Neuropsychologia 78 (2015) 130–141

Contents lists available at ScienceDirect

Neuropsychologia

journal homepage: www.elsevier.com/locate/neuropsychologia

Compositionality and the angular gyrus: A multi-voxel similarity analysis of the semantic composition of nouns and verbs

Christine Boylan [*,1], John C. Trueswell, Sharon L. Thompson-Schill

Department of Psychology, University of Pennsylvania, United States

ARTICLE INFO

Article history:
Received 22 April 2015
Received in revised form
1 October 2015
Accepted 3 October 2015
Available online 18 October 2015

Keywords:
Pattern similarity analysis
Angular gyrus
Anterior temporal lobe
Conceptual combination
Semantic composition

ABSTRACT

The cognitive and neural systems that enable conceptual processing must support the ability to combine (and recombine) concepts to form an infinite number of ideas. Two candidate neural systems for conceptual combination—the left anterior temporal lobe (ATL) and the left angular gyrus (AG)—have been characterized as "semantic hubs" due to both functional and anatomical properties; however, these two regions likely support different aspects of composition. Here we consider two hypotheses for the role of AG in conceptual combination, both of which differ from a putative role for the ATL in "feature-based" combinatorics (i.e., meaning derived by combining concepts' features). Firstly, we examine whether AG is more sensitive to function-argument relations of the sort that arise when a predicate is combined with its arguments. Secondly, we examine the non-mutually exclusive possibility that AG represents information carried on a verb in particular, whether this be information about event composition or about thematic relations denoted uniquely by verbs. We identified voxels that respond differentially to two-word versus one-word stimuli, and we measured the similarity of the patterns in these voxels evoked by (1) pairs of two-word phrases that shared a noun that was an argument, thus sharing function-argument composition (e.g. *eats meat* and *with meat*), in comparison with two-word phrases that shared only a noun, not an argument (e.g., *eats meat* and *tasty meat*); and (2) stimulus pairs that shared only an event (operationalized here as sharing a verb; e.g. *eats meat* and *eats quickly*), in comparison to both of the above. We found that activity patterns in left AG tracked information relating to the presence of an event-denoting verb in a pair of two-word phrases. We also found that the neural similarity in AG voxel patterns between two phrases sharing a verb correlated with subjects' ratings of how similar the meanings of those two verb phrases were. These findings indicate that AG represents information specific to verbs, perhaps event structure or thematic relations mediated by verbs, as opposed to argument structure in general.

1. 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.

1.1. 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 and Poeppel, 2007; Binder and Desai, 2011; Binder et al., 2009; Lau et al., 2008; Pallier et al., 2011; Patterson et al., 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

* Corresponding author. Fax: +1 215 573 9247.
E-mail address: chboylan@sas.upenn.edu (C. Boylan).
[1] Present address: Institute for Research in Cognitive Science, University of Pennsylvania, 3401 Walnut Street, Suite 400A, Philadelphia, PA 19104, USA.

http://dx.doi.org/10.1016/j.neuropsychologia.2015.10.007
0028-3932/

to functional and anatomical patterns that are consistent with multimodal convergence (Binder and 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 and 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; Binder and Desai, 2011; Binney et al., 2012; Felleman and 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 and Van Essen, 1991) and auditory (Rauschecker and 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 and Scott, 2009; Stringer and 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 and 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 and 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 and 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 et al., 2003; Xu, Kemeny et al., 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 and Desai, 2011; Binder et al., 2009; Lau et al., 2008; but cf. Noonan et al. (2013), and Section 4 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, 2003; Wisniewski and 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; Thompson et al., 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 et al., 2008a, 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.

1.2. 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.

1.2.1 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 et al., 2000; Vigliocco et al., 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 et al., 2002; Cappa and Perani, 2003; Damasio and Tranel, 1993; Shapiro et al., 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 et al., 2011; Liljeström et al., 2008; Tranel et al., 2005; Tyler et al., 2001). Also, a study by Bedny et al. (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 and 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 et al., 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 et al., 2013).

Therefore, because we focus on events as spatio*temporal* 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.

1.2.2 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 and Kratzer, 1998; Pylkkänen et al., 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 et al. (2013) found that left ATL is sensitive to verb-argument structure, specifically the subcategorization frames of verbs. Westerlund et al. (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 me*at, 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.

1.3. 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 et al., 2014; Coutanche and Thompson-Schill, 2014; Polyn et al., 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 and Tyler, 2014; Connolly et al., 2012; Fair et al., 2009; Fairhall and Caramazza, 2013; Haxby, 2001; Kriegeskorte, 2008; Weber et al., 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 and 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.

2. Material and methods

2.1. 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.

2.2. Stimuli

2.2.1 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 and Pylkkänen, 2011; 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.

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

Table 1
Stimuli design.

	Argument (argument-type) composition	*Adjunct (adjunct-type) composition*
+verb	*eatsmeat*	*eats quickly*
verb control	*eats fghjl/fghjl eats*	*eats fghjl/fghjl eats*
−verb	*with meat*	*tasty meat*
−verb (noun) control	*meat fghjl/fghjl meat*	*meat fghjl/fghjl meat*

An example set of two-word compositional and one-word non-compositional items sharing a given noun or 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).

2.2.2 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 Section 4), 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).

2.3. 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 s. The critical phrase was immediately followed by either a 9-s fixation cross, during which the subject need only passively view the screen, or a two-word phrase probe presented for 3 s 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.

2.4. 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_0 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 × 0.9 mm × 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 × 2 mm × 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 and 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 and Tournoux, 1988) and voxels were resampled in the process to 3.5 mm × 3.5 mm × 3.5 mm.

2.5. Analysis

2.5.1 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. 1).

For those voxels within a given ROI, multiple regression was used to generate

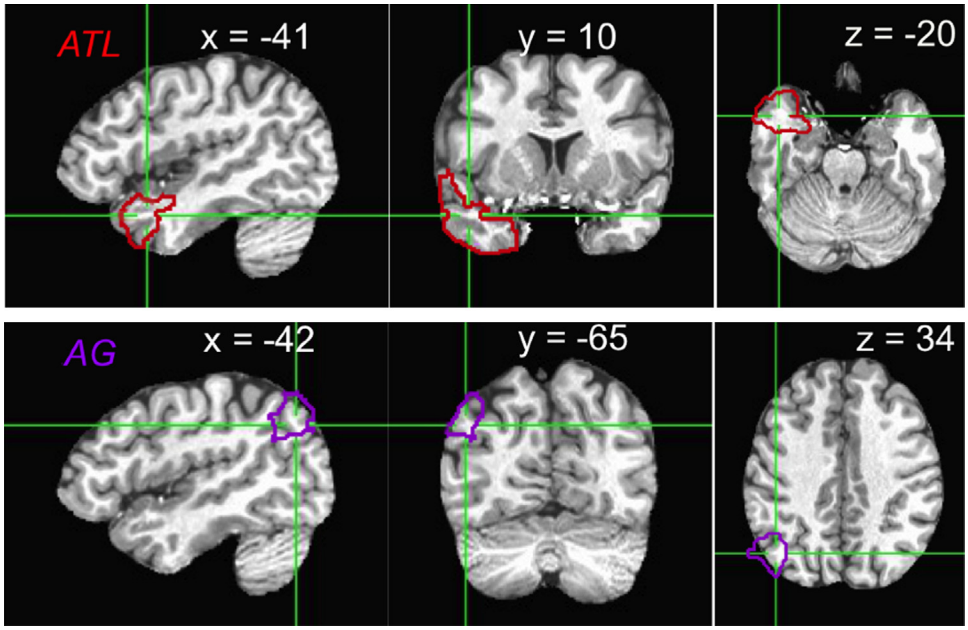

Fig. 1. ROI boundaries of left ATL and left AG.

parameter estimates (β) representing each voxel's activity in each stimulus item condition within subject. Voxels' βs were calculated by convolving all variables with a gamma-variate hemodynamic response function and entering them into a general linear model (GLM) (AFNI; Cox and Jesmanowicz, 1999). Motion estimates were included as regressors of no interest. After implementing our voxel selection criteria (see below), the per-voxel β values were entered into the similarity analysis.

2.5.2 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 Section 4 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 βs for those previously chosen 100 best voxels into vectors of 100 values per condition per ROI per subject (see Fig. 2). 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 A1–A6.

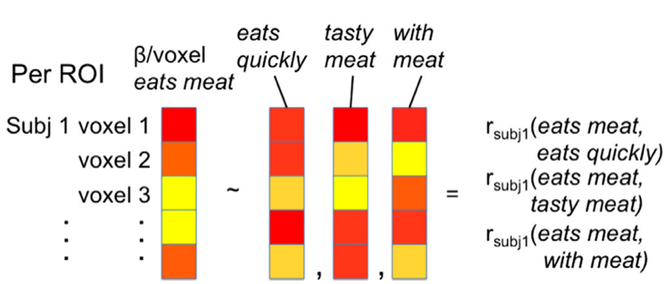

Fig. 2. Diagram of similarity analysis. A vector of the N (50, 100, 200, and 500) best voxels' β values for a given condition (e.g. +verb_arg, or *eats meat*) was correlated against another vector of β values from the same voxel array for a different condition in the same composition set (e.g. +verb_adj, or *eats quickly*).

2.5.3. 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.

3. Results

3.1. 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 Fig. 3, the shared verb correlations

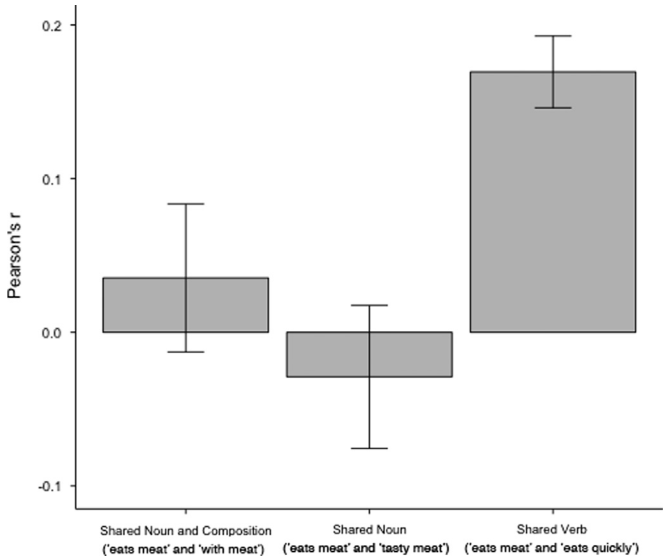

Fig. 3. Pairwise correlations between relevant pairs of word phrases in 100 best voxels in left AG. Error bars are 95% confidence intervals.

($r=0.17$) 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 A1 and A2 in Appendix). These findings favor the hypothesis that the left AG is specialized for verb semantics.

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 A3 and A4 in Appendix for left ATL comparisons). The noun-sharing phrases and the noun+composition-sharing phrases were significantly greater than both chance (Appendix Table A3) and the verb-sharing condition (Appendix Table A4), 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 3×2 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* ($F(2,96)=5.02$, $p=0.008$).

3.2. 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 A5 and A6 of Appendix for AG and ATL correlations across voxel selection criteria).

4. 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 et al., 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 et al., 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.

4.1. 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., 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) morpho-syntactically 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 and 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 et al. (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 et al. (2015) for evidence that both taxonomic and thematic associations engage AG, while ATL

Table A1
Pearson's *r* values for verb-relevant correlations in left AG.

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

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.

† $0.05 < p < 0.1$.
* $p < 0.05$.

Table A2
T-tests of differences between verb- or noun-based correlations in left AG (*p*-values reported).

	50 voxels	100 voxels	200 voxels	500 voxels
Shared verb vs. shared noun	***0.092***†	< ***0.0001****	***0.052***†	***0.034****
Shared verb vs. shared noun + composition	*0.48* ns	< ***0.0001****	*0.23* ns	*0.11* ns
Shared noun vs. shared noun + composition	*0.54* ns	***0.051***†	*0.18* ns	*0.62* ns

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 A1 of Appendix.

† $0.05 < p < 0.1$.
* $p < 0.05$.

Table A3
Pearson's *r* values for verb-relevant correlations in left ATL.

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

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.

† $0.05 < p < 0.1$.
* $p < 0.05$.

Table A4
T-tests of differences between verb- or noun-based correlations in left ATL (*p*-values reported).

	50 voxels (ns)	100 voxels	200 voxels (ns)	500 voxels (ns)
Shared verb vs. shared noun	*0.96*	***0.041******	*0.27*	*0.92*
Shared verb vs. shared noun +composition	*0.91*	***0.0050******	*0.13*	*0.66*
Shared noun vs. shared noun+composition	*0.91*	*0.13* ns	*0.74*	*0.62*

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 A3 of Appendix.

* $p < 0.05$.

Table A5
Pearson's *r* values for correlations of AMT similarity norms with neural similarity scores in left AG split by nouns and verbs.

	50 voxels	100 voxels	200 voxels	500 voxels
Verbs	*−0.10*	***−0.12****	*−0.096*	*−0.14*
Nouns	*−0.016*	*0.023*	*−0.039*	*−0.028*

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) items) 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 items. Table A5 shows those verb- and noun-based correlations in best (unsigned) 50, 100, 200, and 500 voxels in left AG.

* $p < 0.05$.

Table A6
Pearson's *r* values for correlations of AMT similarity norms with neural similarity scores in left ATL split by nouns and verbs.

	50 voxels	100 voxels	200 voxels	500 voxels
Verbs	*−0.013*	*0.041*	*0.038*	*0.026*
Nouns	*−0.090*	*−0.097*	*−0.034*	*−0.048*

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) items) 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 items. Table A6 shows those verb- and noun-based correlations in best (unsigned) 50, 100, 200, and 500 voxels in left 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.

4.2. 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 et al., 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 and 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 and 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 et al., 2005; Wang et al., 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.

4.3. 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 and 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 et al., 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 et al. (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.

Acknowledgments

This research was supported by National Institutes of Health Grants R0I-DC009209 and R01-EY021717 to S.T.-S. and an NSF Graduate Research Fellowship to C.B. The authors wish to thank members of the Thompson-Schill and Trueswell labs, as well as participants in the Common Ground seminars at the Institute for Research in Cognitive Science at the University of Pennsylvania.

Appendix A

See Tables A1–A6.

References

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. Cogn. Neurosci. 1 (1), 44–51.

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. J. Cogn. Neurosci 9 (5), 664–686.

Bedny, M., Caramazza, A., Grossman, E., Pascual-Leone, A., Saxe, R., 2008. Concepts are more than percepts: the case of action verbs. J. Neurosci. 28 (44), 11347–11353.

Bedny, M., Dravida, S., Saxe, R., 2013. Shindigs, brunches, and rodeos: the neural basis of event words. Cogn. Affect. Behav. Neurosci. 14, 891–901.

Bedny, M., Thompson-Schill, S.L., 2006. Neuroanatomically separable effects of imageability and grammatical class during single-word comprehension. Brain Lang. 98 (2), 127–139.

Beeman, M., Friedman, R., Grafman, J., Perez, E., Diamond, S., Lindsay, M., 1994. Summation priming and coarse semantic coding in the right hemisphere. J. Cogn. Neurosci. 6 (1), 26–45.

Bemis, D.K., Pylkkänen, L., 2011. Simple composition: a magnetoencephalography investigation into the comprehension of minimal linguistic phrases. J. Neurosci. 31 (8), 2801–2814.

Bemis, D.K., Pylkkänen, L., 2012. Basic linguistic composition recruits the left anterior temporal lobe and left angular gyrus during both listening and reading. Cereb. Cortex 23, 1859–1873.

Ben Shalom, D., Poeppel, D., 2007. Functional anatomic models of language: assembling the pieces. Neuroscientist 14 (1), 119–127.

Binder, J.R., Desai, R.H., 2011. The neurobiology of semantic memory. Trends Cogn. Sci. 15, 527–536 ⟨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. Cereb. Cortex 19 (12), 2767–2796.

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. J. Cogn. Neurosci. 11 (1), 80–93.

Binder, J.R., Westbury, C., McKiernan, K., Possing, E., Medler, D., 2005. Distinct brain systems for processing concrete and abstract concepts. J. Cogn. Neurosci. 17 (6), 905–917.

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. J. Cogn. Neurosci. 24 (10), 1998–2014.

Bird, H., Howard, D., Franklin, S., 2000. Why is a verb like an inanimate object? Grammatical category and semantic category deficits. Brain Lang. 72 (3), 246–309.

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.

Boylan, C., Trueswell, J.C., Thompson-Schill, S.L., 2014. Multi-voxel pattern analysis of noun and verb differences in ventral temporal cortex. Brain Lang. 137, 40.

Buccino, G., Binkofski, F., Fink, G.R., Fadiga, L., Fogassi, L., Gallese, V., et al., 2001. Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. Eur. J. Neurosci. 13; , pp. 400–404.

Buhrmester, M., Kwang, T., Gosling, S.D., 2011. Amazon's mechanical turk a new source of inexpensive, yet high-quality, data? Perspect. Psychol. Sci. 6 (1), 3–5.

Cappa, S.F., Perani, D., 2003. The neural correlates of noun and verb processing. J. Neurolinguist. 16 (2–3), 183–189.

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.

Chao, L.L., Martin, A., 2000. Representation of manipulable man-made objects in the dorsal stream. NeuroImage 12 (4), 478–484.

Clarke, A., Tyler, L.K., 2014. Object-specific semantic coding in human perirhinal cortex. J. Neurosci. 34 (14), 4766–4775.

Connolly, A.C., Guntupalli, J.S., Gors, J., Hanke, M., Halchenko, Y.O., Wu, Y.-C., et al., 2012. The representation of biological classes in the human brain. J. Neurosci. 32 (8), 2608–2618.

Coutanche, M.N., Thompson-Schill, S.L., 2014. Creating concepts from converging features in human cortex. Cereb. Cortex 25, 2584–2593.

Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput. Biomed. Res. 29 (3), 162–173.

Cox, R.W., Jesmanowicz, A., 1999. Real-time 3D image registration for functional MRI. Magn. Reson. Med. 42 (6), 1014–1018.

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. Front. Hum. Neurosci. 7, 303.

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 Lang. 116 (1), 33–49. http://dx.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. Proc. Natl. Acad. Sci. 90 (11), 4957–4960.

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://dx.doi.org/10.1006/nimg.2001.0775.

Dehaene, S., Piazza, M., Pinel, P., Cohen, L., 2003. Three parietal circuits for number processing. Cogn. Neuropsychol 20 (3–6), 487–506.

Embick, D., Marantz, A., Miyashita, Y., O'Neil, W., Sakai, K.L., 2000. A syntactic specialization for Broca's area. Proc. Natl. Acad. Sci. 97 (11), 6150–6154.

Estes, Z., 2003. Attributive and relational processes in nominal combination. J. Mem. Lang. 48 (2), 304–319.

Fair, D.A., Cohen, A.L., Power, J.D., Dosenbach, N.U.F., Church, J.A., Miezin, F.M., et al., 2009. Functional brain networks develop from a "local to distributed" organization. PLoS Comput. Biol. 5 (5), e1000381. http://dx.doi.org/10.1371/journal.pcbi.1000381.

Fairhall, S.L., Caramazza, A., 2013. Brain regions that represent amodal conceptual knowledge. J. Neurosci. 33 (25), 10552–10558. http://dx.doi.org/10.1523/JNEUROSCI.0051-13.2013.

Felleman, D.J., Van Essen, D.C., 1991. Distributed hierarchical processing in the primate cerebral cortex. Cereb. Cortex 1 (1), 1–47.

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.

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://dx.doi.org/10.1126/science.1106138.

Gainotti, G., 1996. Cognitive and anatomical locus of lesion in a patient with a category-specific semantic impairment for living beings. Cogn. Neuropsychol. 13 (3), 357–390. http://dx.doi.org/10.1080/026432996381944.

Gottlieb, J., 2007. From thought to action: the parietal cortex as a bridge between perception, action, and cognition. Neuron 53 (1), 9–16.

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://dx.doi.org/10.1016/j.neuroimage.2010.06.055.

Haxby, J.V., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science 293 (5539), 2425–2430. http://dx.doi.org/10.1126/science.1063736.

Heim, I., Kratzer, A., 1998. Semantics in Generative Grammar. Blackwell, Malden, MA.

Hoeks, J.C.J., Stowe, L.A., Doedens, G., 2004. Seeing words in context: the interaction of lexical and sentence level information during reading. Cogn. Brain Res. 19 (1), 59–73.

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://dx.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.

Humphries, C., Binder, J.R., Medler, D., Liebenthal, E., et al., 2006. Syntactic and semantic modulation of neural activity during auditory sentence comprehension. J. Cogn. Neurosci 18 (4), 665–679.

Kravitz, D.J., Saleem, K.S., Baker, C.I., Mishkin, M., 2011. A new neural framework for visuospatial processing. Nat. Rev. Neurosci. 12 (4), 217–230.

Kriegeskorte, N., 2008. Representational similarity analysis – connecting the branches of systems neuroscience. Front. Syst. Neurosci. 2, 4. http://dx.doi.org/10.3389/neuro.06.004.2008.

Kuperberg, G.R., 2007. Neural mechanisms of language comprehension: challenges to syntax. Brain Resc 1146, 23–49. http://dx.doi.org/10.1016/j.brainres.2006.12.063.

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://dx.doi.org/10.1016/j.neuroimage.2007.10.009.

Lambon Ralph, M.A., 2014. Neurocognitive insights on conceptual knowledge and its breakdown. Philos. Trans. R. Soc. Lond. B: Biol. Sci. 369 (1634), 20120392. http://dx.doi.org/10.1098/rstb.2012.0392.

Lau, E.F., Phillips, C., Poeppel, D., 2008. A cortical network for semantics:(de) constructing the N400. Nat. Rev. Neurosci. 9 (12), 920–933.

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://dx.doi.org/10.1016/j.neuropsychologia.2015.01.011.

Liljeström, M., Tarkiainen, A., Parviainen, T., Kujala, J., Numminen, J., Hiltunen, J., et al., 2008. Perceiving and naming actions and objects. NeuroImage 41 (3), 1132–1141. http://dx.doi.org/10.1016/j.neuroimage.2008.03.016.

Lin, Y., Michel, J.-B., Aiden, E.L., Orwant, J., Brockman, W., Petrov, S., 2012. Syntactic annotations for the Google Books Ngram Corpus. In: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics, vol. 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.

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 Cogn. Process 28 (8), 1154–1168. http://dx.doi.org/10.1080/01690965.2012.672754.

Musz, E., Thompson-Schill, S.L., 2014. Semantic variability predicts neural variability of object concepts. Neuropsychologia 76, 41–51. http://dx.doi.org/10.1016/j.neuropsychologia.2014.11.029.

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. J. Cogn. Neurosci. 25, 1824–1850. http://dx.doi.org/10.1162/jocn_a_00442.

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

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. Nat. Rev. Neurosci. 8 (12), 976–987.

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://dx.doi.org/10.1126/science.1117645.

Price, C.J., 2010. The anatomy of language: a review of 100 fMRI studies published in 2009. Ann. N. Y. Acad. Sci. 1191 (1), 62–88.

Pylkkänen, L., Brennan, J., Bemis, D.K., 2011. Grounding the cognitive neuroscience of semantics in linguistic theory. Lang. Cogn. Process. 26 (9), 1317–1337.

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. Proc. Natl. Acad. Sci. 98 (2), 676–682. http://dx.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. Nat. Neurosci. 12 (6), 718–724. http://dx.doi.org/10.1038/nn.2331.

Schwartz, M.F., Kimberg, D.Y., Walker, G.M., Brecher, A., Faseyitan, O.K., Dell, G.S., et al., 2011. Neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. Proc. Natl. Acad. Sci. 108 (20), 8520–8524.

Seghier, M.L., 2012. The angular gyrus: multiple functions and multiple subdivisions. 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. J. Neurosci. 30 (50), 16809–16817. http://dx.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, UK.

Shapiro, K.A., Moo, L.R., Caramazza, A., 2006. Cortical signatures of noun and verb production. Proc. Natl. Acad. Sci. 103 (5), 1644–1649. http://dx.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. Cereb. Cortex 19 (10), 2209–2229.

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. J. Cogn. Neurosci. 20 (11), 2037–2057.

Sitnikova, T., Holcomb, P.J., Kuperberg, G.R., 2008b. Neurocognitive mechanisms of human comprehension. In: Shipley, Thomas F., Zacks, Jeffrey M. (Eds.), Understanding Events: From Perception to Action, pp. 639–683.

Stringer, S.M., Rolls, E.T., 2002. Invariant object recognition in the visual system with novel views of 3D objects. Neural Comput. 14 (11), 2585–2596. http://dx.doi.org/10.1162/089976602760407982.

Talairach, J., Tournoux, P., 1988. Co-Planar Stereotaxic Atlas of the Human Brain. Thieme Medical Publishers, Inc., New York.

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. J. Cogn. Neurosci. 19 (11), 1753–1767.

Thompson, C.K., Bonakdarpour, B., Fix, S.F., 2010. Neural mechanisms of verb argument structure processing in agrammatic aphasic and healthy age-matched listeners. J. Cogn. Neurosci. 22 (9), 1993–2011.

Tranel, D., Damasio, H., Damasio, A.R., 1997. A neural basis for the retrieval of conceptual knowledge. Neuropsychologia 35 (10), 1319–1327. http://dx.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 Lang. 92 (3), 288–299. http://dx.doi.org/10.1016/j.bandl.2004.01.011.

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://dx.doi.org/10.1093/brain/124.8.1619.

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. Neurosci. Biobehav. Rev. 35 (3), 407–426. http://dx.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.

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. Hum. Brain Mapp. 31 (10), 1459–1468. http://dx.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://dx.doi.org/10.1016/j.neuropsychologia.2008.12.029.

Westerlund, M., Kastner, I., Kaabi, M. Al, Pylkkänen, L., 2015. The LATL as locus of composition: MEG evidence from English and Arabic. Brain Lang. 141, 124–134. http://dx.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://dx.doi.org/10.1016/j.neuropsychologia.2014.03.001.

Wisniewski, E.J., Love, B.C., 1998. Relations versus properties in conceptual combination. J. Mem. Lang. 38, 177–202.

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://dx.doi.org/10.1016/j.neuroimage.2009.12.036.