ORIGINAL ARTICLE

Contrasting Semantic versus Inhibitory Processing in the Angular Gyrus: An fMRI Study

Gwyneth A. Lewis¹, David Poeppel^{1,2} and Gregory L. Murphy¹

¹Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA and

²Department of Neuroscience, Max-Planck-Institute (MPIEA), Grüneburgweg 14, 60322 Frankfurt, Germany

Address correspondence to Gwyneth A. Lewis, Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA.
Email: gwyneth@nyu.edu

Abstract

Recent studies of semantic memory have focused on dissociating the neural bases of two foundational components of human thought: taxonomic categories, which group similar objects like dogs and seals based on features, and thematic categories, which group dissimilar objects like dogs and leashes based on events. While there is emerging consensus that taxonomic concepts are represented in the anterior temporal lobe, there is disagreement over whether thematic concepts are represented in the angular gyrus (AG). We previously found AG sensitivity to both kinds of concepts; however, some accounts suggest that such activity reflects inhibition of irrelevant information rather than thematic activation. To test these possibilities, an fMRI experiment investigated both types of conceptual relations in the AG during two semantic judgment tasks. Each task trained participants to give negative responses (inhibition) or positive responses (activation) to word pairs based on taxonomic and thematic criteria of relatedness. Results showed AG engagement during both negative judgments and thematic judgments, but not during positive judgments about taxonomic pairs. Together, the results suggest that activity in the AG reflects functions that include both thematic (but not taxonomic) processing and inhibiting irrelevant semantic information.

Key words: angular gyrus, fMRI, taxonomic categories, concepts, thematic relations

Introduction

Semantic representations structure knowledge of real-world entities into coherent categories. Categories have often been divided into types depending on how they are constituted. *Taxonomic categories* group things together that share properties. For example, dogs have 4 legs, fur, and many other features. In contrast, *thematic categories* are defined by extrinsic functional or spatial relations between objects. Dogs and leashes are thematically related because the dog wears a leash during walks. However, thematic pairs do not necessarily share features—dogs and leashes are not both made of nylon and do not both bark. (For more on this distinction, see Markman 1989; Lin and Murphy 2001; Murphy 2001; Estes et al. 2011.)

Researchers dispute whether distinct semantic systems represent taxonomic and thematic knowledge. On the one hand,

thematic and taxonomic relations have different bases—thematic relations group often dissimilar objects based on external, complementary, and spatial links, whereas taxonomic relations group similar objects based on internal, independent, and similarity-based properties (Estes et al. 2011). On the other hand, the same semantic system could encode “internal” features of objects in the “external” events they occur in: Although dogs and leashes may go together to form a thematic category, people must identify its two (taxonomic) components in order to form that category. Thus, thematic relations are to some degree dependent on taxonomic knowledge.

Neurocognitive theories of semantic memory have taken two opposing views of taxonomic versus thematic relations. Hub-and-spoke models propose that verbal and non-verbal information is distributed across modality-specific cortices (the

spokes), and that a transmodal hub located in the anterior temporal lobe (ATL) assembles information from the spokes into concepts. Such *one-system views* propose that the ATL supports both kinds of relations (e.g., Lambon Ralph et al. 2010; Jackson et al. 2015; Lambon Ralph et al. 2016). An alternative view is the *two-systems view*, which proposes an additional hub for thematic relations in the temporoparietal junction (TPJ) (e.g., Schwartz et al. 2011; for reviews see Lambon Ralph et al. 2016; Mirman et al. 2017). Recent experiments have used a range of techniques to test these views. Lesion studies have linked thematic and taxonomic conceptual deficits to distinct brain areas (see below), and semantic priming studies have tested neural dissociations in healthy populations as measured with EEG (e.g., Maguire et al. 2010), fMRI (e.g., de Zubicaray et al. 2013), MEG (e.g., Lewis et al. 2015), and TMS (e.g., Davey et al. 2015). The ATL is widely believed to be involved in taxonomic processing, but the degree to which it is involved in thematic processing is not yet certain. Damage to the ATL harms the detection of associative relations (see Jackson et al. 2015), but Schwartz et al., in their analysis of speech errors from patients with damage to the TPJ, argued that thematic relations were predominantly found in the angular gyrus (AG) area of the TPJ. Our own MEG results (Lewis et al. 2015) suggest that both relations involve the AG. The ATL is clearly central to semantic knowledge; we return to its role in thematic relations in particular in the Discussion. Our study used fMRI to investigate the role of the AG in conceptual processing. We now describe the (somewhat contradictory) literature regarding the AG and related areas.

Divergent evidence comes from studies of various neuropsychological populations, such as those comparing patients with semantic dementia (ATL damage) and semantic aphasia (temporoparietal and/or prefrontal (TP/PF) damage). The latter are most relevant to the current investigation. Jefferies and Lambon Ralph (2006) compared such patients' performance on a battery of semantic tests. The TP/PF patients made taxonomic errors (e.g., calling a dog "seal" or "animal") but also made thematic errors (e.g., calling a dog "leash"). Moreover, their naming accuracy increased with phonemic cues, possibly by increasing activation of the correct label relative to the distractors. Thematic errors may occur because TP/PF patients have trouble navigating attention away from co-activated associate distractors. Jefferies and Lambon Ralph concluded that TP/PF damage deregulates controlled semantic retrieval. (Subgroups of their TP/PF patients did not differ from one another in any meaningful way, suggesting that damage to either frontal or parietal areas can disrupt cognitive control.)

Using a battery of semantic tasks, Noonan et al. (2010) linked parietal damage with deficits in semantic navigation, control over irrelevant information, and retrieval of uncommon meanings. Rather than loss of thematic representations, parietal damage might deregulate semantic processes. Supporting evidence comes from Whitney et al. (2012), who used TMS to simulate virtual lesions in healthy participants during tasks with varying semantic and non-semantic control demands. They found that TMS to parietal areas disrupted top-down semantic selection and non-semantic control. This conclusion was seconded by Thompson et al. (2017), who concluded that semantic control issues were responsible for difficulties in detecting thematic relations in semantic aphasia patients. In sum, parietal areas may support complementary control mechanisms important for semantic retrieval but not semantic processing in particular.

An alternative account proposes that the AG supports spatial and functional relations of events rather than complementary

control mechanisms. Like Jefferies and Lambon Ralph (2006), Schwartz et al. (2011) coded picture naming errors of aphasic patients with language deficits. After isolating each error type, they found that ATL damage predicted taxonomic errors and AG damage predicted thematic errors. Schwartz et al. proposed that the TPJ (specifically, AG) links contextually relevant (event-based) information to the target concept via activation of the correct label (e.g., "dog") as well as a thematically related word (e.g., "leash"). Although rare, thematic errors can occur because damage weakens signal strength for the target concept "dog" relative to its thematic associate "leash." However, one might wonder whether damage to a brain area that represents thematic relations should not reduce rather than increase thematic intrusions, as thematic relations would be less available (Lewis et al. 2015). Also, their statistical technique was designed to distinguish thematic from taxonomic errors, and so it did not reveal areas they might have in common.

Mirman and Graziano's (2012) eye-tracking study of aphasics and control participants tested alternative accounts of Schwartz et al.'s AG result: Temporoparietal damage increases errors in naming but not thematic comprehension, or it impairs cognitive control but not thematic semantic processing (e.g., Noonan et al. 2010). They tracked incidental activation of semantic relations with a passive viewing task during word comprehension. ATL patients fixated longer on taxonomically related objects (possibly reflecting greater difficulty in distinguishing matching objects from taxonomic competitors), while AG patients fixated later and less frequently on thematically related objects. Under Schwartz et al.'s account, AG damage does not decrease thematic processing but rather increases thematic errors (as discussed above), so parietal patients should have looked longer and more frequently at thematic objects. In our view, the Mirman and Graziano results are more consistent with an association between the AG and thematic processing, but these results are part of a pattern in which evidence for both inhibitory mechanisms and thematic relations seem to be found in different studies or using different measures, as we review next.

Data from fMRI studies of semantic representation and control in healthy populations have yet to converge on a consistent interpretation of the AG's role; some results indicate specialization for thematic relations (e.g., Kalénine et al. 2009; de Zubicaray et al. 2013; Davey et al. 2016; Geng and Schnur 2016), and others indicate a general role in semantic control (e.g., Jackson et al. 2015). The studies described next are broadly organized by task paradigm.

Several fMRI studies used a two-alternative forced choice (2AFC) task in which a probe is presented with a related and unrelated foil and concluded that temporoparietal areas are involved in detecting thematic relations. Kalénine et al. (2009) found that thematically related probes and targets increased temporoparietal activity over taxonomically related pairs. Davey et al. (2016) probed specific aspects of semantic knowledge and found greater AG activity during similarity judgments about function than about color.

de Zubicaray et al. (2013) used a picture-word interference naming paradigm, presenting individual objects to be named along with a foil word to be ignored. (The foil is thought to elicit spreading activation among related concepts.) Thematic foils yielded greater activity than taxonomic foils in the left AG and middle temporal gyrus (MTG). Compared with unrelated foils, thematic foils speeded naming and increased activity in the AG, while taxonomic foils delayed naming and decreased activity primarily in the posterior MTG. Decreased posterior MTG

activity could reflect lexical competition, in which taxonomic foils activated shared features between category members that in turn activated lexical competitors. (In contrast, priming benefits may only be observed for thematic distractors, which should not induce lexical competition because they activate conceptually distinct items.)

Finally, evidence for thematic specialization comes from two recent fMRI studies. [Geng and Schnur \(2016\)](#) presented thematic and taxonomic word pairs in an fMRI adaptation paradigm. The paradigm assumes that repeated viewing induces BOLD signal change in brain areas responsible for processing some aspect of the stimulus, namely, featural versus functional aspects. They found that taxonomic pairs elicited ATL adaptation and thematic pairs elicited AG adaptation. [Xu et al. \(2018\)](#) found some evidence that both the ATL and TPJ were involved in both kinds of categorization. However, when the two areas were contrasted, the TPJ was more involved in thematic categorization (and the ATL in taxonomic categorization).

Thus, there is considerable evidence that the AG is involved in processing thematic categories.

Differences in semantic control demands could explain reported dissociations of taxonomic and thematic relations. [Jackson et al. \(2015\)](#) noted that Schwartz et al.'s two-systems account conflicts with findings that TP/PF damaged patients generate both taxonomic and thematic errors, and that both TP/PF and ATL damaged patients show deficits in using thematic relations ([Jefferies and Lambon Ralph 2006](#)). Their fMRI experiment probed semantic and non-semantic control demands with several 2AFC tasks. First, they found greater activity over supplementary motor and frontal areas for taxonomic judgments, and over AG and supramarginal areas (SMG) for thematic judgments. Their conjunction analysis showed that taxonomic > thematic effects overlapped with areas responding to high semantic control trials, suggesting that differences reflect greater executive demands required for more difficult judgments. Lastly, their ROI analysis showed equivalent ATL enhancement for both semantic conditions over the non-semantic task, as well as equivalent AG deactivation for thematic, taxonomic, and non-semantic conditions. Jackson et al. proposed that the ATL encodes not just features of concepts, but also thematic links between associated items that amalgamate as types of features via experience. Although encoded together, features and associations may entail different levels of semantic control if they vary in terms of the range of concepts they link to. In sum, previous reports of neural dissociation may reflect differences in semantic control demands rather than recruitment of distinct semantic hubs.

If the AG represents thematic relations, some of which must be inhibited during retrieval, this could account for our finding of AG engagement with both kinds of relations ([Lewis et al. 2015](#)). Our MEG study tested predictions from the one-system and two-systems views by comparing AG and ATL responses to taxonomic, thematic, and unrelated word pairs. Taxonomic relations selectively involved the ATL, but both relations involved the AG. Specifically, AG activity distinguished both thematic and taxonomic pairs from unrelated pairs over the same time window. The sequentially presented taxonomic pairs (e.g., DOG → SEAL) may have elicited inhibition if the first word (e.g., DOG) activated thematic relations (e.g., *leash*, *bone*, *collar*). Such activation would be irrelevant to processing the taxonomic relation with the second word (e.g., SEAL) because our taxonomic pairs were not also thematically related. Support for this hypothesis comes from [Maguire et al. \(2010\)](#), whose EEG study identified decrease and increase in parietal activity for thematic and taxonomic primes, respectively, over the same

time window. The activity increase could have reflected inhibition of thematic memory traces elicited by the taxonomic prime that are irrelevant for accessing the correct taxonomic category. As such, taxonomic relations may require greater resources in part because thematically related words must be inhibited. However, the results are not necessarily incompatible with thematic processing in parietal areas.

The Present Study

This existing literature reveals a confusing collection of findings, some of which emphasize the AG's importance in semantic inhibition and others of which suggest its importance in the representation or processing of thematic relations. Of course, there is no necessary contradiction between these two sets of findings—the AG could be involved in both, possibly because thematic associates require inhibition in some tasks, to prevent them from being incorrectly processed or produced.

The present fMRI study simultaneously tests inhibition and thematic processing in the AG. Unlike previous imaging studies of conceptual relations, we manipulated inhibition by factorially crossing word pair relation type—thematic, taxonomic, and unrelated—with tasks requiring inhibition either of taxonomic or thematic knowledge.

First, we obtained relatedness ratings for a large set of word pairs and selected pairs such that conditions were carefully matched for relatedness across tasks. This was important because thematic pairs tend to receive higher relatedness ratings than taxonomic pairs ([Lewis et al. 2015](#)), and differences in activation could be caused by differences in difficulty rather than type of processing ([Jackson et al. 2015](#)). Conditions were also matched on linguistic properties such as written frequency, syllables, and word length. We tested the final stimulus set in an event-related fMRI design, manipulating the type of semantic relationship (taxonomic, thematic, filler) and the type of processing required (inhibition, activation) to explore predictions from models of semantic organization. We manipulated inhibition by specifying positive or negative responses for each relation type in two tasks. The taxonomic task trained participants to accept items of the same general kind (e.g., COTTAGE-CASTLE) and to reject thematic and unrelated items. Conversely, the thematic task taught participants to accept items that tend to interact or occur in the same event (e.g., DOG-LEASH) and to reject taxonomic and unrelated word pairs. Thus, each task entailed inhibition of the other relation in some word pairs, allowing us to test whether AG activity reflects inhibitory processes or perhaps inhibition of thematic processes in particular, as some accounts have suggested.

Our notion of inhibition is that when pairs are related in some way, responding that they are unrelated (for that particular task) requires the suppression of activation arising from the relation. The exact neural mechanism underlying this is not known: It could arise from increased activation in an attentional control mechanism that must overcome the unwanted relational activation, or it could involve deactivation of an area due to task requirements. Past results have found deactivation in the AG from thematic relations, albeit not in an inhibitory task ([Humphreys and Lambon Ralph 2017](#)). We did not have a prediction about which mechanism would be found; our method could discover either one. For expository ease, we discuss AG “activation” in our predictions (as we did not, in fact, find deactivation).

The thematic account predicts greater AG activity for thematic pairs, and the inhibition account predicts greater AG

activity when the semantic relationship mismatches the task criteria. Specifically, within the thematic judgment task, the thematic account predicts greater AG activity for accepting thematic than rejecting taxonomic pairs, while the inhibition account predicts the opposite result. The inclusion of the filler (unrelated) condition provides a baseline response for each task to calculate relative differences. In particular, fillers can be rejected without inhibition, since they are not related in any way. Therefore, the difference between the fillers and the related negative pairs (thematic pairs in the taxonomic task and taxonomic pairs in the thematic task) is a measure of inhibitory processing, as participants must reject them even though they are related.

Given our goals, we focused on the AG as a region of interest (ROI). However, because of the ATL's role in taxonomic processing, we also specified that area as an ROI. Unfortunately, the ATL is difficult to image with standard fMRI. For one, "the signal-to-noise ratio diminishes substantially near the temporal poles, owing to their proximity to air filled sinuses (the so-called 'susceptibility artefact')" (Patterson et al. 2007, p. 981). The likelihood of detecting ATL activity is also influenced by the nature of the baseline task and the depth of the field of view (Visser et al. 2010). (See Jackson et al. 2015 for a relevant study that overcame challenges associated with imaging the ATL.) Our standard fMRI experiment was not designed to control such factors, so we did not make strong predictions about ATL results. This is one reason why our previous study used MEG, which could find significant ATL activity (Lewis et al. 2015). For the purpose of the present study, we instead focused on attaining maximal coverage of the AG. We defined the AG by focusing on the ROIs from Schwartz et al. (2011) and Jackson et al. (2015), to allow more direct comparisons to those studies. (Schwartz et al. used the term TPJ to refer to their posterior area, and we followed their practice in Lewis et al. But the ROI actually corresponds more specifically to the AG; the TPJ generally refers to a larger area. Thus, we use the more specific term here.)

Materials and Method

Participants

Eighteen right-handed native English speakers (10 males; mean age = 27) recruited from the New York University community participated in the fMRI experiment. The New York University IRB approved the study, and all participants completed a safety screening form and provided written informed consent. Two participants were excluded from the analysis (one due to inadequate coverage of the parietal lobe, one due to low response accuracy) resulting in a total of 16 included participants.

Stimuli

An initial set of over 400 word pairs was drawn in part from our previous experiment (Lewis et al. 2015), from Battig and Montague (1969), and the University of South Florida Free Association Norms corpus (Nelson et al. 2004). The thematic pairs related spatially (e.g., TABLE-LAMP) or functionally (e.g., THREAD-NEEDLE) but not taxonomically (e.g., not CAT-MOUSE). The taxonomic pairs were basic-level categories (e.g., SEAL-DOG, both are mammals) that did not share an obvious thematic relation (e.g., not CAT-MOUSE). The filler pairs (e.g., PANE-NUT) were unrelated. We next collected ratings from online participants ($n = 32$) who rated the relatedness of each word pair on a 1–7 scale (1 = not at all related and 7 = highly

related). Critically, the ratings instructions did not include specific definitions of taxonomic and thematic relatedness because we did not want to emphasize certain features but rather assess the general strength of word pair relatedness. To ensure participants took the task seriously, we included "sanity check" items, consisting of 10 highly related items (e.g., SQUARE-CIRCLE) and 10 highly unrelated (e.g., SCHEME-MOOSE) items (not used in the experiment). We excluded 3 participants based on their ratings of these items.

The ratings data confirmed that filler pairs were not related ($M = 1.6$, $SD = 1.0$). Additionally, thematic pairs received higher relatedness ratings ($M = 5.6$, $SD = 1.1$), than taxonomic pairs ($M = 4.5$, $SD = 1.3$). Given this disparity, we excluded items from each list to match the two conditions on relatedness (new taxonomic relatedness: $M = 5.2$, $SD = 0.5$; new thematic relatedness: $M = 5.3$, $SD = 0.6$), as well as on psycholinguistic variables to ensure that any effects could not be attributed to differences in the stimuli other than the intended manipulation. The final stimulus set included a total of 320 word pairs: 120 thematic, 120 taxonomic, and 80 controls. Supplementary Table 1 shows a summary of the psycholinguistic variables by condition, and Supplementary Table 2 shows the stimulus pairs and their relatedness values.

Procedure

Prior to each task, participants were trained in making relatedness judgments in the scanner with 30 word pairs (15 thematic, 15 taxonomic). Each response received feedback and an explanation. The taxonomic training instructed participants to respond "yes" to items of the same general kind (e.g., *MELON and BERRY: both are kinds of fruit*), and "no" to items that tend to interact or occur together (e.g., *HEN and CAGE: while hens live in cages, only one is an animal*). The thematic training defined relatedness in opposite terms. The instructions and practice items are provided in the Supplementary Methods 1. Each trial began with a fixation dot presented for 0.5 s, followed by the stimuli (one word above the other) for 1.5 s, presented in black Arial font at size 30 against a light gray background. The trials were pseudo-randomly ordered and were jittered by a variable ITI (2–6 s, with optimal fMRI schedules programmed in Optseq2). Feedback ("Incorrect") was only provided on incorrect trials. Figure 1 shows the trial sequence. There were four 7-min runs, each consisting of 80 items. Each task presented participants with 160 items across two sequential scans. Counterbalancing of task order ensured that half of the participants completed the taxonomic task in runs 1–2 and the thematic task in runs 3–4, and vice versa. No item was repeated within or across tasks.

Imaging and Data Analysis

Scans were conducted with a 3-T Siemens Allegra system with a Nova Medical NM-011 head coil at New York University's Center for Brain Imaging lab. Stimuli were projected on a screen at the back of the scanner (approximately 58 cm from each participant's forehead), and reflected through a mirror above the participant's head. Echo-planar images (EPI) were acquired with a multi-gradient-echo EPI sequence [TR = 2 s, TE = 15 s, FA = 82°, field of view = 192 × 240 mm], slices = 34 [slice thickness = 3 mm, voxel resolution = 3 × 3 × 3 mm, volumes per run = 206]. Slices were oriented off the AC-PC line and adjusted for each individual to ensure maximal coverage of both the ATL and the AG. High-resolution structural images were acquired with a magnetization-prepared rapid-acquisition

Figure 1. Trial sequence in the fMRI experiment.

gradient echo (MPRAGE) (voxel resolution: $1 \times 1 \times 1$ mm) at the end of the session. Magnetic field estimates were acquired with a field map sequence for later use in image distortion correction.

Data were analyzed using FMRI Expert Analysis Tool (FEAT) version 6.00, part of FSL (Smith et al. 2004; Jenkinson et al. 2012) FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl. The first 6 and last 6 volumes were discarded. Preprocessing included MCFLIRT motion correction (Jenkinson et al. 2002), Fourier-space time-series phase-shifting slice-timing correction, brain extraction (BET; Smith 2002), spatial smoothing (8-mm FWHM Gaussian kernel), grand-mean intensity normalization of all volumes, and highpass temporal filtering (Gaussian-weighted least-squares straight line fitting, $\sigma = 50.0$ s). Functional and MPRAGE images were co-registered with FLIRT (Jenkinson and Smith 2001; Jenkinson et al. 2002) and with the standard Montreal Neurological Institute (MNI) brain using FNIRT nonlinear registration (Andersson et al. 2007a, 2007b).

Statistical Analysis

We used FMRIB's Improved Linear Model (FILM) and local autocorrelation correction (Woolrich et al. 2001) to perform general linear modeling of time-series data. Event-related responses corresponding to the onset of taxonomic, thematic, and filler trials were modeled using a double-Gaussian HRF. Responses from incorrect trials (~18% of the data) were included as a variable of no interest; relatedness score and response time (RT) were included as parametric regressors. The model included temporal derivatives. The model employed individual contrasts corresponding to the onset of the taxonomic, thematic, and filler trials, each of which was 2 s in duration (0.5 s fixation + 1.5 s word pair). Initial contrasts compared each condition (taxonomic, thematic, filler) to the time between trials when the stimulus was not present (rest). Each participant's contrast was combined over the two sequential runs from each task in fixed-effects analyses using FLAME (FMRIB's Local Analysis of Mixed Effects; Beckmann et al. 2003; Woolrich et al. 2004; Woolrich 2008). The combined contrasts were analyzed at the group level in mixed effects analysis using FLAME that applied Gaussian Random Field (GRF) theory to correct for multiple comparisons using a cluster-defining threshold of $Z \geq 2.3$ and corrected significance threshold of $P = 0.05$ (Worsley 2001). In light of recent concerns regarding how cluster inference can lead to false positives when using more lenient thresholds, we conducted

additional analyses based on suggestions from the literature for reducing Type 1 errors (Woo et al. 2014; Eklund et al. 2016). Details are provided in the Supplementary Methods 2.

Region-of-Interest Analysis

We performed ROI analyses on activation in a priori motivated cortical areas from Schwartz et al. (2011) and Jackson et al. (2015). We masked ROIs with labels from the subcortical FreeSurfer parcellation of the probabilistic Destrieux Atlas (Destrieux et al. 2010). Based on results of Schwartz et al.'s (2011) lesion study that isolated residual thematic naming errors to AG (BA 39) damage and residual taxonomic naming errors to temporal pole (BA 38) damage, we restricted the ROI analysis to activity in these areas. In addition to mixed effects analyses with these ROIs, we conducted a repeated measures ANOVA on mean signal for the 4 semantic contrasts against fillers. To increase sensitivity for this analysis, we refined our anatomical AG ROI with a meta-analytic map. First, we used NeuroSynth (www.neurosynth.org; Yarkoni et al. 2011) to conduct a meta-analysis across thousands of published fMRI studies to identify peak coordinates from papers that mention the term "semantic" (search performed on 02/28/2018). Next, we used the resulting specificity map (based on 827 studies) to mask our AG ROI. The final refined AG ROI only included voxels ($n = 76$) from the original anatomical ROI with high probability of being related to the term "semantic" (for a similar procedure see Gruber et al. 2014; Ripollés et al. 2016). Individual parameter estimates were then calculated by averaging over voxels in the refined ROI for each of the 4 semantic conditions against filler. We excluded two participants whose parameter estimates exceeded 2 SDs from the mean. Values were submitted to a 2×2 ANOVA to examine effects of word pair (taxonomic/thematic) and judgment (positive/negative). Figure 2 shows the anatomical ROIs and average parameter estimates for each condition against implicit rest.

Results

Behavioral Results

Recall that we identify conditions by the task (Tax or Them) plus positive or negative correct response (P or N), resulting in names such as *TaxP* or *ThemN*. The behavioral data averages are presented in Table 1. A repeated measures ANOVA on

Figure 2. Left: The AG (red) and ATL (blue) ROIs shown on the standard FreeSurfer brain; Right: Parameter estimates of activity (arbitrary units) against implicit rest over all voxels in each ROI by condition and task (L = left, R = right).

Table 1 Average RT and response accuracy by condition

Condition	RT (ms)		Proportion correct	
	M	SD	M	SD
TaxP	982	113	0.84	0.10
ThemP	1002	108	0.84	0.09
TaxN	1055	110	0.77	0.12
ThemN	1091	113	0.75	0.16
Fill1	1101	108	0.80	0.18
Fill2	1104	114	0.88	0.08

Fill1 and Fill2 correspond to filler (unrelated) pairs in the taxonomic and thematic task, respectively.

response accuracy did not reveal a main effect of task ($F(1,15) = 1.06$, $P = 0.32$, $\eta^2 = 0.07$) or word pair ($F(2,30) = 2.04$, $P = 0.15$, $\eta^2 = 0.12$), but there was a significant interaction between the effects of task and word pair ($F(2,30) = 8.40$, $P < 0.001$, $\eta^2 = 0.36$). Negative trials, which required ignoring a presented relation, yielded the numerically lowest average response accuracy. The Tax and Them accuracy scores were very similar in each task, with the exception of higher accuracy for filler trials in the thematic than in the taxonomic task, yielding the significant interaction. We excluded trials with incorrect responses from further analyses.

An ANOVA on RTs did not reveal a significant effect of task ($F(1,15) = 0.17$, $P = 0.68$, $\eta^2 = 0.011$), but there was a significant effect of word pair ($F(2,30) = 37.20$, $P < 0.001$, $\eta^2 = 0.71$) and a significant interaction ($F(2,30) = 43.51$, $P < 0.001$, $\eta^2 = 0.74$). Post-hoc Tukey HSD tests conducted at the $P < 0.05$ significance level revealed that, overall, RTs were slower on thematic than taxonomic trials, and RTs on filler trials were slower than both. Among related conditions, RTs were slower on negative than

positive trials, and on negative thematic than negative taxonomic trials.

Lastly, single-trial analyses showed significant correlations between relatedness score and positive thematic pairs with faster responses for higher relatedness ($r = -0.20$, $P < 0.0001$). In the thematic task, higher relatedness of fillers was correspondingly associated with slower (negative) responses ($r = 0.10$, $P < 0.001$). Other correlations between relatedness score and RT did not reach significance ($P > 0.05$). Note that we removed variance due to RT and relatedness from the fMRI analyses.

Neural Results

Region-of-Interest Results

Contrasts with fillers. Initial ROI analyses contrasted the related items against the unrelated (filler) items. Contrasts with fillers were always with items from the same task such that differences could not be attributed to the task instructions. Second, the fillers and negative related items required the same response such that differences could not be explained by activity for negative judgments in general. Third and most importantly, the fillers were not related, thus, greater activity for positive related items and negative related items will reflect relational processing or inhibition, respectively. Contrasts with negative items tested whether rejecting DOG-SEAL or rejecting HORSE-REINS required more activity than rejecting TOAD-MOP.

The contrasts showed that both thematic and taxonomic negative conditions elicited greater activity than fillers, suggesting that one function reflected in AG activity is the inhibition of irrelevant semantic information, regardless of relation. (The right AG showed a small cluster of significant activation for negative thematic over filler trials. None of the other contrasts with right AG activity were significant.) Contrasts of the positive items against fillers showed that accepting HORSE-REINS elicited

more activity than rejecting TOAD-MOP. Importantly, accepting DOG-SEAL did not elicit more activity than rejecting TOAD-MOP, suggesting that the AG plays a role in representing thematic information in particular. As for the left and right ATL ROIs, none of the contrasts yielded significant clusters. As mentioned in the Method we did not make strong predictions concerning differences in the ATL because the region is difficult to image with fMRI. The remainder of this section focuses on the AG.

Contrasts among related items. Contrasts among related items focused on effects of thematic versus taxonomic items and positive versus negative judgments. Differences between taxonomic and thematic conditions should reflect specialized semantic processing, while differences between negative and positive items should reflect inhibition. Additionally, differences should not be driven by semantic difficulty or psycholinguistic variables because conditions were matched on these measures. Contrasts with positive pairs revealed that accepting thematic pairs (HORSE-REINS) elicited more activity than accepting taxonomic pairs (DOG-SEAL), which suggests that relational processing in the AG is specialized for thematic concepts. Although not significant following cluster correction, both negative thematic and negative taxonomic pairs elicited more activity than positive taxonomic pairs, which further suggest that the AG supports general inhibition and relational processing of thematic information, but not relational processing of taxonomic information. No other contrast yielded significant differences. Figure 3 shows significant activation maps.

To explore whether spatial differences existed between conditions, we conducted an additional ROI analysis to include activity in the AG as well as the supramarginal gyrus and the intraparietal sulcus (IPS). The new label was created by merging our original AG label (*G_pariet_inf-Angular*) with the supramarginal gyrus (*G_pariet_inf-Supramar*) and IPS (*S_intrapariet_and_P_trans*) labels from the same Destrieux et al. (2010) Brain Atlas. This way we could verify whether the peak for inhibition was in fact in the AG—as opposed to IPS, which has been shown to play a role in executive control demands in a variety of tasks (e.g., Humphreys and Lambon Ralph 2014). The results confirmed that significant peak effects were in left AG for the following contrasts: ThemN > Fill, ThemP > Fill, TaxN > Fill. The ThemP > TaxP contrast did not yield significant (corrected) activity. The significant peak voxels were identical to those in the analysis with the original AG ROI. See Supplementary Table 3 for the statistical tests and peak voxel coordinates from both ROI analyses.

To explore the possibility of an interaction between the semantic relation of the word pair (taxonomic/thematic) and the type of judgment (negative/positive), parameter estimates

for each of the 4 contrasts against filler were extracted from the mean signal of the refined AG ROI and submitted to a 2×2 repeated measures ANOVA. Neither factor yielded a significant main effect (word pair: $F(1,13) = 3.81$, $P = 0.073$, $\eta^2 = 0.226$; judgment: ($F(1,13) = 0.47$, $P = 0.504$, $\eta^2 = 0.035$)), but there was a significant interaction between word pair and judgment. ($F(1,13) = 4.78$, $P < 0.05$, $\eta^2 = 0.269$). While negative and positive judgments about thematic pairs yielded similarly stronger responses, negative judgments about taxonomic pairs yielded stronger responses than positive judgments. Supplementary Figure 1 shows the interaction.

In sum, the results of the ROI analysis suggest that greater AG activity is required for accepting and rejecting thematic trials and for rejecting taxonomic trials than for rejecting filler trials or accepting taxonomic trials. This pattern is a mixture of the two hypotheses' predictions, giving evidence of both inhibitory processes (more activation for negative trials) and thematic processing (activation for positive thematic trials).

Whole-Brain Results

We supplemented our ROI analysis with contrasts of activity over the entire brain, to examine the extent to which thematic and inhibitory processes share a neuronal basis. Consistent with the ROI results, both thematic conditions elicited significantly greater left AG activity relative to fillers, while the positive taxonomic condition did not. Inconsistent with the ROI results, left AG activity was not differentiated for negative taxonomic trials relative to filler trials. The whole-brain analysis additionally showed significant clusters in the precuneus and cingulate gyrus for the contrast TaxN > Fill, as well as in the middle frontal gyrus for ThemN > Fill contrast. Figure 4 shows activation maps from the whole-brain analysis and Supplementary Table 4 reports significant clusters.

Finally, conjunction analyses at the whole-brain level were conducted to identify areas generally involved in thematic and inhibitory processing. We employed activity from the contrasts against filler trials to identify areas specifically involved in conceptual inhibition rather than those generally involved in negative judgments. Conjunction analyses were therefore on negative related conditions (ThemN > Fill $\wedge$ TaxN > Fill) and on thematic conditions (ThemN > Fill $\wedge$ ThemP > Fill). The analysis revealed a conjunction for the negative judgments over the AG, cingulate gyrus, and middle frontal gyrus, as well as a conjunction for thematic judgments over AG, cingulate gyrus, and MTG. The conjunction between all 3 conditions > filler included the AG and cingulate gyrus. The Supplementary Figure 2 shows the conjunctions over the left hemisphere (recall that the only significant cluster in the right hemisphere was for the ThemN > Fill effect).

Figure 3. Significant activation in the left AG ROI from contrasts between conditions following cluster correction (voxel inclusion threshold $z = 2.3$, cluster significance threshold $P < 0.05$).

Figure 4. Significant activation from the whole-brain contrasts following cluster correction (voxel inclusion threshold $z = 2.3$, cluster significance threshold $P < 0.05$) shown on the left-hemisphere. Top: Sagittal view; Bottom: Medial view. All clusters remained significant using a $P < 0.01$ threshold with the exception of the AG cluster in the negative thematic > filler contrast.

Discussion

We examined conceptual processing under conditions requiring negative or positive responses to elucidate whether mechanisms in the AG underlie semantic inhibition or retrieval. Our *a priori* motivated ROI analysis focused on the area where damage has been found to predict thematic picture naming errors (Schwartz et al. 2011), and where level of semantic control has been found to modulate activity in neurologically healthy participants (Jackson et al. 2015). Results from our experiment help to integrate two seemingly conflicting accounts of the AG's function, as they indicate AG engagement to (a) reject related negative pairs, as predicted by the inhibition account, and (b) to accept thematically related pairs, as predicted by the thematic account, in the same task. Contrasts in both the whole-brain and ROI analyses with the filler pairs revealed significantly greater activity for negative and positive thematic pairs (rejecting or accepting HORSE-REINS > rejecting TOAD-MOP). Contrasts with positive items revealed significantly greater activity for positive thematic than taxonomic items (accepting HORSE-REINS > accepting DOG-SEAL). Although not significant in the whole-brain analysis, contrasts with fillers in the ROI analysis yielded significantly stronger activity for negative taxonomic items (rejecting DOG-SEAL > rejecting TOAD-MOP). Lastly, contrasts against the positive taxonomic items showed greater activity for both negative taxonomic and thematic items (rejecting DOG-SEAL or rejecting HORSE-REINS > accepting DOG-SEAL), although not significant following cluster correction. None of the contrasts in either analysis revealed significantly greater activity for positive taxonomic items. The ANOVA of signal in areas of the AG related to semantics also found a similar interaction, in which positive taxonomic trials generated less signal than thematic or negative trials.

While we were also interested in replicating taxonomic effects on ATL responses, the ATL is difficult to image because of its susceptibility to signal loss with our procedure (see Patterson et al. 2007). We suspect this is why results of higher-level contrasts of ATL activation were not significant (see Lewis et al. 2015 for a more successful approach using MEG). The majority of the discussion therefore focuses on the AG. Before relating the results to predictions from the thematic and

inhibition account, we consider the stages required for negative and positive semantic judgments.

Stages of Semantic Processing

In the present experiment, participants used thematic and taxonomic criteria to make positive and negative judgments about the relatedness of words. How did they use these criteria and what was the nature of the AG's involvement? As described in the Introduction, knowledge of real-world entities like dogs includes features. While relational processing begins with identification of individual items, subsequent stages may depend on task instructions and the type of semantic relationship. To verify the relatedness of DOG-SEAL, participants might activate conceptual representations of real-world referents, evaluate common features, and then generalize across the features to derive MAMMAL (see Murphy et al. 2012). For featurally dissimilar items like HORSE-REINS, participants might activate knowledge of each item's conventional role, assess whether the roles are complementary, and then retrieve the external representation of a horse-riding scenario. In some cases, thematically related items are directly associated, presumably making the relation easy to retrieve.

How might participants process negative pairs in which the words are related but with the wrong relation for that task? Conceivably, they could activate only correct, task-relevant information, which would preclude the need for inhibition. This seems incompatible with the results (see below). We consider an alternative scenario requiring inhibition based on the claim that thematic relations activate automatically (Sachs et al. 2008; Maguire et al. 2010). Viewing HORSE-REINS may activate knowledge of their complementary relationship, which must then be inhibited in the taxonomic task to evaluate conceptual similarity of the items. Rejecting DOG-SEAL in the thematic task may involve a similar sequence in which they activate their taxonomic category, which must then be inhibited to evaluate thematic relations. This is less certain because the literature does not (to our knowledge) suggest automaticity of taxonomic activation. Indeed, the cognitive literature often characterizes thematic relations as "attractive" wrong answers (Gentner and

Brem 1999), for example, “Thematic relations are intrusive. They are apprehended involuntarily in tasks for which they are irrelevant and even counterproductive” (Estes et al. 2011: 261). Our behavioral results do not reveal whether it was harder to reject taxonomic than filler items in the thematic task, given that responses to fillers were both slower and more accurate.

A different form of inhibition would be to somehow inhibit an entire set of relations, for example, inhibiting thematic information as a whole in the taxonomic task. Such inhibition might not be completely successful, accounting for the difficulty of negative trials. Neurally, this could be accomplished through deactivation of brain areas that compute or represent that kind of information. We next consider the AG’s involvement in semantic processing in light of this analysis.

Evidence for the Thematic Account

Enhanced AG responses for positive and negative thematic items as predicted by the thematic account were identified both in the ROI and whole-brain analyses. Because our thematic pairs did not share features and our taxonomic pairs did not share complementary roles, neither was predicted to link to task-irrelevant information during positive judgments. We therefore interpret greater AG activity for the positive thematic items as indexing semantic processing rather than inhibition. The lack of a significant difference between positive taxonomic and filler items in both analyses also supports the thematic account, although it should be noted that these are different responses and therefore not perfectly comparable.

Evidence for the Inhibition Account

The expectation from the inhibition account for greater AG activity during negative judgments was confirmed—both the ROI and whole-brain analysis showed that negative judgments for related conditions engaged the AG relative to the unrelated conditions. The prediction for greater activity for both negative related items than positive taxonomic items was weakly supported in the whole-brain results, as neither was significant after correction for multiple comparisons. This pattern of results is not consistent with the notion that the entire set of thematic relations was inhibited through AG deactivation. Unlike some previous studies (e.g., Jackson et al. 2015), we did not find AG deactivation at all; our effects were all positive. This could very well be due to procedural differences in the comparison to rest and differences between blocked and event-related designs. However, finding greater activation in negative than in positive trials does not seem consistent with deactivation as a mechanism for this task.

Does the AG Underlie Both Semantic Processing and Inhibition?

The results did not uniformly support either account. First, the prediction from the inhibition account of significant differences between negative and positive taxonomic conditions was not met, nor was the prediction for no significant differences between positive thematic and positive taxonomic conditions. By the same token, some predictions from the thematic account were not met, including greater activity for negative thematic than negative taxonomic pairs. The prediction from the thematic account of no significant difference between negative taxonomic and filler items was also not met. In most cases, the

failures of one view’s predictions correspond to findings of the other view’s predictions.

Our key finding was that both negative judgments and thematic relations engaged the AG, while taxonomic relations per se did not. That is, the AG was more active when rejecting DOG-SEAL and accepting and rejecting HORSE-REINS, but was not more active when accepting DOG-SEAL (compared with fillers). As such, greater activation for negative taxonomic pairs seems related to inhibition rather than semantic processing. Earlier we considered whether participants activated only correct, task-relevant features. If so, AG responses should not have differed between negative judgments about related and filler items because neither pair would have activated task-irrelevant information. Conversely, we interpret the finding of more activity for accepting HORSE-REINS (but not accepting DOG-SEAL) as indicative of semantic processing, since these positive trials do not require inhibition. Finally, results showed that AG areas that were generally active during negative judgments about related items (with the response to unrelated items subtracted out) showed large conjunctions with areas that were generally active for thematic word pairs. This effect is reinforced by contrasts with the filler items, suggesting that the conjunction cannot be attributed to processes underlying negative judgments in general. Based on these results, we propose that the AG supports both semantic inhibition and thematic specialization. We cannot say how broadly this inhibitory role extends, because our study (and most of the studies reviewed in the Introduction) focused on semantic relations. However, some findings suggest that it is involved in attentional control more generally (Whitney et al. 2012; Humphreys and Lambon Ralph 2017).

Recall that Lewis et al. (2015) identified AG involvement in both taxonomic and thematic relations during general relatedness judgments about sequentially presented word pairs (however, that task required participants to verify any relation, not a particular one as in this experiment). As described in the Introduction, involvement of the AG in taxonomic relations (e.g., DOG-SEAL) could have reflected inhibition of thematic relations activated by the first word (e.g., *leash*, *bone*, *collar*) but irrelevant to processing the taxonomic relation. That account seems inconsistent with the present finding that the AG was not more active either when accepting DOG-SEAL or rejecting TOAD-MOP, when the task explicitly required inhibition of thematic associates. A likely explanation for the difference is that Lewis et al. staggered the presentation of word pairs (required for MEG recording), and the present study used simultaneous presentation. With the present procedure, participants may have been unlikely to predict the next word, which was already present on the screen.

Future research should attempt to specify in more detail which aspects of semantic attention or control the AG carries out. Our attentional manipulation had to do with rejecting attractive items that did not match the task requirements. However, other studies have used quite different attentional tasks, such as detecting weak associations or subordinate meanings (Thompson et al. 2017, who also used a task more similar to ours)—see Hoffman et al. (2018) for a review. An important goal is to eventually match a fine-grained analysis of semantic relations and semantic control to neural functions.

We should also point out that participants were significantly slower to reject thematic than taxonomic items, consistent with the idea that thematic relations are particularly “attractive” relations (Estes et al. 2011). That is, saying “no” to thematic pairs is harder than saying “no” to taxonomic pairs. (Recall that our imaging results removed variance due to RTs

prior to analysis.) This difference is not attributable to differences between conditions such as relatedness strength (the variance of which was also removed), familiarity, word length, or other linguistic measures, which were carefully controlled. Moreover, the finding of equivalent RT and accuracy in accepting both thematic and taxonomic items demonstrates that the conditions were effectively balanced. Based solely on the neural results, we can conclude that the AG responds during rejection of taxonomic and thematic items. However, the finding of slower RT when rejecting thematic than taxonomic items suggests a functional difference. Measures of RT may be more sensitive to differences in semantic processing than fMRI.

Role of the ATL in Thematic Relations

Finally, it is worth considering whether our results conflict with arguments that the ATL is also involved in thematic processing (what Jackson et al. 2015 call associative relations). Jackson et al. noted that patients with ATL damage suffer on tests of associative relations like the Camel and Cactus test (devised by Bozeat et al. 2000). They suggest that the concept hub in the ATL connects to thematic as well as featural information. This is likely true. Hoffman et al. (2018) present a recurrent connectionist model that encompasses both taxonomic and associative knowledge in a single network, which is a potential model of the ATL. Our conclusions about AG functions do not contradict such claims. Event knowledge and associations are rich enough that they might involve multiple brain areas. For example, the ATL might represent associations between concepts while the AG represents more detailed relational information about events and scenes, e.g., exactly how a leash interacts with a dog (Price et al. 2015; Williams et al. 2017). Furthermore, harm to the ATL would also disrupt thematic knowledge, because thematic relations rely on taxonomic information (as noted in the Introduction). As information about leashes is degraded, so is the dog-leash connection. In contrast, when the AG is damaged, perhaps it is the specific relation between dog and leash that is lost, even if the component concepts are intact.

Hoffman et al. (2018) report that both semantic dementia and semantic aphasia patients, with damage to the ATL and the frontal-parietal network, respectively, show diminished performance on the Camel and Cactus test. They attribute this to different psychological deficits: damaged conceptual hubs and control processes, respectively. This is compatible with our finding that the AG is involved in inhibitory processing, a type of semantic control. Our results, along with other findings involving the AG reviewed in the introduction, additionally suggest that thematic knowledge is also represented in the AG. This account is speculative but it helps to explain the multiple loci of semantic information that have been identified in lesion and imaging studies.

General Conclusion

The literature has been split on the question of semantic processing and the AG. There is evidence for two different functions. One is that the AG represents thematic relations—or more generally event knowledge of how objects occur together and interact. The other is that the AG serves to select relevant semantic relations and inhibit irrelevant ones. These two claims are not completely incompatible, as both involve the selection (broadly speaking) of particular semantic relations. The existence of multiple studies providing evidence for both accounts already suggests that the AG's semantic functions are

diverse. However, it is difficult to integrate results of all the past work, which has used different measures and paradigms, including naming errors, eyetracking, classification, tests of brain-damaged patients, and electrophysiological and fMRI studies. No doubt some of the differences in results can be explained by details of task requirements.

Our goal was to compare these two accounts within a single experiment and a single task. With our design, the tests of the two accounts were independent, in that we could find evidence for either one. (In some designs, participants are forced to choose between responses, such as thematic and taxonomic groupings, making it impossible to find evidence of both.)

Our chief finding was that semantic judgments requiring either inhibition or thematic processing engaged the AG, suggesting that it underlies both functions. When conditions were equated for relatedness strength and psycholinguistic variables, judgments involving either thematic relations or semantic inhibition engaged this area, which has been simultaneously associated with semantic control (Jackson et al. 2015) and thematic specialization (Schwartz et al. 2011). In sum, we propose that the AG plays a role both in inhibition and thematic activation, depending on the computation required. This might help explain why damage to this area could cause both reduced activation of thematic relations in comprehension tests (Mirman and Graziano 2012)—through loss of thematic relations—and increased production of thematic associates in a naming task (Schwartz et al. 2011)—through loss of semantic inhibition.

Our experiment required that participants give different responses to a given item depending on the task: HORSE-REINS was a positive item under one set of instructions but required a negative response under the other. Keeping track of the current experimental instructions is probably a function of the prefrontal area, as much research in executive control suggests (e.g., Thompson-Schill et al. 1997; Badre et al. 2005). Other research suggests that selecting relevant features in the activated classification task is carried out in the AG or nearby areas. Hanson and Chrysikou (2017) recently proposed that posterior parietal areas are involved in semantic selection. Using a classification task that varied the feature on which items were to be grouped, they discovered that left inferior parietal areas were active in choosing between two concrete features as well as two abstract features (less strongly). In our task, the prefrontal cortex presumably interacted with posterior areas to ensure that the correct kinds of features are used. Future research should examine the role of frontal-parietal circuits in semantic flexibility. People are extraordinarily good at noticing connections of all kinds between items, but they are also very good at ignoring the irrelevant ones. The AG seems to be involved in both of these abilities.

Supplementary Material

Supplementary material is available at *Cerebral Cortex* online.

Funding

This work was supported by National Institutes of Health Grant NIH 2R01DC05660 to D.P.

Notes

We thank Pablo Ripollés for assistance with the masked ROI analysis. *Conflict of Interest*: None declared.

References

- Andersson JLR, Jenkinson M, Smith S 2007a. Non-linear optimisation. FMRIB Analysis Group. Report No.: TR07JA1. Available at: www.fmrib.ox.ac.uk/analysis/techrep/
- Andersson JLR, Jenkinson M, Smith S 2007b. Non-linear registration, aka spatial normalization. FMRIB Analysis Group. Report No.: TR07JA2. Available at: www.fmrib.ox.ac.uk/analysis/techrep/
- Badre D, Poldrack RA, Paré-Blagoev EJ, Insler RZ, Wagner AD. 2005. Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. *Neuron*. 47:907–918.
- Battig WF, Montague WE. 1969. Category norms of verbal items in 56 categories A replication and extension of the Connecticut category norms. *J Exp Psychol G*. 80(3p2):1.
- Beckmann CF, Jenkinson M, Smith SM. 2003. General multilevel linear modeling for group analysis in fMRI. *Neuroimage*. 20: 1052–1063.
- Bozeat S, Lambon Ralph MA, Patterson K, Garrod P, Hodges JR. 2000. Non-verbal semantic impairment in semantic dementia. *Neuropsychologia*. 38:1207–1215.
- Davey J, Rueschemeyer SA, Costigan A, Murphy N, Krieger-Redwood K, Hallam G, Jefferies E. 2015. Shared neural processes support semantic control and action understanding. *Brain Lang*. 142:24–35.
- Davey J, Thompson HE, Hallam G, Karapanagiotidis T, Murphy C, De Caso I, Jefferies E. 2016. Exploring the role of the posterior middle temporal gyrus in semantic cognition: integration of anterior temporal lobe with executive processes. *Neuroimage*. 137:165–177.
- de Zubicaray GI, Hansen S, McMahon KL. 2013. Differential processing of thematic and categorical conceptual relations in spoken word production. *J Exp Psychol Gen*. 142:131–142.
- Destrieux C, Fischl B, Dale A, Halgren E. 2010. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage*. 53:42750.
- Eklund A, Nichols TE, Knutsson H. 2016. Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. *Proc Natl Acad Sci U S A*. 113:7900–7905.
- Estes Z, Golonka S, Jones LL. 2011. Thematic thinking: the apprehension and consequences of thematic relations. In: Ross BH, editor. *Psychology of learning and motivation: advances in research and theory*. Vol. 54. San Diego (CA): Elsevier Academic Press. p. 249–294.
- Geng J, Schnur TT. 2016. Role of features and categories in the organization of object knowledge: evidence from adaptation fMRI. *Cortex*. 78:174–194.
- Gentner D, Brem SK 1999. Is snow really like shovel? Distinguishing similarity from thematic relatedness. *Proceedings of the 21st annual meeting of the Cognitive Science Society*. Mahwah (NJ): Erlbaum. p. 179–184.
- Gruber MJ, Gelman BD, Ranganath C. 2014. States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*. 84:486–496.
- Hanson GK, Chrysikou EG. 2017. Attention to distinct goal-relevant features differentially guides semantic knowledge retrieval. *J Cogn Neurosci*. 29:1178–1193.
- Hoffman P, McClelland JL, Lambon Ralph MA. 2018. Concepts, control, and context: a connectionist account of normal and disordered semantic cognition. *Psychol Rev*. 125(3):293–328.
- Humphreys GF, Lambon Ralph MA. 2014. Fusion and fission of cognitive functions in the human parietal cortex. *Cereb Cortex*. 25:3547–3560.
- Humphreys GF, Lambon Ralph MA. 2017. Mapping domain-selective and counterpointed domain-general higher cognitive functions in the lateral parietal cortex: evidence from fMRI comparisons of difficulty-varying semantic versus visuo-spatial tasks, and functional connectivity analyses. *Cereb Cortex*. 27:4199–4212.
- Jackson RL, Hoffman P, Pobric G, Lambon Ralph MA. 2015. The nature and neural correlates of semantic association versus conceptual similarity. *Cereb Cortex*. 25:4319–4333.
- Jefferies E, Lambon Ralph MA. 2006. Semantic impairment in stroke aphasia versus semantic dementia: a case-series comparison. *Brain*. 129:2132–2147.
- Jenkinson M, Bannister P, Brady M, Smith S. 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*. 17:825–841.
- Jenkinson M, Beckmann CF, Behrens TE, Woolrich MW, Smith SM. 2012. Fsl. *Neuroimage*. 62:782–790.
- Jenkinson M, Smith S. 2001. A global optimisation method for robust affine registration of brain images. *Med Image Anal*. 5:143–156.
- Kalénine S, Peyrin C, Pichat C, Segebarth C, Bonthoux F, Baciuc M. 2009. The sensory-motor specificity of taxonomic and thematic conceptual relations: a behavioral and fMRI study. *Neuroimage*. 44:1152–1162.
- Lambon Ralph MA, Jefferies E, Patterson K, Rogers TT. 2016. The neural and computational bases of semantic cognition. *Nat Rev Neurosci*. 18:42–55.
- Lambon Ralph MA, Sage K, Jones RW, Mayberry EJ. 2010. Coherent concepts are computed in the anterior temporal lobes. *Proc Natl Acad Sci U S A*. 107:2717–2722.
- Lewis G, Poeppel D, Murphy GL. 2015. The neural bases of taxonomic and thematic conceptual relations: an MEG study. *Neuropsychologia*. 68:176–189.
- Lin EL, Murphy GL. 2001. Thematic relations in adults' concepts. *J Exp Psychol Gen*. 130:3–28.
- Lund K, Burgess C. 1996. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behav Res Methods Instrum Comput*. 28:203–208.
- Maguire MJ, Brier MR, Ferree TC. 2010. EEG theta and alpha responses reveal qualitative differences in processing taxonomic versus thematic semantic relationships. *Brain Lang*. 114:16–25.
- Markman EM. 1989. *Categorization and naming in children: problems of induction*. Cambridge (MA): MIT Press.
- Mirman D, Graziano KM. 2012. Damage to temporo-parietal cortex decreases incidental activation of thematic relations during spoken word comprehension. *Neuropsychologia*. 50: 1990–1997.
- Mirman D, Landrigan JF, Britt AE. 2017. Taxonomic and thematic semantic systems. *Psychol Bull*. 143:499–520.
- Murphy GL. 2001. Causes of taxonomic sorting by adults: a test of the thematic-to-taxonomic shift. *Psychon Bull Rev*. 8: 834–839.
- Murphy GL, Hampton JA, Milovanovic GS. 2012. Semantic memory redux: an experimental test of hierarchical category representation. *J Mem Lang*. 67:521–539.
- Nelson DL, McEvoy CL, Schreiber TA. 2004. *The University of South Florida free association, rhyme, and word fragment norms*. *Behav Res Methods Instrum Comput*. 36:402–407.
- Noonan KA, Jeffries E, Corbett F, Lambon Ralph MA. 2010. Elucidating the nature of deregulated semantic cognition in semantic aphasia: evidence for the roles of prefrontal and temporo-parietal cortices. *J Cogn Neurosci*. 22:1597–1613.

- Patterson K, Nestor P, Rogers TT. 2007. Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat Rev Neurosci*. 8:976–987.
- Price AR, Bonner MF, Peele JE, Grossman M. 2015. Converging evidence for the neuroanatomic basis of combinatorial semantics in the angular gyrus. *J Neurosci*. 35:3276–3284.
- Ripollés P, Marco-Pallares J, Alicart H, Tempelmann C, Rodriguez-Fornells A, Noesselt T. 2016. Intrinsic monitoring of learning success facilitates memory encoding via the activation of the SN/VTA-Hippocampal loop. *Elife*. 5:e17441.
- Sachs O, Weis S, Zellagui N, Huber W, Zvyagintsev M, Mathiak K, Kircher T. 2008. Automatic processing of semantic relations in fMRI: neural activation during semantic priming of taxonomic and thematic categories. *Brain Res*. 1218:194–205.
- Schwartz MF, Kimberg DY, Walker GM, Brecher A, Faseyitan OK, Dell GS, Mirman D, Coslett HB. 2011. Neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. *Proc Natl Acad Sci U S A*. 108:8520–8524.
- Smith SM. 2002. Fast robust automated brain extraction. *Hum Brain Mapp*. 17:143–155.
- Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Niazzy RK. 2004. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*. 23:S208–S219.
- Thompson H, Davey J, Hoffman P, Hallam G, Kosinski R, Howkins S, Wooffindin E, Gabbitas R, Jefferies E. 2017. Semantic control deficits impairs understanding of thematic relationships more than object identity. *Neuropsychologia*. 104:113–125.
- Thompson-Schill SL, D'Esposito M, Aguirre GK, Farah MJ. 1997. Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. *Proc Natl Acad Sci U S A*. 94:14792–14797.
- Visser M, Jefferies E, Lambon Ralph MA. 2010. Semantic processes in the anterior temporal lobes: a meta-analysis of the functional neuroimaging literature. *J Cogn Neurosci*. 22:1083–1094.
- Whitney C, Kirk M, O'Sullivan J, Lambon Ralph MA, Jefferies E. 2012. Executive semantic processing is underpinned by a large-scale neural network: revealing the contribution of left prefrontal, posterior temporal, and parietal cortex to controlled retrieval and selection using TMS. *J Cogn Neurosci*. 24:133–147.
- Williams A, Reddigari S, Pykkänen L. 2017. Early sensitivity of left perisylvian cortex to relationality in nouns and verbs. *Neuropsychologia*. 100:131–143.
- Woo CW, Krishnan A, Wager T. 2014. Cluster extent based thresholding in fMRI analyses: pitfalls and recommendations. *Neuroimage*. 91:412–419.
- Woolrich MW. 2008. Robust group analysis using outlier inference. *Neuroimage*. 41:286–301.
- Woolrich MW, Behrens TE, Beckmann CF, Jenkinson M, Smith SM. 2004. Multilevel linear modelling for FMRI group analysis using Bayesian inference. *Neuroimage*. 21:1732–1747.
- Woolrich MW, Ripley BD, Brady M, Smith SM. 2001. Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage*. 14:1370–1386.
- Worsley KJ. 2001. Statistical analysis of activation images. In: Jefferies E, Matthews PM, Smith SM, editors. *Functional MRI: an introduction to methods*. New York (NY): Oxford University Press. p. 251–270.
- Xu Y, Wang X, Wang X, Men W, Gao J-H, Bi Y. 2018. Doctor, teacher, and stethoscope: neural representation of different types of semantic relations. *J Neurosci*. 38:3303–3317.
- Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD. 2011. Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods*. 8:665–670.