Neuropsychologia 100 (2017) 131–143

Contents lists available at ScienceDirect

Neuropsychologia

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

Early sensitivity of left perisylvian cortex to relationality in nouns and verbs

Adina Williams[a,*,1], Samir Reddigari[b,1], Liina Pylkkänen[a,b,c]

[a] Department of Linguistics, New York University, New York, NY 10003, United States
[b] Department of Psychology, New York University, New York, NY 10003, United States
[c] NYU Abu Dhabi Institute, New York University, Abu Dhabi, United Arab Emirates

ARTICLE INFO

Keywords:
Relationality
Magnetoencephalography
Inferior parietal lobe
Superior temporal cortex
Argument structure
Grammatical category

ABSTRACT

The ability to track the relationality of concepts, i.e., their capacity to encode a relationship between entities, is one of the core semantic abilities humans possess. In language processing, we systematically leverage this ability when computing verbal argument structure, in order to link participants to the events they participate in. Previous work has converged on a large region of left posterior perisylvian cortex as a locus for such processing, but the wide range of experimental stimuli and manipulations has yielded an unclear picture of the region's exact role(s). Importantly, there is a tendency for effects of relationality in single-word studies to localize to posterior temporo-parietal cortex, while argument structure effects in sentences appear in left superior temporal cortex. To characterize these sensitivities, we designed two MEG experiments that cross the factors relationality and eventivity. The first used minimal noun phrases and tested for an effect of semantic composition, while the second employed full sentences and a manipulation of grammatical category. The former identified a region of the left inferior parietal lobe sensitive to relationality, but not eventivity or combination, beginning at 170 ms. The latter revealed a similarly-timed effect of relationality in left mid-superior temporal cortex, independent of eventivity and category. The results suggest that i) multiple sub-regions of perisylvian cortex are sensitive to the relationality carried by concepts even in the absence of arguments, ii) linguistic context modulates the locus of this sensitivity, consistent with prior studies, and iii) relationality information is accessed early – before 200 ms – regardless of the concept's event status or syntactic category.

1. Introduction

A central goal for cognitive neuroscience is to understand how humans neurally represent our vast conceptual space. One way to divide this space is by the relationality of concepts: while *woman* and *car* describe properties of an individual, *brother* is a relation between two siblings. Similarly, in the verbal domain, a *laugh* is an action of one person whereas a *greeting* cannot happen without two or more people. In development, preschool-age children already show awareness of the relationality of concepts (Smiley and Brown, 1979; Mirman and Graziano, 2012) and this factor critically influences the syntactic distribution of the words conveying the meanings of concepts. But how is relationality neurally encoded? Here we addressed this by carefully comparing the effects of relationality with the effects of closely correlated factors, such as the eventivity and syntactic category of words.

Within the set of brain regions that are generally thought to comprise the brain's "semantic system" (Binder et al., 2009), one area in particular has risen as a possible candidate for the encoding of relational knowledge: the angular gyrus (AG), which lies in the posterior portion of the inferior parietal lobule (IPL). While the hypothesized roles of the AG cover a broad range of functions within and outside of language (Seghier, 2013), a series of studies has identified it as sensitive to the types of manipulations that a region encoding relations should respond to. In particular, the so-called "argument structure" of a word is a linguistic reflection of its relational structure: event concepts that imply one event participant are typically expressed in language via intransitive one-argument verbs (such as *run*) while events that require two participants typically map onto transitive, two-argument verbs (e.g., *bribe*). Thus manipulating the argument structure of linguistic expressions is a natural way to tap onto relational encoding. Within the literature pursuing this, the left AG (lAG) has shown more activation for verbs with complex argument structures than for verbs with simple argument structures in both aphasic and normal populations (Thompson et al., 2007, 2010; Meltzer-Asscher et al., 2013, 2015). However, while a relationality hypothesis is capable of explaining these findings, it is not the only such hypothesis.

Alternatively, the AG has been proposed to play a unique role in the

* Corresponding author.
E-mail addresses: adinawilliams@nyu.edu (A. Williams), samir.reddigari@nyu.edu (S. Reddigari), liina.pylkkanen@nyu.edu (L. Pylkkänen).
[1] Denotes equal contribution.

http://dx.doi.org/10.1016/j.neuropsychologia.2017.04.029
Received 18 November 2016; Received in revised form 9 April 2017; Accepted 23 April 2017
Available online 24 April 2017
0028-3932/

representation of event concepts (Binder and Desai, 2011). Such a hypothesis has plausible anatomical motivation given that the AG is contiguous with brain regions that contribute to spatio-temporal processing (Kravitz et al., 2011). An event structure hypothesis was also proposed by Boylan et al. (2015), who observed that the left AG tracks information relating to the presence of an event-denoting verb and suggested that this region may represent information specific to verbs. If the AG represents event concepts, then one might expect “bigger” or more complex events to activate it more, explaining why verbs naming events with more participants activate the AG more.

Also relevant to the notions of relationality and eventivity are the results of a PET study by Damasio et al. (2001). Although this experiment did not focus on these semantic features, a number of incidental findings speak to the role of the left IPL in the processing of both. First, naming actions elicited more activity in the lAG compared to naming concrete entities, consistent with the event hypothesis described above. Second, naming actions that involve tools (which tend to be transitive verbs) versus those not involving tools (generally intransitive), elicited more activity in the left supramarginal gyrus (SMG), which lies directly anterior to the lAG. Lastly, a similarly localized effect was reported for the naming of spatial relations compared to concrete entities. While these clusters were centered in the SMG rather than the AG, they offer evidence for a sensitivity to relationality in the left IPL. However, it is important to realize that two of these three contrasts (actions minus entities and spatial relations minus entities) are comparisons between different grammatical categories (verb/noun and preposition/noun, respectively). Such a confound might be eliminated by investigating effects of relationality and eventivity within nouns or verbs alone.

Finally, the AG has been hypothesized as a general site for the “integration of semantic knowledge” (Lau et al., 2008) or “conceptual combination” (Price et al., 2016, 2015; Graves et al., 2010). Given the evidence that single words activate their possible contexts (McDonald and Shillcock, 2001; Baayen et al., 2011), an integration hypothesis could also explain the relationality findings. Crucially, relationality, eventivity, and integration have not been pitted against each other in any single experiment in the extant literature.

Although argument structure, eventivity and integration are all somewhat related, and often correlated, aspects of language, they can be teased apart. In particular, research on the angular gyrus (or the neural correlates of argument structure processing in general) has not yet taken advantage of a class of semantically relational words that are non-eventive, namely relational nouns. Relational nouns name static relations between individuals: *mother* names a relation between a child and their female parent, and *age* names a relation between an individual and the amount of time it has existed. Intuitively, one can appreciate the relationality of a term like *mother* by the fact that one cannot be a mother without being the mother of someone whereas one can be a woman without having a similar attached relation. Research in theoretical semantics has shown that just like the argument structure of verbs, the relational structure of nouns has grammatical consequences (Baker, 1968; Löbner, 1981; Caponigro and Heller, 2007; Nathan, 2006; Barker, 2016). Kinship terms are naturally a prominent class of relational nouns and in fact, a recent voxel-wise modeling of fMRI data revealed a strong clustering of social concepts, populated by kinship terms and other words relating to human social interactions, in the AG (Huth et al., 2016, Supplementary materials).

Of the studies discussed above on the role of the lAG, most employ experimental paradigms where lexical items are presented in isolation, as word pairs, or in minimal phrases – environments that may not capture the way thematic relations or event concepts are processed in the brain when encountered in natural language. In fact, findings from the handful of studies utilizing sentential stimuli to investigate the brain basis of argument structure processing suggest sensitivities in posterior to middle regions of the STS and STG, rather than inferior parietal cortex. Shetreet et al. (2007) report graded activation in the posterior STG for increasing numbers of thematic options taken by verbs in Hebrew sentences, while a practically identical region was shown to respond to violations of unmarked transitivity in German (Grewe et al., 2007). By manipulating argument hierarchy and verb type, Bornkessel et al. (2005) implicated the left posterior STS and STG (but also the left intraparietal sulcus) in thematic role processing. More recently, an fMRI experiment revealed that classifiers trained on simple transitive sentences could reliably distinguish mirror image propositions (i.e. “The truck hit the ball” vs. “The ball hit the truck”) across active and passive constructions, as well as across subjects (Frankland and Greene, 2015). The activity driving the classification was localized to the left mid-superior temporal cortex (lmSTC), with distinct anterior sub-regions able to classify the identity of the agent and patient. Wang et al. (2016) identified an overlapping region whose activity successfully classified participants of animated events as the agent or patient. That this swath of the temporal lobe appears crucial in encoding the thematic roles of an event suggests that the context in which predicates are embedded can modulate where relationality is processed in the brain, and that previously reported lAG effects might not generalize to more natural or contextualized linguistic stimuli.

Finally, relational processing must obviously lie within a broader neural system accomplishing all of semantic processing. In addition to the inferior parietal cortex, this network is thought to consist of many integrative sites, including at least the anterior temporal lobe, ventromedial temporal cortex, dorsomedial prefrontal cortex, ventromedial prefrontal cortex, the inferior frontal gyrus and medial parietal cortex (Binder et al., 2009; Binder, 2016). How this network together orchestrates the many facets of semantic processing that are required for the brain implementation of natural language is still largely unknown. Revealing these dynamics will clearly require measurements with temporal resolution that reflects the speed of language processing, but such data are still a minority within the broader, mostly hemodynamic, literature on semantic processing. However, in our own work we have used magnetoencephalography (MEG) to characterize the spatiotemporal profiles of several specific integrative computations, contributing to our understanding of some of these integrative hubs. In this work, combinatory processing appears to begin in the left anterior temporal lobe (LATL) at ~200–250 ms after the comprehension of a word in its context (Bemis and Pylkkänen, 2011, 2012, 2013; Westerlund et al., 2015). At this point, composition appears to already be clearly semantic (Westerlund and Pylkkänen, 2014; Zhang and Pylkkänen, 2015), though limited to cases in which the meanings of the composing words and the specific combinatory rule are both relatively simple (Ziegler and Pylkkänen, 2016; Poortman and Pylkkänen, 2016); this may be because the comprehension of more complex words has not sufficiently unfolded by 200 ms in order for their meanings to be composed (Ziegler and Pylkkänen, 2016; cf., Binder and Desai, 2011). A later stage of combinatory effects has also been identified, localizing to ventromedial prefrontal cortex (vmPFC) at around 400 ms (Bemis and Pylkkänen, 2011) – this activity clearly responds to least some cases of complex semantic composition where LATL effects are absent, specifically cases of so-called “coercion” in which a syntax-semantics mismatch must be resolved (Brennan and Pylkkänen, 2008, 2010; Pylkkänen and McElree, 2007; Pylkkänen et al., 2009). Crucially, both the LATL and vmPFC combinatory effects have also been elicited during language production, supporting a clearly amodal semantic role for both (Pylkkänen et al., 2014). Finally a set of studies on reference resolution has identified the medial parietal cortex as sensitive to the linking of word meanings to their referents (Brodbecket al, 2016; Brodbeck and Pylkkänen, 2017). These effects have occurred relatively late, at roughly 400 ms after the onset of a reference resolving word, consistent with the idea that reference resolution cannot successfully take place until lexical access has had a chance to properly unfold.

In the context of this prior body of MEG research, a key motivation for the current study was our so far null findings on the AG: though the

AG has been hypothesized as an integrative hub, we have not yet seen evidence for this in our studies (with the exception of one study: Bemis and Pylkkänen, 2012). Thus, the present work sought to find positive MEG evidence for the participation of the AG in semantic processing, exploring the alternative variables of relationality and eventivity as possible modulators of this activity. Given the spatio-temporal dynamics revealed by the prior work, the timing of any such effects will clearly be of interest. For example, while the relational structure of a concept is in some sense a rather deep lexico-semantic property, making a relatively late effect of this variable plausible, a word's argument structure also crucially determines how it combines with the words around it. This could necessitate relatively early processing of relational structure. With the use of MEG, the current work was able to address this both for isolated words (Experiment 1) and for relational concepts within sentences (Experiment 2).

2. Experiment 1

The goal of Experiment 1 was to pit relationality against eventivity, and take the possibility of integration into account. To directly test whether the AG responds more strongly to eventivity or relationality while testing the effect of integration, Experiment 1 contrasted non-eventive high relational and low relational nouns (e.g., *name* vs. *desk*) with event-denoting nouns, which were also divided into two categories on the basis of their argument structure. Specifically, event-denoting nouns describing multiple-participant events, such as *embrace*, were labeled as ‘high relational’ and events describing single-participant events, such as *dance,* as ‘low-relational,’ yielding a crossed design varying both eventivity and (relative) relationality. To test the integration hypothesis, target nouns were then placed in minimal noun phrase contexts, preceded by an adjective, a possessor, or an unpronounceable consonant string. MEG was used to monitor brain activity, allowing us to not only focus our measurement on the critical items, but also to assess the detailed timing of any obtained effects, in contrast to prior hemodynamic literature.

2.1. Methods

2.1.1. Participants

Twenty three right-handed, native English speakers participated in this experiment. All had normal or corrected-to-normal vision and gave informed consent. Two participants were excluded due to excessive and irreducible environmental noise; a third participant was excluded due to response accuracy below 60% on the forced choice task; a fourth participant was excluded due to excessively long response times that were twice as long as the study average response time. As a result, nineteen participants were included in the final factorial data analysis (13 female, average age: 26.7 years).

2.1.2. Stimuli and experimental design

As shown in Fig. 1, this study employed a $2 \times 2 \times 3$ design with the factors relationality (high vs low), eventivity (eventive vs. non-eventive) and combinatorial context (non-combinatorial vs. possessor vs. adjective) fully crossed. The non-combinatorial contexts were created as in prior studies, by inserting unpronounceable consonant strings before the target noun (e.g., Bemis and Pylkkänen, 2011; Westerlund and Pylkkänen, 2014; Zhang and Pylkkänen, 2015; Boylan et al., 2015). The consonant strings were constructed to be length-matched with the average of the length of the possessor and the adjective that were associated with each target noun. Over the course of the experiment, each target noun appeared in each of the three contexts, and each context word appeared four times, once with each type of target noun. Stimulus order was fully random. We employed a forced-choice semantic relatedness task that was designed to hold the attention of the participants and to encourage them to fully process the stimuli for meaning.

Research in theoretical semantics has uncovered several grammatical consequences of the relational structure of nouns. For example, a relational noun can stand in for a question in the object position of a verb, whereas a non-relational noun sounds ill-formed in this position. This is illustrated by the contrast between *I found out Bill's age*, which can be paraphrased roughly as *I found out what Bill's age is*, whereas *I found out Bill's cat*, is hard to make sense of, *cat* being a non-relational noun (Baker, 1968; Löbner, 1981; Caponigro and Heller, 2007; Nathan, 2006; Barker, 2016). Similarly, relational nouns are more comfortable in the *of*-genitive in English than non-relational nouns, as shown by the contrast between *I got to know the brother of my boss* vs. *I got to know the cat of my boss* (Barker, 1995, Barker, 2011).

Target nouns were categorized into their stimulus conditions on the basis of four linguistic tests: two relationality tests and two eventivity tests. We relied on the expert intuitions of 4 trained linguists (2 native speakers and two fluent non-native speakers) and stipulated that at least two linguists needed agree that word in question passed a relevant test for that word to be included in our stimuli. We diagnosed target nouns as high relational if they passed one or both of the relationality tests: the X-of-John Test (e.g., ‘The *daughter / slaughter* of John’ vs. ‘The *policy/laugh* of John’; Barker, 1995; Barker, 2011), and a Concealed Questions Test (e.g., ‘I can guess your *weight’* vs. ‘I can guess your *coat*’; Heim, 1979; Nathan, 2006; Harris et al., 2008). If a target noun was found to be anomalous under both relationality tests, it was binned as low relational. The low relational items selected by our metric corresponded to basic objects (e.g. *chair/policy)*, and unergative event nominalizations (e.g. *laugh / yawn*); both of which lack an internal argument (see Marantz, 2013; Borer, 2012; Postal, 2010; Grimshaw,

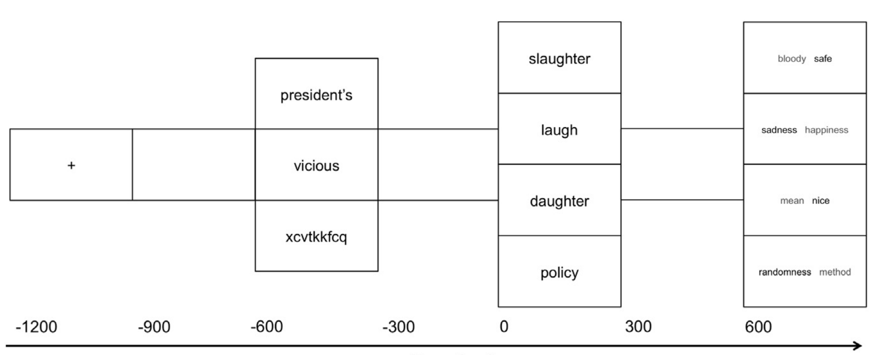

Fig. 1. Trial structure for Experiment 1. Participants viewed two words presented in sequence followed by a forced choice task in which they decided which of the following two elements was a “better fit” to the phrase, given here in green. The first word presented was either a possessor, and adjectival modifier, or an unpronounceable consonant string. The second word was the target that varied by relationality and eventivity. The task either targeted the first word or the second (target) word.

Table 1
Experiment 1 design.

	Possession	Modification	Consonant string
High relational eventive	president's slaughter	vicious slaughter	xcvtkkfcq slaughter
Low relational eventive	president's laugh	vicious laugh	xcvtkkfcq laugh
High relational non-eventive	president's daughter	vicious daughter	xcvtkkfcq daughter
Low relational non-eventive	president's policy	vicious policy	xcvtkkfcq policy

1990).

We selected eventive nominals by relying on two linguistic tests which diagnose the presence of an eventuality. According to (Davidson, 1969), manner adverbials and temporal prepositions can modify events, and therefore be used to diagnose the presence of an event. Two tests were adapted from this insight, which diagnosed whether a noun can be interpreted as a dynamic event: if a noun is eventive, it can be modifiable by a manner adverbial (e.g. '*the president's slaughter*, he did it viciously' vs. '*the president's policy*, he did it viciously'), or by a temporal prepositional phrase ('*the vicious slaughter* was at midnight' vs. '*the vicious daughter* was at midnight'). Thus, if a noun failed both eventivity tests, it was binned as non-eventive.

In our stimuli, there were two potential locations of ambiguity. First, most of the eventive nouns were noun-verb ambiguous, in isolation, though not all (such as the one-place predicate *laughter* and the two-place predicate *slaughter*). If this ambiguity were to affect the processing of relationality, we might expect these words to be processed differently in isolation than in adjectival or prepositional phrase contexts, which heavily biased readers to interpret these words as nouns. Second, high relational eventive nouns in our stimulus set often displayed a thematic ambiguity. For example *president's slaughter* could be interpreted as *the slaughter of the president,* or as *the slaughter by the president*. Since high relational eventive nouns was the only cell that was thematically ambiguous in this way, this design also let us probe whether ambiguities in thematic assignment would drive lAG activity. If this were the case, we might anticipate an interaction between eventivity and relationality.

To test for the sensitivity of the left angular gyrus to semantic integration of two concepts, as predicted by several proposals (e.g., Lau et al., 2008; Price et al., 2015), the critical stimuli were presented both as single words and in combinatory contexts. Two types of combinatory contexts were used: one that served to saturate an argument of the predicate and one that did not. To achieve the former, possessors were used, given that the English s-possessive can notoriously take on many semantic roles. For example, consider the following two instances of the s-possessive: *president's daughter* and *president's laughter*. With the relational noun *daughter*, the s-possessive provides a means to indicate one of the participants of the daughter-relation, namely the parent (the full proposition *Malia is the president's daughter* would serve to saturate both arguments of *daughter*). Similarly, in *president's laughter*, the s-possessive serves to introduce the laugher. Thus, by using the s-possessive as one of our contexts, we were able to have a matched context across the eventive and non-eventive nouns that in both cases served to name (or, more formally, "saturate") one participant of the relation. Consequently, if the AG responded particularly strongly to the saturation of the arguments of relations, as opposed to the sheer presence of relations, our design had the ability to capture this. Therefore, if composition type affects how relationality or eventivity are processed in the Angular Gyrus, we would expect an interaction.

In our stimuli, possessors were formed from animate occupation nouns and the Saxon Genitive (e.g. *president's, pilot's, director's*). We used occupation nouns because they are natural in a minimal, two word paradigm without determiners, and they can saturate the arguments of the high relational non-eventive nouns in the design (e.g. *president's daughter / height*). When a non-relational noun such as *desk* is composed with a possessor, as in *president's desk*, the possessor in a sense introduces a relation, or "coerces" the noun into a relational one (Barker, 1995), such that the resulting interpretation is 'desk associated with the president (in some way).' Thus, if the AG was specifically sensitive to the saturation of arguments, it might respond in a special way to combinations of possessors and non-relational nouns, since the argument slot needed by the possessor is not readily available within the semantics of non-relational nouns. Finally, as a non-saturating combinatory control, we used adjectival modifiers (e.g., *vicious daughter/laughter*). For adjective context words, they uniformly modified target nouns (e.g. *vicious slaughter* vs. *vicious laugh* vs. *vicious daughter* vs. *vicious policy*). The set of adjective-noun phrases contained a range of adjective types, including intersective (e.g. 'happy child', who is both happy and a child) and non-intersective adjectives (a scalar adjective such as 'loud' has a different interpretation depending on whether it is applied to a typically quiet event, such as 'whisper', or a typically loud one such as 'scream'; see Bierwisch, 1989; Kamp and Partee, 1995). In all, our design was a $2 \times 2 \times 3$ manipulation of relationality, eventivity, and context, as depicted in Table 1.

Lexical statistics of the target nouns are summarized in Table 2. Target nouns did not significantly differ in mean reaction time or letter bigram mean frequency (summed frequency divided by the number of successive bigrams, as determined by the English Lexicon Project (Balota et al., 2007; Lund and Burgess, 1996)). Due to the restrictions present on the class of high relational nouns, the stimuli were not perfectly matched for Log Frequency (in the HAL corpus), Number of Letters per word, and Number of Morphemes per word. We take the fact that relational nouns are more common in our study to derive from the empirical fact that relational nouns are more common in natural language (Gentner, 2005; Asmuth and Gentner, 2005). Based on the British National Corpus, Asmuth and Gentner (2005) estimated nearly half of the nouns in the adult vocabulary are relational. To address the inherent lexical differences present in the stimuli, we ran an additional

Table 2
Lexical statistics of Experiment 1 stimuli.

	ELP RT	ELP Acc.	Bigram Freq.	Length	Freq.	No. Morphemes
High relational eventive	649.43 (55.05)	0.97 (0.05)	1836 (846)	6.74 (1.95)	9.41 (1.44)	1.36 (0.63)
Low relational eventive	626.54 (64.79)	0.96 (0.04)	2086 (1186)	4.94 (1.50)	8.77 (1.34)	1.06 (0.24)
High relational non-eventive	624.95 (61.02)	0.98 (0.03)	2032 (858)	6.27 (1.96)	9.72 (1.55)	1.43 (0.65)
Low relational non-eventive	622.46 (58.21)	0.98 (0.03)	1730 (779)	5.22 (1.53)	9.17 (1.42)	1.16 (0.37)
Possessors	644.57 (75.76)	0.98 (0.02)	2405 (1078)	6.65 (1.83)	8.85 (1.52)	1.73 (0.60)
Modifiers	636.98 (73.78)	0.98 (0.03)	1854 (954)	6.20 (2.27)	9.87 (1.44)	1.57 (0.61)

single trial analysis that was able to take them into account (discussed in later sections).

Combinatory context words were matched for the following factors: mean reaction time on lexical decision, length in letters, number of morphemes, bigram frequency of the context word with the following target noun. We note that adjectives were slightly, but significantly more frequent than possessors, and had slightly higher letter bigram means. Keeping the noun and the modifiers constant across the whole paradigm resulted in most of the composed 2-word phrases in our materials being very low frequency (possession-noun had an average of 2.18 instances, modification- noun had an average of 8.98 instances out of 520 million words, $t(98) = 3.1282$, $p = 0.0023$) in the Corpus of Contemporary American English (Davies, 2008).

We chose 50 nouns for of each of the four types of target noun forming a total of 200 unique target nouns and presented each in the two combinatory contexts and the baseline non-combinatory context, resulting in a total of 600 trials. An alternative stimulus list was presented to two of the 19 participants. Specifically, after the first two participants, thirty-two of the 600 target nouns in the alternative list (5.3%; See Additional materials) were noticed to create confounds in the design; they were lexically ambiguous or orthographically anomalous. They were consequently removed for the remaining participants. Qualitative comparisons of behavioral results and brain activity between the two stimuli sets revealed no noticeable differences. Subjects receiving the alternate stimulus list were included in the analysis.

Fig. 1 shows the trial structure. Each trial consisted of a fixation cross, followed by a context word (i.e. possessor, adjective, or none), followed by the target noun. Each item was presented for 300 ms, with an intervening interval of 300 ms of blank screen. In order to verify that participants were paying attention, we added a forced-choice semantic task after every trial. At the end of the trial, two words would appear and participants chose which of the two options—the option presented on the left-side or the option presented on the right-side—was a better fit with the stimulus phrase. The task questions could be related to the target noun, the context word, or the combination of the two. For example, if the participant saw the trial 'xcvtkkfcq daughter' and the options 'male' and 'female', the participant would be expected to pick 'female'. The correct answer to half of the task questions was the left-side option and the correct answer to the other half would be the right-side option.

Trials were divided into ten blocks; each block contained 60 trials with the constraint that no two continuous trials share either the same context word or the same target noun. Stimulus sets were pseudo-randomized and constructed separately for each subject. Between blocks, participants could take a short rest, or continue immediately. The stimuli were presented by PsychToolBox software (Brainard, 1997; Pelli, 1997).

2.1.3. Procedure

A digital model of each subject's head shape was acquired with a FastSCAN laser scanner (Polhemus, Colchester, VT). The coordinates of three anatomical landmarks (the nasion and the left and right tragus) and five reference points were also digitized for purposes of coregistration. Before entering the magnetically shielded room, participants received a verbal explanation of the task and then practiced on a shortened block containing items that were distinct from the test blocks. If the participant answered less than 60% of the practice questions correctly on the practice block, the instructions were repeated verbally and the practice was repeated inside the magnetically shielded room prior to recording. The practice block provided feedback after responses that the participants could use to verify their comprehension, and to clarify task question with the experimenter before the recording began. No feedback was provided during recording.

Participants lay supine in a dimly lit and magnetically shielded room. The positions of the marker coils were measured at the beginning and at the end of the experiment. MEG data were recorded at the New York University facility using a whole-head MEG system (157 axial gradiometer sensors; Kanazawa Institute of Technology, Nonoichi, Japan) with 1000 Hz sampling rate, and filtered during data collection with a 200 Hz low-pass filter. Stimuli were projected onto a screen about 50 cm away from the participants' eyes. The target items and task questions were presented in white, 30-point, lower-case, Courier font on a grey background. In both the practice block and the recording blocks participants responded to the task question by using the index and middle fingers of their left hands to choose the correct option; button presses with the middle finger selected the left-side option, while button presses with the index finger selected the right-side option. The task options would remain on the screen until the button-press. The recording session lasted between 45 min and 1 h.

2.1.4. Data analysis

Data were pre-processed and analyzed with MNE-Python (Gramfort et al., 2013, 2014) and the Eelbrain wrapper (https://pythonhosted.org/eelbrain). Raw data were band-pass filtered offline between 1 and 40 Hz. We extracted epochs from −700 to 600 ms relative to the onset of the target nouns. Data were noise reduced via the Continuously Adjusted Least-Squares Method (Adachi et al., 2001), in the MEG Laboratory software (Yokogawa Electric Corporation and Eagle Technology Corporation, Tokyo, Japan). If one or more channel amplitudes was higher than 3000 fT or lower than −3000 fT, that epoch was identified as containing environmental noise and automatically removed. An average of 13.4% of trials were rejected due to noise or artifacts. Artifact rejection was based on visual inspection for known artefactual noise in the laboratory environment. For example, if the raw activity for individual channels presented the characteristic frontal distribution of a blink, that trial was rejected. Baseline correction was performed using data from the 100 ms before the first word of each phrase.

The position of the subject's head with respect to the MEG sensors was established by feeding current to five marker coils attached to the digitized reference points. The digitized head shape was co-registered to anatomical MRI data. Structural MRIs were available for 10 participants; these scans were acquired in a separate session at the Center for Brain Imaging at New York University (3 T Siemens Allegra scanner with T1-weighted MPRAGE sequences). For the remaining subjects we used the average structural MRI provided by FreeSurfer (http://surfer.nmr.mgh.harvard.edu/), which was linearly scaled to reflect the size of the subject's head. A source space containing 2562 sources per hemisphere was constructed on the cortical mantle of the scaled brain, and the forward solution for each source was calculated with a boundary element method using the inner skull boundary. The inverse operator was calculated with the forward solution and a regularized noise covariance matrix reflecting the channel covariance structure of the pre-stimulus intervals. This inverse operator was applied to the average evoked responses to obtain a time course of minimum norm estimates at each source for each condition. The direction of the current estimates was freely oriented with respect to the cortical surface, and thus all magnitudes were non-negative. The source estimates were then noise-normalized at each source (Dale et al., 2000), generating dynamic statistical parameter maps (dSPM) that were used in statistical analyses.

Our statistical analyses searched for spatio-temporal clusters that were reliably affected by our stimulus manipulation, corrected for multiple comparisons. To form a reasonably large area for the cluster search, we included in our search space a large portion of the left perisylvian cortex spanning the angular gyrus and surrounding regions (Brodmann Areas 39, 40, 1, 2, 3, 5, 7, 19, 22, 41, 42) as isolated from the PALS-B12 atlas (Van Essen, 2005). We then tested the hypotheses that this region is sensitive to relationality and/or eventivity with spatio-temporal cluster-permutation tests based on a $2 \times 2 \times 3$ ANOVA (see Nichols and Holmes, 2002; Maris and Oostenveld, 2007) so that we could fully investigate interactions with combinatory context words.

Since we did not have fine-grained predictions as to the timing of our effects, all tests were performed over a generous time window of 100–600 ms after the onset of the target noun, designed to cover both early and late estimates of lexical access (e.g., Hauk et al., 2004, 2006; Pylkkänen and Marantz, 2003) as well as the timing of previously reported effects of phrasal composition (e.g., Bemis and Pylkkänen, 2011). An F-statistic and a corresponding p-value were calculated at each source-time point pair for every main effect and interaction. Clusters were formed by identifying source-time points contiguous in space and time (at least 10 adjacent sources and a temporal extent of 25 ms) with p-values at or below 0.05, and cluster-level test statistics were calculated by summing the F-values within each cluster. We then repeated this procedure with 10,000 permutations of the condition labels and extracted for each permutation the test statistic of the largest cluster for each effect. These values represent the expected distribution of cluster-level statistics under the null hypothesis that the measured activity is unaffected by the factor in question. We then assigned a p-value to each cluster in the un-permuted data by calculating the proportion of values in the permutation distribution that were larger than the observed test statistic, thereby correcting for multiple comparisons non-parametrically at the cluster level (Maris and Oostenveld, 2007).

To address the concern that significant clusters might be driven by the unmatched lexical factors in the stimuli (i.e. length, frequency, and number of morphemes), we planned a follow-up of brain activity at the trial level. Average dSPM values within significant spatiotemporal clusters were averaged over space and time, and entered as the dependent variable in linear mixed-effects models. We used the lme4 package in R (Bates et al., 2015) to fit a linear mixed-effects model with a fixed effect for the effect of the cluster in question, as well as log frequency, number of morphemes, length, and mean bigram frequency, all taken from the English Lexicon Project (Balota et al., 2007). We included crossed random intercepts for participant and item (Baayen et al., 2008). The outcome variable was logarithmically transformed to approximate normality, and continuous variables were scaled to zero mean and unit variance. Significance was established by testing each estimated parameter against zero using the lmerTest package (Kuznetsova et al., 2016).

2.2. Results

2.2.1. Behavioral results

Accuracy on the semantic relatedness judgment task was high for low relational (83.08% correct) and high relational nouns (82.51% correct). Since the purpose of the task was simply to hold the participants' attention, and because task questions themselves varied in difficulty, we did not analyze the behavioral data further.

2.2.2. MEG results: factorial analysis

The $2 \times 2 \times 3$ spatio-temporal cluster permutation ANOVA performed over all 12 conditions on the search region including the left AG (Brodmann 39) and all contiguous Brodmann areas revealed a distributed spatio-temporal cluster in the left inferior parietal lobe from 170 to 260 ms post target noun onset ($F_{SUM} = 6489.3$, $p = 0.022$), with a main effect of relationality (Fig. 2). No main effects were found for eventivity or combinatory context, nor did we find any interactions. In each context and across eventivity, high relational nouns showed more activation than low relational ones. Individual pairwise *t*-tests on eventive nouns yielded significant differences between high and low relational nouns within possession [$t(18) = 2.80$ $p = 0.012$], modification [$t(18) = 2.81$, $p = 0.012$], and consonant string [$t(18) = 3.10$, $p = 0.006$] contexts. Individual pairwise *t*-test on non-eventive nouns yielded significant differences between high and low relational nouns within possession [$t(18) = 5.43$, $p < 0.001$], and consonant string [$t(18) = 2.23$, $p = 0.039$] contexts, but not within modification [$t(18) = 1.48$, $p = 0.155$].

2.2.3. MEG results: single trial analysis

The single trial analysis was designed to test whether the results of the factorial analysis could be attributed to some of the confounding lexical variables discussed above (e.g. word frequency, length in letters, or number of morphemes per word). None of them significantly predicted brain activity in our linear model, while relationality was significant (estimate = 0.00554, $p = 0.011$). See Table 3 for model coefficients and significance.

2.3. Summary and discussion

Experiment 1 pitted the hypothesis that the left AG is sensitive to relationality against the hypothesis that it is sensitive to eventivity. We found a region of the left inferior parietal lobe that is sensitive to relationality in the linguistic domain (as operationalized through varying the number of arguments nouns take) from 170 to 260 ms, which was verified using a single trial analysis. We found no significant effect of eventivity, suggesting that prior findings linking the AG to eventivity could have been driven by the inherent relationality of verbal predicates. Additionally, although our search region included a wide swatch of left perisylvian cortex (which included the lmSTC), no further clusters were identified. As for the hypothesis about conceptual combination, if it were the case that the LIPL engages in the integration of semantic information, we would expect Experiment 1 to find increased activation when the noun was modified when compared to the noun in isolation (without a modifier). In contrast to this, we found no main effect of context nor any modulation of the relationality effect by contextual considerations (i.e., no interaction). However, recent work shows that the left AG is sensitive to the canonicity of a combination as operationalized using two word co-occurrence, with lAG activity to plausible combinations positively correlating with their log co-occurrence frequency (Price et al., 2015, 2016). In Experiment 1, all our combinations were intuitively plausible but had low frequencies of co-occurrence, so based on the findings in Price et al., (2015, 2016), we might not expect to see an effect of contextual combination in lAG.

To further investigate contextual combination on our dataset, we tested our variables of interest in the right hemisphere as well. We anticipated that perhaps the right AG might contribute to relational, eventive, or combinatory processing of our stimuli, following the findings of Graves et al. (2010). This study found that canonically ordered noun-noun compounds (e.g., *lake house*) activated right AG more than their reversed counterparts (e.g., *house lake*), suggesting a role for the right AG in combinatory processing. When we tested the right hemisphere in this dataset for the same liberal time window (100–600 ms) and search region used for the spatio-temporal cluster on the left, we found no effect for of any of our variables of interest.

The prior literature on semantic processing has found AG activation for vastly different semantic tasks and manipulations. For example, in production studies using picture-word interference, thematic relations (e.g. COW-pasture) have been shown to elicit increased AG activation as compared to categorical (or taxonomic) relations (e.g. COW-rat) (de Zubicaray et al., 2013), a generalization also supported by deficit-lesion (Schwartz et al., 2011), TMS (Davey et al., 2015) and (somewhat more weakly) MEG data (Lewis et al., 2015). If the processing of thematically related items involves activation of a predicate that could relate them (e.g., cows graze on pasture), a region that encodes the relations between predicates and their arguments should be activated. First language acquisition work shows that words referring to relational concepts are acquired slower than those referring to concrete objects and familiar individuals (Gentner, 1978), supporting the idea that a certain level of neural sophistication is required to process relations. If there is a single function underwriting all the disparate findings in the literature—and their might not be; there might be some fine-grained cortical organization our neuroimaging tools cannot detect—it would have to be something very broad. A good candidate explanation could

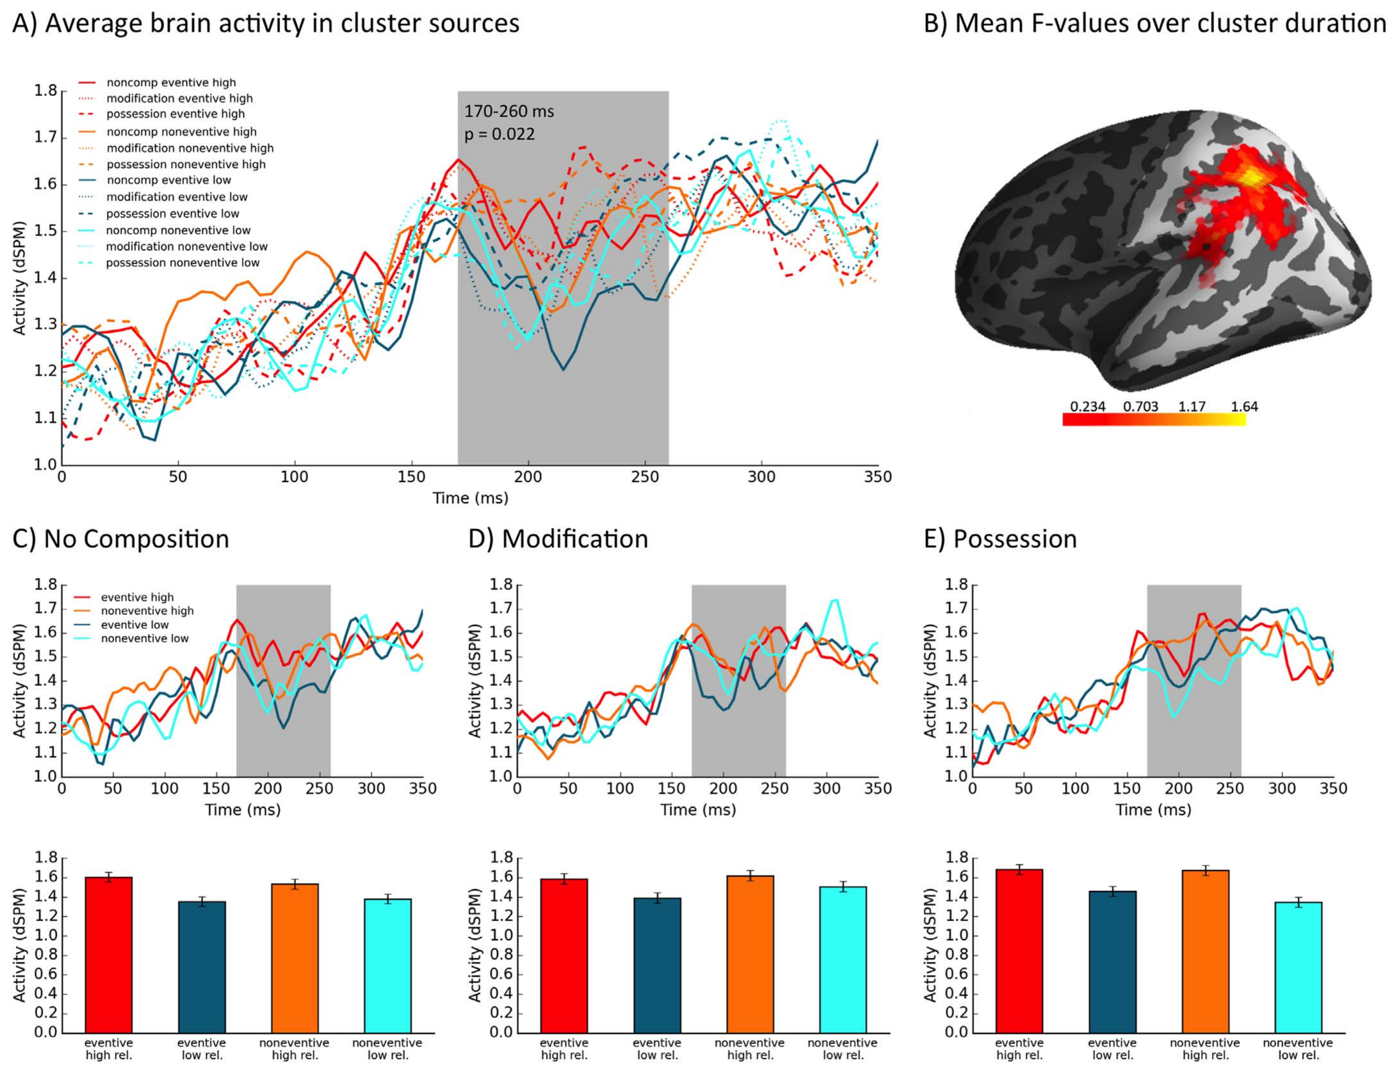

Fig. 2. Results of Experiment 1. A cluster of activity was identified in a left inferior parietal region showing increased activity for high relational nouns. The waveform plot in (A) shows the average brain activity for all conditions over the sources that form the cluster; the grey area delimits the temporal extent of the cluster. For ease of viewing, the waveforms in (A) have been separated by context into three bipartite sub-figures (C-E); the top portion of (C) shows average activation for the consonant string condition; the top portion of (D) shows average activation for the modification condition; the top portion of (E) shows average activation for the possession condition. The bottom portions display bar graphs that show the average activation by context over the sources in the cluster, averaged over the full temporal extent of the cluster; error bars show pooled standard error. The topographic map in (B) shows the F-values over the full spatial extent of the cluster in source space, with yellow portion corresponding to a higher F-value than red portion. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Single trial analysis results, Experiment 1.

	Estimate	Std. Err.	df	*t*	*p*
Intercept	0.00151	0.00959	18.5	0.16	0.876
Relationality	0.00554	0.00216	375.8	2.57	0.011*
Log Freq.	0.00080	0.00109	222.6	0.73	0.464
No. Morphemes	−0.00064	0.00128	198.1	−0.50	0.616
Length	−0.00034	0.00133	208.4	−0.26	0.799
Bigram Freq.	−0.00051	0.00106	240.1	−0.48	0.629

* $p \leq 0.05$.

be relationality, one of the primary high-level semantic abilities we possess.

3. Experiment 2

While Experiment 1 fully crossed eventivity and relationality, it did so only in minimal noun phrases, and not in a verbal or sentential context. Since much of the existing evidence for the lAG's role in processing argument structure comes from isolated verbs and verb phrases (and thus events), a thorough investigation of the neural correlates of event and argument structure processing should manipulate these factors in both nominal and verbal domains, and in a more natural linguistic context. (For a review of the large body of research on noun and verb representation, see Vigliocco et al., 2011.) We therefore designed a second experiment, in which noun-verb ambiguous items varying in their eventivity as nouns and transitivity as verbs were embedded at the beginning of generic sentences, which participants judged as true or false. For example, the transitive verb “bribe” and intransitive verb “riot” are both eventive as nouns, while “hammer” and “garden” are non-eventive nouns. Each word occurred once as a noun and once as a verb (e.g. “to garden”, “a garden”), always at the sentence initial position. The stimulus “A riot results in property damage.” should be judged generally true, while “To bribe someone is ethical” is false. Table 4 summarizes the design. Experiment 2 addressed three questions: i) whether effects of relationality in the lAG are replicated in sentential stimuli, ii) whether these effects are present for both event and non-event concepts, and iii) whether packaging these concepts as nouns rather than verbs modulates effects of relationality. Given the

Table 4
Experiment 2 design.

		Intransitive	Transitive
Eventive as Noun	Noun	a riot	a bribe
	Verb	to riot	to bribe
Non-Eventive as Noun	Noun	a garden	a hammer
	Verb	to garden	to hammer

inclusion of verb phrases in this experiment, we refer to the argument structure variable as "transitivity" rather than "relationality," but consider the manipulation to be conceptually the same.

We address these questions in two brain regions. The first is the lAG, which is strongly motivated by previous findings (Thompson et al., 2007, 2010; Meltzer-Asscher et al., 2013, 2015; Boylan et al., 2015.). The second is the lmSTC as defined by Frankland and Greene (2015), who proposed that this region encodes the arguments of a verb in sentential contexts.

3.1. Methods

3.1.1. Participants

Twenty-six right-handed native English speakers (11 male, mean age 27.8 years; three participants did not report age) participated in the experiment at New York University Abu Dhabi. All participants gave informed consent, and all were included in the analysis.

3.1.2. Stimuli and experimental design

Target stimuli consisted of 124 class-ambiguous words possessing a noun and a verb entry in CELEX2 (Baayen et al., 1995).

Eventivity of a noun was tested with the assumption that eventive nouns can be said to occur at a point in time, while non-eventive nouns cannot. We conducted a norming study on Amazon Mechanical Turk, where participants rated the naturalness of sentences of the form, "The [noun] happened last [time]", where "time" was one of the twelve months, four seasons, or the word "year." For example, "The bribe happened last September" was expected to be rated as natural, while "The hammer happened last year" was not. Each word was rated between 15 and 20 times on a scale from 1 to 7, with 7 being completely natural. An average rating of 4.0 was used as the cutoff between non-eventive and eventive nouns, and there was a significant difference between their average ratings (non-eventive: M = 1.87, eventive: M = 5.75, $t(122) = 28.3$, $p < 0.001$).

As an index of transitivity, we considered the frequency with which a verb occurs with a noun phrase complement according to the smoothed and filtered version of the VALEX subcategorization corpus, and used 0.50 to divide intransitive and transitive verbs (Korhonen et al., 2006). There was a significant difference in this measure between intransitive and transitive verbs (intransitive: M = 0.19, transitive: M = 0.78, $t(122) = 29.8$, $p < 0.001$). It is important to note at this point that the factor of transitivity was decided by the verb form of each stimulus. On the surface, the noun forms (especially of the non-eventive subset, i.e. *a hammer*) are no more relational than the intransitive ones (i.e. *a garden*). Any potential main effects of transitivity must be considered in light of this design.

Conditions were controlled for the following lexical variables: length, log CELEX frequency, verb-noun dominance as measured by the inverse tangent of the frequency ratio [see Tyler et al. (2008)], number of morphemes, orthographic bigram frequency, subcategorization frame entropy [see Linzen et al. (2013)], and lexical decision reaction time as reported in the English Lexicon Project (Balota et al., 2007). Because each word was presented in a noun context (preceded by *a/an*) and a verb context (preceded by *to*), we also controlled for the dominance of the relevant bigrams, again using the inverse tangent of the ratio between the bigram frequencies taken from the Corpus of Contemporary American English (Davies et al., 2008). Table 5 shows the means, standard deviations, and results of an eventivity by transitivity ANOVA on each of the controlled variables. Each target word was embedded at the beginning of two generic sentences, once as a noun in an indefinite determiner phrase (e.g. *a bribe*), and once as a verb in an infinite phrase (e.g. *to bribe*). Half of all sentences were generally true, and the rest generally false. The sentential stimuli were not otherwise controlled for content, as only the brain response to the target word was of interest in this study. Although the internal arguments of transitive verbs were eventually saturated in the sentences, the processing of those arguments is not captured in the activity analyzed. This is an important difference between the present experiment and previous neuroimaging studies of argument structure processing at the sentence level.

3.1.3. Procedure

Headshape digitization was conducted as in Experiment 1. Data for Experiment 2 were collected using a 208-channel axial gradiometer MEG system (Kanazawa Institute of Technology, Kanazawa, Japan) at New York University Abu Dhabi. Data were sampled at 1000 Hz with an online bandpass filter between 0.1 and 200 Hz.

All instructions and stimuli were projected onto a screen 90 cm from the subject's face, and were presented in white 24-point Courier font on a grey background. The experiment was divided into four blocks of 62 trials. Blocks were pseudo-randomized such that each contained roughly equal numbers of words from each condition, and a near equal true-false balance. In order to reduce the presence of eye movement artifacts, each trial began with a screen suggesting the subject blink if necessary. The trial structure was as follows. A fixation cross was displayed for 300 ms in the center of the screen, followed by a blank screen lasting 300 ms. The complete sentence was then displayed word-by-word, each word appearing for 300 ms and separated from the next by a blank screen also lasting 300 ms (Fig. 3). All words were lowercase, and no punctuation was included. Following the last word, a response prompt appeared, and the subject indicated their true/false judgment with a button press.

3.1.4. Data analysis

All subjects' headshapes were coregistered to a scaled FreeSurfer average brain, as no structural MRI data were available. An offline low-pass filter of 40 Hz was applied. Independent component analysis was employed on the filtered data to select and remove components corresponding to artifacts from eye movement, heartbeat, and environ-

Table 5
Lexical statistics and tests of Experiment 2 stimuli.

	ELP RT	Length	Log Freq.	No. Morphemes	NV Dom.	NV Bigram Dom.	SCF Entropy
Eventive Intransitive	668 (81)	5.97 (1.68)	2.51 (0.67)	1.10 (0.30)	0.55 (0.44)	0.47 (0.43)	1.73 (0.33)
Eventive Transitive	678 (70)	6.35 (1.40)	2.41 (0.51)	1.16 (0.45)	0.65 (0.34)	0.51 (0.38)	1.72 (0.42)
Non-Eventive Intransitive	659 (82)	5.52 (1.41)	2.51 (0.61)	1.06 (0.25)	0.43 (0.47)	0.48 (0.46)	1.85 (0.31)
Non-Eventive Transitive	639 (73)	5.90 (1.35)	2.75 (0.50)	1.13 (0.34)	0.51 (0.45)	0.65 (0.42)	1.75 (0.37)
ANOVA *p* val. (Event.)	0.08	0.09	0.10	0.60	0.09	0.34	0.26
ANOVA *p* val. (Trans)	0.74	0.14	0.47	0.30	0.28	0.16	0.42
ANOVA *p* val. (Interaction)	0.28	1.00	0.10	1.00	0.90	0.39	0.45

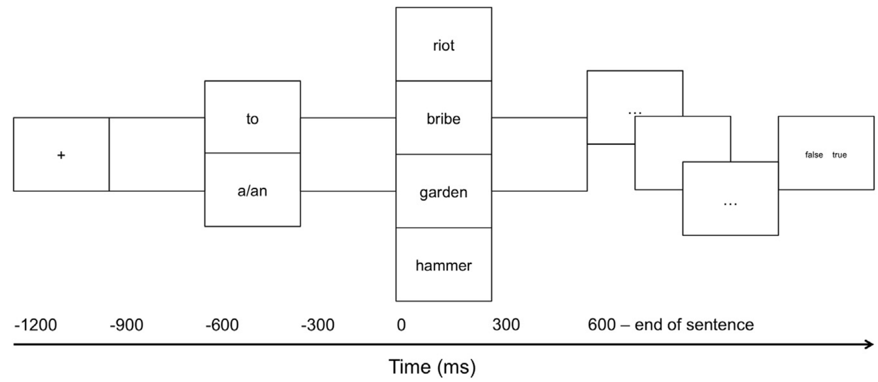

Fig. 3. Trial structure for Experiment 2. Participants read generic sentences beginning with an indefinite determiner phrase or infinitive phrase, the head word of which was the target word varying in its eventivity and relationality. Sentences were presented word-by-word and varied in length. At the end of the sentence, participants responded to the true/false prompt with a button press.

mental noise specific to the MEG device. The data were then epoched around the events of interest, beginning 200 ms prior to the context word, and extending 600 ms after presentation of the target word. Epochs in which the amplitude of any sensor still exceeded 2000 fT after ICA were excluded from analysis (eliminating 2.5% of epochs on average). Remaining epochs were averaged to obtain an evoked response for each category-eventivity-transitivity condition. Source estimation proceeded as in Experiment 1.

We conducted temporal permutation cluster tests in two regions of interest, the lmSTC and the lAG, from 100 to 400 ms after onset of the target word. The lmSTC ROI was defined by taking all sources within 25 mm of the coordinates reported in Frankland and Greene's (2015) first experiment (Talairach [−59, −25, 6], MNI [−63, −35, 3]). The lAG was Brodmann area 39 as defined in the PALS-B12 atlas (Van Essen, 2005). A 2 × 2 × 2 ANOVA with factors category, eventivity, and transitivity was conducted on each time point in each region. Clusters were formed by identifying temporally contiguous timepoints with p-values below 0.05, with a minimum cluster duration of 25 ms. Correction for multiple comparisons was performed via permutation tests as in Experiment 1, but with cluster-level test statistics calculated over temporal rather than spatiotemporal clusters (Maris and Oostenveld, 2007). To correct for running the test in multiple regions, the permutation distribution was generated from the largest cluster in *either* region for each permutation of the data.

3.2. Results

3.2.1. Behavioral results

Behavioral responses for one participant were not recorded because of a button box malfunction. The remaining participants performed at ceiling with mean accuracy of 95.6% on the true/false judgment task.

3.2.2. MEG results

The temporal permutation test on activity within the lmSTC revealed three separate main effects of transitivity: 105–170 ms ($p = 0.005$), 185–230 ms ($p = 0.032$), and 240–305 ms ($p = 0.008$). In each cluster, there was more activity for relational compared to non-relational words. There were no main effects of category or eventivity, nor any interactions.

In the lAG, no significant main effects or interactions were identified. There was a marginally significant 3-way interaction (category × eventivity × transitivity) from 250 to 290 ms ($p = 0.057$). There was no systematic pattern apparent in the mean activity within this cluster for any of the three factors. In the noun context, for the eventive stimuli, relational words elicited more activity than non-relational ones; for the non-eventive stimuli, relational words elicited less activity than their non-relational counterparts. These patterns both reversed in the verb context. Because the interaction was neither significant nor obviously interpretable, we do not attempt to further characterize or unpack this effect (Fig. 4).

All p-values are after correction across ROIs. All effects are visualized in Fig. 4.

3.3. Summary and discussion

Experiment 2 revealed a robust sensitivity to the relationality of the target word, independent of its event status or grammatical category, in the lmSTC. This effect manifested as three nearly contiguous clusters spanning from 105 ms to 305 ms after the onset of the target, within which relational concepts elicited more brain activity than non-relational ones. This finding is in line with previous fMRI studies which localize various effects of argument structure processing in sentences to the lateral temporal cortex, and not the lAG (Frankland and Greene, 2015; Bornkessel et al., 2005; Grewe et al., 2007; Shetreet et al., 2007; Ben-Shachar et al., 2003). However, the wide variety of manipulations and tasks in these studies (number of thematic options, argument animacy violations, etc.), as well as the inconsistent localization (middle and posterior regions of the left MTG, STS and STG) make it difficult to clarify the exact roles of these regions.

Frankland and Greene (2015) propose that the lmSTC encodes higher-order sentential meaning, given that the patterns of activity in this region could distinguish between mirror image propositions. In particular, they claim the region computes the sentence's agent and patient, or “who did what to whom,” a hypothesis supported by Wang et al. (2016) using nonlinguistic stimuli. Such a computation is an implausible explanation for the relationality effects observed in the lmSTC in the present study, as we analyzed brain activity in response to words whose arguments are not saturated. The present effect of relationality, then, suggests that the lmSTC is also involved in a lower level of argument structure processing than that described by Frankland and Greene.

While relationality affected lmSTC activity, we found no effects of eventivity (consistent with Experiment 1) nor of syntactic category. One factor relevant for the absence of syntactic category effects is that all our targets were in fact noun-dominant – perhaps this might have somehow mitigated against positive effects of category. However, Tyler et al. (2008) found that noun-verb dominance (defined with the same metric employed here) did not modulate brain activity evoked by ambiguous words, either in isolation or embedded in noun and verb phrases. We therefore consider it unlikely the noun-dominance of these stimuli obscured a category effect in either ROI.

4. General discussion

The present experiments identified two left perisylvian regions that are sensitive to the relationality of words, but not their event status. The

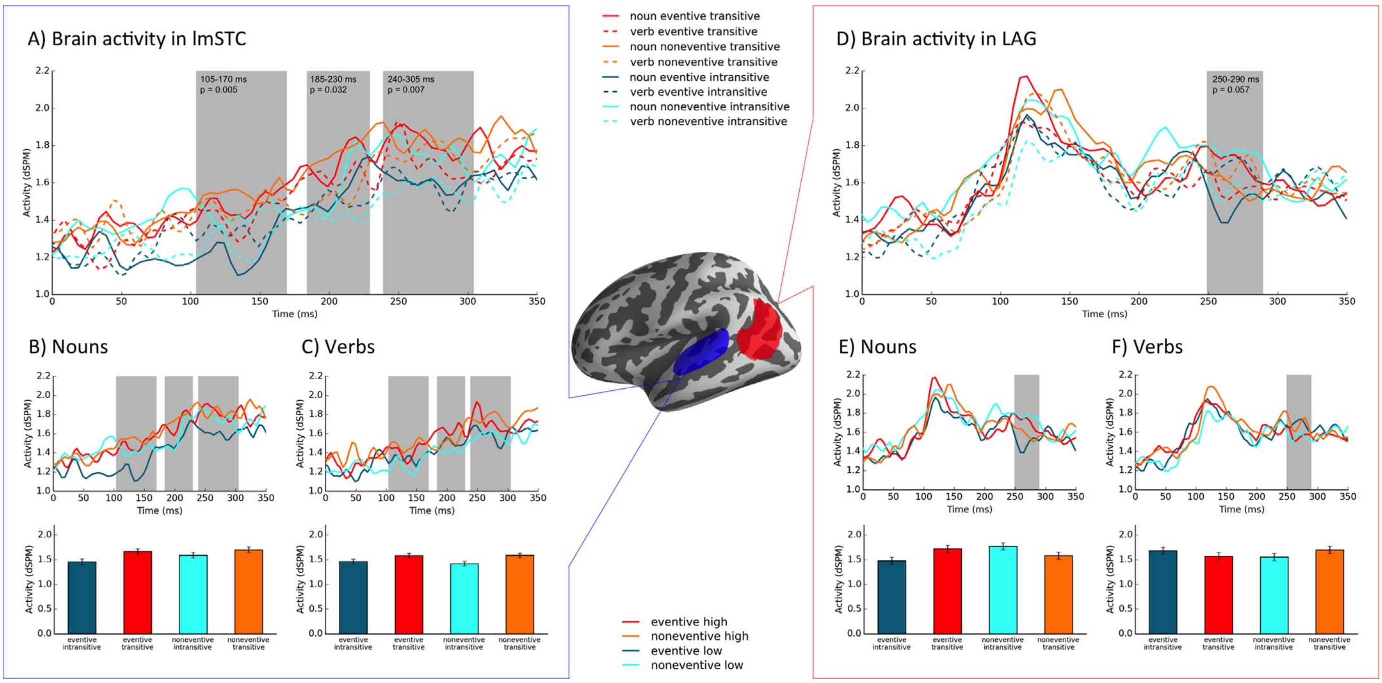

Fig. 4. Results of Experiment 2. Three main effects of relationality were identified in the lmSTC time course analysis. A trending 3-way interaction was found in the lAG (BA39).

first study used minimal noun phrases, and found the effect was independent of combinatory context. Experiment 2 used full sentences, and showed the effect to independent of grammatical category.

Our results in both brain regions are consistent with lesion studies: patients with posterior perisylvian lesions have trouble processing argument structure violations (McCann and Edwards, 2002; Kim and Thompson, 2000), and are less sensitive to argument structure in online sentence processing tasks (Shapiro and Levine, 1990). Of particular relevance to Experiment 2 are the findings of Wu et al. (2007), who showed that damage to lateral superior temporal cortex was associated with decreased comprehension of sentences with thematic information (*The circle kicked the square*) but not locative information (*The circle is above the square*). These posterior sensitivities further accord with models of language processing under which syntactic aspects of predicates activate portions of the posterior language network (Ben-Shachar et al., 2003; Hadar et al., 2002), as well as models which take lexical entries to be encoded in the mind with syntactic information about their use (e.g., grammatical class or subcategorization frame; Hagoort, 2003; Levelt, 1999; Bock and Levelt, 1994; Roelofs, 1992, 1993).

While previous studies implicating the superior temporal cortex in argument structure processing are limited to verbs, in Experiment 2 we found an effect of relationality that does not interact with category; the pattern is the same across the stimuli's nominal and verbal forms. This is of particular interest for the non-eventive noun phrases (e.g. *a hammer*, *a garden*). Although the verb *hammer* requires thematic roles (agent and patient), there is no relevant thematic information for the noun phrase *a hammer*. That these words still elicit more activity than the noun forms of intransitive verbs suggests that within a sentential context i) every potential meaning of a noun-verb ambiguous word is activated in lmSTC automatically, regardless of the category cued by context, and ii) the thematic information associated with eventive interpretation of a word is available even when that word is interpreted as non-eventive.

4.1. Effects of linguistic context on argument structure processing

A key finding of the present studies is that the two experiments revealed early effects of relationality in different brain regions: Experiment 1 in the left inferior parietal lobe, and Experiment 2 in the left mid superior temporal cortex. Though this is a between-experiment observation, the divergence is in line with evidence from fMRI studies discussed above, in which the argument structure of single words modulates activity in the LIPL, and argument structure manipulation in sentences affects activity in left middle and posterior temporal cortex. (See localization of argument structure effects from selected studies in Fig. 5.) The finding is also compatible with the model proposed by Thompson and Melzer-Asscher (2014), which describes a similar division of labor between inferior parietal and superior temporal regions. An advantage of the design of Experiment 2 is that the target word is always encountered in a minimal phrase, preceded by *a/an* or *to*. In this way it is analogous to the design of Experiment 1, and comparable to the single word presentations of lexical or semantic decision studies. Here, that phrase is embedded in a sentence, the entirety of which must be comprehended to complete the task. This embedding appears to be sufficient to recruit the lmSTC in argument structure processing, consistent with the hemodynamic studies that used sentential stimuli.

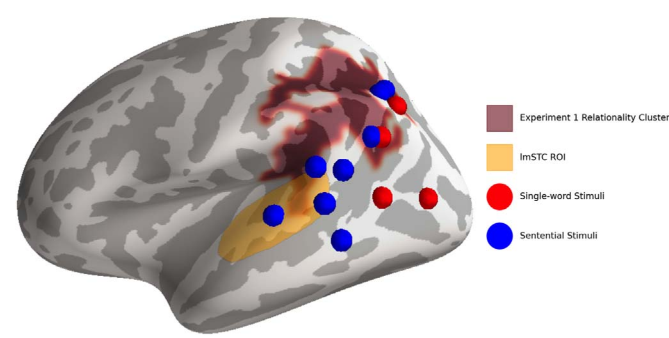

Fig. 5. Localization of argument structure effects in selected studies and present experiments. Single-word studies: Thompson et al. (2007, 2010); Meltzer-Asscher et al., (2013, 2015). Sentence studies: Ben-Shachar et al. (2003); Bornkessel et al. (2005); Grewe et al. (2007); Shetreet et al. (2007); Frankland and Greene (2015).

Such a consistent discrepancy suggests that the effects in these two regions may be driven by distinct processes. What computation is the lmSTC performing for words within sentences, but not for those in phrases or in isolation? The LIPL may process smaller units of meaning, while lmSTC may process relationality as it pertains to full propositions. For example, noun phrases can refer to individuals (e.g., president's daughter), or events (e.g., president's slaughter), but they do not convey the as much information as full propositions (e.g., an assault is a crime).

4.2. Timing of the relationality effects

In both experiments, the effects of relationality began strikingly early: 170 ms and 105 ms in Experiments 1 and 2, respectively. Such timing is unexpected if the time window between 100 and 200 ms after word onset is associated with processing of a word's visual form (Pylkkänen et al., 2002; Solomayak and Marantz, 2010; Simon et al., 2012; Laszlo et al., 2012; Fruchter and Marantz, 2015). However, early lexical effects have been reported in other domains. For example, word frequency effects (taken as evidence of lexical access) have been reported as early as 110 ms (Hauk and Pulvermüller, 2004; Hauk et al., 2006; Penolazzi et al., 2007). Likewise, sensitivities to semantic properties have been shown prior to 170 ms (Segalowitz and Zheng, 2009; Hauk et al., 2012; Lewis and Poeppel, 2014). To our knowledge, the present studies provide the first evidence that argument structure information is accessed before 200 ms. The timing of these effects suggests comprehenders glean relational information from nouns and verbs very early, which might imply that information about argument structure and conceptual relations may be prerequisite for further semantic processing.

4.3. Eventivity hypothesis

The spatiotemporal analysis in Experiment 1 found no main effect of or interaction with eventivity in the lAG or any other region, offering no evidence in support of Binder and Desai's (2011) proposal that the lAG is involved in event representation. An event structure hypothesis was also proposed by Boylan et al. (2015), who observed that the presence of an event-denoting verb activates the left AG, suggesting that this region may represent verb-specific information. These authors acknowledged however that their findings could be accounted for both in terms of eventivity and in terms of the relational/thematic structure of specific events. It also stands in contrast to the findings of Bedny et al. (2014) who reported more activity for event nouns relative to object nouns in the left MTG, left SPL, and left SFG; however, they did not address relationality as a potential confound in their noun stimuli. A relationality hypothesis could go towards explaining their data, because eventualities are often relational predicates that are predominantly realized in language as verbs (e.g. Bedny et al., 2014; Frawley, 1992; Langacker, 1987; Talmy, 1975). In Experiment 2, the time course analysis on the lAG found a marginally significant interaction between eventivity, category, and relationality; however, the pattern of activity did not suggest any systematic difference between eventive and non-eventive words. The present studies indicate that the brain's sensitivity to relationality (at least in the lAG and lmSTC) is not specific to event concepts.

4.4. Domain generality

If the LIPL does subserve as integral a cognitive function as relationality, this function might be detectable in non-linguistic domains as well. One of the main non-linguistic areas of cognitive science investigating the role of the AG is in the study of mathematics (Grabner et al., 2007, Pyke et al., 2015, Anderson et al., 2014). A large body of literature explores the role of the angular gyrus in arithmetic problem solving (Dehaene et al., 2003; Grabner et al., 2007; Bemis and Pylkkanen, 2013, Pyke et al., 2015, Anderson et al., 2014), as well as in the processing of non-singular semantics (Domahs et al., 2012). The shared ingredient between these disparate cognitive domains might be the capacity to encode a functional relation between multiple entities in the world, be it a linguistic correspondence that holds between the participants in an event, a thematic relationship that holds between pairs of individuals (e.g. a cow and its pasture) based on world knowledge, or a mathematical mapping between two numbers. It is possible that all of these competencies have a shared underlying cognitive component that sets up functional mappings between symbolic or mental representations.

5. Conclusion

In summary, the two experiments presented in this paper show that regions in the left perisylvian cortex display a sensitivity to relationality independent of eventivity. Furthermore, we found no sensitivity to combinatory context in the LIPL, as well as no sensitivity to grammatical category in the lmSTC. Both effects of relationality were isolated at relatively early time windows (with effects beginning prior to 200 ms), suggesting that argument structure information is extracted from lexical items early, and may be in the service of preparing for later semantic computations. When taken together, the two experiments suggest that the localization of relationality effects depends on the semantic context in which stimuli are presented: left IPL for phrasal stimuli, and lmSTC for propositional (sentential) stimuli.

Funding

This research was supported by the National Science Foundation grant BCS-1221723 (L. Pylkkänen) and grant G1001 from the NYUAD Institute, New York University Abu Dhabi (L. Pylkkänen).

Additional materials

The source estimates and stimulus lists for both experiments are available at: http://www.psych.nyu.edu/pylkkanen/lab/relationality_materials.html.

Acknowledgments

Special thanks to Becky Laturnus and Yohei Oseki for help during the stimulus design for Experiment 1, as well as to Paul Del Prato and Miriam Lautner for help with coding the stimulus presentation script for Experiment 1. We thank Jessica Mow for data collection for Experiment 2.

References

Adachi, Y., Shimogawara, M., Higuchi, M., Haruta, Y., Ochiai, M., 2001. Reduction of non-periodic environmental magnetic noise in MEG measurement by continuously adjusted least squares method. IEEE Trans. Appl. Supercond. 11 (1), 669–672.
Anderson, J.R., Lee, H.S., Fincham, J.M., 2014. Discovering the structure of mathematical problem solving. NeuroImage 97, 163–177.
Asmuth, J., Gentner, D., 2005. Context sensitivity of relational nouns. In: Proceedings of the twenty-seventh annual meeting of the cognitive science society. pp. 163–168.
Baayen, R.H., Piepenbrock, R., Gulikers, L., 1995. The CELEX lexical database (CD-ROM). Linguistic Data Consortium. University of Pennsylvania, Philadelphia, PA.
Baayen, R.H., Davidson, D.J., Bates, D.M., 2008. Mixed-effects modeling with crossed random effects for subjects and items. J. Mem. Lang. 59 (4), 390–412.
Baayen, R.H., Milin, P., Đurđević, D.F., Hendrix, P., Marelli, M., 2011. An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. Psychol. Rev. 118 (3), 438.
Baker, C.L., 1968. Indirect Questions in English. University of Illinois.
Balota, D.A., Yap, M.J., Hutchison, K.A., Cortese, M.J., Kessler, B., Loftis, B., Treiman, R., 2007. The English lexicon project. Behav. Res. Methods 39 (3), 445–459.
Barker, C., 1995. Possessive descriptions. CSLI, Stanford, CA.
Barker, C., 2011. Possessives and relational nouns. Semantics: An International Handbook of Natural Language Meaning, 2, 1109-1130.
Barker, C., 2016. Why relational nominals make good concealed questions. Lingua.
Bates, D., Mächler, M., Bolker, B., Walker, S., 2015. Fitting linear mixed-effects models using lme4. J. Stat. Softw. 67 (1), 1–48.
Bedny, M., Dravida, S., Saxe, R., 2014. Shindigs, brunches, and rodeos: The neural basis of event words. Cognit. Affect. Behav. Neurosci. 14 (3), 891–901.
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 (bhs170).
Bemis, D.K., Pylkkanen, L., 2013. Combination across domains: an MEG investigation into the relationship between mathematical, pictorial, and linguistic processing. Front. Psychol. 3, 583.
Ben-Shachar, M., Hendler, T., Kahn, I., Ben-Bashat, D., Grodzinsky, Y., 2003. The neural reality of syntactic transformations evidence from functional magnetic resonance imaging. Psychol. Sci. 14 (5), 433–440.
Bierwisch, M., 1989. The semantics of gradation. (261) In: Bierwisch, M., Lang, E. (Eds.), Dimensional Adjectives 71 Springer, Berlin.
Binder, J.R., 2016. *In defense* of abstract conceptual representations. Psychon. Bull. Rev. 23 (4), 1096–1108.
Binder, J.R., Desai, R.H., 2011. The neurobiology of semantic memory. Trends Cogn. Sci. 15 (11), 527–536.
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.
Bock, J.K., Levelt, W., 1994. Language production: grammatical encoding. In: Gernsbacher, Morton Ann (Ed.), Handbook of Psycholinguistics. Academic Press, San Diego, pp. 945–984.
Borer, H., 2012. In the event of a nominal. Theta Syst. Argum. Struct. Interface 103–149.
Bornkessel, I., Zysset, S., Friederici, A.D., von Cramon, D.Y., Schlesewsky, M., 2005. Who did what to whom? The neural basis of argument hierarchies during language comprehension. NeuroImage 26 (1), 221–233.
Boylan, C., Trueswell, J.C., Thompson-Schill, S.L., 2015. Compositionality and the angular gyrus: a multi-voxel similarity analysis of the semantic composition of nouns and verbs. Neuropsychologia 78, 130–141.
Brainard, D.H., 1997. The psychophysics toolbox. Spat. Vision. 10, 433–436.
Brennan, J., Pylkkänen, L., 2008. Processing events: behavioral and neuromagnetic correlates of aspectual coercion. Brain Lang. 106 (2), 132–143.
Brennan, J., Pylkkänen, L., 2010. Processing psych verbs: behavioural and MEG measures of two different types of semantic complexity. Lang. Cogn. Process. 25 (6), 777–807.
Brodbeck, C., Pylkkänen, L., 2017. Language in context: characterizing the comprehension of referential expressions with MEG. NeuroImage 147, 447–460.
Brodbeck, C., Gwilliams, L., Pylkkänen, L., 2016. Language in Context: MEG Evidence for Modality-General and-Specific Responses to Reference Resolution. eneuro, 3(6), ENEURO-0145.
Caponigro, I., Heller, D., 2007. The non-concealed nature of free relatives: Implications for connectivity in specificational sentences. Direct Compositionality, Ed. Barker, C., Jacobson, P., pp. 237–263.
Dale, A.M., Liu, A.K., Fischl, B.R., Buckner, R.L., Belliveau, J.W., Lewine, J.D., Halgren, E., 2000. Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. Neuron 26 (1), 55–67.
Damasio, H., Grabowski, T.J., Tranel, D., Ponto, L.L., Hichwa, R.D., Damasio, A.R., 2001. Neural correlates of naming actions and of naming spatial relations. NeuroImage 13 (6), 1053–1064.
Davey, J., Cornelissen, P.L., Thompson, H.E., Sonkusare, S., Hallam, G., Smallwood, J., Jefferies, E., 2015. Automatic and controlled semantic retrieval: tms reveals distinct contributions of posterior middle temporal gyrus and angular gyrus. J. Neurosci. 35 (46), 15230–15239.
Davidson, D., 1969. The individuation of events. In: Essays in honor of Carl G. Hempel (pp. 216–234). Springer.
Davies, M., 2008. The Corpus of Contemporary American English. BYE, Brigham Young University.
Dehaene, S., Piazza, M., Pinel, P., Cohen, L., 2003. Three parietal circuits for number processing. Cogn. Neuropsychol. 20 (3–6), 487–506.
Domahs, F., Nagels, A., Domahs, U., Whitney, C., Wiese, R., Kircher, T., 2012. Where the mass counts: common cortical activation for different kinds of nonsingularity. J. Cogn. Neurosci. 24 (4), 915–932.
Frankland, S.M., Greene, J.D., 2015. An architecture for encoding sentence meaning in left mid-superior temporal cortex. Proc. Natl. Acad. Sci. USA 112 (37), 11732–11737.
Frawley, W., 1992. Linguistic Semantics. Erlbaum, Hitlsdale, NJ.
Fruchter, J., Marantz, A., 2015. Decomposition, lookup, and recombination: meg evidence for the full decomposition model of complex visual word recognition. Brain Lang. 143, 81–96.
Gentner, D., 1978. On relational meaning: the acquisition of verb meaning. Child Dev. 988–998.
Gentner, D., 2005. The development of relational category knowledge. Build. Object Categ. Dev. Time 245–275.
Grabner, R.H., Ansari, D., Reishofer, G., Stern, E., Ebner, F., Neuper, C., 2007. Individual differences in mathematical competence predict parietal brain activation during mental calculation. NeuroImage 38 (2), 346–356.
Gramfort, A., Luessi, M., Larson, E., Engemann, D.A., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., Hämäläinen, M., 2013. MEG and EEG data analysis with MNE-Python. Front. Neurosci. 7, 267.
Gramfort, A., Luessi, M., Larson, E., Engemann, D.A., Strohmeier, D., Brodbeck, C., Hämäläinen, M.S., 2014. MNE software for processing MEG and EEG data. NeuroImage 86, 446–460.
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.
Grewe, T., Bornkessel-Schlesewsky, I., Zysset, S., Wiese, R., von Cramon, D.Y., Schlesewsky, M., 2007. The role of the posterior superior temporal sulcus in the processing of unmarked transitivity. NeuroImage 35 (1), 343–352.
Grimshaw, J., 1990. Argument Structure. The MIT Press, Cambridge, Massachusetts.
Hadar, U., Palti, D., Hendler, T., 2002. The cortical correlates of verb processing: Recent neuroimaging studies. Brain Lang. 83, 175–176.
Hagoort, P., 2003. How the brain solves the binding problem for language: a neurocomputational model of syntactic processing. NeuroImage 20, S18–S29.
Harris, J., Pylkkänen, L., McElree, B., Frisson, S., 2008. The cost of question concealment: eye-tracking and MEG evidence. Brain Lang. 107 (1), 44–61.
Hauk, O., Pulvermüller, F., 2004. Neurophysiological distinction of action words in the fronto-central cortex. Human. Brain Mapp. 21 (3), 191–201.
Hauk, O., Davis, M.H., Ford, M., Pulvermüller, F., Marslen-Wilson, W.D., 2006. The time course of visual word recognition as revealed by linear regression analysis of ERP data. NeuroImage 30 (4), 1383–1400.
Hauk, O., Coutout, C., Holden, A., Chen, Y., 2012. The time-course of single-word reading: evidence from fast behavioral and brain responses. NeuroImage 60 (2), 1462–1477.
Heim, I., 1979. Concealed questions. In Semantics from different points of view (pp. 51–60). Springer.
Huth, A.G., de Heer, W.A., Griffiths, T.L., Theunissen, F.E., Gallant, J.L., 2016. Natural speech reveals the semantic maps that tile human cerebral cortex. Nature 532 (7600), 453–458.
Kamp, H., Partee, B., 1995. Prototype theory and compositionality. Cognition 57 (2), 129–191.
Kim, M., Thompson, C.K., 2000. Patterns of comprehension and production of nouns and verbs in agrammatism: implications for lexical organization. Brain Lang. 74 (1), 1–25.
Korhonen, A., Krymolowski, Y., Briscoe, T., 2006. A large subcategorization lexicon for natural language processing applications. In: Proceedings of LREC (Vol. 6).
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.
Kuznetsova, A., Brockhoff, P.B., Christensen, R.H.B., 2016. Package lmerTest.
Langacker, R.W., 1987. Foundations of Cognitive Grammar: Theoretical Prerequisites. 1 Stanford University Press, Stanford, California.
Laszlo, S., Stites, M., Federmeier, K.D., 2012. Won't get fooled again: an event-related potential study of task and repetition effects on the semantic processing of items without semantics. Lang. Cogn. Process. 27 (2), 257–274.
Lau, E.F., Phillips, C., Poeppel, D., 2008. A cortical network for semantics: (de) constructing the N400. Nat. Rev. Neurosci. 9 (12), 920–933.
Levelt, W.J., 1999. Models of word production. Trends Cogn. Sci. 3 (6), 223–232.
Lewis, G., Poeppel, D., 2014. The role of visual representations during the lexical access of spoken words. Brain Lang. 134, 1–10.
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.
Linzen, T., Marantz, A., Pylkkänen, L., 2013. Syntactic context effects in single word recognition: an MEG study. Ment. Lex. 8 (2), 117–139.
Löbner, S., 1981. Intensional verbs and functional concepts: more on the" rising temperature" problem. Linguist. Inq. 471–477.
Lund, K., Burgess, C., 1996. Producing high-dimensional semantic spaces from lexical co-occurrence. Behav. Res. Methods, Instrum. Comput. 28 (2), 203–208.
Marantz, A., 2013. Verbal argument structure: events and participants. Lingua 130, 152–168.
Maris, E., Oostenveld, R., 2007. Nonparametric statistical testing of EEG-and MEG-data. J. Neurosci. Methods 164 (1), 177–190.
McCann, C., Edwards, S., 2002. Verb problems in fluent aphasia. In *Brain and Language* (Vol. 83,pp. 42–44). Academic Press Inc Elsevier Science 525 B St, Ste 1900, San Diego, Ca 92101-4495 USA.
McDonald, S.A., Shillcock, R.C., 2001. Rethinking the word frequency effect: the neglected role of distributional information in lexical processing. Lang. Speech 44 (3), 295–322.
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. Lang. Cogn. Process. 28 (8), 1154–1168.
Meltzer-Asscher, A., Mack, J.E., Barbieri, E., Thompson, C.K., 2015. How the brain processes different dimensions of argument structure complexity: evidence from fMRI. Brain Lang. 142, 65–75.
Mirman, D., Graziano, K.M., 2012. Individual differences in the strength of taxonomic versus thematic relations. J. Exp. Psychol.: General. 141 (4), 601.
Nathan, L.E., 2006. On the Interpretation of Concealed Questions. Massachusetts Institute of Technology.
Nichols, T.E., Holmes, A.P., 2002. Nonparametric permutation tests for functional neuroimaging: a primer with examples. Human Brain Mapp. 15 (1), 1–25.
Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat. Vision. 10 (4), 437–442.
Penolazzi, B., Hauk, O., Pulvermüller, F., 2007. Early semantic context integration and lexical access as revealed by event-related brain potentials. Biol. Psychol. 74 (3), 374–388.
Poortman, E.B., Pylkkänen, L., 2016. Adjective conjunction as a window into the LATL's contribution to conceptual combination. Brain Lang. 160, 50–60.
Postal, P.M., 2010. Edge-Based Clausal Syntax: A Study of (Mostly) English Object Structure. MIT Press, Cambridge, Massachusetts.
Price, A.R., Bonner, M.F., Peelle, J.E., Grossman, M., 2015. Converging evidence for the neuroanatomic basis of combinatorial semantics in the angular gyrus. J. Neurosci. 35 (7), 3276–3284.
Price, A.R., Peelle, J.E., Bonner, M.F., Grossman, M., Hamilton, R.H., 2016. Causal evidence for a mechanism of semantic integration in the angular gyrus as revealed by high-definition transcranial direct current stimulation. J. Neurosci. 36 (13), 3829–3838.
Pyke, A., Betts, S., Fincham, J.M., Anderson, J.R., 2015. Visuospatial referents facilitate

the learning and transfer of mathematical operations: Extending the role of the angular gyrus. Cogn. Affect. Behav. Neurosci. 15 (1), 229–250.
Pylkkänen, L., Marantz, A., 2003. Tracking the time course of word recognition with MEG. Trends Cogn. Sci. 7 (5), 187–189.
Pylkkänen, L., McElree, B., 2007. An MEG study of silent meaning. J. Cogn. Neurosci. 19 (11), 1905–1921.
Pylkkänen, L., Stringfellow, A., Marantz, A., 2002. Neuromagnetic evidence for the timing of lexical activation: an MEG component sensitive to phonotactic probability but not to neighborhood density. Brain Lang. 81 (1), 666–678.
Pylkkänen, L., Martin, A.E., McElree, B., Smart, A., 2009. The anterior midline field: coercion or decision making? Brain Lang. 108 (3), 184–190.
Pylkkänen, L., Bemis, D.K., Elorrieta, E.B., 2014. Building phrases in language production: an MEG study of simple composition. Cognition 133 (2), 371–384.
Roelofs, A., 1992. A spreading-activation theory of lemma retrieval in speaking. Cognition 42 (1–3), 107–142.
Roelofs, A., 1993. Testing a non-decompositional theory of lemma retrieval in speaking: retrieval of verbs. Cognition 47 (1), 59–87.
Schwartz, M.F., Kimberg, D.Y., Walker, G.M., Brecher, A., Faseyitan, O.K., Dell, G.S., Coslett, H.B., 2011. Neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. Proc. Natl. Acad. Sci. USA 108 (20), 8520–8524.
Segalowitz, S.J., Zheng, X., 2009. An ERP study of category priming: evidence of early lexical semantic access. Biol. Psychol. 80 (1), 122–129.
Seghier, M.L., 2013. The angular gyrus multiple functions and multiple subdivisions. Neuroscientist 19 (1), 43–61.
Shapiro, L.P., Levine, B.A., 1990. Verb processing during sentence comprehension in aphasia. Brain Lang. 38 (1), 21–47.
Shetreet, E., Palti, D., Friedmann, N., Hadar, U., 2007. Cortical representation of verb processing in sentence comprehension: number of complements, subcategorization, and thematic frames. Cereb. Cortex 17 (8), 1958–1969.
Simon, D.A., Lewis, G., Marantz, A., 2012. Disambiguating form and lexical frequency effects in MEG responses using homonyms. Lang. Cogn. Process. 27 (2), 275–287.
Smiley, S.S., Brown, A.L., 1979. Conceptual preference for thematic or taxonomic relations: a nonmonotonic age trend from preschool to old age. J. Exp. Child Psychol. 28 (2), 249–257.
Solomyak, O., Marantz, A., 2010. Evidence for early morphological decomposition in visual word recognition. J. Cogn. Neurosci. 22 (9), 2042–2057.
Talmy, L., 1975. Semantics and syntax of motion. Syntax Semant. 4, 181–238.
Thompson, C.K., Melzer-Asscher, A., 2014. Neurocognitive mechanisms of verb argument structure processing. In: Bachrach, A., Roy, I., Stockall, L. Structuring the argument. John Benjamins, (141–168).
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.
Tyler, L.K., Randall, B., Stamatakis, E.A., 2008. Cortical differentiation for nouns and verbs depends on grammatical markers. J. Cognit. Neurosci. 20 (8), 1381–1389.
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.
Van Essen, D.C., 2005. A population-average, landmark-and surface-based (PALS) atlas of human cerebral cortex. Neuroimage 28 (3), 635–662.
Wang, J., Cherkassky, V.L., Yang, Y., Chang, K.K., Vargas, R., Diana, N., Just, M.A., 2016. Identifying thematic roles from neural representations measured by functional magnetic resonance imaging. Cogn. Neuropsychol. 33 (3–4), 257–264.
Westerlund, M., Pylkkänen, L., 2014. The role of the left anterior temporal lobe in semantic composition vs. semantic memory. Neuropsychologia 57, 59–70.
Westerlund, M., Kastner, I., Al Kaabi, M., Pylkkänen, L., 2015. The LATL as locus of composition: meg evidence from English and Arabic. Brain Lang. 141, 124–134.
Wu, D.H., Waller, S., Chatterjee, A., 2007. The functional neuroanatomy of thematic role and locative relational knowledge. J. Cogn. Neurosci. 19 (9), 1542–1555.
Zhang, L., Pylkkänen, L., 2015. The interplay of composition and concept specificity in the left anterior temporal lobe: an MEG study. NeuroImage 111, 228–240.
Ziegler, J., Pylkkänen, L., 2016. Scalar adjectives and the temporal unfolding of semantic composition: an MEG investigation. Neuropsychologia 89, 161–171.
de Zubicaray, G.I., Hansen, S., McMahon, K.L., 2013. Differential processing of thematic and categorical conceptual relations in spoken word production. J. Exp. Psychol.: Gen. 142 (1), 131.