

Relational vs. attributive interpretation of nominal compounds differentially engages
angular gyrus and anterior temporal lobe

Christine Boylan, John C. Trueswell, and Sharon L. Thompson-Schill

Department of Psychology, University of Pennsylvania

Corresponding author: Christine Boylan
Present Address: Institute for Research in Cognitive Science
University of Pennsylvania
3401 Walnut Street, Suite 400A
Philadelphia, PA 19104, USA
Phone: (001) 630-346-2989
Fax: (001) 215-573-9247
Email: chboylan@sas.upenn.edu

ABSTRACT

The angular gyrus (AG) and anterior temporal lobe (ATL) have been found to respond to a number of tasks involving combinatorial processing. In this study, we investigate the conceptual combination of nominal compounds, and ask whether ATL/AG activity is modulated by the *type* of combinatorial operation applied to a nominal compound. We compare relational and attributive interpretations of nominal compounds and find that ATL and AG both discriminate these two types, but in distinct ways. While right AG demonstrated greater positive task-responsive activity for relational compounds, there was a greater negative deflection in the BOLD response in left AG for relational compounds. In left ATL, we found an earlier peak in subjects' BOLD response curves for attributive interpretations. In other words, we observed dissociations in both AG and ATL between relational and attributive nominal compounds, with regard to magnitude in the former and to timing in the latter. These findings expand on prior studies that posit roles for both AG and ATL in conceptual processing generally, and in conceptual combination specifically, by indicating possible functional specializations of these two regions within a larger conceptual knowledge network.

1. INTRODUCTION

Language's infinite generative capacity allows us to produce utterances ranging from the prosaic, as in "Close the door," to the ridiculous, as in "Hold the newsreader's nose squarely, waiter, or friendly milk will countermand my trousers" (Stephen Fry, *A Bit of Fry and Laurie*). Less ridiculous, but no less novel, sentences are uttered every day, and the ability of a reader or listener to understand such novel sentences, the propositional meanings of which cannot be retrieved from memory, requires a compositional algorithm that takes word meanings and combines them in such a way as to produce a more complex meaning. The neural substrates of this compositional algorithm remain elusive, in part because we still lack consensus on a typology of compositional operations.

Many approaches to the study of composition have benefitted from a clear distinction between syntactic and semantic composition (Kuperberg, Sitnikova, & Lakshmanan, 2008; Pyllkänen & McElree, 2006). Studies on Jabberwocky sentences demonstrate that subjects parse Jabberwocky phrases into syntactic hierarchical constituents similar to their natural language counterparts, even without knowing what the phrase means (e.g. "the mouse that eats our cheese" vs. the Jabberwocky version "the couse that rits our treeve" (example from Pallier, Devauchelle, & Dehaene, 2011)). Studies of complement coercion also suggest that syntactic and semantic argument structures are not isomorphic: evidence from behavioral, eye-tracking, and electrophysiological measures demonstrate a processing cost where semantic material unexpressed in the syntax must be inserted in order to coerce a coherent argument structure; e.g. "The man began the book" is interpreted as "The man began [reading/writing] the book" via implicit insertion of some event information (Baggio, Choma, van Lambalgen, & Hagoort, 2010; Kuperberg, Choi,

Cohn, Paczynski, & Jackendoff, 2009; McElree, Pyllkkänen, Pickering, & Traxler, 2006; Pyllkkänen & McElree, 2007; Traxler, McElree, Williams, & Pickering, 2005).

In this study, we proceed one step further, and suggest that within the domain of semantic composition, there is evidence for a distinction between two basic combinatorial operations, *even when syntax is held constant*. Specifically, we investigate the case of noun-noun compounds, in which the structure is always a modifier noun followed by a head noun (e.g. *mountain lake*, where the syntax dictates this is a lake (in the mountains), not a mountain (in a lake)). Noun-noun compounds are a tractable case of minimal composition and are particularly interesting because the first noun – the modifier noun – can be either “attributive” (as in *zebra clam*, where *zebra* denotes the attribute “striped”) or “relational” (as in *mountain lake* where “mountain” is not an attribute but an object bearing a spatial relation with “lake”). Attributive combinations are similar to predicating combinations, which can be paraphrased as “a [noun] that is [adjective],” such as *red ball* (“a ball that is red”). Non-predicating combinations, on the other hand, cannot be paraphrased in such a way: e.g. *tennis ball* is not “a ball that is tennis,” but rather is “a ball for playing tennis” (Downing, 1977; Gagné & Shoben, 1997; Levi, 1978). Attributive noun-noun compounds can be paraphrased as “a [head noun] that is [modifier noun]-like”, as in *zebra clam* – “a clam that is zebra-like” (likely “a clam that is striped”). Relational noun-noun compounds are more complex in that they are non-predicating, and derive their meaning from some extrinsic predicating relation (e.g. “a ball *for playing* tennis”) (Levi, 1978; Murphy, 1990).

We find evidence that relational and attributive interpretations of noun compounds differentially engage two regions of the brain otherwise broadly implicated in semantic

composition: the angular gyrus (AG) and the anterior temporal lobe (ATL). Below, we discuss how the distinction between relational and attributive combination may shed light on the functional differences between these two putative neural “hubs” of semantic composition.

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

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

high intra-temporal connectivity (Binney et al., 2012). Such neuroanatomical sequestration may be conducive to representing abstract, modality-invariant semantics. Thus, ATL is a prime candidate for attributive semantic composition.

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

In addition to the ATL, researchers have also ascribed the label “semantic hub” to the AG, as it lies at the junction between temporal, parietal, and occipital lobes and thus receives a confluence of auditory, somatosensory, spatial, and visual inputs. Conceptual combination studies of the sort described above (Bemis & Pylkkänen, 2012) have

demonstrated involvement of both left AG and left ATL, and several studies implicate bilateral AG in the contrast between well-formed sentences on the one hand and word lists, pseudowords, or scrambled sentences on the other (Bavelier et al., 1997; Bottini et al., 1994; Humphries, Binder, Medler, & Liebenthal, 2007; Humphries, Binder, Medler, Liebenthal, & others, 2006). AG also shows greater activity for semantic violations vs. congruent well-formed sentences (Kang, Constable, Gore, & Avrutin, 1999), particularly semantic incongruities violating verb-argument structure and thematic constraints rather than real-world knowledge (Kuperberg et al., 2008; Newman, Pancheva, Ozawa, Neville, & Ullman, 2001), and also for connected discourse vs. unrelated sentences (Fletcher et al., 1995; Homae, Yahata, & Sakai, 2003; Xu, Kemeny, Park, Frattali, & Braun, 2005). This broad profile of effects has led some to suggest that the AG may play a potentially domain-general role in semantic information integration structured around events.

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

picking out an attribute of the modifier noun. It is likely that these stimuli were probing semantic thematic relations in particular rather than combinatorial semantics in general.

1.2 Relational vs. attributive interpretation of nominal compounds

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

It is worth noting here that the terminology "relational" and "attributive" bear the misfortune of being both very common and denoting very different qualities depending on the theoretical framework. For instance, there are certain non-deverbal "relational" nouns, like *sister*, *boss*, *edge*, *height*, etc., that seem to take implicit semantic arguments (Partee & Borschev, 2003), but we do not consider such relational nouns here. However, "relational" and "attributive" eschew the perhaps overly syntactic typology of, say, "predicating vs. non-predicating," in describing a phenomenon we argue arises from conceptual-semantic affordances; that is, *robin hawk* is interpretable as "a hawk that

hunts robins” only because of the association between *hawk* and *hunting*, even though the “hunting” relation is unstated (cf. *robin chicken*, which would be unlikely to invite such a predator-prey relation). Therefore, we will refer to these nominal compound types as “relational” and “attributive.”

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

combined (e.g. LOCATED is a more apt relation than ABOUT when construing the compound *mountain lake*). This model stipulates that the RESEMBLE relation (that is, the attributive interpretation) is largely dispreferred.

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

would interfere with comprehension of a relational target. Estes (2003) did not find such an interaction, and interpreted these findings as inconsistent with CARIN. However, this null result is not particularly damning to the CARIN model, and evidence arbitrating whether attributive and relational compound interpretation involve distinct processes remains equivocal.

1.3 Relational and attributive semantics in the brain

The distinction between property- and relation-based semantic processes also appears in the neuropsychological semantic memory literature. Here, one abiding question has been: Do relational and attributive conceptual combinations arise from neuroanatomically separable components of the semantics, or are they subsumed by the same combinatorial operation? If the latter, are these operations hierarchically disposed in some way: that is, are relational interpretations preferred over attributive interpretations, or vice-versa?

In the semantic memory literature, the distinction between so-called taxonomic and thematic associations serve as a parallel to what we have described as attributive and relational associations, respectively. The literature on taxonomic and thematic associations does not relate explicitly to nominal compound interpretation, or even conceptual combination *per se*, and the experimental paradigms focus rather on speech errors, priming, and judgment of semantic association between words that are not part of a compositional phrase. However, we argue that the semantic ontology of “properties” vs. “relations” underlying the distinction between taxonomic and thematic associations

also defines the distinction between attributive and relational interpretations of compounds, and thus both literatures should be explored in tandem.

Just as the conceptual combination literature includes theories in which attributive and relational combination are part of the same process and not categorically distinct (as in CARIN), so does the semantic memory literature accommodate models whereby taxonomic and thematic relations are not or need not be distinguished. For instance, under many connectionist frameworks, thematic relations are reified into features, such that the concept *dog* might be linked not only to nodes for “furry,” “warm-blooded,” and “loyal,” but also nodes for explicitly relation-based facts, such as “bears live young,” “is led on a leash,” and “chews/buries bones” (Rogers & McClelland, 2004). Another account holds that taxonomic properties are the constituents of thematic relations: for instance, the thematic relation *chew*, between, say, *dog* and *bone*, is supervenient on knowledge of the properties of *dog* and *bone*, so that thematic processing cannot occur without taxonomic processing (Lewis, Poeppel, & Murphy, 2015). This account is supported by an MEG priming study finding that activity in left ATL was only sensitive to property-based taxonomic associations, while both taxonomic and thematic associations predicted activity in the left temporoparietal junction (TPJ), inclusive of AG (Lewis et al., 2015). Both these accounts contrast with a model whereby thematic relations constitute a qualitatively different level of representation from taxonomic features or properties, where the unit of representation is the event rather than the feature (Boylan, Trueswell, & Thompson-Schill, 2015).

Another study, however, found evidence of a double dissociation between taxonomic and thematic semantic errors in left ATL and left AG, respectively (Schwartz

et al., 2011). Speakers' semantic errors can be divided into either taxonomic category errors (that is, uttering an incorrect word, but one which has commensurate features, such as when "apple" is named as "pear") or thematic relation errors (that is, uttering "dog" when "bone" was intended, reflecting the thematic relation between "dog" and "bone"). Schwartz et al. (2011) examined the taxonomic and thematic errors produced by 86 post-stroke aphasics in a picture-naming task and conducted voxel-based lesion-symptom mapping (VLSM) on each error type separately (with shared variance between error types regressed out). Taxonomic errors were mapped to left ATL lesions, while thematic errors were localized to left AG. This double dissociation between ATL and AG supports the view that the ATL and AG support categorically distinct semantic computations, corresponding to property-based and relation-based operations, respectively.

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

We find evidence for (1) a single dissociation in bilateral AG showing more task-responsive activity for relational compounds than attributive compounds, and (2) a timing difference in ATL, specifically an earlier ATL response to attributive compounds than

relational compounds. This order is not consistent with that put forth in the CARIN model, which predicts that relational combinations are processed prior to attributive combinations, but is consistent with a model where the distinction between attributive and relational processing could be more fluid in ATL.

2. MATERIAL AND METHODS

2.1 Participants

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

2.2 Stimuli

2.2.1 Stimulus Design

We drew our nominal compound stimuli from two studies investigating the effects of attributive- and relational-biased compounds (Estes, 2003; Wisniewski & Love, 1998). Of the stimuli used in these studies, we chose the 64 most attributive-biased and 64 most relational biased items according to a norming study we conducted via Amazon Mechanical Turk (Buhrmester, Kwang, & Gosling, 2011). Subjects (n=17, none of whom participated in the fMRI study) were asked to write their interpretations of each noun compound and indicate their familiarity with the noun compound on a 1-5 Likert scale, and three independent coders designated these interpretations as either attributive or relational. Coder agreement was over 95%, and where coders' designations diverged, the

primary author's designation was used (Boylan). Criteria for qualitative definitions of relational and attributive compounds were taken from Estes (2003) and Wisniewski & Love (1998). Relational and attributive items were matched on unigram (log) frequency ($M_{\text{attributive}} = 1.07$, s.d. = 0.57; $M_{\text{relational}} = 1.15$, s.d. = 0.60; SUBTLEXus; bigram frequencies were not possible because many compounds were infrequent or novel), compound length ($M_{\text{attributive}} = 11.98$, s.d. = 2.46; $M_{\text{relational}} = 12.13$, s.d. = 2.84), compound familiarity (Likert scale 1-5, from 17-subject norming study on Amazon Mechanical Turk: $M_{\text{attributive}} = 2.29$, s.d. = 0.51; $M_{\text{relational}} = 2.40$, s.d. = 0.49), and word association (latent semantic analysis (LSA) over general reading at 1st year college level; $M_{\text{attributive}} = 0.12$, s.d. = 0.15; $M_{\text{relational}} = 0.14$, s.d. = 0.12) (Brysbaert & New, 2009; Laham, 1998). Compounds ranged from relatively unfamiliar or novel (least familiar attributive compound: *tiger paper*; least familiar relational compound: *pine mushroom*) to relatively common (most familiar attributive compound: *bullet train*; most familiar relational compound: *wood stove*).

We found that the attributive modifier nouns (first nouns in the two-word compounds) were less frequent than the head nouns ($M_{\text{attributive_modifier}} = 0.82$, s.d. = 0.45; $M_{\text{attributive_head}} = 1.32$, s.d. = 0.58) while the converse was true for relational compounds ($M_{\text{relational_modifier}} = 1.24$, s.d. = 0.61; $M_{\text{relational_head}} = 1.08$, s.d. = 0.58). Ramifications for the attributive compounds' "modifier markedness" are addressed in the Discussion section below.

2.2.2 Norming for relational-attributive bias

In order to measure the variability in the interpretations of our noun-noun compound stimuli, and thus the degree of bias toward attributive and relational

combination, we combined responses from the Mechanical Turk stimulus norming survey with responses taken from a survey of our fMRI subjects after they left the scanner. Both surveys asked subjects to describe what they thought each noun compound meant, with the slightly amended instruction to the fMRI participants that they write down the interpretations they had entertained while viewing the stimuli inside the scanner. Responses from a total of 35 subjects per item (17 from Mechanical Turk norming, 18 from fMRI subjects) were coded as either attributive or relational (see 2.2.1 for coding practices).

Figure 1 shows the distribution of attributive and relational bias by item, order-ranked from unanimously attributive interpretations to unanimously relational interpretations. The average “relational bias” for an item categorically labeled as relational in the Estes and Wisniewski & Love studies was 89.6% (SD = 0.14); likewise, the average “attributive bias” for an attributive item was 94.5%, (SD = 0.10). Despite having identified 128 items from the Mechanical Turk survey as relatively biased towards either relational or attributive meanings, additional responses from fMRI subjects confirmed that some items were much more ambiguous than others. For instance, while the item “cow parsnip” was originally designated as relational-biased (example interpretation: “a parsnip fed to a cow”), as was “pine mushroom” (e.g. “mushroom that grows on pine trees”), additional responses indicated these items were equally likely to have attributive readings, where interpretations such as “a parsnip shaped like a cow,” and “mushroom that looks like a pine cone” were offered for “cow parsnip” and “pine mushroom,” respectively.

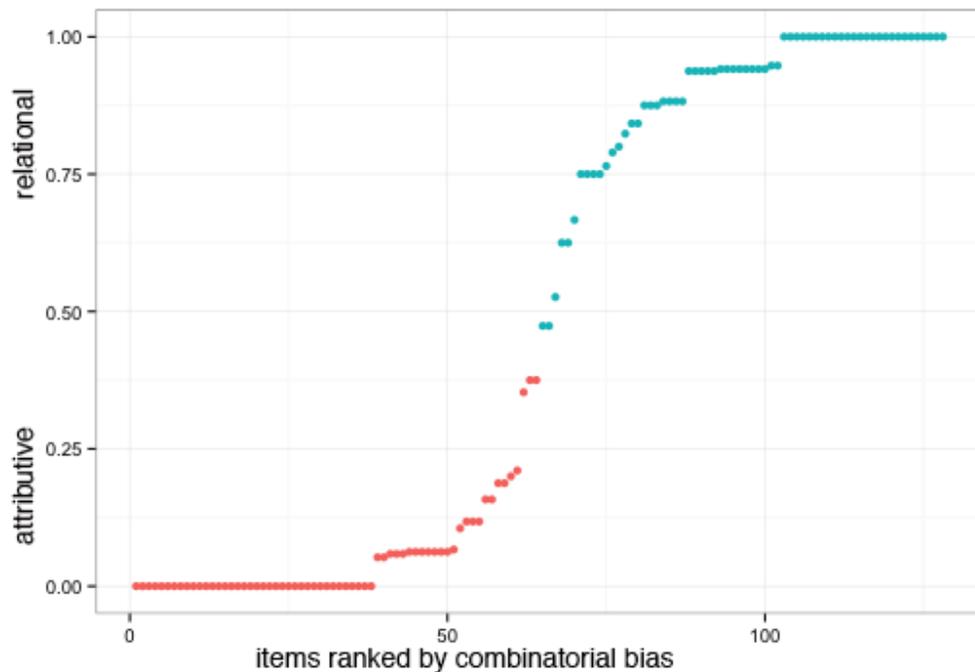


Figure 1. Plot of relational-attributive combinatorial bias (normed on 35 participants, inclusive of 18 fMRI subjects). Figure 1 shows the distribution of attributive and relational combinatorial bias by item, order-ranked from unanimously attributive (relational bias = 0) interpretations to unanimously relational (relational bias = 1) interpretations. Blue indicates those items marked as relational in Estes (2003) and Wisniewski & Love (1998) studies and red indicates those marked as attributive in those studies. This aggregate combinatorial bias measure was included as a continuous covariate in one model, while the individual subject response (attributive/relational) was treated as a categorical variable in a model of individual BOLD timecourses.

2.3 Experimental Task and Design

The subject's task on each trial was to read two simultaneously centrally presented words constituting a nominal compound and indicate by button press (1) when they had decided on a coherent meaning for the compound and (2) whether a subsequent "probe" matched

the meaning they had in mind. The nominal compound was presented for 2 seconds, and was immediately followed by a fixation cross, on screen for 6 or 8 seconds, during which time the subject need only passively view the screen after pressing the button to indicate they had arrived at an interpretation. This fixation period was followed by a probe phrase (2 seconds), which either matched the modal interpretation for a given item (as determined in the Mechanical Turk survey; e.g. “a prickly carpet” for “cactus carpet”) or was a dispreferred interpretation (e.g. “a carpet on which a cactus stands” for “cactus carpet”), where one out of eight trials had a dispreferred probe. The probe phrase served as a prompt for the subject to indicate by button press whether the probe matched their initial interpretation. We included the probe (not split by attributive/relational) as a separate event in our model, and imposed a relatively long (8-10 seconds) and jittered separation between compound and probe, in order to avoid probe processing effects from bleeding into the compound interpretation event. This was followed by a 6-8-second fixation-cross ITI. The ordering of the attributive and relational events was scheduled using optseq2 (<http://surfer.nmr.mgh.harvard.edu/optseq>). The entire experiment consisted of 8 runs of 16 trials each. After leaving the scanner, subjects were given a questionnaire which asked them to write down what they had thought each nominal compound meant when they had viewed them in the scanner.

2.4 Image acquisition

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

anatomic localizer; (2) a high-resolution, T1-weighted sequence for localization of fMRI activity in standard stereotactic space; (3) T2*-weighted images from 8 experimental runs; (4) a B₀ field map sequence for subsequent geometric unwarping of T2*-weighted images.

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

2.5 Analysis

2.5.1 Image analysis and ROIs

FMRI data were pre-processed offline using the AFNI (Cox & Jesmanowicz, 1999) software package. The first four volumes of each functional run were removed so as to allow the signal to reach steady-state magnetization. Functional images were slice-time corrected, and a motion correction algorithm employed in AFNI registered all volumes to a mean functional volume. Images were then unwrapped via B₀ field maps (using FSL software; <http://www.fmrib.ox.ac.uk/fsl>) to reduce non-linear magnetic field distortions. We applied a high-pass filter of 0.01 Hz on each run to remove low frequency trends. Functional data were registered to the individual subject's anatomical MRI. Transient spikes in the signal were removed using AFNI's 3dDespike.

Our *a priori* ROIs were left and right anterior temporal lobes and left and right angular gyri, which we delimited using AFNI's CA_ML_18_MNIA atlas. Our anterior temporal ROIs spanned labels “left/right temporal pole” and “left/right medial temporal

pole”, while our angular gyrus ROIs circumscribed only the atlas’s “left/right angular gyrus” ROI (see Figure 2).

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

To identify task-activated voxels for inclusion in further analysis, we first conducted a GLM with FIR regressors for task and fixation ITI. The task TENT series was time-locked to the onset of the nominal compound, and the ITI TENT began at the onset of the fixation ITI (10-12 seconds after the onset of the nominal compound). To investigate effects of relational-attributive combinatorial bias in each ROI, we used a model with covariates for task and jittered probe event, where task TENTs were again time-locked to the nominal compound presentation and probe TENTs synced to the onset of the probe question (8-10 seconds post compound onset), along with a continuous covariate for the relational-attributive bias (see Figure 1) of the noun compound in each trial. Head movement and global signal were included as covariates of no interest in both models.

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

3. RESULTS

3.1 Task-responsive voxels in anatomical ROIs

In a group-level contrast targeting bilateral AG and bilateral ATL, we found several clusters of voxels with a reliable ($p < 0.01$, uncorrected) activation difference between task (comprehension of nominal compound) and ITI fixation baseline, where the task condition collapsed attributive and relational trials together (Figure 2). Two distinct clusters of activity were revealed in left AG, where one cluster was positively activated for task relative to baseline (cluster centroid: [-45 -54 38] Talairach coordinates) and another more posterior, inferior cluster was more active during baseline relative to task (cluster centroid: [-39 -67 35] Talairach coordinates). This motivated us to treat positively and negatively activated task-responsive voxels as distinct functional subregions within the anatomical left AG ROI. Other clusters in right AG and left ATL

were largely positively task-responsive, and so did not prompt any functional division between above- and below-fixation task activation (see outlined ROIs in Figure 2). No significant task-responsive clusters survived even a liberal threshold in right ATL, and so this anatomical region was not analyzed further.

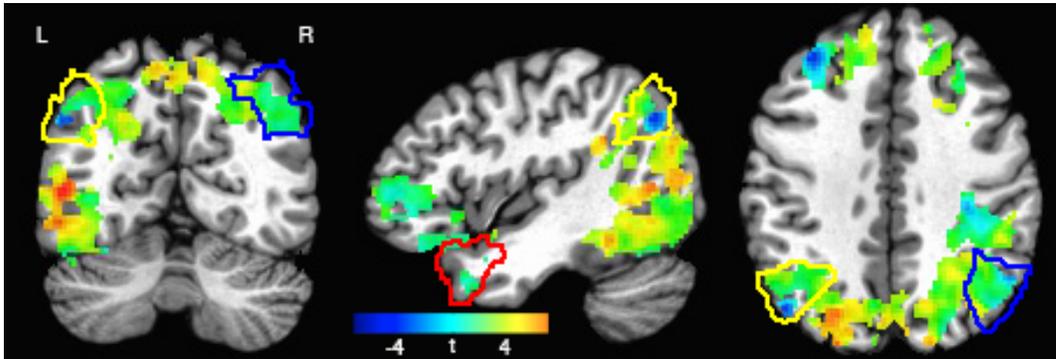


Figure 2: Whole-brain task-responsive voxel activity (task vs. ITI fixation baseline). ROIs circumscribed: left ATL (red), left AG (yellow), and right AG (blue) at Talairach coordinates [-45 -57 38].

3.2 Combinatorial bias predicts activity in left and right AG

3.2.1 Categorical analysis

For each subject, we identified the 50 most positive task-responsive voxels in each of our three ROIs: right AG, left AG, and left ATL. We also identified the 50 most negative task-responsive voxels in left AG, for a total of four functional ROIs. The location and distribution of these top 50 task-responsive voxels varied greatly across subjects for each ROI, precluding clear anatomical delineation of positive vs. negative task-responsive regions within left AG for a given subject, even though group-level clusters suggest an anatomical divide along PGa/PGp (Noonan, Jefferies, Visser, &

Lambon Ralph, 2013; see Discussion). Based on subjects' post-scanning survey responses, we were able to code each trial as attributive or relational at the individual subject level. Using this categorical coding scheme for each fMRI subject, we found no significant difference between attributive and relational conditions in positive task-responsive left ATL voxels ($t(34) = 0.26$, $p = 0.80$) or positive task-responsive left AG voxels ($t(34) = 0.18$, $p = 0.86$). We found significantly more positive activity for the relational condition relative to the attributive in right AG positive task-responsive voxels ($t(34) = 2.89$, $p = 0.0067$). Looking at negative task-responsive activity in left AG, we found significantly greater negative activation for the relational condition relative to the attributive (i.e. more positive activity in the attributive condition; $t(34) = 2.12$, $p = 0.042$).

3.2.2 Continuous analysis

Because we found a gradient of bias in our compound norming study, ranging from highly relational (all subjects indicating a relational interpretation) to ambiguous (some subjects indicating relational interpretations, others attributive interpretations) to highly attributive, we also examined the effect of relational-attributive bias (referred to henceforth as simply “combinatorial bias,” see Fig. 1) on BOLD signal amplitude. This combinatorial bias measure was the average relational/attributive rating across both the norming study responses and the fMRI subjects' post-scanning survey for a combined total of 35 ratings per item. This analysis was conducted within each of the same four 50-voxel ROIs (positive task-responsive voxels in left ATL, right ATL, left AG, and also negative task-responsive voxels in left AG) in each subject.

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

It is interesting to note that while we expected to find evidence of increased competition for the more ambiguous compounds (that is, the compounds near 0.5 bias in Figure 1), we failed to find differential activation related to increased ambiguity in either our ATL and AG ROIs or in a targeted analysis of left inferior frontal gyrus (left IFG). Neither did we see an effect of either attributive/relational interpretation or degree of ambiguity first button-press reaction time (the participant's button press indicating convergence on an interpretation before being shown the probe). Thus we do not find evidence for an effect of ambiguity *per se*, but rather the degree of relational or attributive bias along the combinatorial bias gradient.

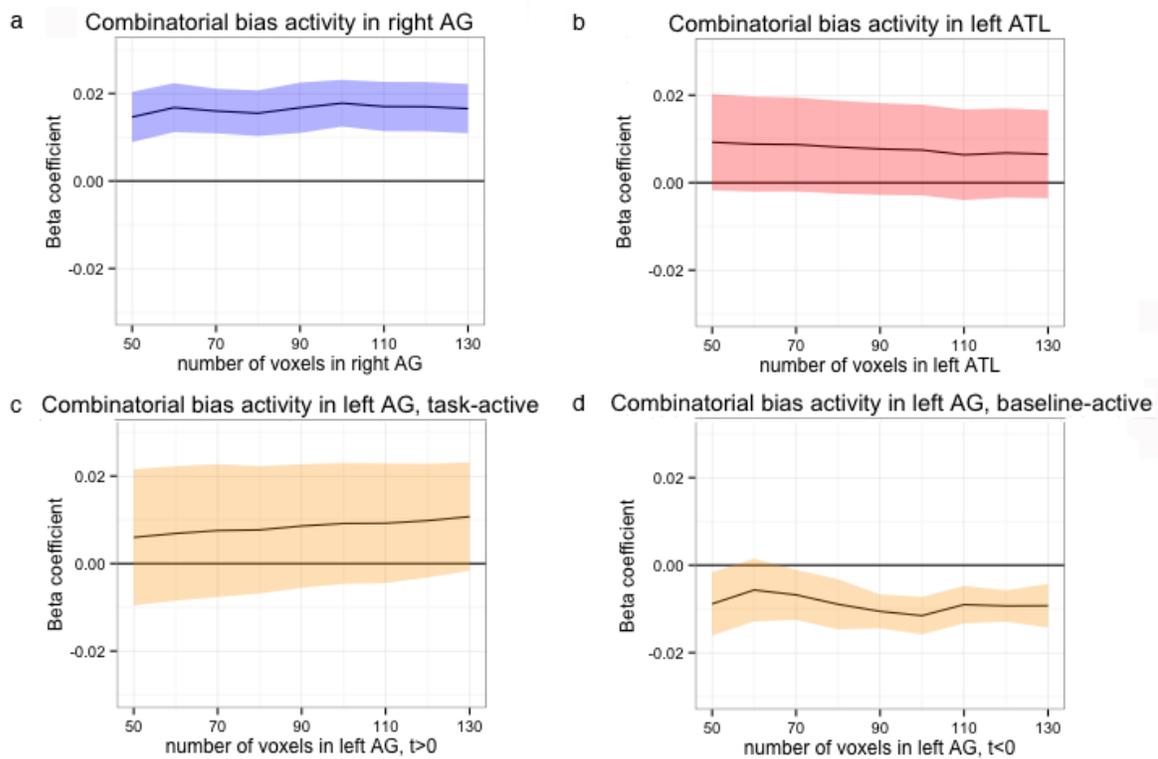


Figure 3. Beta coefficients from subject-wise parametric analysis of voxel activity by combinatorial bias (see Fig. 1) across a range of voxels within (a) right AG, (b) left ATL, (c) left AG, positively task-responsive voxels only, and (d) left AG, negatively task-responsive voxels only. Voxels chosen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

3.3 BOLD timecourse of attributive interpretation differs by ROI

Finding evidence that combinatorial bias predicts the magnitude of the response in both left and right AG, we then asked whether relational and attributive combination effects might also show distinct BOLD timecourses across ROIs. In particular, we were interested in whether the timecourse of BOLD activity in ATL would reveal a distinction

between attributive and relational processing that was not apparent in the previous analysis.

Using the same task-responsive voxel selection criteria as above, we compared the timecourse of BOLD activity for relational and attributive combinations in the 50 most task-responsive voxels in right AG, left AG (negatively task-responsive voxels only), and left ATL. In this analysis, we treated relational and attributive combination categorically, coding each trial condition based on individual subjects' responses in a post-scan survey.

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

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

distinction between attributive and relational combination in addition to activity in bilateral AG.

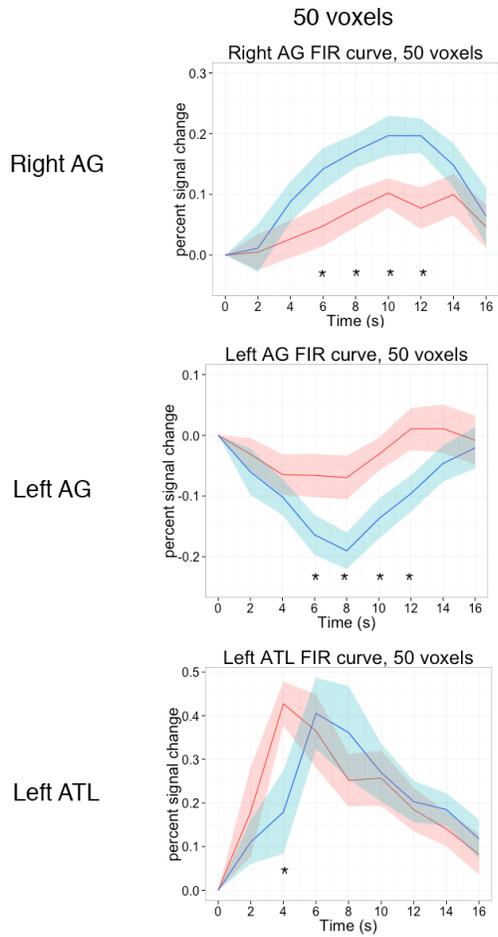


Figure 4. Peristimulus FIR curves of attributive (red) vs. relational (blue) activation in each ROI. FIR curves shown are from the 50 most positively task-responsive voxels for the task-vs.-fixation contrast in each anatomical ROI, except for left AG, where only the 50 most negatively task-responsive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

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

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

IV. DISCUSSION

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

These combined findings support an account whereby relational and attributive operations are not dissociable by a coarse neuroanatomical divide, but rather are encoded differently in different regions. The multiple, potentially redundant, instantiations of a

combinatorial relational-attributive code across the brain provide compelling evidence that the relational-attributive distinction is a productive one. While these data suggest that bilateral AG is more engaged in computing relational combination, comparison with the left ATL profile does not constitute a true double dissociation between relational and attributive processes. While we do not find evidence to support a full dual-route, parallel process model of relational and attributive processing, we find that these combinatorial operations may be instantiated in overlapping networks across ATL and AG. The time course of attributive and relational BOLD response in the left ATL suggests that attributive processing may be facilitated relative to relational processing, rather than attributive “relations” being the interpretations of last resort (as predicted by the CARIN model). This is broadly consistent with certain claims that attributive feature extraction and re-combination must occur before relational information can be processed. The earlier ATL response peak for attributive compounds may also be facilitated by what we call “modifier markedness”: the fact that these attributive compounds contain modifier nouns that are less frequent than the subsequent head noun, and perhaps therefore more salient.

4.1 Angular gyrus and thematic relations

4.1.1 AG sensitivity to thematic relations as verbs

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

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

While left AG has been implicated in the detection of syntactic errors (Embick, Marantz, Miyashita, O’Neil, & Sakai, 2000; G. R. Kuperberg et al., 2003), it is also involved in the detection of semantic incongruities, particularly when the semantics of an argument violates the semantic constraints of a verb (Kang et al., 1999; Kuperberg et al., 2008; Newman et al., 2001). This suggests that the sensitivity of AG to thematic roles and verb structure is not limited to the syntactic composition alone, but also to the semantic content (Pallier et al., 2011). It may be that left AG acts as an interface between semantic memory and syntactic structure, mapping semantic-thematic relations onto structural constraints surrounding verbs and their arguments. Indeed, electrophysiological and neuroimaging studies support an overlap between (morpho-)syntactic and semantic-thematic verb violations. Kuperberg et al. (2008) compared three different types of verb violations: (1) semantic–thematically violated verbs (e.g. “at breakfast the eggs would eat”) (2) morphosyntactically violated verbs (e.g. “at breakfast

the boys would eats”) and (3) real-world violations (e.g. “at breakfast the boys would plant”). They found that, unlike real-world violations, both semantic-thematic and morpho-syntactic violations elicited activity in a frontal/inferior parietal/basal ganglia network. This was consistent with previous electrophysiological findings showing that semantic-thematic and syntactic violations evoked a P600 event-related potential (ERP) that is distinct from the N400 produced by real-world incongruities (Hoeks, Stowe, & Doedens, 2004; Kuperberg, 2007; Kuperberg et al., 2003). Kuperberg et al. (2008) suggested that this frontal/AG/basal ganglia activity, as well as the P600 activity reported in previous studies to the same types of semantic violations, reflected prolonged attempts to combine the structural constraints of the verb with semantic properties of the Agent NP argument. This description is also consistent with previous hypotheses suggesting such activity might underlie an attempt to relate the goal-directed action denoted by the verb with properties of the Agent (Buccino et al., 2001; Chao & Martin, 2000; Damasio et al., 2001; Fogassi et al., 2005).

4.1.2 Lateralized effects of thematic relations in AG

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

implicit verb/event; in the “dog-bone” case, “chews” or “buries,” etc. That is, thematic knowledge is precisely knowledge of verbs and their arguments: that is, in conceiving of the relation, one is accessing the very semantic information that identifies the verb, or that the verb denotes. We speculate that the verb may be the minimal linguistic expression of the fundamental thematic and event-based concepts that AG subserves.

It may also be that the degree to which a thematic relation is grammaticalized – for instance, whether it is spelled out in an explicit verb – accounts for the subtly different profiles of activity between right and left AG, and even within left AG itself. Graves and colleagues' (2010) study of nominal compounds vs. their non-attested reversals found BOLD activity in right AG, but not left, increased for familiar nominal compounds like *lake house* (as compared with *house lake*). The authors offer an extended version of Beeman and colleagues' (1994) right hemisphere “coarse semantic coding hypothesis” to account for how noun-noun compounds might engage right AG but not left AG. They describe the dynamics of left and right AG by way of attractor networks, where such a network settles into an attractor basin when it optimizes the error space in the mapping between inputs (words) and outputs (meanings). 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. With less discrete minima, right AG would support more extensive overlap in meanings. Graves and colleagues propose that this increased “overlap” in right AG is what supports conceptual combination.

We suggest these attractor dynamics may explain a more subtle difference than that between the representation of words/concepts (left AG) and their combinatorial

instances (right AG). After all, if the distinction between left and right AG representations were a matter of “lexicalized meanings” vs. “compositional meanings”, it becomes difficult to defend that the Graves et al. compounds were the latter and not the former. Familiar relational stimuli in the Graves et al. study were items like “rock star” and “lake house”: arguably highly lexicalized, if not idiomatic. Instead, we suggest that both left and right AG support conceptual combination, but the greater “semantic reach” afforded by wide, shallow attractor basins in right AG may accommodate access to the extrinsic relation necessary for interpreting compounds like *dog bone* that lack the explicit (morpho)syntactic information about the relation (i.e. “a bone that a dog chews on”). Where explicit syntactic information detailing the agent, theme, and relation between dog and bone (“chews on”) is lacking, right AG is able to impute the extrinsic relation. While left AG relies on explicit morphosyntactic cues to arrive at highly constrained meanings, right AG can find implicit relational content in the process of determining a coherent meaning between two nouns. This would account for why verb-based minimal composition, where the thematic relation is grammaticalized and explicit, would be derived in the narrow attractor basins of left AG (Boylan et al., 2015), while the more fluid meanings afforded by nominal compounds are derived in the wider semantic net of right AG. That is, the left AG may be more engaged in syntactic-based relational combination, while the right AG mediates more semantic-based relational combination, particularly of novel nominal compounds as in our stimuli.

4.1.3 AG and the default mode network

The effect of combinatorial bias we found in right AG accords with the Graves et al. (2010) nominal compound effect in that region, and the distinction between relational

and attributive interpretations adds another dimension to the characterization of right AG as a combinatorial hub. However, our nominal compound effect in negatively task-activated voxels in left AG is rather more novel.

It has been noted that left AG is functionally heterogeneous, and recent work has begun to map this heterogeneity to subregions within left AG and surrounding areas. Of particular interest is the role of certain regions of the AG in the “default mode network” (DMN) (Raichle et al., 2001), a network of regions showing more positive activation in the absence of a task. Left AG’s involvement in the default network has invited the contention that differential activation in this region indexes task-unrelated difficulty effects independent of the content or representation of the stimuli. While certain regions of AG do indeed appear to index difficulty effects, the baseline-active (“deactivated”) regions of AG do not follow this pattern. 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 executive semantic control conditions, and was characterized as allocating attention to semantic representations in a task-dependent and goal-driven manner. However, dorsal AG’s role in semantic control was contrasted with mid AG, the activity of which was modulated by the semantic representational content of stimuli *even when matched on task demands*. Crucially, mid AG was associated with the DMN, showing more positive activation in the absence of a task. Noonan et al. (2013) note that

centers of activation in the putatively functionally distinct regions of dorsal AG and mid AG lie at Talairach coordinates [-41 -55 45] and [-39 -65 30], respectively. These align well with the two clusters of activation we observed for the task-vs.-baseline contrast in left AG: a more dorsal cluster of positive activation at [-45 -54 38] and a more posterior, inferior cluster of more negative activation at [-39 -67 35]. We find that the latter cluster aligns both anatomically and functionally with the left mid AG region of the DMN. Unlike the more dorsal cluster, the left mid AG cluster showed a combinatorial bias effect. Several studies also find stimulus-dependent activity unrelated to task difficulty in this region: some found this region showed more negative activity to abstract items than to concrete items (*prima facie* similar to our finding that relational compounds elicited more negative activity relative to attributive items in left mid AG) (Binder, Westbury, McKiernan, Possing, & Medler, 2005; Noonan et al., 2013; Wang, Conder, Blitzer, & Shinkareva, 2010). Other studies found this region responded more negatively to perceptual decisions than semantic decisions, while both stimuli elicited a negative divergence from baseline (as did both types of nominal compound stimuli relative to baseline in our study) (Seghier, Fagan, & Price, 2010). Outside the domain of language, Grabner, et al. (2013) found greater (more positive, less “deactivated”) left mid AG activity in more difficult arithmetic tasks, inconsistent with the model characterizing this region as reflecting task difficulty, and consistent with the hypothesis that this region supports symbol-referent mapping. Given that we did not find evidence that relational compounds were more difficult to interpret than attributive compounds, our differential activity in the left mid AG is unlikely to be an effect of task difficulty. Thus, our findings are consistent with a model whereby left AG task-negative activity reflects categorical

differences in the semantic content of the stimuli, rather than differences in executive semantic control required to complete the task.

4.1.4 A stimulus-specific negative BOLD response?

While we characterize the left mid AG and right AG effects as similarly indexing a greater magnitude of activation for relational interpretations, this is based on the absolute value of activation relative to baseline. It is entirely possible that the left mid AG effect should rather be characterized as an “attributive > relational” effect when considered as a local vascular and/or neural suppression. Proposed hemodynamic mechanisms for negative BOLD responses include “blood stealing” or “blood sharing,” whereby blood is diverted to active regions and away from the site of the negative BOLD response. However, these purely vascular explanations are not sufficient to predict the negative BOLD response profiles found in our study and several visual studies (Bressler, Spotswood, & Whitney, 2007; Shmuel et al., 2002; Smith, Williams, & Singh, 2004). Shmuel et al. (2006), while showing that negative BOLD activity is correlated with decreased neuronal activity measured via simultaneous electrical recordings, also demonstrated that the local decrease in neuronal activity predicted the spatiotemporal properties of the negative BOLD pattern better than surrounding positive BOLD activity did. This suggests that local neural suppression drives the negative BOLD response.

Moreover, recent studies suggest that such neural suppression may be more stimulus-specific than most attentional accounts predict. It is well known that attention directed to a given location corresponds with an increase in neural activity at that region (even in the absence of visual/auditory stimulation) and a decrease in neural activity at unattended regions. This predicts that attentional neural suppression will occur over

relatively broad swathes of (unattended) sensory space, however (Bressler et al., 2007). Recent studies show that negative BOLD responses to stimuli are spatially (retino- and tonotopically) constrained and highly tuned to certain properties of the stimulus (Bressler et al., 2007; Linke, Vicente-Grabovetsky, & Cusack, 2011). Thus, it could be that neural suppression itself carries stimulus-specific information. On the other hand, for a reduction in signal to carry such stimulus-specific information, it may be that the firing trace of center-surround responses in a sub-voxel population of neurons is dominated by the surround suppression, even though the most narrowly tuned neurons are being excited (Bressler et al., 2007; Linke et al., 2011; Müller & Kleinschmidt, 2004).

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

4.2 Anterior temporal lobe and semantic composition

While bilateral AG appears to index the relational-attributive distinction by way of different overall activation for the duration of the task, left ATL may derive both relational and attributive combinations, but with a temporal offset in activation peak for the relational combination as compared with the attributive. It is unclear whether the

relational meaning is delayed, perhaps due to the need to access an extrinsic relation linking the two nouns, or the attributive reading is facilitated. Attributive compound interpretation may be facilitated for a number of reasons. For one, it may be that the modifier noun (the first noun) is salient by way of being less frequent. We find that the attributive modifier nouns (first nouns in the two-word compounds) are less frequent than the head nouns, while the converse is true for relational compounds. This “modifier markedness” on the attributive compounds may consist not only in the frequency disparity between the modifier and head nouns within an attributive compound, but also in the fact that the modifier noun in attributive compounds is significantly less frequent than the modifier noun in relational compounds ($p < 0.001$). Modifier markedness may serve as a predictive cue in biasing towards attributive interpretation, and this early cue may account for the earlier peak in the left ATL BOLD response to attributive compounds. Indeed, several studies have found the left ATL to be the locus of semantically mediated predictive facilitation (Fruchter, Linzen, Westerlund, & Marantz, 2015; Lau, Weber, Gramfort, Hämäläinen, & Kuperberg, 2014). However, the facilitation we observe for the attributive “modifier markedness” is likely not semantic: attributive nouns were no more or less semantically associated than relational nouns (association measured by LSA). It is also unclear how “early” the frequency-based cue on the modifier is: the modifier is only infrequent *relative* to a downstream head noun, or to other modifier nouns in relational compounds. Methods more sensitive to the neural timecourse of nominal compound interpretation may elucidate which lexical and semantic factors contribute to facilitating attributive readings.

These findings nonetheless present an intriguing alternative to a strict dual-process hypothesis for attributive and relational combination. The temporal delay in computing relational combinations in left ATL is consistent with a theory whereby feature extraction and attribute-based composition occur prior to relation-based, thematic information, but are, at least at some level, derived by the same combinatorial process. In such a model, “properties” are predicates, and thematic relations can be reified into features of object concepts such that “walks on a leash” is as much a property of *dog* as simple adjectives like “furry” or “loyal.” Since language is flexible in this regard, able to express the verb-/event-based relational concept “bears live young” in the adjective “viviparous,” and to convert between verbs (*relations* between nouns) and participles (properties/*attributes* of nouns), etc., it is perhaps unsurprising that there be a means of converting relational and attributive concepts along a single dimension, where incorporation of an extrinsic relation is simply an added “feature”. While relational and attributive combinations may be qualitatively distinguished in AG, it might be that these differences are collapsed in left ATL. Indeed, in a study directly comparing how conceptual combination and object-concept specificity engage left ATL, Westerlund & Pykkänen (2014) concluded that combination and specificity effects in left ATL likely arise from a single feature-binding operation.

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

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

4.3 Limitations, future directions, and conclusions

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

concepts in wider, shallower attractor basins. While this model of attractor dynamics is an intriguing one, it remains to be directly tested in right and left AG. These results invite further study to pinpoint how right and left AG differ with regard to semantic combination.

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

Rather surprisingly, we did not find evidence of increased competition for the more ambiguous compounds in our ATL, AG, or IFG targeted ROIs. This negative finding invites careful comparison between what we call the “combinatorial bias” of a noun compound along a perhaps simplified, one-dimensional relational-attributive axis, and the more general question of ambiguity and reference resolution, particularly concerning the role of left IFG. While the current study had a range of combinatorial biases, further research is needed to target both relational/attributive ambiguity as well as multiple ambiguities at a finer granularity than attributive vs. relational (e.g. noun compounds with multiple possible attributive readings, not just an attributive vs. relational ambiguity). Our failure to find an effect of ambiguity is troublesome for the CARIN model in particular, which is about competition between different relations. The interaction between ambiguity and the relational-attributive distinction has not been explored, and it is possible that, contrary to CARIN and other models, there may even be

an “ambiguity advantage” effect. Analogous to the finding from lexical decision studies that words with multiple related meanings (polysemes) show a processing advantage where words with a few *unrelated* meanings (homonyms) do not show such an advantage, it is possible that nominal compounds with multiple related attributive meanings (or multiple related relational meanings) show a facilitation effect relative to compounds with either one meaning or a few unrelated meanings (Rodd, Gaskell, & Marslen-Wilson, 2004).

Finally, while our aim was to determine a productive typology of conceptual combination where attributive and relational interpretations are distinct, we did not delineate a typology at any finer grain than attributive vs. relational. Our post-scanning survey indicated that a given compound might have range of possible interpretations, sometimes constrained to relational meanings, sometimes to attributive, and sometimes ambiguous enough to invite both relational and attributive meanings. Yet the heterogeneity within the semantic space of attributive vs relational interpretations may differ, such that the space of possible relational meanings might be dilated relative to the space of attributive meanings, or vice versa. This problem of defining a productive typology of possible predicates and relations has been a pressing issue in the conceptual combination for the past 40 decades (some prominent models and approaches to be found in e.g. Downing, 1977; Gagné & Shoben, 1997; Levi, 1978; Murphy, 1990), but is not yet resolved. Further study is warranted, in which the stimuli are chosen to better sample the range and heterogeneity of attributive and relational items than we did here.

ACKNOWLEDGMENTS

This research was supported by National Institutes of Health Grants R01-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

item	type	compound
1	attributive	tiger paper
2	attributive	warehouse brain
3	attributive	spider flowers
4	attributive	gymnast squirrel
5	attributive	grease fish
6	attributive	octopus chair
7	attributive	skunk cigar
8	attributive	balloon pregnancy
9	attributive	blimp belly
10	attributive	catapult promotion
11	attributive	inferno radiator
12	attributive	medicine music
13	attributive	rocket sprinter
14	attributive	junkyard desk
15	attributive	razor insult
16	attributive	surgeon fish
17	attributive	cardinal flower
18	attributive	needle grass
19	attributive	ice marriage
20	attributive	volcano pimple
21	attributive	parasite citizen

22	attributive	mosquito fly
23	attributive	zebra clam
24	attributive	cactus carpet
25	attributive	turtle jogger
26	attributive	boomerang lie
27	attributive	megaphone talker
28	attributive	kangaroo rat
29	attributive	willow oak
30	attributive	pillow lips
31	attributive	piranha lawyer
32	attributive	snow goose
33	attributive	butter grip
34	attributive	leopard lizard
35	attributive	porcupine fish
36	attributive	molasses traffic
37	attributive	dinosaur computer
38	attributive	magnet smile
39	attributive	mule deer
40	attributive	barrel cactus
41	attributive	umbrella tree
42	attributive	gem idea
43	attributive	feather luggage
44	attributive	sponge memory

45	attributive	lemon paint
46	attributive	vampire insect
47	attributive	shark politician
48	attributive	sedative voice
49	attributive	mirror pond
50	attributive	funnel tornado
51	attributive	thunder applause
52	attributive	leech boyfriend
53	attributive	soldier ant
54	attributive	twig legs
55	attributive	mushroom cloud
56	attributive	silk hair
57	attributive	elbow macaroni
58	attributive	sandpaper skin
59	attributive	pen knife
60	attributive	handlebar moustache
61	attributive	hourglass body
62	attributive	iron fist
63	attributive	pie chart
64	attributive	bullet train
65	relational	pine mushroom
66	relational	orchard oriole
67	relational	milk virus

68	relational	lounge beer
69	relational	copper horse
70	relational	patio cigarette
71	relational	water celery
72	relational	cow parsnip
73	relational	daisy extract
74	relational	rodeo magazine
75	relational	sailing hat
76	relational	honey soup
77	relational	floor television
78	relational	winter mushroom
79	relational	glass rose
80	relational	acorn woodpecker
81	relational	harbor porpoise
82	relational	prisoner graffiti
83	relational	nose sound
84	relational	bowling sweater
85	relational	pasta jar
86	relational	battle theory
87	relational	servant scandal
88	relational	mustard spot
89	relational	gender comedy

90	relational	motorcycle documentary
91	relational	bacon tongs
92	relational	finger nerve
93	relational	cracker tray
94	relational	rugby shoes
95	relational	microwave sandwich
96	relational	telescope observation
97	relational	onion tears
98	relational	pancake spatula
99	relational	concrete fountain
100	relational	burrito stain
101	relational	council mandate
102	relational	storage hutch
103	relational	mountain snake
104	relational	sugar tea
105	relational	jungle bird
106	relational	ocean nausea
107	relational	city riots
108	relational	employee vote
109	relational	scalpel incision
110	relational	gambling policy
111	relational	sympathy bouquet

112	relational	room cabinets
113	relational	party dance
114	relational	chest rash
115	relational	conference chairman
116	relational	gold flakes
117	relational	doctor testimony
118	relational	baseball injury
119	relational	grill steak
120	relational	song book
121	relational	quartz clock
122	relational	lapel microphone
123	relational	raisin cereal
124	relational	sign post
125	relational	desk stapler
126	relational	flower season
127	relational	cream sauce
128	relational	wood stove

REFERENCES

- Baggio, G., Choma, T., van Lambalgen, M., & Hagoort, P. (2010). Coercion and compositionality. *Journal of Cognitive Neuroscience*, *22*(9), 2131–2140.
- Baron, S. G., Thompson-Schill, S. L., Weber, M., & Osherson, D. (2010). An early stage of conceptual combination: Superimposition of constituent concepts in left anterolateral temporal lobe. *Cognitive Neuroscience*, *1*(1), 44–51.
<https://doi.org/10.1080/17588920903548751>
- Bavelier, D., Corina, D., Jezzard, P., Padmanabhan, S., Clark, V. P., Karni, A., ... Rauschecker, J. P. (1997). Sentence reading: A functional MRI study at 4 Tesla. *Journal of Cognitive Neuroscience*, *9*(5), 664–686.
- Beeman, M., Friedman, R., Grafman, J., Perez, E., Diamond, S., & Lindsay, M. (1994). Summation priming and coarse semantic coding in the right hemisphere. *Cognitive Neuroscience, Journal of*, *6*(1), 26–45.
- Bemis, D. K., & Pylkkänen, L. (2011). Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience*, *31*(8), 2801–2814.
- Bemis, D. K., & Pylkkänen, L. (2012). Basic Linguistic Composition Recruits the Left Anterior Temporal Lobe and Left Angular Gyrus During Both Listening and Reading. *Cerebral Cortex*, *23*, 1859–1873.
<https://doi.org/10.1093/cercor/bhs170>
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*.

- Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, *19*(12), 2767–2796.
- Binder, J. R., Westbury, C., McKiernan, K., Possing, E., & Medler, D. (2005). Distinct brain systems for processing concrete and abstract concepts. *Cognitive Neuroscience, Journal of*, *17*(6), 905–917.
- Binney, R. J., Parker, G. J. M., & Lambon Ralph, M. A. (2012). Convergent Connectivity and Graded Specialization in the Rostral Human Temporal Lobe as Revealed by Diffusion-Weighted Imaging Probabilistic Tractography. *Journal of Cognitive Neuroscience*, *24*(10), 1998–2014.
- 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. (2015). Compositionality and the angular gyrus: A multi-voxel similarity analysis of the semantic composition of nouns and verbs. *Neuropsychologia*, *78*, 130–141.
<https://doi.org/10.1016/j.neuropsychologia.2015.10.007>
- Bressler, D., Spotswood, N., & Whitney, D. (2007). Negative BOLD fMRI Response in the Visual Cortex Carries Precise Stimulus-Specific Information. *PLoS ONE*, *2*(5). <https://doi.org/10.1371/journal.pone.0000410>
- Brysbaert, M., & New, B. (2009). Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new

- and improved word frequency measure for American English. *Behavior Research Methods*, 41(4), 977–990. <https://doi.org/10.3758/BRM.41.4.977>
- Buccino, G., Binkofski, F., Fink, G. R., Fadiga, L., Fogassi, L., Gallese, V., ... Freund, H.-J. (2001). Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. *European Journal of Neuroscience*, 13(2), 400–404. <https://doi.org/10.1111/j.1460-9568.2001.01385.x>
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science*, 6(1), 3–5. <https://doi.org/10.1177/1745691610393980>
- Chao, L. L., & Martin, A. (2000). Representation of Manipulable Man-Made Objects in the Dorsal Stream. *NeuroImage*, 12(4), 478–484. <https://doi.org/10.1006/nimg.2000.0635>
- Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, 24(2), 299–349.
- Coutanche, M. N., & Thompson-Schill, S. L. (2014). Creating Concepts from Converging Features in Human Cortex. *Cerebral Cortex*, 25, 2584–2593. <https://doi.org/10.1093/cercor/bhu057>
- Cox, R. W., & Jesmanowicz, A. (1999). Real-time 3D image registration for functional MRI. *Magnetic Resonance in Medicine*, 42(6), 1014–1018.
- 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. <https://doi.org/10.1006/nimg.2001.0775>

- Downing, P. (1977). On the Creation and Use of English Compound Nouns. *Language*, 53(4), 810. <https://doi.org/10.2307/412913>
- Embick, D., Marantz, A., Miyashita, Y., O'Neil, W., & Sakai, K. L. (2000). A syntactic specialization for Broca's area. *Proceedings of the National Academy of Sciences*, 97(11), 6150–6154. Retrieved from <http://www.pnas.org/content/97/11/6150.short>
- Estes, Z. (2003). A tale of two similarities: comparison and integration in conceptual combination. *Cognitive Science*, 27(6), 911–921. <https://doi.org/10.1016/j.cogsci.2003.01.001>
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral 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. <https://doi.org/10.1126/science.1106138>
- Fruchter, J., Linzen, T., Westerlund, M., & Marantz, A. (2015). Lexical Preactivation in Basic Linguistic Phrases. *Journal of Cognitive Neuroscience*, 1–24. https://doi.org/10.1162/jocn_a_00822
- Gagné, C. L. (2000). Relation-Based Combinations Versus Property-Based Combinations: A Test of the CARIN Theory and the Dual-Process Theory of

- Conceptual Combination. *Journal of Memory and Language*, 42(3), 365–389.
<https://doi.org/10.1006/jmla.1999.2683>
- Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(1), 71–87.
<https://doi.org/10.1037/0278-7393.23.1.71>
- Grabner, R. H., Ansari, D., Koschutnig, K., Reishofer, G., & Ebner, F. (2013). The function of the left angular gyrus in mental arithmetic: Evidence from the associative confusion effect. *Human Brain Mapping*, 34(5), 1013–1024.
<https://doi.org/10.1002/hbm.21489>
- 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.
<https://doi.org/10.1016/j.neuroimage.2010.06.055>
- Henson, R. N. A., Shallice, T., Josephs, O., & Dolan, R. J. (2002). Functional Magnetic Resonance Imaging of Proactive Interference during Spoken Cued Recall. *NeuroImage*, 17(2), 543–558. <https://doi.org/10.1006/nimg.2002.1229>
- Hoeks, J. C. J., Stowe, L. A., & Doedens, G. (2004). Seeing words in context: the interaction of lexical and sentence level information during reading. *Cognitive Brain Research*, 19(1), 59–73.
<https://doi.org/10.1016/j.cogbrainres.2003.10.022>
- Hoffman, P., Binney, R. J., & Lambon Ralph, M. A. (2015). Differing contributions of inferior prefrontal and anterior temporal cortex to concrete and abstract

- conceptual knowledge. *Cortex*, 63, 250–266.
<https://doi.org/10.1016/j.cortex.2014.09.001>
- 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. [https://doi.org/10.1016/S1053-8119\(03\)00272-6](https://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.
<https://doi.org/10.1016/j.neuroimage.2007.03.059>
- Humphries, C., Binder, J. R., Medler, D., Liebenthal, E., & others. (2006). Syntactic and semantic modulation of neural activity during auditory sentence comprehension. *Cognitive Neuroscience, Journal of*, 18(4), 665–679.
- Kang, A. M., Constable, R. T., Gore, J. C., & Avrutin, S. (1999). An Event-Related fMRI Study of Implicit Phrase-Level Syntactic and Semantic Processing. *NeuroImage*, 10(5), 555–561.
- Kuperberg, G. R. (2007). Neural mechanisms of language comprehension: Challenges to syntax. *Brain Research*, 1146, 23–49.
<https://doi.org/10.1016/j.brainres.2006.12.063>
- Kuperberg, G. R., Choi, A., Cohn, N., Paczynski, M., & Jackendoff, R. (2009). Electrophysiological Correlates of Complement Coercion. *Journal of Cognitive Neuroscience*, 22(12), 2685–2701. <https://doi.org/10.1162/jocn.2009.21333>

- Kuperberg, G. R., Holcomb, P. J., Sitnikova, T., Greve, D., Dale, A. M., & Caplan, D. (2003). Distinct patterns of neural modulation during the processing of conceptual and syntactic anomalies. *Journal of Cognitive Neuroscience, 15*(2), 272–293.
- 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. <https://doi.org/10.1016/j.neuroimage.2007.10.009>
- Laham, D. (1998). *Latent Semantic Analysis at CU Boulder*. Retrieved from <http://lsa.colorado.edu/>
- Lambon Ralph, M. A. (2014). Neurocognitive insights on conceptual knowledge and its breakdown. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 369*(1634), 20120392. <https://doi.org/10.1098/rstb.2012.0392>
- Lau, E. F., Weber, K., Gramfort, A., Hämäläinen, M. S., & Kuperberg, G. R. (2014). Spatiotemporal Signatures of Lexical–Semantic Prediction. *Cerebral Cortex, bhu219*. <https://doi.org/10.1093/cercor/bhu219>
- Levi, J. N. (1978). *The Syntax and Semantics of Complex Nominals*. Academic Press.
- Lewis, G. A., Poeppel, D., & Murphy, G. L. (2015). The neural bases of taxonomic and thematic conceptual relations: An MEG study. *Neuropsychologia, 68*, 176–189. <https://doi.org/10.1016/j.neuropsychologia.2015.01.011>
- Linke, A. C., Vicente-Grabovetsky, A., & Cusack, R. (2011). Stimulus-specific suppression preserves information in auditory short-term memory.

- Proceedings of the National Academy of Sciences*, 108(31), 12961–12966.
<https://doi.org/10.1073/pnas.1102118108>
- McElree, B., Pylkkänen, L., Pickering, M. J., & Traxler, M. J. (2006). A time course analysis of enriched composition. *Psychonomic Bulletin & Review*, 13(1), 53–59. <https://doi.org/10.3758/BF03193812>
- Meltzer-Asscher, A., Schuchard, J., den Ouden, D.-B., & Thompson, C. K. (2013). The neural substrates of complex argument structure representations: Processing “alternating transitivity” verbs. *Language and Cognitive Processes*, 28(8), 1–15. <https://doi.org/10.1080/01690965.2012.672754>
- Müller, N. G., & Kleinschmidt, A. (2004). The attentional “spotlight”s’ penumbra: center-surround modulation in striate cortex. *Neuroreport*, 15(6), 977–980.
- Murphy, G. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language*, 29, 259–288.
- Newman, A. J., Pancheva, R., Ozawa, K., Neville, H. J., & Ullman, M. T. (2001). An Event-Related fMRI Study of Syntactic and Semantic Violations. *Journal of Psycholinguistic Research*, 30(3), 339–364.
<https://doi.org/10.1023/A:1010499119393>
- Noonan, K. A., Jefferies, E., Visser, M., & Lambon Ralph, M. A. (2013). Going beyond Inferior Prefrontal Involvement in Semantic Control: Evidence for the Additional Contribution of Dorsal Angular Gyrus and Posterior Middle Temporal Cortex. *Journal of Cognitive Neuroscience*, 1824–1850.
https://doi.org/10.1162/jocn_a_00442

- Pallier, C., Devauchelle, A. D., & Dehaene, S. (2011). Cortical representation of the constituent structure of sentences. *Proceedings of the National Academy of Sciences*, *108*(6), 2522–2527. Retrieved from <http://www.pnas.org/content/108/6/2522.short>
- Partee, B. H., & Borschev, V. (2003). Genitives, relational nouns, and argument-modifier ambiguity. *Modifying Adjuncts*, *4*, 67–112.
- Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, *8*(12), 976–987.
- Pylkkänen, L., & McElree, B. (2006). The syntax-semantics interface: On-line composition of sentence meaning. *Handbook of Psycholinguistics*, *2*, 537–577.
- Pylkkänen, L., & McElree, B. (2007). An MEG Study of Silent Meaning. *Journal of Cognitive Neuroscience*, *19*(11), 1905–1921. <https://doi.org/10.1162/jocn.2007.19.11.1905>
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682. <https://doi.org/10.1073/pnas.98.2.676>
- Rauschecker, J. P., & Scott, S. K. (2009). Maps and streams in the auditory cortex: nonhuman primates illuminate human speech processing. *Nature Neuroscience*, *12*(6), 718–724. <https://doi.org/10.1038/nn.2331>
- Rodd, J. M., Gaskell, M. G., & Marslen-Wilson, W. D. (2004). Modelling the effects of semantic ambiguity in word recognition. *Cognitive Science*, *28*(1), 89–104.

- Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. MIT Press.
- Schwartz, M. F., Kimberg, D. Y., Walker, G. M., Brecher, A., Faseyitan, O. K., Dell, G. S., ... Coslett, H. (2011). Neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. *Proceedings of the National Academy of Sciences*, *108*(20), 8520–8524.
- Seghier, M. L. (2012). The angular gyrus: multiple functions and multiple subdivisions. *The Neuroscientist*, *19*(1), 43–61.
- Seghier, M. L., Fagan, E., & Price, C. J. (2010). Functional Subdivisions in the Left Angular Gyrus Where the Semantic System Meets and Diverges from the Default Network. *Journal of Neuroscience*, *30*(50), 16809–16817.
<https://doi.org/10.1523/JNEUROSCI.3377-10.2010>
- Shmuel, A., Augath, M., Oeltermann, A., & Logothetis, N. K. (2006). Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nature Neuroscience*, *9*(4), 569–577.
<https://doi.org/10.1038/nn1675>
- Shmuel, A., Yacoub, E., Pfeuffer, J., Van de Moortele, P.-F., Adriany, G., Hu, X., & Ugurbil, K. (2002). Sustained negative BOLD, blood flow and oxygen consumption response and its coupling to the positive response in the human brain. *Neuron*, *36*(6), 1195–1210.
- Shoben, E. J., & Gagné, C. L. (1997). Thematic relations and the creation of combined concepts. In *Creative thought: An investigation of conceptual structures and*

- processes* (pp. 31–50). Washington, D.C., US: American Psychological Association.
- Smith, A. T., Williams, A. L., & Singh, K. D. (2004). Negative BOLD in the visual cortex: evidence against blood stealing. *Human Brain Mapping, 21*(4), 213–220. <https://doi.org/10.1002/hbm.20017>
- Staresina, B. P., Fell, J., Do Lam, A. T. A., Axmacher, N., & Henson, R. N. (2012). Memory signals are temporally dissociated in and across human hippocampus and perirhinal cortex. *Nature Neuroscience, 15*(8), 1167–1173. <https://doi.org/10.1038/nn.3154>
- Stringer, S. M., & Rolls, E. T. (2002). Invariant Object Recognition in the Visual System with Novel Views of 3D Objects. *Neural Computation, 14*(11), 2585–2596. <https://doi.org/10.1162/089976602760407982>
- Thompson, C. K., Bonakdarpour, B., Fix, S. C., Blumenfeld, H. K., Parrish, T. B., Gitelman, D. R., & Mesulam, M.-M. (2007). Neural correlates of verb argument structure processing. *Journal of Cognitive Neuroscience, 19*(11), 1753–1767.
- Thompson, C. K., Bonakdarpour, B., & Fix, S. F. (2010). Neural mechanisms of verb argument structure processing in agrammatic aphasic and healthy age-matched listeners. *Journal of Cognitive Neuroscience, 22*(9), 1993–2011.
- Traxler, M. J., McElree, B., Williams, R. S., & Pickering, M. J. (2005). Context effects in coercion: Evidence from eye movements. *Journal of Memory and Language, 53*(1), 1–25. <https://doi.org/10.1016/j.jml.2005.02.002>
- Wang, J., Conder, J. A., Blitzer, D. N., & Shinkareva, S. V. (2010). Neural representation of abstract and concrete concepts: A meta-analysis of neuroimaging studies.

- Human Brain Mapping*, 31(10), 1459–1468.
<https://doi.org/10.1002/hbm.20950>
- Warren, B. (1978). Semantic patterns of noun-noun compounds. In *Acta Universitatis Gothoburgensis. Gothenburg Studies in English Goteborg* (Vol. 41, pp. 1–266).
- Westerlund, M., & Pylkkänen, L. (2014). The role of the left anterior temporal lobe in semantic composition vs. semantic memory. *Neuropsychologia*, 57, 59–70.
<https://doi.org/10.1016/j.neuropsychologia.2014.03.001>
- Wisniewski, E. J. (1996). Construal and Similarity in Conceptual Combination. *Journal of Memory and Language*, 35(3), 434–453.
<https://doi.org/10.1006/jmla.1996.0024>
- Wisniewski, E. J., & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In *Understanding word and sentence. Advances in psychology* (pp. 241–284). Oxford, England: North Holland.
- Wisniewski, E. J., & Love, B. C. (1998). Relations versus Properties in Conceptual Combination. *Journal of Memory and Language*, 38, 177–202.
- 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.