zone” receiving heavy traffic from the ventral visual object identification pathway and auditory “what” pathway, supporting its involvement in compositional operations over features of object-concepts (Rauschecker & Tian, 2000). The AG is surrounded by the dorsal spatial attention networks, the posterior temporal regions involved in motion perception, and the anterior parietal regions involved in representing action (Kravitz et al., 2011). This supports the conditions of AG involvement in action and event representation. While AG may have originated as a dorsal “where/how” convergence zone of spatial, goal-oriented, and action information, it may have been co-opted by language to represent increasingly abstract relational information. These relations might be learned by tracking co-occurrence statistics (“dog” often co-occurs with “bone”; “eggs” often co-occur with “breakfast), and/or these thematic relations are learned part and parcel of hierarchical structures arising in natural language syntax. Thus, the emergence of event and argument structure in thought and language may have been an

extension of the already extant dorsal pathways underpinning action and goal understanding.

### **The lateralization puzzle**

Though the fronto-temporal language network appears to have become strongly left-lateralized, it is less clear the degree to which right and left AG diverged with regard to processing events and representing thematic relations. Chapter 2 offers evidence that lateralization in AG did not necessarily isolate language to one hemisphere, but that the division of labor across bilateral AG is more subtle. That left AG might be more attuned to grammaticalized relation information (that is, information encoded directly on the verb) is supported by the fact that left AG is sensitive to the relatively fine-grained level of verb argument structure, and not other types of argument structure (Chapter 3). The fact that we see activity in right AG reflecting the distinction between non-grammaticalized relational vs. attributive interpretations of nominal compounds suggests that right AG also represents thematic relations between concepts (Chapter 2). However, while left AG may operate within the more constrained limits of grammatically explicit argument structure, right AG may pick up the slack in computing thematic associations between conceptual relations that are not explicitly realized in the morphosyntax.

### **The AG and verb/event semantics**

Chapters 2 and 3 provided evidence that AG tracks relational information potentially tuned to *verb* semantics. Taken in combination with work showing that verb valency – the number of arguments a verb can take – also modulates activity in AG, this invites the question as to whether AG is also sensitive to number information on a subject

NP argument. Does the singular/plural number on a subject NP and verb modulate activity in AG? If so, could observing patterns of AG and surrounding parietal activity elucidate the nature of plural representation when the event could be either collective or distributive (that is, when the subject NP could be either a group or a set of atomic individuals)? A recent fMRI study by Boiteau, Bowers, Nair, & Almor (2014) found that left AG was more active in response to plural subjects as compared with singular subjects. When comparing conjoined NP sentences (e.g. *Jeremy and Lucy did some work on the house*) with unconjoined NPs (*Jeremy did some work with Lucy on the house*), they found that conjoined NPs elicited more activity in right AG compared with unconjoined NPs. This is consistent with studies reviewed in chapter 4 indicating that conjoined NPs are more likely to prompt plural anaphors like *they*, and as such may be “more plural” than unconjoined NPs (Moxey et al., 2011). Given this finding, AG activity might prove a tractable means of measuring the gradient along which collective and distributive plurality might vary. A targeted fMRI analysis of AG activity might also be another means of assessing whether the cardinality of *The boys* is more or less plural than *John and Bill* (see chapter 4).

## **Conclusion**

This dissertation introduced and motivated a framework by which to begin a systematic typology of semantic composition. Using a range of methods, from fMRI multi-voxel pattern analysis to eye-tracking, we were able to measure both the “where,” and possibly also the “when,” of semantic composition. Of particular interest is the role of AG in function-/relation-based semantic composition. It remains to be seen how left

and right AG compare with respect to representing events and thematic relations grammatically or otherwise. The study of the role of bilateral AG as a semantic combinatorial hub benefits from comparison with another established “semantic hub,” the left ATL, and continuing to study these regions as a pair is well motivated both theoretically and empirically.

Finally, it is worth mentioning that the ATL and AG are at the center of a debate over whether representations in these areas are multi-modal or amodal. Though this dissertation did not arbitrate this particular debate, it should nonetheless make clear the utility of modality-invariant representations in a system, such as language, that must be both receptive and expressive in multiple modalities. While distributed semantic knowledge could be instantiated in modality-specific, multimodal, or amodal areas, regions involved specifically in compositional semantics are, more or less by necessity, multimodal or amodal. After all, if we were constantly composing meanings that carried with them fully intact sensorimotor and emotional simulations, it is not clear that we would be able to process spoken language at the rate of 3-4 words per second (Binder & Desai, 2011). Ultimately, it is exactly this disposition to compose which allowed the brain to achieve the abstract symbolic system that is language.

APPENDIX

Table 1: Pearson's r values for verb-relevant correlations in left AG

	50 voxels	100 voxels	200 voxels	500 voxels
<b>shared noun AND</b>	<i>0.027</i>	<i>0.035</i>	<i>0.019</i>	<i>0.018</i>
<b>shared composition</b>				
<b>type</b>				
<i>(eats meat ~ with meat)</i>				
<b>shared noun</b>	<i>0.010</i>	<i>-0.029</i>	<i>-0.0097</i>	<i>0.0036</i>
<i>(eats meat ~ tasty</i>				
<i>meat)</i>				
<b>shared verb</b>	<i>0.044</i>	<i>0.17 *</i>	<i>0.056 †</i>	<i>0.059 *</i>
<i>(eats meat ~ eats</i>				
<i>quickly)</i>				

†:  $0.05 < p < 0.1$ ; \*:  $p < 0.05$

Table 1: Correlations in left AG between the argument-saturated (*eats meat*) item with (a) an argument-saturated phrases *without* a verb (*with meat*), (b) a phrase sharing only a noun, but not composition type or verb (*tasty meat*), and (c) a phrase sharing a verb but constituting adjunct phrase (*eats quickly*) without argument saturation. Pairwise



comparisons are reported using the best 50, 100, 200, and 500 voxels with highest (unsigned) t-statistics from the composition -vs.-non-composition contrast in left AG.

Table 2: T-tests of differences between verb- or noun-based correlations in left AG (p-values reported)

	<b>50 voxels</b>	<b>100 voxels</b>	<b>200 voxels</b>	<b>500 voxels</b>
<b>shared verb vs. shared noun</b>	<i>0.092 †</i>	<i>&lt; 0.0001 *</i>	<i>0.052 †</i>	<i>0.034 *</i>
<b>shared verb vs. shared noun +composition</b>	<i>0.48 ns</i>	<i>&lt; 0.0001 *</i>	<i>0.23 ns</i>	<i>0.11 ns</i>
<b>shared noun vs. shared noun + composition</b>	<i>0.54 ns</i>	<i>0.051 †</i>	<i>0.18 ns</i>	<i>0.62 ns</i>

†:  $0.05 < p < 0.1$ ; \*:  $p < 0.05$

Table 2: T-tests of pairwise differences between verb-sharing phrase correlations, noun-sharing phrase correlations, and noun+composition phrase correlations using the best 50, 100, 200, and 500 voxels with highest (unsigned) t-statistics from the composition -vs.-non-composition contrast in left AG. Correlations being contrasted are those between the argument-saturated verb phrase “eats meat” and those phrases delineated in Table 1 of Appendix.

Table 3: Pearson's r values for verb-relevant correlations in left ATL

	50 voxels	100 voxels	200 voxels	500 voxels
<b>shared noun AND</b>	0.030	<b>0.13 *</b>	0.071†	0.026
<b>shared composition</b>				
<b>type</b>				
<i>(eats meat ~ with meat)</i>				
<b>shared noun</b>	0.026	<b>0.088 *</b>	0.052	0.042
<i>(eats meat ~ tasty</i>				
<i>meat)</i>				
<b>shared verb</b>	0.027	0.023	0.012	0.037
<i>(eats meat ~ eats</i>				
<i>quickly)</i>				

†:  $0.05 < p < 0.1$ ; \*:  $p < 0.05$

Table 3: Correlations in left ATL between the argument-saturated (*eats meat*) item with (a) an argument-saturated phrases *without* a verb (*with meat*), (b) a phrase sharing only a noun, but not composition type or verb (*tasty meat*), and (c) a phrase sharing a verb but constituting adjunct phrase (*eats quickly*) without argument saturation. Pairwise comparisons are reported using the best 50, 100, 200, and 500 voxels with highest (unsigned) t-statistics from composition -vs.-non-composition contrast in left ATL.

Table 4: T-tests of differences between verb- or noun-based correlations in left ATL (p-values reported)

	<b>50 voxels</b>	<b>100 voxels</b>	<b>200 voxels</b>	<b>500 voxels</b>
<b>shared verb vs. shared noun</b>	<i>0.96 ns</i>	<b><i>0.041 *</i></b>	<i>0.27 ns</i>	<i>0.92 ns</i>
<b>shared verb vs. shared noun +composition</b>	<i>0.91 ns</i>	<b><i>0.0050 *</i></b>	<i>0.13 ns</i>	<i>0.66 ns</i>
<b>shared noun vs. shared noun + composition</b>	<i>0.91 ns</i>	<i>0.13 ns</i>	<i>0.74 ns</i>	<i>0.62 ns</i>

†:  $0.05 < p < 0.1$ ; \*:  $p < 0.05$

Table 4: T-tests of pairwise differences between verb-sharing phrase correlations, noun-sharing phrase correlations, and noun+composition phrase correlations using the best 50, 100, 200, and 500 voxels with highest (unsigned) t-statistics from composition -vs.-non-composition contrast in left ATL. Correlations being contrasted are those between the argument-saturated verb phrase “eats meat” and those phrases delineated in Table 3 of Appendix.

Table 5: Pearson's r values for correlations of AMT similarity norms with neural similarity scores in left AG split by nouns and verbs

	<b>50 voxels</b>	<b>100 voxels</b>	<b>200 voxels</b>	<b>500 voxels</b>
<b>verbs</b>	-0.10	<b>-0.12 *</b>	-0.096	-0.14
<b>nouns</b>	-0.016	0.023	-0.039	-0.028

†:  $0.05 < p < 0.1$ ; \*:  $p < 0.05$

Table 5: Comparisons of correlations between similarity norms from AMT survey of pairwise phrase similarity (inclusive of both compositional (e.g. *eats meat*) and non-compositional (e.g. *meat*) phrases) and neural similarity norms calculated by averaging correlations between all pairs of phrases sharing a verb and all pairs of phrases sharing a noun (including both compositional and non-compositional phrases). Table 5 shows those verb- and noun-based correlations in best (unsigned) 50, 100, 200, and 500 voxels in left AG.

Table 6: Pearson's r values for correlations of AMT similarity norms with neural similarity scores in left ATL split by nouns and verbs

	<b>50 voxels</b>	<b>100 voxels</b>	<b>200 voxels</b>	<b>500 voxels</b>
<b>verbs</b>	<i>-0.013</i>	<i>0.041</i>	<i>0.038</i>	<i>0.026</i>
<b>nouns</b>	<i>-0.090</i>	<i>-0.097</i>	<i>-0.034</i>	<i>-0.048</i>

†:  $0.05 < p < 0.1$ ; \*:  $p < 0.05$

Table 6: Comparisons of correlations between similarity norms from AMT survey of pairwise phrase similarity (inclusive of both compositional (e.g. *eats meat*) and non-compositional (e.g. *meat*) phrases) and neural similarity norms calculated by averaging correlations between all pairs of phrases sharing a verb and all pairs of phrases sharing a noun (including both compositional and non-compositional phrases. Table 6 shows those verb- and noun-based correlations in best (unsigned) 50, 100, 200, and 500 voxels in left ATL.

## BIBLIOGRAPHY

- Albrecht, J. E., & Clifton, C. (1998). Accessing singular antecedents in conjoined phrases. *Memory & Cognition*, 26(3), 599–610. Retrieved from <http://link.springer.com/article/10.3758/BF03201166>
- Albright, A. (2009). Feature-based generalisation as a source of gradient acceptability\*. *Phonology*, 26(01), 9. <http://doi.org/10.1017/S0952675709001705>
- Baggio, G., Choma, T., van Lambalgen, M., & Hagoort, P. (2010). Coercion and compositionality. *Journal of Cognitive Neuroscience*, 22(9), 2131–2140. Retrieved from <http://www.mitpressjournals.org/doi/abs/10.1162/jocn.2009.21303>
- Barker, C. (1992). Group Terms in English: Representing Groups as Atoms. *Journal of Semantics*, 9(1), 69–93.
- Baron, S. G., Thompson-Schill, S. L., Weber, M., & Osherson, D. (2010). An early stage of conceptual combination: Superimposition of constituent concepts in left anterolateral temporal lobe. *Cognitive Neuroscience*, 1(1), 44–51. <http://doi.org/10.1080/17588920903548751>
- Barr, D. J. (2008). Analyzing “visual world” eyetracking data using multilevel logistic regression. *Journal of Memory and Language*, 59(4), 457–474. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0749596X07001015>
- Bavelier, D., Corina, D., Jezzard, P., Padmanabhan, S., Clark, V. P., Karni, A., ... Rauschecker, J. P. (1997). Sentence reading: A functional MRI study at 4 Tesla. *Journal of Cognitive Neuroscience*, 9(5), 664–686. Retrieved from <http://www.mitpressjournals.org/doi/abs/10.1162/jocn.1997.9.5.664>

- Bedny, M., Caramazza, A., Grossman, E., Pascual-Leone, A., & Saxe, R. (2008). Concepts are more than percepts: the case of action verbs. *The Journal of Neuroscience*, 28(44), 11347–11353.
- Bedny, M., Dravida, S., & Saxe, R. (2013). Shindigs, brunches, and rodeos: The neural basis of event words. *Cognitive, Affective, & Behavioral Neuroscience*, 14, 891–901. <http://doi.org/10.3758/s13415-013-0217-z>
- Bedny, M., & Thompson-Schill, S. L. (2006). Neuroanatomically separable effects of imageability and grammatical class during single-word comprehension. *Brain and Language*, 98(2), 127–139. <http://doi.org/10.1016/j.bandl.2006.04.008>
- Beeman, M., Friedman, R., Grafman, J., Perez, E., Diamond, S., & Lindsay, M. (1994). Summation priming and coarse semantic coding in the right hemisphere. *Cognitive Neuroscience, Journal of*, 6(1), 26–45. Retrieved from [http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=6793678](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6793678)
- Bemis, D. K., & Pykkänen, L. (2011a). Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience*, 31(8), 2801–2814.
- Bemis, D. K., & Pykkänen, L. (2011b). Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience*, 31(8), 2801–2814.
- Bemis, D. K., & Pykkänen, L. (2012). Basic Linguistic Composition Recruits the Left Anterior Temporal Lobe and Left Angular Gyrus During Both Listening and Reading. *Cerebral Cortex*, 23, 1859–1873. <http://doi.org/10.1093/cercor/bhs170>
- Bennett, M. R. (1974). *Some Extensions of a Montague Fragment of English*. UCLA.



- Ben Shalom, D., & Poeppel, D. (2007). Functional Anatomic Models of Language: Assembling the Pieces. *The Neuroscientist*, *14*(1), 119–127.  
<http://doi.org/10.1177/1073858407305726>
- Berent, I., Pinker, S., Tzelgov, J., Bibi, U., & Goldfarb, L. (2005). Computation of semantic number from morphological information. *Journal of Memory and Language*, *53*(3), 342–358. <http://doi.org/10.1016/j.jml.2005.05.002>
- Berko, J. (1958). *The child's learning of English morphology*. Radcliffe College.
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*. Retrieved from  
<http://www.sciencedirect.com/science/article/pii/S1364661311002142>
- Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, *19*(12), 2767–2796. Retrieved from  
<http://cercor.oxfordjournals.org/content/19/12/2767.short>
- Binder, J. R., Frost, J. A., Hammeke, T. A., Bellgowan, P. S. F., Rao, S. M., & Cox, R. W. (1999). Conceptual Processing during the Conscious Resting State: A Functional MRI Study. *Journal of Cognitive Neuroscience*, *11*(1), 80–93.  
<http://doi.org/10.1162/089892999563265>
- Binder, J. R., Westbury, C., McKiernan, K., Possing, E., & Medler, D. (2005). Distinct brain systems for processing concrete and abstract concepts. *Cognitive Neuroscience, Journal of*, *17*(6), 905–917. Retrieved from  
[http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=6788595](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6788595)

- Binney, R. J., Parker, G. J. M., & Lambon Ralph, M. A. (2012). Convergent Connectivity and Graded Specialization in the Rostral Human Temporal Lobe as Revealed by Diffusion-Weighted Imaging Probabilistic Tractography. *Journal of Cognitive Neuroscience*, 24(10), 1998–2014. Retrieved from [http://www.mitpressjournals.org/doi/abs/10.1162/jocn\\_a\\_00263](http://www.mitpressjournals.org/doi/abs/10.1162/jocn_a_00263)
- Bird, H., Howard, D., & Franklin, S. (2000). Why is a verb like an inanimate object? Grammatical category and semantic category deficits. *Brain and Language*, 72(3), 246–309.
- Boersma, P., & Weenink, D. (2015). Praat: doing phonetics by computer (Version 5.4.12). Retrieved from <http://www.praat.org/>
- Boiteau, T. W., Bowers, E., Nair, V. A., & Almor, A. (2014). The neural representation of plural discourse entities. *Brain and Language*, 137, 130–141. <http://doi.org/10.1016/j.bandl.2014.08.003>
- Bottini, G., Corcoran, R., Sterzi, R., Paulesu, E., Schenone, P., Scarpa, P., ... Frith, C. D. (1994). The role of the right hemisphere in the interpretation of figurative aspects of language. *Brain*, 117, 1241–1253. Retrieved from <http://brain.oxfordjournals.org/content/brain/117/6/1241.full.pdf>
- Boylan, C., Trueswell, J. C., & Thompson-Schill, S. L. (submitted). Compositionality and the angular gyrus: a multi-voxel similarity analysis of the semantic composition of nouns and verbs.
- Boylan, C., Trueswell, J. C., & Thompson-Schill, S. L. (2014). Multi-voxel pattern analysis of noun and verb differences in ventral temporal cortex. *Brain and Language*, 137, 40–49. <http://doi.org/10.1016/j.bandl.2014.07.009>

- Bressler, D., Spotswood, N., & Whitney, D. (2007). Negative BOLD fMRI Response in the Visual Cortex Carries Precise Stimulus-Specific Information. *PLoS ONE*, 2(5). <http://doi.org/10.1371/journal.pone.0000410>
- Brysbaert, M., & New, B. (2009). Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*, 41(4), 977–990. <http://doi.org/10.3758/BRM.41.4.977>
- Buccino, G., Binkofski, F., Fink, G. R., Fadiga, L., Fogassi, L., Gallese, V., ... Freund, H.-J. (2001). Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. *European Journal of Neuroscience*, 13(2), 400–404. <http://doi.org/10.1111/j.1460-9568.2001.01385.x>
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science*, 6(1), 3–5. <http://doi.org/10.1177/1745691610393980>
- Cappa, S. F., & Perani, D. (2003). The neural correlates of noun and verb processing. *Journal of Neurolinguistics*, 16(2-3), 183–189. [http://doi.org/10.1016/S0911-6044\(02\)00013-1](http://doi.org/10.1016/S0911-6044(02)00013-1)
- Cappa, S. F., Sandrini, M., Rossini, P. M., Sosta, K., & Miniussi, C. (2002). The role of the left frontal lobe in action naming rTMS evidence. *Neurology*, 59(5), 720–723. <http://doi.org/10.1212/WNL.59.5.720>
- Chandler, D. (2007). *Semiotics: the basics*. London: Routledge.

- Chao, L. L., & Martin, A. (2000). Representation of Manipulable Man-Made Objects in the Dorsal Stream. *NeuroImage*, *12*(4), 478–484.  
<http://doi.org/10.1006/nimg.2000.0635>
- Clarke, A., & Tyler, L. K. (2014). Object-Specific Semantic Coding in Human Perirhinal Cortex. *Journal of Neuroscience*, *34*(14), 4766–4775.  
<http://doi.org/10.1523/JNEUROSCI.2828-13.2014>
- Clifton, C., & Frazier, L. (2012). Interpreting conjoined noun phrases and conjoined clauses: Collective versus distributive preferences. *The Quarterly Journal of Experimental Psychology*, *65*(9), 1760–1776.  
<http://doi.org/10.1080/17470218.2012.667425>
- Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y.-C., ... Haxby, J. V. (2012). The Representation of Biological Classes in the Human Brain. *The Journal of Neuroscience*, *32*(8), 2608–2618.  
<http://doi.org/10.1523/JNEUROSCI.5547-11.2012>
- Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, *24*(2), 299–349. Retrieved from  
<http://www.sciencedirect.com/science/article/pii/S0364021300000203>
- Coutanche, M. N., & Thompson-Schill, S. L. (2014). Creating Concepts from Converging Features in Human Cortex. *Cerebral Cortex*, *25*, 2584–2593.  
<http://doi.org/10.1093/cercor/bhu057>
- Cox, R. W. (1996). AFNI: Software for Analysis and Visualization of Functional Magnetic Resonance Neuroimages. *Computers and Biomedical Research*, *29*(3), 162–173. <http://doi.org/10.1006/cbmr.1996.0014>

- Cox, R. W., & Jesmanowicz, A. (1999). Real-time 3D image registration for functional MRI. *Magnetic Resonance in Medicine*, 42(6), 1014–1018. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.43.253&rep=rep1&type=pdf>
- Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. *Natural Language Parsing*, 320–358.
- Crepaldi, D., Berlingeri, M., Cattinelli, I., Borghese, N. A., Luzzatti, C., & Paulesu, E. (2013). Clustering the lexicon in the brain: a meta-analysis of the neurofunctional evidence on noun and verb processing. *Frontiers in Human Neuroscience*, 7, 303. <http://doi.org/10.3389/fnhum.2013.00303>
- Crepaldi, D., Berlingeri, M., Paulesu, E., & Luzzatti, C. (2011). A place for nouns and a place for verbs? A critical review of neurocognitive data on grammatical-class effects. *Brain and Language*, 116(1), 33–49. <http://doi.org/10.1016/j.bandl.2010.09.005>
- Damasio, A. R., & Tranel, D. (1993). Nouns and verbs are retrieved with differently distributed neural systems. *Proceedings of the National Academy of Sciences*, 90(11), 4957–4960. Retrieved from <http://www.pnas.org/content/90/11/4957>
- Damasio, H., Grabowski, T. J., Tranel, D., Ponto, L. L. B., Hichwa, R. D., & Damasio, A. R. (2001). Neural Correlates of Naming Actions and of Naming Spatial Relations. *NeuroImage*, 13(6), 1053–1064. <http://doi.org/10.1006/nimg.2001.0775>

- Dehaene, S., Piazza, M., Pinel, P., & Cohen, L. (2003). Three parietal circuits for number processing. *Cognitive Neuropsychology*, 20(3-6), 487–506. Retrieved from <http://www.tandfonline.com/doi/abs/10.1080/02643290244000239>
- De Saussure, F. (1916). *Cours de Linguistique Generale*. Paris: Payot.
- Dikker, S., & Pylkkanen, L. (2011). Before the N400: Effects of lexical-semantic violations in visual cortex. *Brain and Language*.
- Dikker, S., Rabagliati, H., Farmer, T. A., & Pylkkanen, L. (2010). Early Occipital Sensitivity to Syntactic Category Is Based on Form Typicality. *Psychological Science*, 21(5), 629–634. <http://doi.org/10.1177/0956797610367751>
- Downing, P. (1977). On the Creation and Use of English Compound Nouns. *Language*, 53(4), 810. <http://doi.org/10.2307/412913>
- Embick, D., Marantz, A., Miyashita, Y., O’Neil, W., & Sakai, K. L. (2000). A syntactic specialization for Broca’s area. *Proceedings of the National Academy of Sciences*, 97(11), 6150–6154. Retrieved from <http://www.pnas.org/content/97/11/6150.short>
- Estes, Z. (2003a). A tale of two similarities: comparison and integration in conceptual combination. *Cognitive Science*, 27(6), 911–921. <http://doi.org/10.1016/j.cogsci.2003.01.001>
- Estes, Z. (2003b). Attributive and relational processes in nominal combination. *Journal of Memory and Language*, 48(2), 304–319.
- Fair, D. A., Cohen, A. L., Power, J. D., Dosenbach, N. U. F., Church, J. A., Miezin, F. M., ... Petersen, S. E. (2009). Functional Brain Networks Develop from a “Local

- to Distributed” Organization. *PLoS Computational Biology*, 5(5), e1000381.  
<http://doi.org/10.1371/journal.pcbi.1000381>
- Fairhall, S. L., & Caramazza, A. (2013). Brain Regions That Represent Amodal Conceptual Knowledge. *Journal of Neuroscience*, 33(25), 10552–10558.  
<http://doi.org/10.1523/JNEUROSCI.0051-13.2013>
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1(1), 1–47. Retrieved from  
<http://cercor.oxfordjournals.org/content/1/1/1.1.short>
- Ferreira, F., & Henderson, J. M. (1991). Recovery from misanalyses of garden-path sentences. *Journal of Memory and Language*, 30(6), 725–745.  
[http://doi.org/10.1016/0749-596X\(91\)90034-H](http://doi.org/10.1016/0749-596X(91)90034-H)
- Fletcher, P. C., Happe, F., Frith, U., Baker, S. C., Dolan, R. J., Frackowiak, R. S. J., & Frith, C. D. (1995). Other minds in the brain: a functional imaging study of “theory of mind” in story comprehension. *Cognition*, 57(2), 109–128. Retrieved from <http://www.sciencedirect.com/science/article/pii/001002779500692R>
- Fogassi, L., Ferrari, P. F., Gesierich, B., Rozzi, S., Chersi, F., & Rizzolatti, G. (2005). Parietal Lobe: From Action Organization to Intention Understanding. *Science*, 308(5722), 662–667. <http://doi.org/10.1126/science.1106138>
- Frazier, L., Pacht, J. M., & Rayner, K. (1999). Taking on semantic commitments, II: collective versus distributive readings. *Cognition*, 70(1), 87–104. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0010027799000025>
- Friederici, A. D., Rüschemeyer, S. A., Hahne, A., & Fiebach, C. J. (2003). The role of left inferior frontal and superior temporal cortex in sentence comprehension:

- localizing syntactic and semantic processes. *Cerebral Cortex*, 13(2), 170–177.  
Retrieved from <http://cercor.oxfordjournals.org/content/13/2/170.short>
- Gagné, C. L. (2000). Relation-Based Combinations Versus Property-Based Combinations: A Test of the CARIN Theory and the Dual-Process Theory of Conceptual Combination. *Journal of Memory and Language*, 42(3), 365–389.  
<http://doi.org/10.1006/jmla.1999.2683>
- Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(1), 71–87.  
<http://doi.org/10.1037/0278-7393.23.1.71>
- Gainotti, G. (1996). Cognitive and Anatomical Locus of Lesion in a Patient with a Category-specific Semantic Impairment for Living Beings. *Cognitive Neuropsychology*, 13(3), 357–390. <http://doi.org/10.1080/026432996381944>
- Garrod, S., & Sanford, A. J. (1988). Thematic subjecthood and cognitive constraints on discourse structure. *Journal of Pragmatics*, 12, 519–534.
- Gordon, P. C., Hendrick, R., Ledoux, K., & Yang, C. L. (1999). Processing of Reference and the Structure of Language: An Analysis of Complex Noun Phrases. *Language and Cognitive Processes*, 14(4), 353–379.  
<http://doi.org/10.1080/016909699386266>
- Gottlieb, J. (2007). From Thought to Action: The Parietal Cortex as a Bridge between Perception, Action, and Cognition. *Neuron*, 53(1), 9–16.  
<http://doi.org/10.1016/j.neuron.2006.12.009>



- Grabowski, T. J., Damasio, H., Tranel, D., Ponto, L. L. B., Hichwa, R. D., & Damasio, A. R. (2001). A role for left temporal pole in the retrieval of words for unique entities. *Human Brain Mapping, 13*(4), 199–212. <http://doi.org/10.1002/hbm.1033>
- Graves, W. W., Binder, J. R., Desai, R. H., Conant, L. L., & Seidenberg, M. S. (2010). Neural correlates of implicit and explicit combinatorial semantic processing. *NeuroImage, 53*(2), 638–646. <http://doi.org/10.1016/j.neuroimage.2010.06.055>
- Hagoort, P. (2005). On Broca, brain, and binding: a new framework. *Trends in Cognitive Sciences, 9*(9), 416–423. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1364661305002093>
- Hausser, R. (1974). *Quantification in an Extended Montague Grammar*. University of Texas, Austin.
- Haxby, J. V. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science, 293*(5539), 2425–2430. <http://doi.org/10.1126/science.1063736>
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Malden, MA: Blackwell.
- Heim, I., Lasnik, H., & May, R. (1991). Reciprocity and plurality. *Linguistic Inquiry, 22*(1), 63–101. Retrieved from <http://www.jstor.org/stable/10.2307/4178708>
- Henson, R. N. A., Shallice, T., Josephs, O., & Dolan, R. J. (2002). Functional Magnetic Resonance Imaging of Proactive Interference during Spoken Cued Recall. *NeuroImage, 17*(2), 543–558. <http://doi.org/10.1006/nimg.2002.1229>

- Hodges, J. R., Graham, N., & Patterson, K. (1995). Charting the progression in semantic dementia: Implications for the organisation of semantic memory. *Memory*, 3(3-4), 463–495.
- Hoeks, J. C. J., Stowe, L. A., & Doedens, G. (2004). Seeing words in context: the interaction of lexical and sentence level information during reading. *Cognitive Brain Research*, 19(1), 59–73. <http://doi.org/10.1016/j.cogbrainres.2003.10.022>
- Hoffman, P., Binney, R. J., & Lambon Ralph, M. A. (2015). Differing contributions of inferior prefrontal and anterior temporal cortex to concrete and abstract conceptual knowledge. *Cortex*, 63, 250–266. <http://doi.org/10.1016/j.cortex.2014.09.001>
- Holm, S. (1979). A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.
- Homae, F., Yahata, N., & Sakai, K. L. (2003). Selective enhancement of functional connectivity in the left prefrontal cortex during sentence processing. *NeuroImage*, 20(1), 578–586. [http://doi.org/10.1016/S1053-8119\(03\)00272-6](http://doi.org/10.1016/S1053-8119(03)00272-6)
- Humphries, C., Binder, J. R., Medler, D. A., & Liebenthal, E. (2007). Time course of semantic processes during sentence comprehension: An fMRI study. *NeuroImage*, 36(3), 924–932. <http://doi.org/10.1016/j.neuroimage.2007.03.059>
- Humphries, C., Binder, J. R., Medler, D., Liebenthal, E., & others. (2006). Syntactic and semantic modulation of neural activity during auditory sentence comprehension. *Cognitive Neuroscience, Journal of*, 18(4), 665–679. Retrieved from [http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=6794476](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6794476)

- Kang, A. M., Constable, R. T., Gore, J. C., & Avrutin, S. (1999). An Event-Related fMRI Study of Implicit Phrase-Level Syntactic and Semantic Processing. *NeuroImage*, *10*(5), 555–561.
- Kravitz, D. J., Saleem, K. S., Baker, C. I., & Mishkin, M. (2011). A new neural framework for visuospatial processing. *Nature Reviews Neuroscience*, *12*(4), 217–230. Retrieved from <http://www.nature.com/nrn/journal/v12/n4/abs/nrn3008.html>
- Kriegeskorte, N. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*, 4. <http://doi.org/10.3389/neuro.06.004.2008>
- Kuperberg, G., McGuire, P. K., Bullmore, E. T., Brammer, M., Rabe-Hesketh, S., Lythgoe, D. J., ... others. (2000). Common and distinct neural substrates for pragmatic, semantic, and syntactic processing of spoken sentences: an fMRI study. *Cognitive Neuroscience, Journal of*, *12*(2), 321–341. Retrieved from [http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=6787815](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6787815)
- Kuperberg, G. R. (2007). Neural mechanisms of language comprehension: Challenges to syntax. *Brain Research*, *1146*, 23–49. <http://doi.org/10.1016/j.brainres.2006.12.063>
- Kuperberg, G. R., Choi, A., Cohn, N., Paczynski, M., & Jackendoff, R. (2009). Electrophysiological Correlates of Complement Coercion. *Journal of Cognitive Neuroscience*, *22*(12), 2685–2701. <http://doi.org/10.1162/jocn.2009.21333>
- Kuperberg, G. R., Sitnikova, T., & Lakshmanan, B. M. (2008). Neuroanatomical distinctions within the semantic system during sentence comprehension: Evidence

- from functional magnetic resonance imaging. *NeuroImage*, 40(1), 367–388.  
<http://doi.org/10.1016/j.neuroimage.2007.10.009>
- Lambon Ralph, M. A. (2014). Neurocognitive insights on conceptual knowledge and its breakdown. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 369(1634), 20120392. <http://doi.org/10.1098/rstb.2012.0392>
- Lau, E. F., Phillips, C., & Poeppel, D. (2008). A cortical network for semantics:(de) constructing the N400. *Nature Reviews Neuroscience*, 9(12), 920–933. Retrieved from <http://www.nature.com/nrn/journal/v9/n12/abs/nrn2532.html>
- Levi, J. N. (1978). *The Syntax and Semantics of Complex Nominals*. Academic Press.
- Lewis, G. A., Poeppel, D., & Murphy, G. L. (2015). The neural bases of taxonomic and thematic conceptual relations: An MEG study. *Neuropsychologia*, 68, 176–189.  
<http://doi.org/10.1016/j.neuropsychologia.2015.01.011>
- Liljeström, M., Tarkiainen, A., Parviainen, T., Kujala, J., Numminen, J., Hiltunen, J., ... Salmelin, R. (2008). Perceiving and naming actions and objects. *NeuroImage*, 41(3), 1132–1141. <http://doi.org/10.1016/j.neuroimage.2008.03.016>
- Linke, A. C., Vicente-Grabovetsky, A., & Cusack, R. (2011). Stimulus-specific suppression preserves information in auditory short-term memory. *Proceedings of the National Academy of Sciences*, 108(31), 12961–12966.  
<http://doi.org/10.1073/pnas.1102118108>
- Link, G. (1983). The logical analysis of plurals and mass terms: A lattice-theoretical approach. In *Meaning, use and interpretation of language* (pp. 302–323). Berlin, Germany: De Gruyter.

- Lin, Y., Michel, J.-B., Aiden, E. L., Orwant, J., Brockman, W., & Petrov, S. (2012). Syntactic Annotations for the Google Books Ngram Corpus. *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics, 2: Demo Papers*.
- Linzen, T., Marantz, A., & Pylkkanen, L. (2013). Syntactic context effects in visual word recognition: An MEG study. *The Mental Lexicon*, 8(2), 117–139. Retrieved from <http://www.ingentaconnect.com/content/jbp/ml/2013/00000008/00000002/art00001>
- Luke, K.-K., Liu, H.-L., Wai, Y.-Y., Wan, Y.-L., & Tan, L. H. (2002). Functional anatomy of syntactic and semantic processing in language comprehension. *Human Brain Mapping*, 16(3), 133–145. <http://doi.org/10.1002/hbm.10029>
- McElree, B., Pylkkänen, L., Pickering, M. J., & Traxler, M. J. (2006). A time course analysis of enriched composition. *Psychonomic Bulletin & Review*, 13(1), 53–59. <http://doi.org/10.3758/BF03193812>
- Meltzer-Asscher, A., Schuchard, J., den Ouden, D.-B., & Thompson, C. K. (2013). The neural substrates of complex argument structure representations: Processing “alternating transitivity” verbs. *Language and Cognitive Processes*, 28(8), 1–15. <http://doi.org/10.1080/01690965.2012.672754>
- Moxey, L. M., Sanford, A. J., Sturt, P., & Morrow, L. I. (2004). Constraints on the formation of plural reference objects: The influence of role, conjunction, and type of description. *Journal of Memory and Language*, 51(3), 346–364. <http://doi.org/10.1016/j.jml.2004.06.001>

- Moxey, L. M., Sanford, A. J., Wood, A. I., & Ginter, L. M. N. (2011). When do we use “they” to refer to two individuals? Scenario mapping as a basis for equivalence. *Language and Cognitive Processes, 26*(1), 79–120.  
<http://doi.org/10.1080/01690961003701865>
- Müller, N. G., & Kleinschmidt, A. (2004). The attentional “spotlight”s’ penumbra: center-surround modulation in striate cortex. *Neuroreport, 15*(6), 977–980.  
Retrieved from  
[http://journals.lww.com/neuroreport/Fulltext/2004/04290/The\\_attentional\\_spotlight\\_s\\_penumbra\\_9.aspx](http://journals.lww.com/neuroreport/Fulltext/2004/04290/The_attentional_spotlight_s_penumbra_9.aspx)
- Murphy, G. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language, 29*, 259–288.
- Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*(1), 70. Retrieved from  
<http://psycnet.apa.org/journals/xlm/11/1/70/>
- Musz, E., & Thompson-Schill, S. L. (2014). Semantic variability predicts neural variability of object concepts. *Neuropsychologia, 76*, 41–51.  
<http://doi.org/10.1016/j.neuropsychologia.2014.11.029>
- Newman, A. J., Pancheva, R., Ozawa, K., Neville, H. J., & Ullman, M. T. (2001). An Event-Related fMRI Study of Syntactic and Semantic Violations. *Journal of Psycholinguistic Research, 30*(3), 339–364.  
<http://doi.org/10.1023/A:1010499119393>

- Ni, W., Constable, R., Mencl, W., Pugh, K., Fulbright, R., Shaywitz, S., ... Shankweiler, D. (2000). An event-related neuroimaging study distinguishing form and content in sentence processing. *Cognitive Neuroscience, Journal of*, 12(1), 120–133. Retrieved from [http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=6788751](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6788751)
- Noonan, K. A., Jefferies, E., Visser, M., & Lambon Ralph, M. A. (2013). Going beyond Inferior Prefrontal Involvement in Semantic Control: Evidence for the Additional Contribution of Dorsal Angular Gyrus and Posterior Middle Temporal Cortex. *Journal of Cognitive Neuroscience*, 1–26. [http://doi.org/10.1162/jocn\\_a\\_00442](http://doi.org/10.1162/jocn_a_00442)
- Pallier, C., Devauchelle, A. D., & Dehaene, S. (2011). Cortical representation of the constituent structure of sentences. *Proceedings of the National Academy of Sciences*, 108(6), 2522–2527. Retrieved from <http://www.pnas.org/content/108/6/2522.short>
- Papagno, C., & Capitani, E. (2001). Slowly progressive aphasia: a four-year follow-up study. *Neuropsychologia*, 39(7), 678–686. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0028393201000070>
- Partee, B. H., & Borschev, V. (2003). Genitives, relational nouns, and argument-modifier ambiguity. *Modifying Adjuncts*, 4, 67–112. Retrieved from [http://books.google.com/books?hl=en&lr=&id=5fUgAAAAQBAJ&oi=fnd&pg=PA67&dq=%22Barbara+H.+Partee+and+Vladimir%22+%22Recent+proposals+by+Jensen+and+Vikner+\(1994\),+Vikner+and%22+%22some+\(possibly+metaphorical\)+sense+express+possession%3B+and+in+each%22+%22genitives+able+to+get+argumental+interpretations+without+actually%22+&ots=aeMPO0RrOS&sig=kKxztDgQwLjemU1PbG5PIOGHFc0](http://books.google.com/books?hl=en&lr=&id=5fUgAAAAQBAJ&oi=fnd&pg=PA67&dq=%22Barbara+H.+Partee+and+Vladimir%22+%22Recent+proposals+by+Jensen+and+Vikner+(1994),+Vikner+and%22+%22some+(possibly+metaphorical)+sense+express+possession%3B+and+in+each%22+%22genitives+able+to+get+argumental+interpretations+without+actually%22+&ots=aeMPO0RrOS&sig=kKxztDgQwLjemU1PbG5PIOGHFc0)

- Patson, N. D., & Ferreira, F. (2009). Conceptual plural information is used to guide early parsing decisions: Evidence from garden-path sentences with reciprocal verbs. *Journal of Memory and Language*, 60(4), 464–486.  
<http://doi.org/10.1016/j.jml.2009.02.003>
- Patson, N. D., & Warren, T. (2010). Evidence for distributivity effects in comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(3), 782.
- Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, 8(12), 976–987. Retrieved from  
<http://www.nature.com/nrn/journal/v8/n12/abs/nrn2277.html>
- Polyn, S. M., Natu, V. S., Cohen, J. D., & Norman, K. A. (2005). Category-Specific Cortical Activity Precedes Retrieval During Memory Search. *Science*, 310(5756), 1963–1966. <http://doi.org/10.1126/science.1117645>
- Price, C. J. (2010). The anatomy of language: a review of 100 fMRI studies published in 2009. *Annals of the New York Academy of Sciences*, 1191(1), 62–88. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1111/j.1749-6632.2010.05444.x/full>
- 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. Retrieved from  
<http://www.tandfonline.com/doi/abs/10.1080/01690965.2010.527490>
- Pylkkänen, L., & McElree, B. (2006). The syntax-semantics interface: On-line composition of sentence meaning. *Handbook of Psycholinguistics*, 2, 537–577.



Retrieved from

<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.85.5305&rep=rep1&type=pdf>

Pylkkänen, L., & McElree, B. (2007). An MEG Study of Silent Meaning. *Journal of Cognitive Neuroscience*, *19*(11), 1905–1921.

<http://doi.org/10.1162/jocn.2007.19.11.1905>

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682.

<http://doi.org/10.1073/pnas.98.2.676>

Rauschecker, J. P., & Scott, S. K. (2009). Maps and streams in the auditory cortex: nonhuman primates illuminate human speech processing. *Nature Neuroscience*, *12*(6), 718–724. <http://doi.org/10.1038/nn.2331>

Roberts, C. (1987). *Modal Subordination, Anaphora, and Distributivity*. University of Massachusetts Amherst.

Rogers, T. T., Hocking, J., Noppeney, U., Mechelli, A., Gorno-Tempini, M. L., Patterson, K., & Price, C. J. (2006). Anterior temporal cortex and semantic memory: reconciling findings from neuropsychology and functional imaging. *Cognitive, Affective, & Behavioral Neuroscience*, *6*(3), 201–213. Retrieved from

<http://link.springer.com/article/10.3758/CABN.6.3.201>

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., & Patterson, K. (2004). Structure and Deterioration of Semantic Memory: A

- Neuropsychological and Computational Investigation. *Psychological Review*, 111(1), 205–235. <http://doi.org/10.1037/0033-295X.111.1.205>
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. MIT Press.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8(3), 382–439.
- Rosenfelder, I., Fruehwald, J., Evanini, K., & Yuan, J. (2011). *FAVE (Forced Alignment and Vowel Extraction) Program Suite*. Retrieved from <http://fave.ling.upenn.edu>
- Scha, R. (1984). Distributive, collective and cumulative quantification. In *Truth, Interpretation, and Information* (pp. 131–158).
- Schwartz, M. F., Kimberg, D. Y., Walker, G. M., Brecher, A., Faseyitan, O. K., Dell, G. S., ... Coslett, H. (2011). Neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. *Proceedings of the National Academy of Sciences*, 108(20), 8520–8524. Retrieved from <http://www.pnas.org/content/108/20/8520.short>
- Seghier, M. L. (2012). The angular gyrus: multiple functions and multiple subdivisions. *The Neuroscientist*, 19(1), 43–61.
- Seghier, M. L., Fagan, E., & Price, C. J. (2010). Functional Subdivisions in the Left Angular Gyrus Where the Semantic System Meets and Diverges from the Default Network. *Journal of Neuroscience*, 30(50), 16809–16817. <http://doi.org/10.1523/JNEUROSCI.3377-10.2010>
- Seuren, P. A. M. (2009). *Language in Cognition : Language From Within Volume I: Language From Within*. Oxford University Press.

- Shapiro, K. A., Moo, L. R., & Caramazza, A. (2006). Cortical signatures of noun and verb production. *Proceedings of the National Academy of Sciences*, *103*(5), 1644–1649. <http://doi.org/10.1073/pnas.0504142103>
- Shehzad, Z., Kelly, A. M. C., Reiss, P. T., Gee, D. G., Gotimer, K., Uddin, L. Q., ... Milham, M. P. (2009). The Resting Brain: Unconstrained yet Reliable. *Cerebral Cortex*, *19*(10), 2209–2229. <http://doi.org/10.1093/cercor/bhn256>
- Shmuel, A., Augath, M., Oeltermann, A., & Logothetis, N. K. (2006). Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nature Neuroscience*, *9*(4), 569–577. <http://doi.org/10.1038/nn1675>
- Shmuel, A., Yacoub, E., Pfeuffer, J., Van de Moortele, P.-F., Adriany, G., Hu, X., & Ugurbil, K. (2002). Sustained negative BOLD, blood flow and oxygen consumption response and its coupling to the positive response in the human brain. *Neuron*, *36*(6), 1195–1210. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0896627302010619>
- Shoben, E. J., & Gagné, C. L. (1997). Thematic relations and the creation of combined concepts. In *Creative thought: An investigation of conceptual structures and processes* (pp. 31–50). Washington, D.C., US: American Psychological Association.
- Sitnikova, T., Holcomb, P. J., Kiyonaga, K. A., & Kuperberg, G. R. (2008a). Two neurocognitive mechanisms of semantic integration during the comprehension of visual real-world events. *Journal of Cognitive Neuroscience*, *20*(11), 2037–2057.

- Sitnikova, T., Holcomb, P. J., & Kuperberg, G. R. (2008b). Neurocognitive mechanisms of human comprehension. *Understanding Events: From Perception to Action*, 639–683.
- Smith, A. T., Williams, A. L., & Singh, K. D. (2004). Negative BOLD in the visual cortex: evidence against blood stealing. *Human Brain Mapping*, 21(4), 213–220. <http://doi.org/10.1002/hbm.20017>
- Staresina, B. P., Fell, J., Do Lam, A. T. A., Axmacher, N., & Henson, R. N. (2012). Memory signals are temporally dissociated in and across human hippocampus and perirhinal cortex. *Nature Neuroscience*, 15(8), 1167–1173. <http://doi.org/10.1038/nn.3154>
- Stringer, S. M., & Rolls, E. T. (2002). Invariant Object Recognition in the Visual System with Novel Views of 3D Objects. *Neural Computation*, 14(11), 2585–2596. <http://doi.org/10.1162/089976602760407982>
- Syrett, K., & Musolino, J. (2013). Collectivity, Distributivity, and the Interpretation of Plural Numerical Expressions in Child and Adult Language. *Language Acquisition*, 20(4), 259–291. <http://doi.org/10.1080/10489223.2013.828060>
- Tabor, W., & Hutchins, S. (2004). Evidence for Self-Organized Sentence Processing: Digging-In Effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 431–450. <http://doi.org/10.1037/0278-7393.30.2.431>
- Talairach, J., & Tournoux, P. (1988). *Co-Planar Stereotaxic Atlas of the Human Brain*. New York: Thieme Medical Publishers, Inc.
- Thompson, C. K., Bonakdarpour, B., Fix, S. C., Blumenfeld, H. K., Parrish, T. B., Gitelman, D. R., & Mesulam, M.-M. (2007). Neural correlates of verb argument

structure processing. *Journal of Cognitive Neuroscience*, 19(11), 1753–1767.

Retrieved from

<http://www.mitpressjournals.org/doi/abs/10.1162/jocn.2007.19.11.1753>

Thompson, C. K., Bonakdarpour, B., & Fix, S. F. (2010). Neural mechanisms of verb argument structure processing in agrammatic aphasic and healthy age-matched listeners. *Journal of Cognitive Neuroscience*, 22(9), 1993–2011. Retrieved from <http://www.mitpressjournals.org/doi/abs/10.1162/jocn.2009.21334>

Tranel, D., Damasio, H., & Damasio, A. R. (1997). A neural basis for the retrieval of conceptual knowledge. *Neuropsychologia*, 35(10), 1319–1327.

[http://doi.org/10.1016/S0028-3932\(97\)00085-7](http://doi.org/10.1016/S0028-3932(97)00085-7)

Tranel, D., Martin, C., Damasio, H., Grabowski, T. J., & Hichwa, R. (2005). Effects of Noun-Verb Homonymy on the Neural Correlates of Naming Concrete Entities and Actions. *Brain and Language*, 92(3), 288–299.

<http://doi.org/10.1016/j.bandl.2004.01.011>

Traxler, M. J., McElree, B., Williams, R. S., & Pickering, M. J. (2005). Context effects in coercion: Evidence from eye movements. *Journal of Memory and Language*, 53(1), 1–25. <http://doi.org/10.1016/j.jml.2005.02.002>

Tyler, L. K., Russell, R., Fadili, J., & Moss, H. E. (2001). The neural representation of nouns and verbs: PET studies. *Brain*, 124(8), 1619–1634.

<http://doi.org/10.1093/brain/124.8.1619>

Tyler, L. K., Stamatakis, E. A., Bright, P., Acres, K., Abdallah, S., Rodd, J. M., & Moss, H. E. (2004). Processing objects at different levels of specificity. *Journal of*

- Cognitive Neuroscience*, 16(3), 351–362. Retrieved from  
<http://www.mitpressjournals.org/doi/abs/10.1162/089892904322926692>
- Vigliocco, G., Vinson, D. P., Druks, J., Barber, H., & Cappa, S. F. (2011). Nouns and verbs in the brain: A review of behavioural, electrophysiological, neuropsychological and imaging studies. *Neuroscience & Biobehavioral Reviews*, 35(3), 407–426. <http://doi.org/10.1016/j.neubiorev.2010.04.007>
- Vilberg, K. L., & Rugg, M. D. (2008). Memory retrieval and the parietal cortex: A review of evidence from a dual-process perspective. *Neuropsychologia*, 46(7), 1787–1799. <http://doi.org/10.1016/j.neuropsychologia.2008.01.004>
- Wang, J., Conder, J. A., Blitzer, D. N., & Shinkareva, S. V. (2010). Neural representation of abstract and concrete concepts: A meta-analysis of neuroimaging studies. *Human Brain Mapping*, 31(10), 1459–1468. <http://doi.org/10.1002/hbm.20950>
- Weber, M., Thompson-Schill, S. L., Osherson, D., Haxby, J., & Parsons, L. (2009). Predicting judged similarity of natural categories from their neural representations. *Neuropsychologia*, 47(3), 859–868.  
<http://doi.org/10.1016/j.neuropsychologia.2008.12.029>
- 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. <http://doi.org/10.1016/j.bandl.2014.12.003>
- Westerlund, M., & Pylkkänen, L. (2014). The role of the left anterior temporal lobe in semantic composition vs. semantic memory. *Neuropsychologia*, 57, 59–70.  
<http://doi.org/10.1016/j.neuropsychologia.2014.03.001>

- Wisniewski, E. J. (1996). Construal and Similarity in Conceptual Combination. *Journal of Memory and Language*, 35(3), 434–453. <http://doi.org/10.1006/jmla.1996.0024>
- Wisniewski, E. J., & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In *Understanding word and sentence. Advances in psychology* (pp. 241–284). Oxford, England: North Holland.
- Wisniewski, E. J., & Love, B. C. (1998). Relations versus Properties in Conceptual Combination. *Journal of Memory and Language*, 38, 177–202.
- Wojcik, E. H., & Saffran, J. R. (2015). Toddlers encode similarities among novel words from meaningful sentences. *Cognition*, 138, 10–20.  
<http://doi.org/10.1016/j.cognition.2015.01.015>
- Xu, J., Kemeny, S., Park, G., Frattali, C., & Braun, A. (2005). Language in context: emergent features of word, sentence, and narrative comprehension. *Neuroimage*, 25(3), 1002–1015.
- Yee, E., Drucker, D. M., & Thompson-Schill, S. L. (2010). fMRI-adaptation evidence of overlapping neural representations for objects related in function or manipulation. *NeuroImage*, 50(2), 753–763. <http://doi.org/10.1016/j.neuroimage.2009.12.036>