Neural Mechanisms of Conceptual Relations

by

Gwyneth A. Lewis

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Psychology
New York University
January, 2017

David Poeppel and Gregory L. Murphy

ACKNOWLEDGMENTS

First, I would like to thank my advisors, David Poeppel and Gregory Murphy. David, thank you for your wisdom and for mentoring me through every aspect of my graduate career. I will always be grateful for all the extra support and encouragement that you gave me. Greg, thank you for being unfailingly generous with your time, your advice, and your dry humor. I'm very thankful to my committee members–Alec Marantz, Jon Freeman, and Marjorie Rhodes. Alec, thank you especially for your support and advice throughout my time as a lab manager and later as a PhD student. I'm thankful to the Poeppel Lab for years of good memories—Luc Arnal, Florencia Assaneo, Amy Belfi, Greg Cogan, Ido Davidesco, Suzanne Dikker, Nai Ding, Keith Doelling, Adeen Flinker, Saskia Haegens, Anna Kasdan, Taehwan Kim, Jean-Remi King, Ning Mei, Lucia Melloni, Johanna Rimmelle, Jess Rowland, Xiangbin Teng, Xing Tian, Jeffrey Walker, Jean Zarate, and Elana Zion-Golumbic. Suz, thanks especially for the extra moral support in the weeks leading up to my defense. I'm thankful to the Marantz Lab, especially Meera Al-Kaabi, Kevin Chun-Hsien Hsu, Allyson Ettinger, Joseph Fruchter, Laura Gwilliams, Tal Linzen, and Olla Solomyak. Thanks also to the Pylkkänen lab, especially Doug Bemis, Christian Brodbeck, Paul del Prato, and Rebecca Egbert. Thanks to Keith Sanzenbach for training me to run fMRI participants, and to Jeffrey Walker for helping me run MEG participants. Lastly, I'm thankful to my NYU professors for their invaluable instruction over the years, especially David Heeger, Brian McElree, and Liina Pylkkänen

The work in this dissertation was supported by the National Science Foundation (grants DGE-1342536 and BCS-1128769) and the National Institute of Health (grant 2R01DC05660).

ABSTRACT

An over-arching goal in neurolinguistic research is to characterize the neural bases of semantic representation. A particularly relevant goal concerns whether we represent features and events (a) together in a generalized semantic hub or (b) separately in distinct but complementary systems. While the left anterior temporal lobe (ATL) is strongly emphasized in representing both feature-based (taxonomic) knowledge and event-based (thematic) knowledge, recent evidence suggests that the temporal parietal junction (TPJ) plays a unique role in thematic semantics. The primary goal of this dissertation was to identify and characterize the neural mechanisms that support taxonomic and thematic semantics, and the general goal was to shed further light on neural stages of word comprehension. We conducted two magnetoencephalography (MEG) experiments to identify neural indices of visual representations (Chapter 1) and to examine ATL vs. TPJ involvement in taxonomic and thematic semantics (Chapter 2), respectively. We also conducted a functional magnetic resonance imaging (fMRI) experiment to characterize the role of the TPJ in thematic inhibition vs. thematic semantics (Chapter 3). The three experiments employed semantic judgment tasks, equated stimulus conditions on linguistic and psycholinguistic variables, and supplemented analyses with continuous variables as more sensitive hallmarks of lexical access.

Chapter 1 demonstrated that initial stages of spoken word recognition involve contact with visual representations of features associated with the real-world referent. The early timing of the effect suggests that sensory aspects of meaning are not necessarily a product of lexical activation during speech recognition. Chapter 2 demonstrated ATL selectivity for taxonomic relations, and moderate TPJ selectivity for both taxonomic and thematic relations. Results for the TPJ could reflect either inhibition of irrelevant information or conceptual processing. Chapter 3 tested these possibilities by requiring inhibition of the opposite relation in two semantic judgment tasks. Results of this experiment indicate that the TPJ plays a role both in thematic semantics *and* in inhibitory processing when the conceptual computation requires it.

In sum, this dissertation focuses on topics pertaining to neural encoding of words with respect to form and meaning. Across three neurolinguistic experiments, we addressed (1) contributions of visual representations during lexical access, (2) ATL and TPJ selectivity for thematic vs. taxonomic concepts, and (3) TPJ inhibition vs. specialization for thematic concepts, respectively.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

Chapter 2

Chapter 3

LIST OF TABLES

Chapter 1

Chapter 2

Chapter 3

LIST OF APPENDICES

GENERAL INTRODUCTION

Results from neurolinguistic research show that word comprehension recruits distinct brain areas during orthographic, morphological, and semantic processing stages, respectively (e.g., Lewis, Solomyak, & Marantz, 2011; Simon, Lewis, & Marantz, 2012). While word recognition has been argued to involve contact with lexical representations in the posterior middle temporal gyrus (pMTG) (e.g., Hickok & Poeppel, 2007), the temporal dynamics of activation, competition, and access stages of sound-meaning representations are not well understood. One particularly nebulous issue concerns the hypothesized role of perceptual representations in meaning-based resolution. For example, at what point in recognition do words activate imagery of the real-world referent? When "strawberry" activates the color red, does this activation contribute to recognition or is it merely as a by-product of lexical access?

The focus of Chapter 1 is on the temporal role of visual cortices and the pMTG in lexical access. We conducted an MEG experiment to examine the timing effects of perceptual, phonemic, and lexical variables on cortical responses during lexical decision stages. Results revealed that word imageability (an index of the extent to which a particular meaning can be perceptually experienced) modulated visual cortices early in recognition, while written word frequency modulated pMTG responses later in recognition. Based on the timing of effects of these among other linguistic variables, we propose a flowchart of sublexical and lexical stages of spoken word recognition in which sensory aspects of meaning become activated prior to selection of the lexical representation.

Another issue of primary focus in this dissertation concerns how we process different forms of meaning. When we extract the meaning of a word like DOG, we also extract representations of its features and attributes, like *cute, furry, has a tail*, etc. We use these representations to group objects with overlapping features and attributes into *taxonomic* categories, which formally refer to groupings based on overlapping features, attributes, or meanings across two or more items from the same category that do not necessarily interact or co-occur, like DOG and SEAL) (Estes, Golonka, & Jones, 2012; Lin & Murphy, 2001). In addition to features, we also extract representations of contextually relevant objects. We use these representations to group objects that tend to co-occur or interact in the same event into *thematic* categories, which formally refer to groupings based on complementary spatial or functional relations between categorically distinct objects, like DOG and LEASH. Whether taxonomic and thematic knowledge are represented in distinct semantic systems remains an open question for semantic memory researchers. On the one hand, thematic and taxonomic relations have very different properties in that thematic relations are external, complementary, and link dissimilar objects, whereas taxonomic relations are internal, independent, and link objects based on feature similarity (Estes et al., 2012). On the other hand, it is conceivable that the same semantic system encodes 'internal' features of objects in the 'external' events they occur in.

Such questions pertain to wider debates about semantic dissociation across disciplines including philosophy, linguistics, and neuroscience. Within cognitive neuroscience, researchers have investigated the underlying mechanisms of taxonomic and thematic relations with a range of neuroimaging techniques. Lesion studies have linked thematic and taxonomic conceptual deficits to distinct lesion locations (Schwartz et al., 2011) and studies of neurologically intact subjects have reported dissociated neural correlates of taxonomic and thematic relations with EEG (e.g., Maguire, Brier, & Ferree, 2010), fMRI (e.g., Zubicaray et al., 2013), MEG (e.g., Lewis, Poeppel, & Murphy, 2015), and TMS (Davey et al., 2015). Findings from such studies generally converge to suggest that the ATL supports taxonomic concepts, and (to a weaker extent) that the TPJ represents thematic concepts.

The primary focus of Chapter 2 is on the putative role of the ATL and TPJ in taxonomic and thematic relations, respectively. We conducted an MEG experiment with a semantic judgment paradigm to contrast effects of taxonomically, thematically, and unrelated (filler) word pairs on ATL and TPJ activity. We found that while the ATL showed selective involvement in taxonomic relations, the TPJ showed involvement in both kinds of relations. The primary goal in Chapter 3 was to build on results from Chapter 2 to elucidate the role of the TPJ in conceptual processing. We conducted an fMRI experiment with two kinds of semantic judgment tasks to examine what kinds of processes accompany conceptual relations, such as relational inhibition. Inhibition was measured by requiring either negative or positive responses to taxonomic and thematic word pairs. In one task, participants made semantic judgments based on taxonomic criteria of relatedness, which required inhibition of thematic relations. In a second task, participants made judgments based on thematic criteria, which required inhibition of taxonomic relations. In sum, Chapters 2 and 3 focused on ATL and TPJ responses to taxonomically vs. thematically related word pairs to shed light on the neural mechanisms of conceptual relations.

CHAPTER 1: THE ROLE OF VISUAL REPRESENTATIONS DURING THE LEXICAL ACCESS OF SPOKEN WORDS[1]

Gwyneth Lewis and David Poeppel

1.1 Abstract

Do visual representations contribute to spoken word recognition? We examine, using MEG, the effects of sublexical and lexical variables at superior temporal (ST) areas and the posterior middle temporal gyrus (pMTG) compared with that of word imageability at visual cortices. Embodied accounts predict early modulation of visual areas by imageability – concurrently with or prior to modulation of pMTG by lexical variables. Participants responded to speech stimuli varying continuously in imageability during lexical decision with simultaneous MEG recording. We employed the linguistic variables in a new type of correlational time course analysis to assess trial-by-trial activation in occipital, ST, and pMTG regions of interest (ROIs). The linguistic variables modulated the ROIs during different time windows. Critically, visual regions reflected an imageability effect prior to effects of lexicality on pMTG. This surprising effect supports a view on which sensory aspects of a lexical item are not a consequence of lexical activation.

1.2 Introduction

Models of Lexical Processing. Neurocognitive models of lexical access suggest the participation of distinct neural regions in the activation of, competition between, and selection of

---

[1] Published as: Lewis, G., & Poeppel, D. (2014). The role of visual representations during the lexical access of spoken words. *Brain and Language*, *134*, 1-10.

basic sound-meaning representations. In the visual domain, studies of lexical access have determined that certain MEG components are sensitive to the orthographic, morphological, and semantic features of words during different time windows. For example, the frequency of a word’s adjacent letter strings (bigram frequency) modulates occipital activation at ~100 ms post word onset (Solomyak & Marantz, 2010). Around 50 ms later, the morphological transition probability of words (the probability of the whole word form, given the stem) modulates responses in the fusiform gyrus (Lewis, Solomyak, & Marantz, 2011; Solomyak & Marantz, 2010). At ~300 ms, properties of the whole word form modulate a superior temporal response (Lewis, Solomyak, & Marantz, 2011; Pylkkänen & Marantz, 2003; Simon, Lewis, & Marantz; Solomyak & Marantz, 2009; Solomyak and Marantz, 2010). How do the stages of spoken word recognition compare with those involved in visual word recognition?

Embodied perspectives on language processing focus less on activation and competition and more on the role of perceptual (or sensorimotor) representations in, for example, (lexical) semantic access. Strong theories of embodiment view semantic knowledge as grounded in perceptual experience rather than in the relationships between words (Bickhard, 2008). Semantic access is thought to require perceptual simulation and directly engage areas of the brain that are active while perceiving the referent in the real world (Gallese & Lakoff, 2005). Weak-embodiment theories view lexical-semantic access as only moderately dependent on the participation of sensory and motor systems. On such models, semantics may be grounded in sensory and motor information but may also be accessed from higher-level representations (Meteyard & Vigliocco, 2008). In opposition to embodied-based accounts, abstract, symbolic theories view semantic knowledge as derived from correspondences between internal symbols and their extensions to objects in the real world (Fodor & Pylyshyn, 1988). New data could shed light on these theories and disambiguate among some of these predictions.

Recent empirical findings. Results from fMRI studies of visual perception and mental imagery suggest that the same occipital regions active while perceiving objects are similarly active while mentally 'simulating' visual images of objects (Ganis et al., 2004). Evidence that occipital (visual) regions are involved in simulating perceptual visual features during language comprehension also comes from recent fMRI experiments. One study showed, for example, that occipital regions processed shape information of sounds, wherein the stimulus impact sound of an object (such as a ball bouncing) modulated occipital activation when the hearer's instructions were to name the shape (e.g., *round*) rather than the material (e.g., *rubber*) of the object generating the sound (James et al., 2011).

Motivation of the current experiment. Based on such findings, we assume that the visual cortex is at least possibly active during spoken word recognition. Whether and when such activation contributes to meaning-based resolution remains controversial. In previous work, we found that the meaning-based resolution of *visual* words can be verified at around 300 ms post-stimulus onset (Simon, Lewis, & Marantz, 2012). This is reflected in the modulation of a superior temporal response (the MEG M350, comparable to the N400/N400m of Helenius et al., 2007) by the meaning-entropy (semantic ambiguity) of visually presented words. An absence of earlier semantic effects does *not* mean, however, that lexical resolution (selection of the appropriate representation) does not begin much earlier. But can one diagnose lexical resolution and perceptual simulation at earlier stages, and on which brain regions should one focus?

Of particular interest, in the functional anatomic sense, is the posterior middle temporal gyrus (pMTG), which previous work implicates as an indicator of lexical access in spoken word recognition (Hickok & Poeppel, 2007). While traditional accounts of verbal comprehension emphasize the role of Wernicke's (superior temporal) area in speech processing, there is considerable evidence that the MTG plays a central role in lexical processing (see reviews in, e.g., Dronkers et al. 2004; Hickok & Poeppel, 2007; Lau et al. 2008). Evidence that the pMTG is a critical node in the language comprehension network comes from lesion studies that find that, compared with patients with lesions in Wernicke's (superior temporal) and Broca's areas, patients with lesions to pMTG demonstrate poor performance in comprehending and naming single words. The pMTG may therefore link conceptual information to lexical representations (Dronkers et al., 2004). Further evidence comes from a study of connectivity profiles of brain areas within the language comprehension network, which determined that the MTG connectivity pattern is extensively integrated with areas of the network previously found to be critical to sentence comprehension (Turken & Dronkers, 2011). Results from neuroimaging show that MTG activation increases as a function of speech intelligibility (Davis & Johnsrude, 2003) and is also modulated by increasing semantic ambiguity (Rodd, Davis, & Johnsrude, 2005). The MTG therefore provides an ideal testing ground for the study of the processes leading to lexical access of spoken words.

To identify lexical access, we examine responses of the pMTG to phonological neighborhood density (ND), which is based on the number of words that differ from the word by one phoneme (Luce & Pisoni, 1998). We employed this metric based on the theory that words are organized based on their phonetic similarity to other words. As an example, the word *bat* has a dense neighborhood because it phonetically resembles many other words such as bate, ban, bit, etc. We recognize words within dense neighborhoods more slowly because high neighborhood density words activate more lexical representations than do low density items, which entails greater competition among entries (Vitevitch & Luce, 1998). We also examine responses to whole word form (surface) frequency (SF), which has previously been shown to modulate middle and superior temporal areas that are involved in lexical access during the later stages of word recognition (Lewis, Solomyak, & Marantz, 2011; Simon, Lewis, & Marantz, 2012). We also employ concreteness, which is another metric of the sensory information attached to the word. Concrete words (e.g., *apple*) are thought to be more easily encoded and retrieved than abstract words (e.g. *freedom*). For example, concrete words are recognized significantly faster and omitted significantly less from recall memory (Holmes & Langford, 1976). Concrete words additionally induce more negative N400s, which may be because concrete words evoke more sensory information attached to the representation (West & Holcomb, 2000). While concreteness and imageability strongly correlate, we include concreteness as a measure because imageability ratings are based only on visual aspects of the item (e.g., instructions require the rater to evoke a mental image of the item (Paivo et al., 1968).

1.3 Method

Stimuli. We accessed all 1,324 monosyllabic nouns with imageability ratings from the MRC Psycholinguistic Database (Coltheart, 1981). The exclusion procedure removed items with the following characteristics:

- Multiple Part-of-Speech (POS) classes (e.g., yawn that is both a noun and a verb) based on the coding in the English Lexicon Project (ELP) (Balota et al., 2007).

- More than one morpheme, as coded by the ELP.
- Lexical decision accuracy below 70%, as coded by the ELP.
- Shared phonology with orthographically different word(s) (e.g., cent and scent), as based on homophony coding in the program Neighborhood Watch (NW). Because imageability rating tasks are based on visual words, it was necessary to exclude orthographically different items with identical phonologies to ensure that the subject accessed the correct meaning of the word.
- Heteronymy, where orthographically identical words have multiple meanings but different phonologies (e.g., sow, which refers to a female pig or the act of planting by seed), which was determined by accessing the number of dictionary headwords from the Wordsmyth Online Dictionary (Parks, Kennedy, & Broquist, 1998). While imageability ratings tasks make explicit the part of speech class an item belongs to, there are some instances where words have multiple meanings under the same speech class (e.g., the noun yard may denote a unit of measurement or an area of ground). We also excluded homonyms, where orthographically identical words have multiple meanings and the same phonology (e.g., bank, which might refer to a river bank or a financial institution), based on the number of dictionary headwords. We did not remove items where the alternative meaning was obscure, as in mare, which refers obscurely to “a large flat dark area on the moon or Mars.”
- Items outside the range of 3-4 phonemes (83% were within this range).
- Unusual consonant-vowel sequences.

The exclusion procedure reduced the set to 287 items. To increase the stimulus set, we included all 423 nouns not already in the MRC set from the Cortese & Fugett (2004) corpus, which includes imageability ratings for 3,000 monosyllabic words. After applying the same exclusion criteria to these items, we further winnowed down this new list to just 113 items (closely matched to items from the other corpus in terms of phonemic and orthographic frequency and length), for a total of 400 total target items. Imageability ratings in both corpora fall between 100 (lowest-imageability) and 700 (highest-imageability). For example, the imageability of *whim* = 180, *hag* = 400, and *goose* = 690. Our stimulus imageability ranged from 160-690.

Variables. One way of isolating acoustic, phonemic, lexical, and visual effects is to employ the variables in a regression model. By regressing each variable onto the competing variables, one can ensure that any early cortical activation in visual cortices cannot be attributed to phonemic, phonetic, or lexical effects. The variables are described below.

Cohort Competition. This was the ratio of a word's CELEX frequency to its cohort-size, multiplied by 100, so that lower values indicate greater competition (as calculated in Zhuang et al., 2011)

Phonological neighborhood density. This measures the density of phonological neighbors that differ from the word by one phoneme (Luce & Pisoni, 1998). Measures were accessed from the English Lexicon Project (Balota et al., 2007).

Imageability. The measure is based on ratings from the MRC online database and the Cortese and Fugett corpus.

Variable	Mean	SD
Biphone Frequency	2.50	0.66
Cohort Competition	6.85	18.68
Concreteness	527.28	95.04
Duration (ms)	371.82	48.37
Entropy	4.01	1.24
Imageability	493.71	107.76
Length	4.41	0.83
Number of Phonemes	3.51	0.50
Phonological Density	13.76	8.56
Surface Frequency	7.97	1.61

Table 1: Properties for the target items.

Analysis. We followed the same procedure for MEG data processing for source space analyses described in Lewis et al. (2011). Noise reduction with software MEG160 (Yokogawa Electric Corporation and Eagle Technology Corporation, Tokyo, Japan) and data from three MEG reference sensors involved the Continuously Adjusted Least-Squares Method (Calm; Adachi, Shimogawara, Higuchi, Haruta, & Ochiai, 2001). Further processing of the noise reduced data was in MNE (MGH/HMS/MIT Athinoula A. Martinos Center for Biomedical Imaging, Charleston, MA). We reconstructed each subject's structural MRI using FreeSurfer routines (CorTechs Lab Inc., La Jolla, CA). Reconstructions were used to estimate cortically constrained minimum-norm solutions of the MEG data. The forward solution (magnetic field estimates at each MEG sensor) was estimated from a source space of 5124 activity points with a boundary-element model (BEM) method. The inverse solution (an estimate of the temporal and spatial distribution of the MEG data) was calculated from the forward solution. Data was then converted into dynamic statistical parameter map (dSPM) values (Dale, et al., 2000).

Regions of interest. We morphed each subject's cortex to a standard FreeSurfer brain to visualize grand average activation across subjects. We defined the occipital ROI around visible peak activation in the occipital area. Anatomical FreeSurfer labels, including the superior temporal gyrus (STG), superior temporal sulcus (STS), and middle temporal gyrus (MTG), constrained selection of ROIs based on peaks in grand average activation. Figure 1 shows the ROIs.

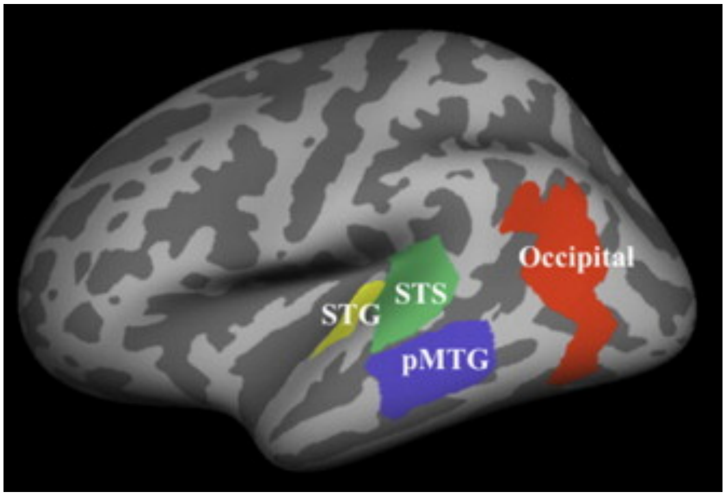

Figure 1: Functionally defined anatomically constrained ROIs on the inflated brain (STS = superior temporal sulcus, STG = superior temporal gyrus, pMTG = posterior middle temporal gyrus). Data within ROIs were employed in correlational analyses with stimulus variables. The ROIs include, roughly, BA areas 42, 22, 21, and 19 (STG, STS, pMTG, and Occipital, respectively).

MNE routines morphed labels back to individual subject brains, and grand average ROI activation within each subject's label was employed in trial-by-trial correlational analyses with the stimulus variables (including imageability, biphone frequency, cohort competition, neighborhood density, entropy, and surface frequency). We focused on the left hemisphere because neurophysiological evidence suggests that speech and language perception is lateralized here, however, we do acknowledge that this may depend on the technique, as hemodynamic and electrophysiological imaging data has indicated that processing may be more bilateral in nature (Price, 2012; Schirmer, Fox, & Grandjean, 2012; Turkeltaub & Coslett, 2010).

Time-course analysis. Our analysis examined effects of the stimulus variables on millisecond level neural activation as the speech played. Specifically, we correlated millisecond level activation within the ROIs with the various stimulus variables. A multiple comparisons correction (Maris & Oostenveld, 2007) was performed on temporal clusters of the point-by-point regressions that were significant prior to correction at the $p < .05$ significance level. An $\sum r$ statistic was constructed by summing coefficients of temporally continuous effects. We tested the significance of the statistic with Monte-Carlo *p*-values. First, we computed a correlation wave by permuting the random variable 10,000 times and then calculated the $\sum r$ statistic at significant clusters at each of the 10,000 permutations. A distribution of $\sum r$ values was constructed from the highest $\sum r$ value at the individual permutations. We defined our Monte-Carlo *p*-value based on the ratio of new values that were higher than the initial statistic.

1.4 Results

Neural data. We report here the significant findings. Each subject displayed the typical auditory M100 response. Contour maps and the grand average waveform for all subjects and trials of the raw MEG data at the M100 response are shown in Figure 2. Early occipital activation was primarily positive (outgoing from the cortex), while peak activation within the STG, STS, and MTG labels was negative (ingoing toward the cortex). Figure 2 also presents the labels along with the grand average time courses of activation.

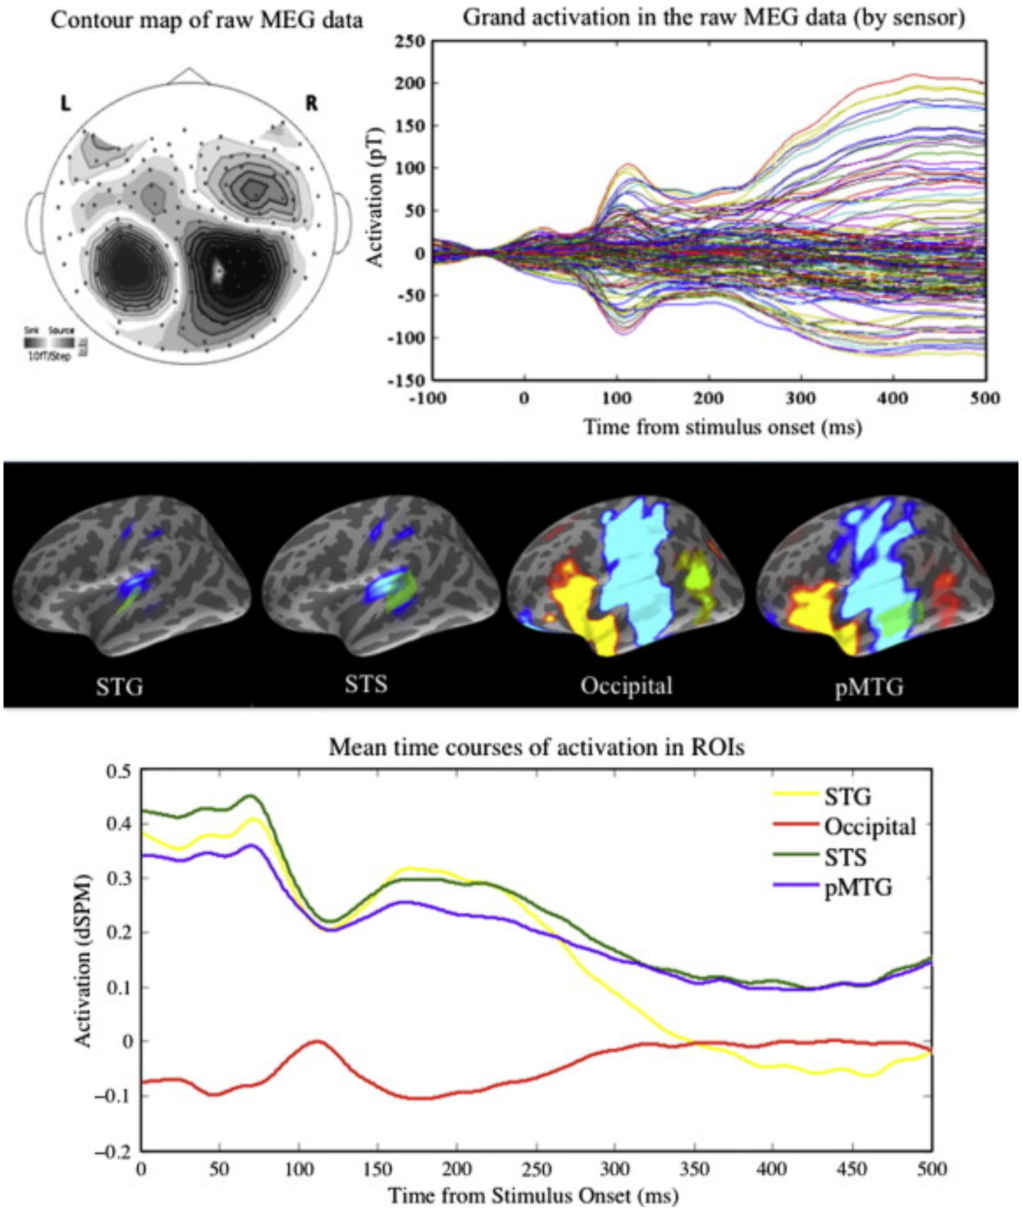

Figure 2: Contour maps and waveform views of the grand average MEG data. Top left: Contour map of the grand average auditory M100 component. Top right: Grand average waveforms of the raw MEG data by sensor averaged over all trials and subjects. Center: Regions of interest (faint

green blobs) and average activation displayed on the standard inflated brain in MNE. Bottom: Average time course of activation within ROIs in arbitrary dSPM units.

Neurophysiological timing results. We investigated the millisecond-by-millisecond, trial-by-trial activation within each subject's ROI in a mixed effects model analysis with subjects and items as random factors. The variables were residual values from linear regressions that removed effects from other variables. Figure 3 displays the correlations and Table 2 provides a summary of the significant correlations.

*Superior temporal gyrus***.** Activation significantly correlated with token biphone frequency over the 160-191 ms time window (Σr =1.4463 for 31 time points, p = .04 following correction for multiple comparisons (CMC) over the 1-200 ms window time window), and also over the 217-255 time window (Σr =1.7320 for 39 time points, p = .03 following CMC over the 200-500 ms window time window), with higher values of biphone frequency resulting in stronger activation.

*Superior temporal sulcus***.** Activation significantly correlated with cohort competition from 255-276 ms (Σr =1.0677 for 21 time points, p = .02 following CMC over the 150-300ms time window), where higher competition had an inhibitory effect on activation (note that lower values denote higher competition). The linear mixed effects model analysis of entropy examined the effect of the millisecond entropy values on each millisecond of STS activation. The analysis identified a large cluster of significant t-values between ~250-280 ms post stimulus onset (significance threshold of $t > 1.96$, $p < .05$).

Posterior middle temporal gyrus. Phonological neighborhood density significantly modulated activation over the 327-347 ms time window but was just at the significance threshold following CMC ($\Sigma r = .91$ for 21 time points, $p = .05$). Activation significantly correlated with surface frequency between 415-442 ms ($\Sigma r = 1.287$ for 28 time points, $p = .04$ following CMC over 200-500ms), with higher values yielding stronger activation.

Occipital. Activation significantly correlated with residual imageability over the 161-191 ms time window ($\Sigma r = 1.241$ for 31 time points, $p < .05$ following CMC over 100-300 ms), with higher values resulting in stronger activation. We included additional variables in the model to rule out alternative plausible explanations for the effect on occipital activation. First, visual activation could signify contact with a visual word form rather than semantic content. To rule this out, we included the words' bigram frequencies. Second, the age of acquisition (AoA) of a concept may be earlier for more imageable words. Previous work shows an association between AoA and visual activation (e.g., Ellis et al., 2006). We acquired AoA ratings from the Cortese and Khanna corpus (2008). Third, occipital activation could be attributed to acoustic features that happen to correlate with imageability, such as duration. Sound symbolism is the idea that certain units of sound share something in common. We thus coded words by their phonetic descriptors (stop-plosive, fricative, nasal, affricative, glide, lateral, and rhotic).

We first ran correlations with the new variables and the imageability variable. Imageability correlated significantly with AoA, with higher imageability associated with lower AoA ($r = -.4149$, $p < .001$). This suggests that more imageable words are learned at an earlier age. However, a correlation with AoA and *residual* imageability was not significant ($p > .05$). We included the new variables in correlational analyses with occipital activation. None of the variables were found to significantly modulate activity following CMC ($p > .05$). Additionally, including the new variables in a regression model with imageability did not affect the significance of the correlation between imageability and occipital activation.

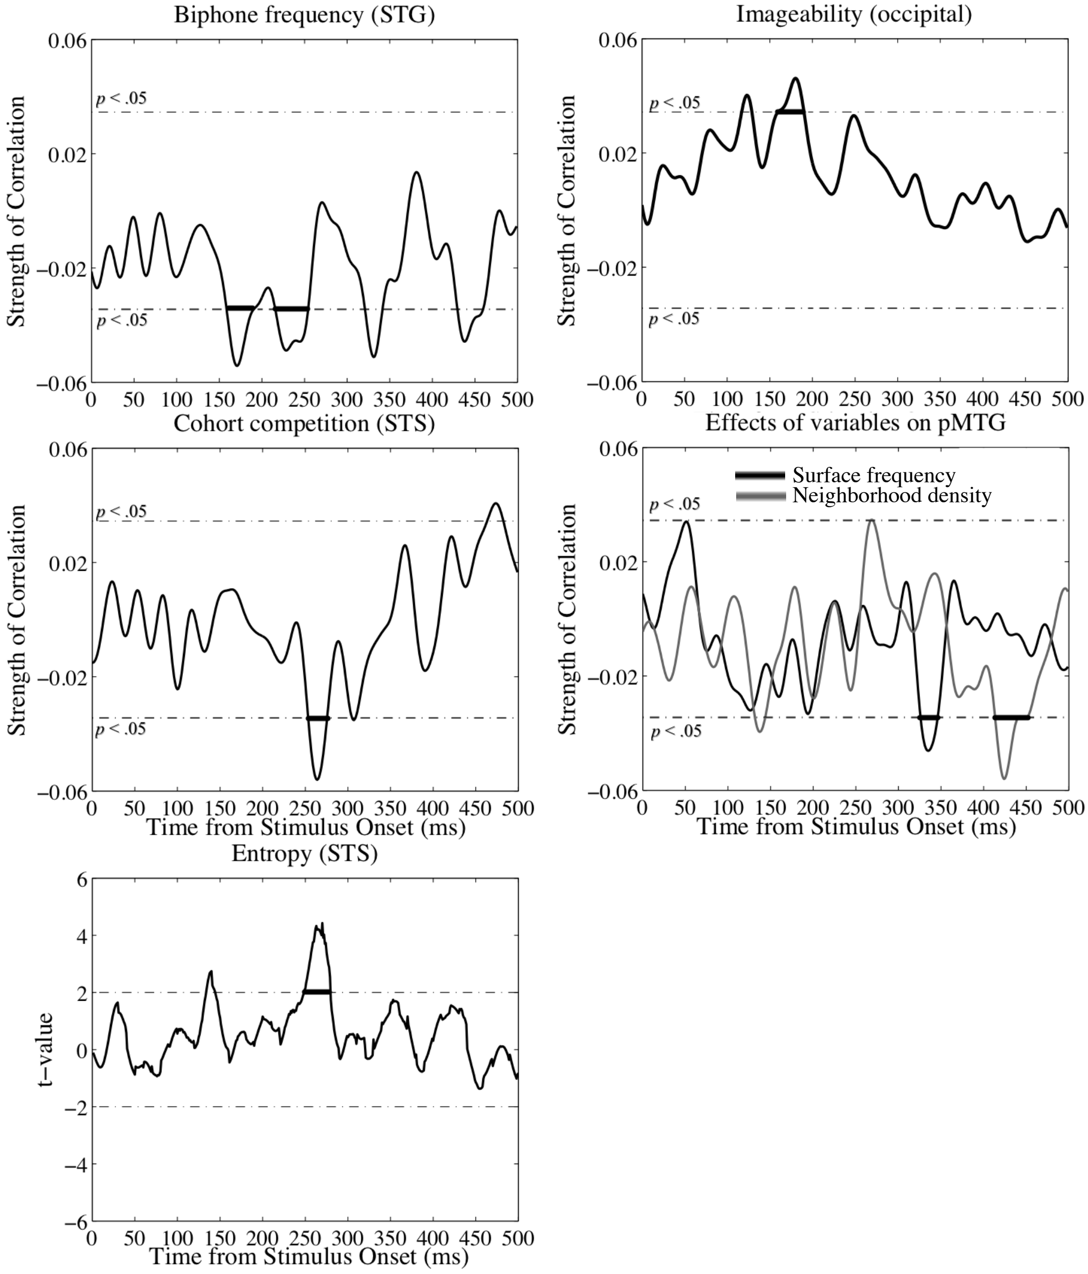

Figure 3: Effects of stimulus variables on ROI activation plotted over time. The $p < .05$ significance level is indicated by the dotted line. Bold lines identify clusters that survived CMC.

Notes: STG = superior temporal gyrus, STS = superior temporal sulcus, pMTG = posterior middle temporal gyrus.

ROI	Variable	*p*	*r*	Time window (ms)
STG	Biphone Frequency	0.04	1.45	160-191
STG	Biphone Frequency	0.03	1.73	217-255
Occipital	Imageability	0.03	1.24	100-300
STS	Cohort Competition	0.02	1.06	255-276
pMTG	Neighborhood Density	0.05	0.91	327-347
pMTG	Surface Frequency	0.04	1.28	415-442

Table 2: Significant correlations between ROI activation and stimulus variables.

STG = superior temporal gyrus, STS = superior temporal sulcus, pMTG = posterior middle temporal gyrus.

1.5 Discussion

This study focused on the temporal organization of the mapping from sound to meaning in lexical processing. We found that perceptual and lexical variables modulated different brain regions during different time windows. Importantly, and somewhat counter-intuitively, visual regions were maximally sensitive to imageability early on in speech processing, prior to effects of cohort competition and surface frequency, typical lexical-level effects. Token biphone frequency modulated STG activation from ~160-190 ms and from ~215-255 ms. While the effect at this region might be predicted, the direction of the effect is surprising as higher frequency was not predicted to result in stronger activation. During the very same time window, imageability modulated occipital activation. The direction of this effect has two interpretations: 1) imageable words have a *stronger* (single) visual representation, or, 2) imageable words have *more* visual representations. We base the latter explanation on the timing of the effect, which occurs prior to lexical access during the activation of multiple competitors. Given the temporal overlap of the imageability and biphone effects, we must assume that incoming sound automatically results in contact with low-level representations of the sound and associated visual properties.

As predicted, cohort competition modulated STS activation prior to lexical access (between ~255-275 ms). Based on the direction of the effect, we hypothesize that greater competition inhibits activation. We additionally found that higher entropy facilitated recognition, presumably because when uncertainty is high (e.g., more competitors), we devote fewer resources in accessing the representation (Ettinger, Linzen, & Marantz, 2014).

Surface frequency modulated pMTG activation between ~415-440 ms. The presence of lexical effects during only this later stage is supportive of "late access" models of lexical resolution. The direction of the effect (higher frequency yielded greater activation) is counter intuitive yet consistent with previous MEG results (e.g., Lewis, Solomyak, & Marantz, 2011; Simon, Lewis, & Marantz, 2012; Solomyak & Marantz, 2009). As stated earlier, we do acknowledge that failure to find earlier lexical effects does not necessarily mean that lexical resolution has not already begun. Figure 4 depicts the stages of spoken word recognition based on the results reported here.

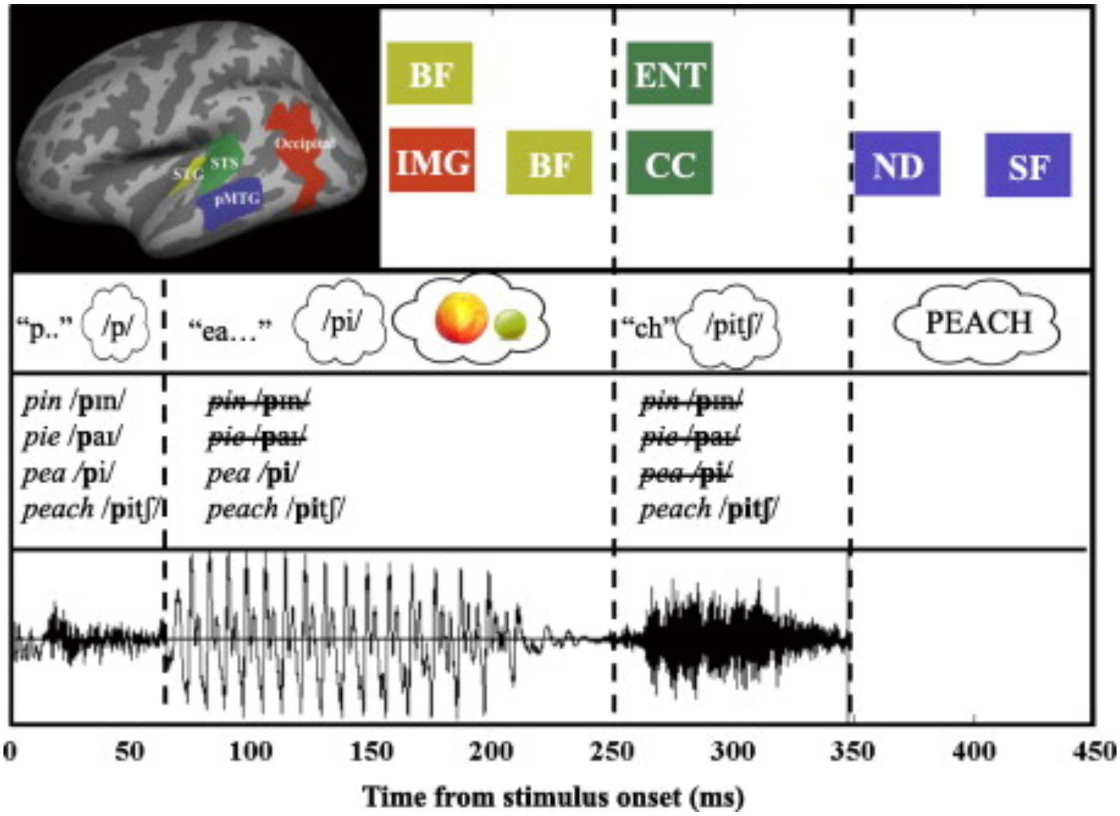

Figure 4: Proposed flowchart of sublexical and lexical stages in spoken word recognition(BF = biphone frequency, IMG = imageability, ENT = entropy, CC = cohort competition, ND = neighborhood density, SF = surface frequency).

In the model, the incoming speech waveform (bottom panel) activates segmental sound representations (middle panel) at STG and visual representations at visual regions. Before the recognition point (in this example, at ∼350 ms), phonemes activate a cohort of competitors at STS. After selection of the representation, the item's phonological family becomes activated at the pMTG. After these processes are complete, the lexical entry is selected at the pMTG.

Although the words were fairly short (mean of ~370 ms), it is unlikely that subjects reached the uniqueness point of the items by as early as 160 ms (where the effect of imageability emerged). Based on these findings, we must conclude that perceptual representations became activated before selection of the lexical entry. We hypothesize that the incoming sound simultaneously activates both segmental sound representations at the STG as well as the associated visual representations at occipital regions. Immediately prior to the recognition point of the word, the phonemic representations activate a cohort of competitors at the STS. Once the winning competitor has been selected, we then activate representations of the item's phonological family. Only after completing these processes do we activate and select the lexical entry. There is tension in these findings as to the extent of their ecological validity. Single spoken word recognition may operate at a much slower pace in the absence of contextual information (compared to, for instance, comprehension of a conversation).

While the imageability effect indicates early contact with visual representations, we cannot conclude from these data whether visual representations causally contribute to lexical access. Because we employed residual imageability values from a regression model including biphone frequency, surface frequency, onset phoneme frequency, and other variables, we can at least assume that early cortical activation in visual cortices cannot be attributed to phonemic, phonetic, or lexical effects. Similar to Zhuang et al. (2011), we found that imageability correlated positively with activation early on in recognition, suggesting that words that are more "semantically rich" (where richness is simply a measure of featural distinctiveness) will activate more perceptual representations, as indicated by stronger activation for such words. Also consistent with Zhuang et al. (2011), we found that higher imageability led to faster and more accurate responses. Both Zhuang et al. and Tyler et al. (2000) found that higher imageability contributed to recognition only when items came from a large cohort of competitors, and therefore concluded that we more easily recognize lower competition words based on their phonemic properties rather than the semantic properties of their cohorts. We, however, found that imageability contributed to recognition across the board, with higher imageability resulting in faster responses to both high and low competition words. This pervasive effect (which partialled out confounding involvement of sublexical and lexical variables) at least indicates that early activation of perceptual representations plays *some* role in computing the meaning of a word, although the nature of this role remains to be understood.

CHAPTER 2: THE NEURAL BASES OF TAXONOMIC AND THEMATIC CONCEPTUAL RELATIONS IN THE BRAIN[2]

Gwyneth Lewis, David Poeppel, and Gregory L. Murphy

2.1 Abstract

Converging evidence from behavioral and neuroimaging studies of human concepts indicate distinct neural systems for taxonomic and thematic knowledge. A recent study of naming in aphasia found involvement of the anterior temporal lobe (ATL) during taxonomic (feature-based) processing, and involvement of the temporoparietal junction (TPJ) during thematic (function-based) processing. We conducted an online magnetoencephalography (MEG) study to examine the spatio-temporal nature of taxonomic and thematic relations. We measured participants' brain responses to words preceded by either a taxonomically or thematically related item (e.g., COTTAGE → CASTLE, KING → CASTLE). In a separate experiment we collected relatedness ratings of the word pairs from participants. We examined effects of relatedness and relation type on activation in ATL and TPJ regions of interest (ROIs) using permutation t-tests to identify differences in ROI activation between conditions as well as single-trial correlational analyses to examine the millisecond-by-millisecond influence of the stimulus variables on the ROIs. Taxonomic relations strongly predicted ATL activation, and both kinds of relations influenced the TPJ. Our results further strengthen the view of the ATL's importance to

---

[2] Published as: Lewis, G., Poeppel, D., & Murphy. G. L. (2015). The neural bases of taxonomic and thematic conceptual relations: An MEG study. *Neuropsychologia, 68,* 176-189.

taxonomic knowledge. Moreover, they provide a nuanced view of thematic relations as involving taxonomic knowledge.

2.2 Introduction

Since at least Inhelder and Piaget (1964), taxonomic knowledge has been a major focus of the study of human concepts. Many concepts can be structured into taxonomies, in which specific, concrete categories nest within more general superordinate categories: for example, schnauzer-dog-mammal-vertebrate, dining room table-table-furniture-artifact-object, birthday party-party-social event. Categories that fall within a common superordinate (e.g., all parties) tend to share properties. Furthermore, the nested quality of many such concepts allows one to draw inferences, such as assuming that schnauzers breathe and give birth to live young, even if one has never encountered a schnauzer. Taxonomic categories are generally similarity-based, that is, they have *shared attributes*. Dogs tend to have four legs, bark, have fur, be pets, and eat meat.

For many years, cognitive psychologists considered taxonomic concepts "real" concepts and other forms of grouping to be the result of immature conceptual systems. For example, children form groupings such as putting a woman with a car because the woman drives the car (e.g., Smiley & Brown, 1979). Such *thematic categories* are not based on similarity (i.e., shared features) but on extrinsic relations between two objects. Later research discovered that adults also form thematic categories if the task is structured correctly (Estes, Golonka, & Jones, 2011; Lin & Murphy, 2001; Murphy, 2001). Some adult subjects even prefer thematic to taxonomic categories when forced to make a choice.

Ellen Markman (1989) made the important observation that thematic relations are not just a primitive form of concept but are an important part of conceptual knowledge. If you want to know what goes on top of a birthday cake, it does not do any good to know about cakes or desserts in general—you have to know that candles go with cake. If you see candles, cake, and balloons, you can infer that a birthday party is taking place. Such relations comprise an important part of our knowledge of events and situations but are theoretically separate from taxonomic knowledge in that taxonomic categories tell us the properties of a set of objects, whereas thematic knowledge tells us how other categories relate to that set (Murphy, 2010).

Theoretically, it is unclear whether we store thematic information as part of the same neural network as taxonomic categories. On the one hand, thematic categories have a very different basis. Dogs and leashes may go together (and some people classify them as the same kind of thing), but they do not share properties. If associations to features represent taxonomic categories, this would not seem to include the extrinsic relations thematic categories are based on. On the other hand, the constituents of thematic relations *are* taxonomic concepts: It is the taxonomic concepts of dogs and leashes that occur in the dog-and-leash thematic concept, so it would not be surprising if the thematic knowledge were mixed in with the taxonomic knowledge. That is, the concept of dog is linked to a node representing four legs and to another node representing walking on a leash.

Alternatively, thematic knowledge may be part of a different knowledge system, namely, our knowledge of events and situations. At a birthday party one lights candles on a cake, sings “Happy Birthday,” gives presents, and so on, in a particular order. Thus, links between candles and cake may not be part of the representation of those taxonomic categories *per se* but could instead be embedded in our event knowledge. Studies of the neural representation of these categories could help to distinguish these possibilities.

Neural responses in taxonomic and thematic associations. The nature of concepts and their neural representation and processing is often investigated in the context of semantic categorization tasks, in which participants determine whether a word falls into a particular category or which of two words goes with a target word. Recent fMRI experiments provide somewhat conflicting results regarding the brain regions involved in thematic and taxonomic categorization. In one study, participants viewed a target and two choice pictures and selected the one picture most related to the target (Kalénine et al., 2009). Stimuli included manipulable and non-manipulable artifacts and natural objects. Half of the related items related taxonomically, half thematically. They found taxonomic processing involved bilateral occipital regions (especially for non-manipulable natural objects like animals) whereas thematic processing involved temporal and parietal (visuo-motor) regions (especially for manipulable artificial objects like tools). Kalénine et al. concluded that taxonomic relations probably rely on perceptual processes while thematic relations rely on event/action processing, perhaps related to object manipulation. Because subjects viewed three pictures during the task, activation in visual areas is expected. The question remains whether such activation reflects conceptual processes or the perception of pictorial similarity. The Kalénine et al. study offers valuable information about conceptual processing of picture naming, but it is unclear whether such occipital activation would be found without pictures—that is, whether taxonomic knowledge is primarily visual or whether such activation mainly reflects the visual similarity of taxonomically related pictures.

Additionally, Carlson, Simmons, Kriegeskorte, and Slevc (2013) did not find occipital activation correlating with semantic properties when viewing single pictures.

Sachs et al. (2008) attempted to identify neural correlates of taxonomic and thematic relations using a similar choice task with words. In a "biased" condition, a taxonomically or thematically related word plus an unrelated preceded a target word. Subjects chose the related option. In contrast to Kalénine et al. (2009), Sachs et al. found similar recruitment of occipital, inferior frontal, and middle temporal brain regions for both relations.

An important region for representing conceptual knowledge is the anterior temporal lobe (ATL), though primarily studied with taxonomic concepts (e.g., Rogers et al., 2004). A recent, large aphasia study by Schwartz et al. (2011) specifically contrasted taxonomic and thematic relations. They analyzed a database of naming errors in 86 patients with various lesion locations to identify which locations corresponded with different error types. When shown a picture of a dog, aphasics with naming problems might give the name of another animal (e.g., *cat*), which would be a taxonomic error. Less often, a thematic error might be produced (e.g., *leash* or *bone*). To isolate the two error types, Schwartz et al. regressed out the mutual variance of taxonomic and thematic errors. Their results suggest two left-hemisphere areas predict such errors: ATL lesions uniquely predicted taxonomic errors whereas temporoparietal junction (TPJ) lesions uniquely predicted thematic errors. Jefferies and Lambon Ralph (2006) also reported naming errors for patients with either semantic dementia resulting from ATL degeneration or aphasia resulting from temporal or frontal damage (or both). They reported that the latter group's errors were "associative" responses 27% of the time, such as squirrel -> *nuts*, glass -> *ice*, and lorry -> *diesel*. These are thematic relations. The ATL group's errors were such responses only 1% of the time and taxonomic responses the rest of the time.

Results from an eye-tracking experiment further support a distinction between the representation of taxonomic and thematic knowledge (Mirman & Graziano, 2012). Aphasics with predominantly ATL or posterior lesions heard individual spoken words, each followed by a screen presenting four images: a target item, a taxonomically or thematically related item, and two distractors. Taxonomic trials yielded longer fixations in the ATL patients, whereas thematic trials led to reduced, delayed fixations in the posterior lesion patients. Taken together, the data from Kalénine et al. (2009), Schwartz et al. (2011), and Mirman and Graziano (2012) suggest a distinction between the neural implementations of taxonomic and thematic knowledge.

According to a prominent distributed-only model of conceptual knowledge, semantic knowledge is represented in a distributed fashion in modality-specific sensory-motor brain regions (Martin, 2007; Martin & Chao, 2001). As an example, our knowledge of *dog* includes its typical shape, the sound of its bark, and its gait. Such features may be represented in cortical areas involved in vision, audition, and motion processing, respectively. An alternative view, known as the *distributed-plus-hub* (AKA *spoke-and-hub*) model, agrees that while specific features of conceptual representations are indeed stored near perceptual or motor areas, semantic knowledge requires a general, supra-modal mechanism that can generalize across similar concepts differing in some featural way (e.g., poodles and beagles). The ATL may be the “hub” that assimilates features from modality specific regions (the spokes). As evidence, while injury to one of the “spokes” typically results in a specific deficiency, damage to the ATL results in more general disabilities, namely, semantic dementia (Jefferies & Lambon Ralph, 2006; Patterson et al., 2006). The role of the ATL in coordinating featural information strongly implies its importance in processing taxonomic concepts that are represented based on associations with their features. A recent magnetoencephalography (MEG) study of noun specificity suggests such involvement of the ATL in the interaction between concepts and their features (Westerlund & Pylkkänen, 2014). This study found an interaction effect at the ATL, wherein activation differed based on the featural specificity of nouns (e.g., *fish* vs. *trout*) and whether they included a modifier (e.g., *spotted fish*).

Schwartz et al.'s (2011) result for taxonomic errors complement other findings involving the ATL (including those summarized above), and their results for thematic errors involving parietal regions generally agree with findings in Kalénine et al. (2009). However, such lesion data can be difficult to interpret. If lesions in the ATL cause taxonomic errors, does that mean the ATL represents taxonomic information? On the one hand, disruption to taxonomic knowledge causes errors. On the other hand, disruption to the taxonomic network should arguably lead aphasic patients to produce more thematic responses, as taxonomic responses become less available. Similarly, if the TPJ and surrounding areas represent thematic relations, one might expect a lesion there to prevent people from providing thematic responses rather than to increase them. Schwartz et al. explain their results as a function of greater noise in the processes controlled by the damaged area. We will consider their account in more detail after presenting our own results.

In short, the Schwartz et al. (2011) study offers an impressive analysis of a large database of language production, concluding that the ATL is critical for taxonomic concepts and the TPJ for thematic concepts. However, online data of taxonomic and thematic processing in intact subjects (where the issue of lesion of effects do not arise) would enhance their interpretation.

The present study. Our experiment further examines the neuroanatomical and neurophysiological differences for thematic vs. taxonomic relations using MEG. We used word stimuli to avoid any incidental visual activation during taxonomic judgments. We also used a simple relatedness task so participants would not have to make difficult judgments about the type of conceptual relation involved, thereby minimizing the decision component of the task. In our study, words pairs appeared individually in sequence, and participants responded as to whether they were related. Some pairs related taxonomically, such as *cottage-castle*, and others related thematically, such as *king-castle*. Foils had no apparent relationship. We recorded neural and behavioral responses to the second word and examined the effects of category type. We analyzed the MEG data constrained with structural MRIs, which enabled us to examine precise timing of effects arising from the different stimuli without sacrificing much spatial resolution.

Standard analyses of choice response time (RT) (Ratcliff, 1978) suggest the following general framework for interpreting this task: When presented with two words, people retrieve information from their conceptual representations to identify overlapping features or shared relations. With retrieval of sufficiently strong positive information, a "Yes" response is made. The more strongly related two words are, the faster people can respond, due to more (and more salient) information linking the two items. Negative responses occur either when enough time passes without retrieving sufficient information to link the two words, or after retrieval of information indicating a lack of connection between them.[3] Negative RTs are typically slower

---

[3] We suspect the latter source of negative responses can seldom be used in the present task. If we had tested only taxonomic relations, it would be possible to retrieve clearly disconfirming information (e.g., one item has leaves and one flies; therefore they are not in the same taxonomic category). However, as thematic relations are very diverse, sometimes linking very different kinds of items, such featural clash cannot effectively serve to identify unrelated word pairs.

than positive RTs because positive responses can be generated as soon as linking information is retrieved, whereas negative responses require waiting for the failure of such retrieval.

The MEG signal is often proportional to the amount of computation required to perform a task. For example, priming a stimulus typically leads to reduction in MEG signal (e.g., Pylkkänen et al., 2006). Although the particular MEG profile of our task is not known, we expected that highly related items would generate a smaller MEG signal and that unrelated items would generate a larger signal, as they require a longer retrieval and comparison process. The results showed that this pattern was generally but not always found.

Results from previous MEG studies suggest spatially and temporally distinct neural stages in visual word recognition. Upon word presentation, these stages begin with orthographic feature detection in occipital regions at around 100 ms, followed by morphological decomposition in inferior temporal regions at around 150 ms, and finally lexical access in superior temporal regions at around 300 ms (e.g., Lewis, Solomyak, & Marantz, 2011; Simon, Lewis, & Marantz, 2012). Little is known, however, about the time course of conceptual relations in the brain. Examination of the millisecond-by-millisecond effects of thematic and taxonomic relations on ROI activation could therefore contribute to our understanding of the temporal and spatial nature of the mechanisms of conceptual relations in visual word recognition.

In sum, our study focuses on the role of the ATL and TPJ during processing of taxonomically vs. thematically related word pairs. If these regions represent distinct systems of semantic knowledge, we should find differential neural activation to taxonomic vs. thematic stimuli. Lastly, we examined taxonomic and thematic effects on posterior occipital regions, which Kalénine et al. (2009) found to be involved in taxonomic processing. These regions may be less involved in identifying taxonomic relations with word stimuli.

2.3 Method

Participants. The MEG experiment included 17 right-handed native English speakers (8 males) from the New York University community with normal or corrected to normal vision. Of the 17 participants, 13 had structural MRI data sets available, which we later used for source localization.

Stimuli. The test stimulus set consisted of 300 primes and 150 targets. Half of the primes related thematically to the targets (e.g., *king-castle*), while the other half related taxonomically to the targets (e.g., *cottage-castle*). The target stimuli could be loosely termed as belonging to one of the following six groups (25 in each): animate, clothing, food, tools/objects, dishes/household, and transportation. We also generated 300 primes and 150 unrelated target filler items (e.g., *nutmeg-reflex*). In constructing the thematic pairs, we took care to avoid items pairs sharing salient taxonomic categories, e.g., *mouse-cat*. The Appendix lists the test stimuli.

Variables. On average, the words were similar across conditions in terms of length, surface (written word) frequency, bigram (adjacent letter) frequency, number of syllables, number of morphemes, number of phonemes, as well as normative behavioral data including mean naming accuracy as reported by the English Lexicon Project (Balota et al., 2007). Table 1 reports the linguistic properties of the stimuli.

Word Type	Length	SF	BGF	NPhon	NSyll	NMorph
Test Targets	5.6(1.7)	8.0(1.7)	7.4(0.5)	4.6(1.6)	1.7(0.7)	1.3(0.5)
Filler Targets	5.6(1.4)	8.1(1.4)	7.4(0.5)	4.8(1.3)	1.9(0.8)	1.3(0.5)
Taxonomic Primes	5.9(1.8)	7.6(1.7)	7.3(0.6)	4.8(1.6)	1.7(0.7)	1.3(0.5)
Thematic Primes	5.3(1.6)	8.5(1.9)	7.4(0.5)	4.2(1.3)	1.5(0.6)	1.3(0.5)
Filler1 Primes	5.6(1.4)	8.0(1.4)	7.4(0.6)	4.8(1.4)	1.8(0.6)	1.4(0.5)
Filler2 Primes	5.7(1.3)	7.9(1.5)	7.4(0.5)	4.9(1.4)	1.9(0.7)	1.3(0.5)

Table 1: Means and standard deviations of the stimulus properties by word type and condition. *Note: Length = letter length; SF = surface (written) word frequency; BGF = log bigram frequency; NPhon = number of phonemes; NSyll = number of syllables; NMorph = number of morphemes.

We carefully matched the stimuli in terms of these properties to ensure any effects could be attributed to the category condition rather than lexical properties of the words. Later, we regressed these and other properties onto the response data to further ensure the validity of any effects. However, it is important to keep in mind that the same test targets served in the thematic and taxonomic conditions, so any differences between them cannot be explained by lexical effects.

Relatedness ratings. In a separate experiment, we obtained relatedness ratings for the stimulus pairs from 38 native English speakers via Mechanical Turk. We excluded 9 participants based on "sanity-check" items (10 highly related and 10 highly unrelated item pairs, e.g., *square-circle* and *scheme-moose*) separate from the experimental stimuli to test whether participants rated the items seriously. We excluded participants when the difference between ratings of highly related and highly unrelated items fell below a criterion value. Instructions for the relatedness task are included in the Appendix. Results of the questionnaire revealed, on average, a slightly higher relatedness score for thematic pairs (*M* = 5.63, *SD* = 0.40) than for taxonomic pairs (*M* = 4.47, *SD* = 0.24). However, both average scores well exceeded that of the filler pairs (*M* = 1.82, *SD* = 0.34). Rather than select a subset of stimuli with equal ratings in the two related conditions, we used each item's relatedness as a predictor of MEG signal and RT in single-trial analyses.

Procedure and recording. We used Matlab (MathWorks, Inc., Natick, MA, USA) with Psychtoolbox helper scripts to present the stimuli. Participants viewed the prime-target pairs over the course of six blocks, with each block consisting of 100 trials. In each block, participants responded to 50 filler pairs, 25 taxonomic pairs, and 25 thematic pairs, with randomized presentation. We shifted the block order for each participant, and ordered the blocks so that the same prime did not appear twice within the first three sequential blocks. Participants completed a practice session before the actual experiment. They read that they would see word pairs in sequence and should decide whether the pairs were related or unrelated. Instructions explicitly mentioned that related words might be the same kinds of things (like *velcro* and *zipper*) or have some relation (like *pants* and *zipper*—zippers open pants). They responded yes or no with their index and middle fingers and were asked to be as fast and accurate as possible. The task lasted approximately half an hour. Figure 1 shows the task sequence. The prime word appeared for 300 ms, followed by a variable interval of .5 - 1 s. The target then appeared for 400 ms, so that word exposure did not vary across conditions. The next trial began after the response. Participants lay supine during the experiment while a whole-head MEG system (157 axial gradiometer sensors) (Kanazawa Institute of Technology, Kanazawa, Japan) acquired their neural data. Recording parameters were the following: 1000 Hz sampling rate, 60 Hz band-pass filter, DC high-pass filter). Structural MRIs for 13 of the participants were acquired from a separate session at the Center for Brain Imaging at New York University (3T Siemens Allegra scanner with T1-weighted MPRAGE sequences).

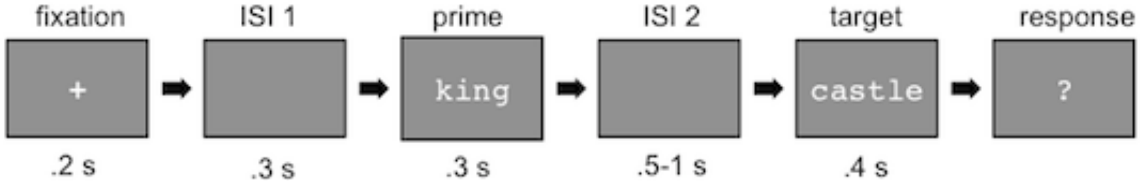

Figure 1: Trial sequence and event durations.

Analysis. The source space analysis closely resembled that described in Lewis et al. (2011) and Lewis and Poeppel (2014). First, we noise-reduced the data in MEG160 (Yokogawa Electric and Eagle Technology Corporation, Tokyo, Japan) using reference sensor data and the Continuously Adjusted Least-Squares Method (CALM; Adachi, Shimogawara, Higuchi, Haruta, & Ochiai, 2001). We next imported the data into MNE (MGH/HMS/MIT Martinos Center for Biomedical Imaging, Charleston, MA) and reconstructed the MRI data sets with FreeSurfer (CorTechs Lab Inc., La Jolla, CA), which we used to calculate minimum-norm solutions. For the four participants without structural MRIs, we employed a standard FreeSurfer brain aligned to each participant's fiducial data. We calculated estimates of the magnetic field at each sensor (the forward solution) from 5124 points of activity using the boundary-element model (BEM) method. We used the forward solution to estimate the spatio-temporal distribution of the MEG data (the inverse solution). We used a free orientation (unconstrained in relation to the cortical surface) to compute the inverse solution. Data were signed, where negative values indicate activity directed downward and positive values indicate activity directed upward with respect to the head based coordinate space. Finally, data were transformed into dynamic statistical parameter map (dSPM) values (Dale et al., 2000) and retained only components normal to the cortical surface. These analytic procedures (in particular, the use of signed, free orientation minimum norm estimates) follow those used in previous MEG studies of visual and spoken word recognition (e.g., Ettinger, Linzen, & Marantz, 2014; Fruchter, Stockall, & Marantz, 2013; Lewis & Poeppel, 2014; Lewis, Solomyak, & Marantz, 2011; Simon, Lewis, & Marantz, 2012).

Regions of interest. Findings from Schwartz et al. (2011) and Kalénine et al. (2009) motivated our interest in four brain regions potentially involved in taxonomic and thematic conceptual processes. For the ATL ROI, this included Brodmann area (BA) regions reported in Schwartz et al. as containing the most significant percentage of voxels associated with taxonomic errors. We therefore created the ATL ROI by merging the following anatomical ROIs from a BA parcellation: BA 38 (33% of error in Schwartz et al.), anterior BA 21 (27% of error), and anterior BA 20 (25% of error). For the TPJ ROI, we merged BA ROIs in posterior regions accounting for the greatest number of voxels associated with thematic errors in Schwartz et al. These included LH BA 39 (35%), posterior LH BA 22 (15%), and LH BA 41 and LH BA 42 (combined 10%). Lastly, we included two additional ROIs encompassing the cuneus and lingual gyrus (LH BA 18 and RH BA 18), reported in Kalénine et al. (2009) as particularly activated in taxonomic categorizations. The average activation within the ATL ROI trended negative between 0-370 ms post target onset (peaking at 160 ms and 290 ms) and then trended positive until 540 ms (peaking at 430 ms). Average TPJ activation was positive from 50-100 ms (peaking at 90 ms) and then remained negative until around 630 ms (peaking at 150 ms and 390 ms). The brain regions and their average time-courses of activation are depicted in Figure 2.

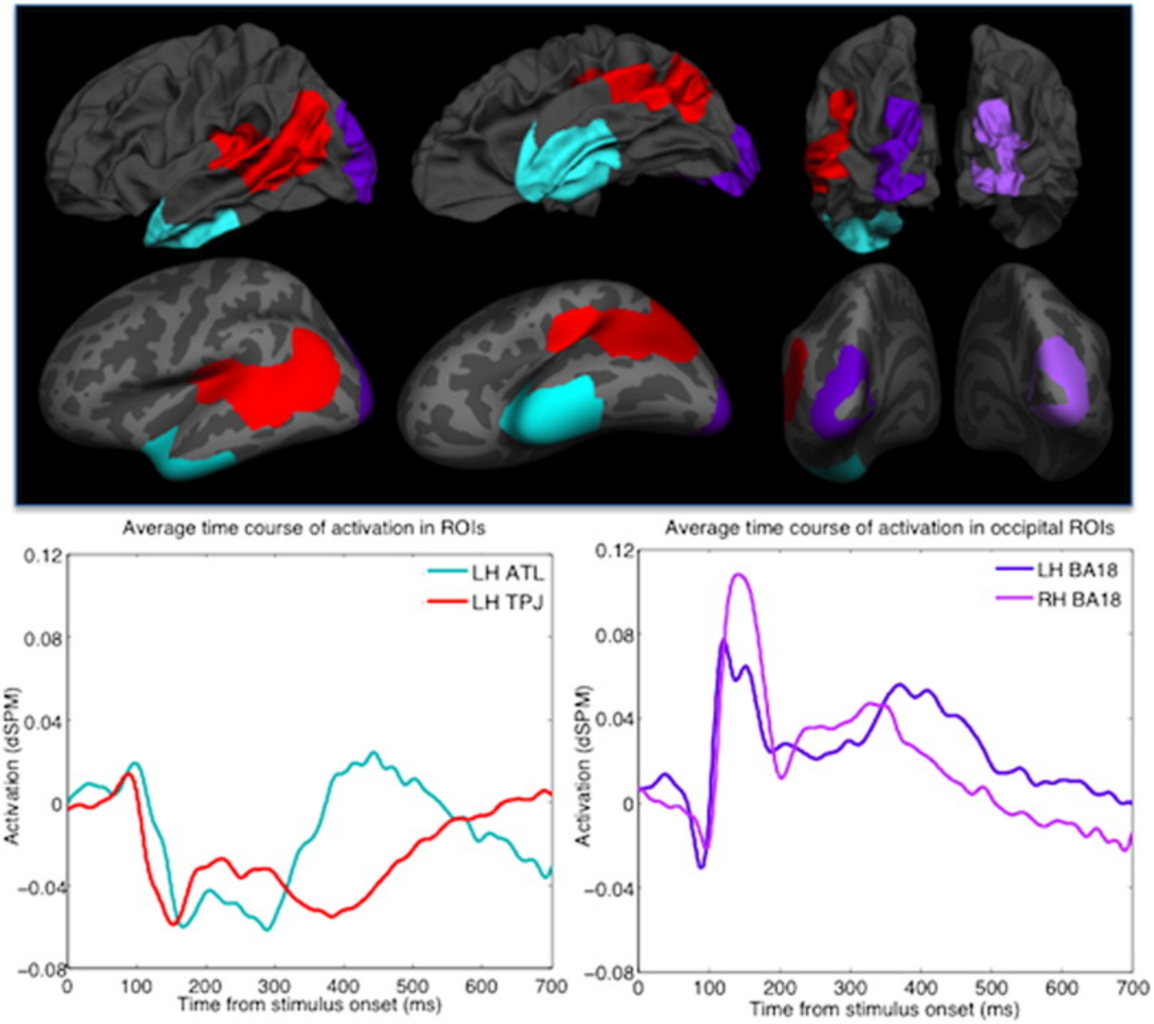

Figure 2: The ROIs presented on the standard pial surface and inflated surface. The ROIs are shown in lateral, ventral, and occipital views (from left to right). We applied a baseline correction of 100 ms pre-stimulus onset to the average time courses of activation within the ROIs (bottom).

Behavioral exclusions**.** We applied the same criteria for excluding behavioral data as in Solomyak, Lewis, & Marantz (2010), Lewis et al. (2011), and Lewis and Poeppel (2014). First, we excluded trials with incorrect responses and/or RTs exceeding 5 s. This removed approximately 9% of the data. After determining average response times in each condition, we converted each participant's RT data into z-scores and excluded trials in excess of three standard deviations above or below a given participant's mean. This procedure removed an additional 2% of the data.

Permutation t-test analyses**.** We conducted two-tailed permutation *t*-tests to compare differences in ROI activation between taxonomic vs. thematic trials, thematic vs. filler trials, and taxonomic vs. filler trials. We performed the analysis on source space activation over the 200-700[4] ms window, relative to the target word onset. In each test, we set criteria for selecting significant clusters of activation (10 sequential time points with a conservative[5] component test threshold *p*-value of $p = .1$). For each cluster of activation, a test statistic was constructed from the sum of *t*-values within a given cluster to identify the largest cluster statistic. Next, the procedure repartitioned each participant's data via random assignment to different conditions, and identified the largest cluster statistic at each of the 10,000 permutations, resulting in a distribution of the largest cluster statistics. Comparisons between the original data and the largest cluster statistics created estimates of each ROI's *p*-value, corresponding to the average proportion of permutation statistics greater than the original statistic.

---

[4] We corrected over 200-700 ms in the permutation t-test, and over 200-600 ms in the correlational analysis, the difference arising from computer memory issues (the correlational analysis in Matlab could only handle data up to this value). These time windows are liberal, as many previous similar analyses correct only over, e.g., 200-500 ms and sometimes less.

[5] Previous studies employed a threshold *p*-value of $p = .3$ (e.g., Bemis & Pylkkänen, 2011; Westerlund & Pylkkänen, 2014).

Finally, we applied a false discovery rate (FDR) controlling technique (Benjamini & Hotchberg, 1995; Benjamini & Yekutieli, 2001; Genovese et al. 2002). This involved ordering the *p*-values in descending order for comparison with the ratio of the index of the ordered p-value to the total number of tests, multiplied by the FDR (here, .05). The *p*-values less than this statistic were deemed significant.

Single-trial analyses. The permutation *t*-test analyses, which we performed on averaged data, provide a general indication of taxonomic vs. thematic processing differences. We supplemented our analyses with single-trial time-course correlations. This analysis technique, which correlates a given stimulus variable with millisecond-by-millisecond ROI activation across all trials and participants, provides a more in-depth examination of the neural responses implicated in conceptual relations. We primarily focused on ATL and TPJ responses to the following dummy coded variables: Taxonomic vs. thematic (Tax-Them), taxonomic vs. filler (Tax-Fill), and thematic vs. filler (Them-Fill). If a brain region is involved in computing a particular kind of relation, its activation might correlate with the strength of a pair's relatedness. We therefore included relatedness in separate analyses for each condition. First, we analyzed ROI activation in regressions to remove any effects of length, bigram frequency, surface frequency, number of morphemes, number of syllables, presentation number (whether first or second time viewing the target), and response time. We next applied a multiple-comparisons correction procedure (Maris & Oostenveld, 2007) to clusters of activation significant at the $p < .05$ level before correction. Next, we computed an $\sum r$ statistic from the summed coefficients of significant contiguous effects and then tested the significance of this statistic with Monte-Carlo *p*-values. To do so, we created 10,000 permutations of the random variable to generate a correlation wave. Next, we calculated the $\sum r$ statistic at each permutation and constructed a new distribution of $\sum r$ values from the highest $\sum r$ value at each permutation. Finally, we based our Monte-Carlo *p*-value on the proportion of values greater than the original statistic.

2.4 Results

Our analyses examined behavioral and neural responses to the relatedness scores of each condition and to the dummy coded variables Tax-Fill, Tax-Them, and Them-Fill. First, we report results of our behavioral analysis, which correlated RTs across trials with the variables of interest. We then present results of our permutation *t*-tests analysis, which provided estimates of the average differences in activation between conditions in each ROI. Finally, we report findings from our single-trial correlational analyses, which revealed whether, when, and to what extent the stimulus variables modulated ATL and TPJ responses.

Behavioral results. On average, thematic targets yielded the fastest responses ($M = 918$, $SD = 497$), filler targets (with "unrelated" responses) resulted in the slowest responses ($M = 952$, $SD = 528$), and taxonomic RTs fell in between ($M = 939$, $SD = 517$). Correlations with RT showed significantly faster responses to taxonomic targets than fillers ($p = .003$, $r = -.03$), significantly faster responses to thematic than fillers ($p < .0001$, $r = -.11$), and significantly faster responses to thematic than taxonomic, even though these RTs were measured on the identical target words ($p < .0001$, $r = .08$). Analysis of the RTs from "positive" trials (thematic and taxonomic) showed a significant correlation with relatedness score, with more related items yielding faster responses ($p < .0001$, $r = -.21$). As expected, RTs from "negative" trials correlated in the opposite direction, with higher relatedness yielding slower RTs ($p = .0017$, $r = .05$).

Participants responded least accurately to taxonomic items (83% correct), most accurately to fillers (95% correct), and somewhere in between to thematic items (88%). These results seem to reflect a slight bias to respond “unrelated.” The faster RTs of thematic over taxonomic trials are to be expected, given that the former were more highly related. A regression including relatedness score of each pair along with its condition determined there was no longer a significant difference between conditions when this difference was controlled for. Thus, taxonomic judgments were not slower *per se*.

Processing speed of taxonomic and thematic judgments differs depending on whether the object is an artifact like *hammer* or a natural kind like *cherry*. Kalénine et al. (2009) found speeded thematic judgments for artifacts and speeded taxonomic judgments for natural kinds. We therefore coded the data into natural kinds and artifacts (counting items from the animate and food groups as natural, although some of the latter are not technically natural kinds, because they do not exist independently in nature; Putnam, 1973) and ran separate correlations with the taxonomic and thematic RTs. As in Kalénine et al. (2009), we found taxonomic responses were significantly faster for natural kinds than artifacts ($p < .02$, $r = -.05$). This effect was then stronger after regressing relatedness score onto RT ($p < .0001$, $r = -.08$). Thematic RT was not significantly correlated. The main goal of this study was not to compare the speed of the two related conditions but rather to investigate their computation in the brain. In the next section we address that issue.

Neural results

Permutation t-tests. We began by analyzing averaged data in each ROI with two-tailed permutation *t*-tests to determine whether contiguous clusters of activation differed significantly between conditions. The analyses compared activation for taxonomic vs. thematic, taxonomic vs. filler, and thematic vs. filler trials. We report the activation means in dSPM units (Dale et al. 2000). Figure 3 shows significant results of the permutation *t*-tests.

*Anterior temporal***.** The two-tailed permutation *t*-tests identified a significant cluster of ATL activity (296-614 ms, *p* = .0079) for taxonomic vs. filler, with stronger amplitude for filler items than taxonomic items (-.671 filler mean [.192 sd] v. -.223 taxonomic mean [.192 sd]). The analysis additionally identified a significant cluster of ATL activity (305-554 ms, *p* = .0093) for thematic vs. filler, with amplitude again stronger for filler items (-.771 filler mean [.137 sd] v. -.302 thematic mean [.137 sd]). The test did not identify significant clusters for taxonomic vs. thematic. As mentioned earlier, stronger activation for the unrelated items likely reflects more processing (as supported by the behavioral results, where filler items yielded the slowest RTs).

*Temporoparietal***.** Analysis of thematic vs. filler pairs identified a contiguous cluster of activation (567-656 ms, *p* = .07, 10,000 permutations) just above significance following FDR correction (-.148 filler mean [.071 sd] v. -.03 thematic mean [.052 sd]). The analysis did not show an effect of taxonomic vs. thematic or taxonomic vs. filler.

Posterior occipital (LH and RH BA 18). The permutation *t*-test did not identify any significant clusters of activation in either LH BA 18, RH BA 18, or in a merger of the two. The posterior occipital regions therefore may not play a role in taxonomic vs. thematic judgment outside of picture naming tasks (as in Kalénine et al., 2009).

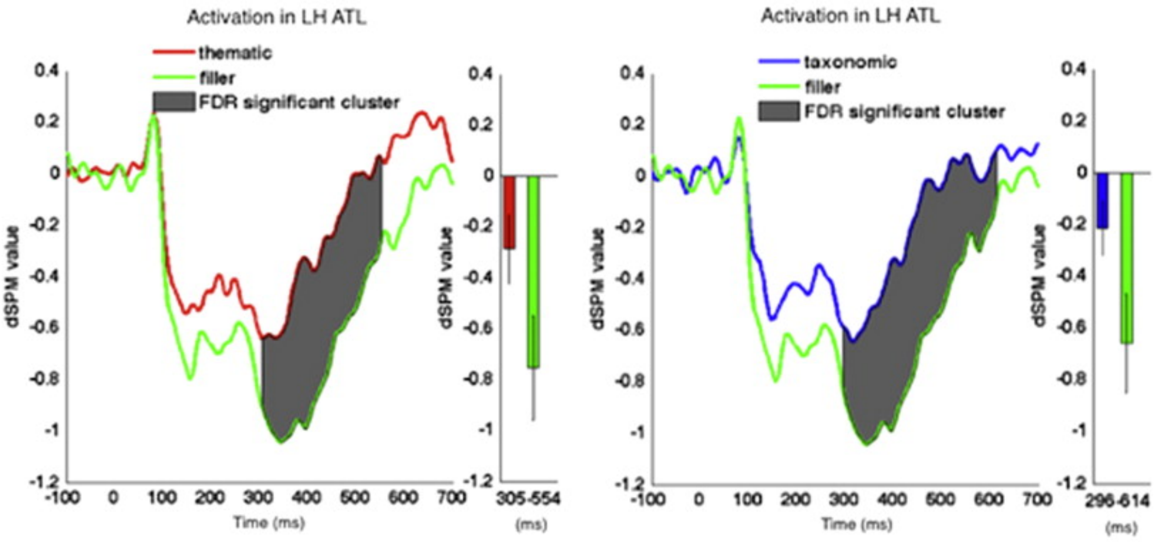

Figure 3: Results of the permutation *t*-tests plotted relative to onset of the second word. Shaded regions signify significantly different clusters of activation between conditions after FDR correction. In both graphs, activation was significantly greater for filler trials. The charts to the right of each plot show the average activation within the significant cluster and the error bars show SEMs.

Single-trial correlations. While the permutation *t*-tests provide a rough estimate of activation differences between conditions, single-trial analyses can provide further evidence that conceptual properties modulate neural activity. We next report the temporal influence of the dummy variables (e.g., taxonomic vs. thematic) as well as the relatedness scores on ATL and TPJ responses. All *p*-values reported next were corrected for multiple comparisons (CMC) over the 200-600 ms time window. Table 2 provides a summary of the results, namely, time windows over which a particular variable correlated significantly with activation as well as the *p*-values following CMC. Figures 4 and 5 show the time course correlation plots.

Anterior temporal. First, we found a correlation between ATL activation and Tax-Fill and Tax-Them, but *not* with Them-Fill. The Tax-Fill effect occurred over 291-415 ms, with stronger activation for filler items, as in the earlier analyses (Σr = 4.54 for 125 time points, p = .004). The Tax-Them effect was between 308-356 ms (Σr = 1.89 for 49 time points, p = .046), with taxonomic pairs generating more activity than thematic. Correlations between relatedness score and activation within the various conditions revealed that the relatedness score of only the taxonomic items modulated the ATL response, with higher relatedness yielding stronger activation between 419-502 ms (Σr = 4.17 for 84 time points, p = .009). Again, neither thematic nor filler relatedness significantly modulated ATL activation following CMC. In short, the ATL seemed correlated with taxonomic judgments but not thematic ones.

Temporoparietal. We found less specific correlations at the TPJ. Activation correlated both with Them-Fill and Tax-Fill. The Them-Fill effect occurred between 356-599 ms (Σr = 10.71 for 244 time points, $p < .0001$) and the Tax-Fill effect was between 327-549 ms (Σr = 11.53 for 223 time points, $p < .0001$). The correlation with Tax-Them was not significant following CMC. We also found that both thematic and filler relatedness modulated the TPJ response, with the thematic effect occurring between 309-355 ms (Σr = 2.57 for 47 time points, p = .037), and the filler effect occurring over 365-472 ms (Σr = 4.16 for 108 time points, p = .002). Taxonomic relatedness did not significantly correlate with TPJ activation. Note that activation within the TPJ ROI was negative throughout the 150-630 ms time window, so the positive correlations with TPJ activation suggest an inhibitory effect. Conversely, as activation within the ATL trended positive from 290-540 ms, inhibitory effects are indicated by negative correlations.

	ATL ROI		TPJ ROI	
Variable	***r***	**range**	***r***	**range**
Tax-Fill	4.53**	291-415	11.53***	327-549
Them-Fill	0.09	597-600	10.71***	356-599
Tax-Them	1.89*	308-356	0.84	263-284
Taxonomic relatedness	4.17**	419-502	1.18	481-503
Thematic relatedness	-	-	2.56*	308-355
Filler relatedness	0.691	200-218	4.16**	365-472

Table 2: Single-trial analysis: Significant values and clusters following Monte-Carlo CMC.

The *r*-values are sums across significant clusters. *Note.* *$p < .05$, **$p < .01$, ***$p < .001$.

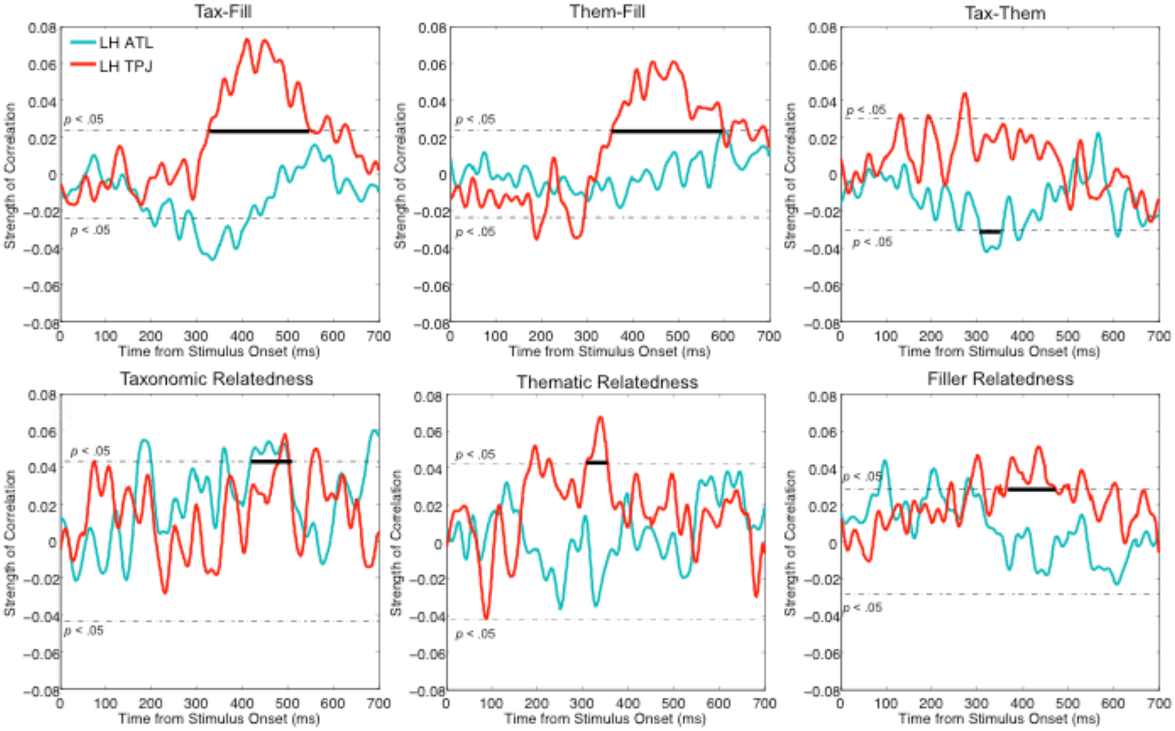

Figure 4: Effects of stimulus variables on ATL and TPJ activation plotted over time. The *p* <.05 significance level (prior to CMC) is denoted by the dotted line. Bold lines identify temporal clusters that survived the Monte–Carlo CMC (10,000 permutations). Positive TPJ correlations indicate an inhibitory effect as activation was negative during these time windows, while positive correlations with the ATL indicate a facilitatory effect as ATL activation was positive during these time windows.

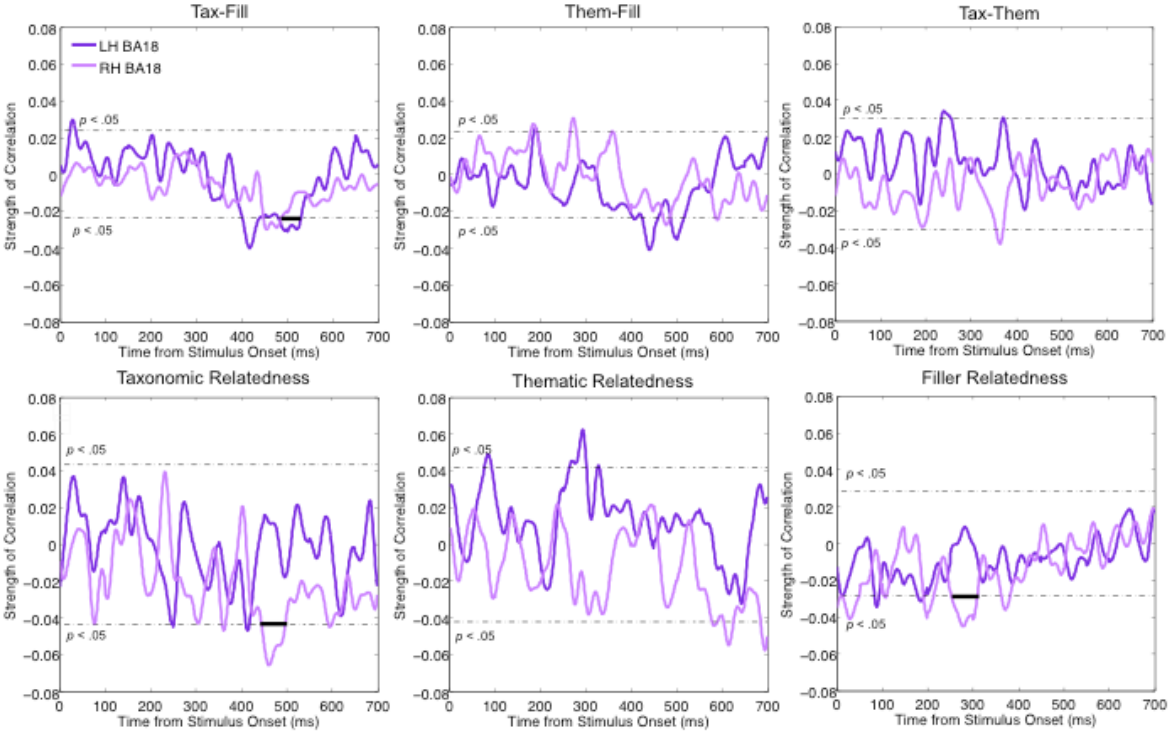

Figure 5: Effects of stimulus variables on LH and RH BA 18 activation plotted over time. The p <.05 significance level (prior to CMC) is denoted by the dotted line. Bold lines identify temporal clusters that survived the Monte–Carlo CMC (10,000 permutations). Positive correlations indicate a facilitatory effect because activation was positive during these time windows.

2.5 Discussion

This experiment focused on the spatio-temporal nature of neural responses implicated in different kinds of conceptual relations. Our main goal was to test hypotheses from Schwartz et al. (2011) regarding the involvement of the ATL and TPJ in taxonomic and thematic processing. We additionally tested the involvement of posterior occipital regions (LH and RH BA 18) in conceptual processing based on Kalénine et al.'s (2009) findings of occipital involvement in taxonomic classification with picture stimuli. In general, our experiment provides converging evidence of the ATL's involvement in processing taxonomic conceptual relations, as found by Schwartz et al. (2011). Our results suggest involvement of the TPJ in both relations. The posterior occipital areas did not appear to be involved in either relation in our task.

Taxonomic specificity in the ATL. Our results imply a distinct role of the ATL in conceptual processing, in agreement with Schwartz et al. (2011). We found modulation of ATL activation by the different conditions beginning as early as 300 ms, which aligns nicely with the timing of lexical access effects in temporal regions found in previous MEG visual word recognition studies (e.g., Lewis et al., 2011; Simon & Lewis, 2012).

First, Tax-Fill modulated the ATL response, with unrelated items resulting in stronger activation. The direction of this effect presumably reflects greater effort spent in attempting to identify a relation among unrelated stimuli, related to the longer RTs for unrelated responses. In contrast, when items are very similar, identification of relatedness requires comparatively little information retrieval and comparison. We also found a distinction for Tax-Them, with taxonomic items yielding stronger ATL activity. While stronger activation signifies greater processing difficulty, it also corresponds with the number of activated features (Lewis & Poeppel, 2014). Another way to put this is that thematic information simply did not involve the ATL: We did not find effects of thematic vs. filler, thematic relatedness, or filler relatedness on ATL activation. Additionally, we found ATL activation uniquely correlated with taxonomic relatedness. Thus, the Tax-Them difference here probably reflects the lack of thematic computations in the ATL. Permutation *t*-tests identified a significant difference in activation for taxonomic vs. filler trials, with stronger activation for unrelated items (compatible with results of the single trial analysis). The test also revealed a significant difference for thematic vs. filler with, again, stronger activation for filler trials.

Thematic processes in the TPJ? Our results suggest a less specific role of the TPJ in conceptual processing. While did not find significant differences between conditions in the permutation *t*-test analysis of the TPJ (other than a suggestive difference for thematic vs. filler), we did find numerous effects in the single trial correlational analysis not specific to thematic processing. We found correlations between TPJ activation and Them-Fill as well and with Tax-Fill during roughly the same time windows (in both correlations, unrelated items generated stronger activation). Unlike the ATL, the TPJ did not distinguish between taxonomic vs. thematic items. We additionally found that the TPJ correlated with both thematic and filler relatedness, but not with taxonomic relatedness during any time window.

Although the finding of taxonomic activation (the Tax-Fill effect) in TPJ seems to conflict with Schwartz et al.'s findings, it is not surprising that this region is involved in conceptual representation and processing. In their meta-analysis of 120 imaging studies, Binder, Desai, Graves, and Conant (2009) identified the angular gyrus (roughly BA 39, contained in our TPJ) as the area most often found to be activated in semantic processing tasks. (They also identified the lateral and ventral temporal cortex as a popular area in such tasks.) Schwartz et al. correlated brain areas with residual errors of the two types after taking into account shared errors; this possibly made it difficult to detect any regions that consistently caused both types of error. Indeed, their TPJ patients made both taxonomic and thematic substitutions, as all but two of their patients made more taxonomic than thematic errors. We address naming errors in general on the next page.

Occipital processing. We did not find any differences between conditions in any analysis involving the posterior occipital ROIs, in contrast to Kalénine et al.'s (2009) results with pictures. Conceptual processing is not one monolithic thing—it differs with context, stimulus, and task. Multiple studies are necessary to discover how conceptual information is processed under different circumstances. We suspect that Kalénine et al.'s results reflect the perceptual similarity of their taxonomically related pictures. For example, fruit tend to be smooth, with rounded shapes and to have stems; animals tend to have legs and heads; furniture tend to have flat perpendicular surfaces. In contrast, thematically related items such as fruit and a basket or a sleeping child and a bed (taken from Kalénine et al.'s Figure 1) generally do not share many perceptual properties.

It seems likely, given our results and those of other studies on conceptual processing (e.g., Binder et al., 2009; Carlson et al., 2013), that the occipital lobe does not directly store featural information about taxonomic categories but is involved in perceptual comparisons used to identify the category of viewed objects and pictures. In studies involving words, such processing will not be in evidence.

Relation to naming errors. Throughout we have compared our results to those of Schwartz et al.'s study of naming errors in aphasia. In the Introduction, we raised the issue of how lesions in an area relate to errors. According to a simplistic analysis, if a brain region does X, then one would expect to find people with lesions in that region doing less of X. However, Schwartz et al. found that damage in the ATL led to more taxonomic errors, whereas damage in the TPJ led to more thematic errors, concluding that they are involved in taxonomic and thematic processing, respectively.

The argument for the ATL is fairly straightforward. As Schwartz et al. suggest (and also see Patterson et al., 2007), when the ATL is damaged, information about category features is lost, causing errors in labeling. If someone loses the information that poodles have curly fur, then he or she might well label a poodle with the name of a straight-haired dog, such as collie, which matches the picture in the preserved features.

It is not as easy to explain thematic errors, however. If the TPJ contains information about thematic relations (as our data suggest), and if damage to that region causes noise in the system, as Schwartz et al. argue, why would aphasic patients provide thematically related names? Surely if their ATL is intact, they would give the correct name; if it is slightly damaged, it would be more correct to give a taxonomic response. There is a sense in which thematic responses are clearly not the right kind of answer—"bone" is not an appropriate name for a picture of a dog. One potential answer is that the TPJ is involved in regulating language use, in addition to its role in conceptual processing. Noonan, Jefferies, Corbett, and Lambon Ralph (2009) argue that this region is part of a language regulation system (including prefrontal cortex) that coordinates sentence processing and semantic activation. In particular, when reading ambiguous words, it would help to assess thematic information to aid disambiguation (e.g., *money-bank* vs. *fishing-bank*), but speakers do not want to produce thematic information instead of the target word, so such associates must be inhibited in production. Noonan et al. propose that a number of aphasic symptoms following damage to the TPJ could be caused by failure to inhibit related semantic information.

Although this notion seems useful in explaining thematic naming errors, our own results suggest thematic knowledge is also processed in the TPJ, given that our task did not involve production or sentence comprehension but merely judgments of relatedness. Semantic control does not seem to be a critical part of that task. Therefore, we propose that the TPJ is involved in detecting thematic relations and also in some form of attentional regulation, perhaps related to controlling the use of those relations in language processing. Thematic errors may derive primarily from the loss of control aspect, in which thematically related names are incorrectly produced, especially when the correct name is not immediately forthcoming.

Summary. Our finding of differential modulation of the ATL and TPJ for taxonomic vs. thematic items in general corroborates Schwartz et al.'s (2011) proposal for two distinct systems for representing different kinds of semantic knowledge. Our results suggest that the ATL represents taxonomic knowledge, while the TPJ (to a lesser degree) seems to represent thematic knowledge. These computations occurred during similar time windows (beginning as early as 300 ms) at about the same time as the beginning of lexical access as observed in previous visual word recognition studies (Lewis et al., 2011; Simon et al., 2012). The specificity of the TPJ for thematic knowledge, however, is less clear, as the TPJ responded to Them-Fill and thematic relatedness, to Tax-Fill, and also filler relatedness. Also, unlike for the ATL, we did not find differences in TPJ activation between conditions. This pattern of results suggests that the ATL plays a strong role in taxonomic judgments, while the TPJ's role appears more generic.

In the Introduction, we raised the question as to whether taxonomic and thematic information are represented together. The answer seems to be both yes and no. ATL activation seemed closely predicted by taxonomic relations, with no influence of thematic relations in the single-trial analyses. The importance of the ATL to taxonomic knowledge is also clear from studies of semantic dementia (e.g., Lambon Ralph et al., 2001; Mesulam et al., 2009). However, the TPJ showed influences of both kinds of relations. As we remarked earlier, thematic relations may be based on events and situations, but the knowledge of those events is about taxonomic categories: what kinds of entities are present and how the entities interact with one another. Therefore there must be some representation that links co-occurring taxonomic categories. That is, understanding the thematic relation of dogs to bones requires one to represent the taxonomic categories of dogs and bones. Consistent with Schwartz et al.'s (2011) findings, those connections seem to be made in the TPJ.

CHAPTER 3: THE ROLE OF THE TEMPORAL PARIETAL JUNCTION IN CONCEPTUAL RELATIONS

Gwyneth Lewis, David Poeppel, and Gregory L. Murphy

3.1 Abstract

The present study examines the neural processes that accompany conceptual relations. Previous neuroimaging results implicate the anterior temporal lobe (ATL) in processing feature-based (taxonomic) concepts and the temporal parietal junction (TPJ) in processing event-based (thematic) concepts (e.g., Schwartz et al., 2011). In recent work, we confirmed ATL selectivity for taxonomic relations but also found TPJ involvement in both kinds of relations (Lewis, Poeppel, & Murphy, 2015). Building on this result, we conducted an fMRI experiment to test whether the TPJ reflects conceptual inhibition vs. comceptual activation. We trained participants to make semantic relatedness judgments based on taxonomic vs. thematic criteria in two tasks (i.e., one task required inhibition of thematic relations and activation of taxonomic relations, while the other task required the opposite responses). While TPJ responses both to positive thematic pairs indicated conceptual activation, responses to negative thematic pairs indicated both conceptual activation and inhibition. Finally, responses to negative (but not positive) taxonomic pairs were associated with conceptual inhibition. Results of this experiment suggest that the TPJ plays a multifaceted role in conceptual processing that includes inhibition and activation of conceptual relations depending on the type of conceptual relation and the type of response required.

3.2 Introduction

As described in the general introduction, the role of the temporal parietal junction (TPJ) in conceptual relations remains an unresolved topic in semantic memory research. One important objective for researchers is to determine whether taxonomic and thematic relations are represented by the same semantic system (the *one-system view*; Jackson, Hoffman, Pobric, & Lambon Ralph, 2015) or two distinct but complementary systems (the *two-systems view*; e.g., Schwartz et al., 2011). While both agree that the anterior temporal lobe (ATL) represents taxonomic knowledge, the two-systems view argues for an additional specialized semantic hub for thematic relations in the TPJ (e.g., Maguire, Brier, & Ferree, 2010; Mirman & Graziano, 2012; Schwartz et al., 2011; Vivas et al., 2016).

Evidence for the two-systems view comes from lesion data linking specific conceptual naming deficits with damage to the ATL and to the TPJ. One of the first of such studies found that stroke aphasics with TPJ damage tended to make thematic errors in picture naming (e.g., labeling a dog "leash") at a rate of 27%, while semantic dementics with ATL damage tended to make taxonomic errors (e.g., labeling a dog "seal") at a rate of 99% (Jefferies & Lambon Ralph, 2006). Rather than attributing deficits in stroke aphasia to loss of thematic representations, Noonan, Jefferies, Corbett, and Lambon Ralph (2010) hypothesized that damage to parietal or frontal areas may reflect problems in regulating semantic control processes during semantic retrieval (p. 1598). Across several tests of semantic control, they associated parietal damage with the following deficits: (1) poor semantic navigation—patients could access semantically proximal but not distant meanings (e.g., CAP-HAT but not CAP-SOCK), (2) poor inhibition of competitors—patients could inhibit weak but not strong semantic competitors (e.g., NEAT-

TIDY but not HAPPY-CHEERFUL), and (3) poor cognitive control over irrelevant information—patients could access more dominant meanings of polysemous words without the help of a cue (e.g., fire – hot*, I lit a fire*) but not less dominant meanings (e.g., fire – rifle, *Fire at will*). This is further supported by results from a TMS experiment that varied semantic and non-semantic control demands (Whitney, Kirk, O'Sullivan, Lambon-Ralph, & Jefferies, 2012). Processes in the TPJ were shown to be domain independent in that parietal areas showed involvement both in (top-down) semantic selection and non-semantic control. These findings suggest that the TPJ supports complementary control mechanisms important for ultimate semantic retrieval but not directly involved in semantic processing.

This account is challenged by Schwartz et al. (2011), who conducted an analysis of picture naming errors of patients with damage to the ATL or TPJ. Picture naming data came from lesion patients who viewed and labeled pictures of objects. Schwartz et al. coded naming errors as taxonomic or thematic. Thematic errors included substituting the label of a complementary object from a different category (e.g., naming a dog a "leash"). Taxonomic errors included substituting a category coordinate label (e.g., naming a dog a "seal," a subordinate label (dog → "hound"), or a superordinate label (dog → "mammal"). As in Jefferies and Lambon Ralph (2006), taxonomic labels dominated the naming errors of both kinds of patients' responses. This is not a surprising finding given that substituting a taxonomic label (e.g., seal for dog) seems more appropriate than substituting a thematic label (e.g., leash for dog), especially in picture naming. When, however, Schwartz et al. regressed out shared variance between error scores to isolate taxonomic from thematic error types, they found that TPJ damage predicted thematic errors and ATL damage predicted taxonomic errors. Schwartz proposed that the ATL and TPJ likely reflect different semantic roles for events and objects, in which the ATL extracts featural information and the TPJ extracts spatial/functional relations. Schwartz et al. theorized that if the TPJ represents cognitive-control mechanisms rather than thematic specialization (e.g., Noonan et al., 2010; Whitney et al., 2012), then TPJ damage should have led to off-task "non-noun" or "off-task" naming errors, in which patients retrieve feature-specific labels (e.g. DOG → "it goes woof"). On the contrary, "off-task" naming errors were associated with ATL but not TPJ damage. This suggests that the TPJ reflects thematic semantics rather than top-down control mechanisms. As such, Schwartz et al. explained their findings as follows: During picture naming of an item like DOG, the ATL signals for the correct label ("dog") and for taxonomically related concepts ("seal"). Damage to the ATL may lead to loss of features such as "four legs," thus making "dog" confusable with "seal." As for the TPJ, if its role is to link contextually relevant information to the target concept, it should normally signal for "dog" and weakly for thematically related concepts such as "leash." Although thematic errors are extremely rare, damage to the TPJ would increase the likelihood of thematic substitutions by reducing the signal strength of "dog" relative to "leash."

Later work by Mirman and Graziano (2012) tested two alternative accounts of Schwartz et al.'s results: (1) TPJ lesions leads to errors in word-production but not thematic comprehension, or (2) TPJ lesions damage cognitive control but not thematic semantic processing (e.g., Noonan et al., 2010). To test these alternative accounts, Mirman and Graziano (2012) compared eye-fixations of neurologically intact subjects to those of aphasics with predominately ATL or posterior lesions during a visual world paradigm experiment (e.g., Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995). During the experiment, participants heard a target word followed by a screen with four images: an object that matched the target word, an object that was taxonomically or thematically related to the target word, and two distractor objects. To minimize cognitive and semantic control demands, participants were instructed to passively view the objects. This enabled the experimenters to track incidental activation of thematic and taxonomic relations during word-comprehension (as opposed to word-production as in Schwartz et al., 2011). First, it was found that relative to controls, parietal patients showed delayed and decreased fixations for thematically (but not taxonomically) related objects. Second, relative to controls, ATL patients showed longer fixations on taxonomically (but not thematically) related objects, which could reflect longer-lasting competition between taxonomic representations. This finding of neural dissociation for taxonomic and thematic relations in a semantic comprehension task with minimal cognitive control demands bolsters Schwartz et al.'s (2011) account of two distinct semantic systems.

In recent work, Jackson et al. (2015) noted three phenomena in the literature that contradict Schwartz et al.'s account of distinct semantic systems: TPJ damaged patients tend to make both taxonomic and thematic errors, TPJ and ATL damaged patients both show deficits in using thematic relationships, and semantic neural dissociations may reflect cognitive but not representational differences. To test whether results of Schwartz et al. (2011) could be accounted for by fundamental differences in semantic control demands, Jackson et al. (2015) conducted an fMRI experiment that manipulated the type of semantic relation (taxonomic vs. thematic), and the level of semantic control (hard vs. easy). Participants completed the following tasks: (1) A semantic judgment, which presented participants with taxonomically or thematically related pairs plus an unrelated foil; (2) a semantic judgment task that varied control demands by presenting taxonomically related pairs like BARREL and BOX with a weakly related foil like PLUM (low demand) or a highly related foil like SEAT (high demand); and (3) a non-semantic letter-matching task (high vs. low demand) that served as baseline for comparison with data from the semantic tasks. These tasks allowed them to test predictions from the two-systems view for TPJ involvement in thematic relations and ATL involvement in taxonomic relations, the one-system view of ATL involvement in both relations, and a mixed account of Schwartz et al. in which the TPJ's involvement in thematic relations can be accounted for by semantic control demands. Contrasts between thematic and taxonomic trials from the semantic judgment task showed greater activation for thematic items across parietal areas (including supramarginal gyrus and angular gyrus). Greater activation for taxonomic trials was found in the frontal gyrus, including the precentral gyrus (as found in Sachs et al., 2008). Whole-brain analyses were conducted with response time (RT) and level of semantic control to identify areas modulated by semantic difficulty. Conjunction analyses with these areas and the taxonomic > thematic contrasts were in almost complete overlap over frontal and supplementary motor areas, suggesting that the taxonomic > thematic difference could be attributed to greater executive demands for more difficult taxonomic judgments. In general, it was found that when RT was included as a parametric regressor, no differential activation for taxonomic vs. thematic relation was found in any brain region. Finally, regions-of-interest (ROI) analyses based on coordinates from Schwartz et al. (2011) showed indiscriminate ATL activation for both kinds of semantic relations, but not for the non-semantic task conditions. The TPJ, however, showed indiscriminate deactivation for both kinds of semantic relations as well as for the non-semantic task conditions. As such, it may be that the TPJ is indiscriminately involved in cognitive but not necessarily semantic control demands.

In sum, results of Jackson et al. (2015) suggest that differing levels of control required for taxonomic vs. thematic relations could account for previous reports of neural dissociations. The authors proposed that the ATL encodes not just the features of concepts, but also the thematic links between spatially or functionally associated items. These associations can then be thought of as types of "features" that become integrated through experience. Although encoded together, they may entail different levels of semantic control because features and associations should vary in terms of the range of concepts they link to, as well as the way in and frequency with which they are experienced.

In response to Jackson et al. (2015), Vivas et al. (2016) argued that if the same unified system represents features and associations, then parietal damage should lead to deficits in using both kinds of relations. By using a forced-choice naming task in which aphasics and neurologically intact patients indicated one of two items that went best with a third, along with a free-choice picture-sorting task in which patients organized pictures based on personal preference for taxonomic or thematic criteria. First, controls showed significantly lower taxonomic scores relative to thematic scores and longer response latencies for taxonomic pairs relative to thematic pairs in both tasks, suggesting that taxonomic relations are more complex or require more activation. Second, patients with parietal lesions had significantly lower thematic scores in both verbal and non-verbal modalities. Assuming that parietal damage predicts loss of semantic control (e.g., Noonan et al., 2010) the authors argued that they should have found worse performance with both relations, particularly in using taxonomic relations, which were apparently more difficult.

Neural suppression and enhancement. Experiments using fMRI to study semantic priming have identified thematic-taxonomic neural dissociations. Such experiments examine effects of semantically related vs. unrelated probe-target stimuli on changes in the BOLD response at the onset of the target item. Such studies interpret reduced activity (neural suppression) and increased activity (neural enhancement) as reflecting less and more effortful processing, respectively. One of the first of such experiments designed to directly test thematic-taxonomic dissociation employed auditory word pairs that related taxonomically or thematically in a lexical decision task (Kotz, Cappa, von Cramon, & Friederici, 2002). It was found that taxonomic relative to thematic pairs led to greater activation in right precuneus and cuneus, areas, suggesting that taxonomic pairs required more effort (individual contrasts between unrelated and related conditions were not reported). Later work used fMRI and a category construction paradigm (e.g., Lin & Murphy, 2001) to investigate whether thematic and taxonomic relations require different cerebral processing demands (Sachs et al., 2008a). In the task, participants chose which of two words went best with a target word. A balanced condition included thematic and taxonomic choices (e.g. CAR → GARAGE/BUS), a biased thematic condition included a thematic choice and a distractor (e.g. CAR → GARAGE/ERASER), and a biased taxonomic condition presented a taxonomic choice and a distractor (e.g. CAR → BUS/ERASER). It was found that the biased taxonomic and thematic conditions recruited similar brain areas in the left hemisphere, particularly in parietal areas (precentral and angular gyrus). Contrasts of each choice in the balanced condition showed that taxonomic choices additionally recruited the left thalamus (implicated in rule-based processing), right middle frontal gyrus (implicated in abstract and verb processing), and left precuneus (implicated in mental imagery). The authors propose that taxonomic categories differ from thematic categories as follows: they are less associated and require more effort for activation, they require assessment of more abstract relationships, and they require mental imagery processing to assess the match between items. This, however, does not explain why no differences in terms of response time or brain activation were identified between the two biased conditions. It may be that differences were only observed in the balanced condition for taxonomic choices because thematic relations are more "automatic." Thus, when a taxonomic choice is to be made, some of these brain areas may be involved in inhibitory processes, i.e., inhibiting competing thematic relations in order to process the taxonomic relation. This is plausible because these areas were not recruited when participants chose between taxonomic and unrelated options.

Work by Sachs et al. (2008b) used event-related fMRI with automatic priming in which a word (e.g., LEASH or SEAL) appeared in rapid succession after a prime (e.g., DOG). The brief stimulus-onset-asynchrony was intended to reduce conscious relational processing to minimize effects of task-related processing strategies during lexical decision. While thematic priming was associated with reduced activity (including the right precuneus and frontal gyrus), taxonomic priming was associated with increased activity in the right precuneus. Sachs et al. (2008b) speculated that the finding of greater activation in right precuneus for taxonomic stimuli could account for behavioral response disadvantages typically found for taxonomic stimuli relative to thematic stimuli (e.g., lower accuracy and longer response times). Recall that Sachs et al. (2008a) proposed that enhancement in left precuneus for taxonomic stimuli could reflect an additional imagery processing stage for taxonomic items. Here, however, the authors interpreted greater precuneus activation as reflecting more effortful processing due to lower salience of taxonomic relative to thematic categories (e.g., *button-coat* is less salient than *jacket-coat*).

Further investigation of dissociation employed a controlled-priming paradigm (Sachs et al., 2011). The use of a longer SOA between the prime and target was manipulated to engage strategic cognitive processes during lexical decision, such as response inhibition. Thematic primes enhanced left inferior frontal activity, while taxonomic primes suppressed responses in right middle frontal gyrus. Contrasts between conditions showed enhancement for taxonomic items in the right insular lobe again suggesting greater effort for taxonomic stimuli.

While these semantic priming studies of neurologically intact participants demonstrate neural dissociability of thematic and taxonomic relations, none of these dissociations involved the ATL or TPJ as predicted by Schwartz et al. (2011). One semantic priming study to report a parietal dissociation for thematic relations was by Kircher, Sass, Sachs, and Krachs (2009). Their study was designed to test whether thematic representations exist for verbal vs. non-verbal information (e.g., words vs. pictures) or whether semantic processes ultimately converge. Using fMRI, they contrasted effects of thematic and unrelated pairs in unimodal and cross-modal priming conditions on brain responses. In both priming conditions, thematic stimuli led to response suppression in parietal areas relative to unrelated pairs, suggesting that this region represents an amodal semantic system specialized for thematic information, independent of modality.

Maguire et al. (2010) used EEG to demonstrate a taxonomic-thematic dissociation in oscillatory activity over parietal areas. Increases in alpha power have been associated with semantic processes relevant for attention necessary for identifying words relative to non-words (e.g., Bastiaansen, Oostenveld, Jensen, & Hagoort, 2008), while increases in theta power have been associated with engagement of episodic memory processes (Klimesch, Schimke, &

Schwaiger, 1994). In an auditory lexical decision EEG experiment, Maguire et al. (2010) examined the scalp topography of EEG power changes for taxonomically vs. thematically related words. First, they found that increases in theta for thematic relations over frontal and parietal areas, indicating that thematic processing engages memory resources. The second and more interesting finding was that of decreases in alpha for thematic priming that mirrored increases in alpha for taxonomic priming. That is, decreases in alpha for thematic priming was of the same magnitude and over the same temporal window as the increases in alpha for taxonomic priming. While alpha increase for taxonomic relations could reflect the need for greater processing resources (e.g., Kotz et al., 2002; Sachs et al., 2008a; 2008b), the authors proposed that it might instead reflect inhibition of thematic memory traces. That is, the first word may have automatically activated thematic links that were irrelevant for retrieving the taxonomic relation with the second word, and were thus inhibited.

If the TPJ represents thematic relations, some of which must be inhibited during retrieval, this could explain our recent finding of an association between TPJ activity with both kinds of relations (Lewis et al., 2015). In an MEG experiment with semantic priming, we contrasted effects of related and unrelated word pairs, as well as strength of taxonomic and thematic relatedness on the ATL and TPJ responses. Schwartz et al.'s (2011) predictions of selective ATL and TPJ involvement in taxonomic and thematic relations (respectively) were partly confirmed in that ATL activity correlated with strength of taxonomic relatedness and distinguished taxonomic from unrelated pairs, and TPJ activity correlated with thematic relatedness and distinguished thematic from unrelated pairs. However, TPJ activity additionally distinguished taxonomic from unrelated pairs over the same time window. While this could reflect direct involvement of the

TPJ in both thematic and taxonomic processing, it could also indicate suppression of competing thematic representations. Although inhibition was not explicitly required, it is possible that the differentiation we observed at the TPJ for taxonomic pairs (e.g., DOG-COW) reflected inhibition of thematic relations (e.g., *leash, bone, collar*) that were initially activated by the first word but required inhibition in order to process the relations.

Results from a recent fMRI experiment suggest the ubiquitous finding of thematic facilitation (e.g., faster RT and higher accuracy) over taxonomic items may be attributable to differences in the mechanisms they involve (de Zubicaray, Hansen, & McMahon, 2013). Using a picture-word interference naming paradigm, de Zubicaray et al. (2013) manipulated taxonomic and thematic relatedness of distractor words to identify facilitation and semantic interference effects in naming latencies and neural activity. Relative to control (unrelated) distractors, thematic distractors facilitated naming and reduced middle temporal gyrus (MTG) activity, while taxonomic distractors delayed naming and decreased activity both in middle MTG and posterior MTG. It was proposed that the differential finding of decreased pMTG only for taxonomic relations could reflect lexical competition that thematic pairs do not induce, wherein only the taxonomic items involve activation of shared features of category members, which elicits a cohort of lexical competitors, and thus interference. In sum, while both thematic and taxonomic relations involve conceptual processing stages, priming is only observed for thematic relations because the conceptual processing stage necessarily induces interference in the taxonomic case. Lastly, thematic items increased activity in the left angular gyrus more than taxonomic items did. This could reflect an additional mechanism for thematic items involved in processing order information given that thematic relations link objects that occur in the same event.

The present experiment. The present fMRI experiment tests competing predictions for inhibition vs. thematic processing, which our MEG experiment (Chapter 2) did not address. Unlike previous functional imaging studies of conceptual relations, we examined whether neural inhibition could be manipulated factorially by crossing word-pair relation type–thematic, taxonomic, and unrelated (fillers)–with tasks requiring inhibition of either taxonomic or thematic knowledge. First, we conducted an online survey of participants who rated the relatedness of thematic, taxonomic, and unrelated word pairs. We then used the ratings to select word pairs for the main experiment such that conditions were matched for relatedness across tasks. This was of particular importance because thematic pairs tend to have higher relatedness ratings than taxonomic pairs (Lewis et al., 2015). We additionally balanced the conditions in terms of linguistic properties such as written frequency, syllables, and word length. We employed the final stimulus set in an event-related fMRI experiment to explore predictions from models of semantic organization. We manipulated the following factors: type of semantic relationship (taxonomic vs. thematic vs. filler), and type of process required (inhibition vs. conceptual). We required inhibition by requiring either "positive" or "negative" responses for each relation type in two tasks. In one task, participants were trained to make taxonomic judgments, wherein words of the same general kind (e.g., COTTAGE-CASTLE) should receive positive relatedness judgments (TaxYes) and thematic and filler (unrelated) pairs should receive negative judgments (ThemNo and Fill1). In a separate task, participants were instructed to make thematic judgments in which word pairs denoting a thematic relationship (e.g., DOG-LEASH) should receive positive relatedness judgments (ThemYes), while taxonomic and unrelated word pairs should receive negative judgments (TaxNo and Fill2). Thus, each task entailed inhibition of the opposition relation. This allowed us to test whether TPJ activity differs as a function of the type of relation that must be inhibited and by the type of relation that must be processed.

The inclusion of the filler condition in each task gave us a baseline response for each task with which we can contrast relative differences between tasks. For between-task contrasts, the thematic account predicts greater TPJ activity for accepting and rejecting thematic pairs over accepting and rejecting taxonomic pairs, respectively. The inhibition account predicts greater TPJ activity for rejecting thematic and rejecting taxonomic pairs over accepting thematic and accepting taxonomic pairs, respectively. For within task contrasts, the thematic account predicts greater TPJ activity for accepting thematic than rejecting taxonomic, while the inhibition account predicts the opposite result. Table 1 shows predictions from the thematic account and inhibition account.

			Task	
			Taxonomic	Thematic
Word Pair	Taxonomic	SEAL and DOG	**X** **X**	**X** ✔
	Thematic	LEASH and DOG	✔ ✔	✔ **X**
	Unrelated	NAIL and MUD	**X** **X**	**X** **X**

Table 1: Predicted TPJ activity under thematic account (red) and inhibition account (blue). The check-marks denote increased TPJ activity; the X’s denote no increase in TPJ activity.

3.3 Method

Participants. Eighteen right-handed native English speakers (ten males; mean age = 27) recruited from the New York University community participated in the fMRI experiment. Participants completed a safety screening form and provided written informed consent. The University Committee on Activities Involving Human Subjects (UCAIHS) of New York University approved the study. Two subjects were excluded from the analysis (one due to inadequate coverage of the parietal lobe, one due to low response accuracy).

Materials. The fMRI experiment employed a 2 x 3 factorial design with repeated measures. Task-type (taxonomic or thematic) and word-pair condition (taxonomic, thematic, control) were varied within subjects. Stimuli were drawn in part from our previous experiment (Lewis et al., 2015), from Battig & Montague (1969), and the University of South Florida Free Association Norms corpus (Nelson, McEvoy, & Schreiber, 1998). Thematic pairs could relate spatially (e.g., TABLE-LAMP) or functionally (e.g., THREAD-NEEDLE) but not taxonomically (e.g., CAT-MOUSE). Taxonomic pairs consisted of basic-level category members (e.g., SEAL-DOG, both are mammals) that did not share an obvious thematic relation (e.g., CAT-MOUSE). Filler pairs (e.g., PANE-NUT) were unrelated.

Relatedness ratings from online Qualtrics participants (n = 32) were acquired similarly as in Lewis et al. (2015). Participants rated the relatedness of 400 word pairs on a 1-7 scale (1 = not at all related and 7 = highly related). Critically, the instructions did not include 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. “Sanity-check” items, which were included to ensure participants took the task seriously, were 10 highly related (e.g. SQUARE-

CIRCLE), and 10 highly unrelated (SCHEME-MOOSE) word pairs. Three participants were excluded because the differences between their mean highly related and highly unrelated sanity check ratings were too low. On average, thematic pairs were the most related (*M* = 5.6, mean *SD* = 1.1, *SD* of the mean = .73), followed by taxonomic pairs (*M* = 4.5, mean *SD* = 1.3, *SD* of the mean = .79), followed by fillers pairs (*M* = 1.6, mean *SD* = 0.96, *SD* of the mean = .48).

The final stimulus set included a total of 320 word pairs: 120 thematic, 120 taxonomic, and 80 controls. The conditions were matched for psycholinguistic properties of word pairs across tasks to ensure that any effects could not be attributed to differences in the stimuli other than the extended manipulation. Table 2 shows the properties by task and condition.

	Relatedness		Concreteness		Length		Frequency		Syllables	
Word Pairs	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Taxonomic Task										
Taxonomic	5.2	0.5	4.7	0.2	5.7	1.3	7.8	1.2	1.7	0.6
Thematic	5.3	0.6	4.7	0.3	5.9	1.2	8.1	1.1	1.6	0.5
Thematic Task										
Taxonomic	5.2	0.5	4.7	0.4	5.9	1.5	7.9	1.4	1.8	0.5
Thematic	5.3	0.6	4.7	0.3	5.6	1.4	8.1	1.1	1.8	0.6
Filler Pairs (in each task)										
Control	1.5	0.3	4.6	0.3	5.4	1.0	7.8	1.0	1.7	0.6

Table 2: Stimulus Properties by Task and Condition (Means and SDs Across Pairs).

*Note: Frequency is measured as the log transformed values of written frequency norms based on 131 million words in the Hyperspace Analogue to Language (HAL) corpus (Lund & Burgess, 1996).

Experimental tasks. The experiment had two main tasks: a taxonomic and a thematic task, each occurring over two sequential two runs. Before each task, participants read instructions for relatedness judgment criteria. The taxonomic task defined relatedness in terms of items that are the same general kind of thing (e.g., MELON and BERRY: both are kinds of fruit) but do not necessarily interact or occur together (e.g., HEN and CAGE: while hens live in cages, only one is an animal). The thematic task instructions defined relatedness in terms of things that tend to occur or interact together (e.g., BEE and HIVE: bees tend to live in hives) but do not necessarily share features (e.g., BARN and CONDO: while both are kinds of buildings, barns and condos do not tend to interact). After reading the instructions, participants practiced making the appropriate kind of judgment with 30 pairs (15 thematic, 15 taxonomic). Each practice was completed in the scanner. Participants received feedback and explanations on each practice trial to ensure they understood the task. Instructions and practice items for each task are shown in the Appendix. Immediately after each practice, participants went through two sequential runs of the experimental trials. Word pairs were presented vertically with one word above the other to reduce saccadic eye movements. The words were separated with the word “and” so that the words would not be processed as compounds. Figure 1 shows a depiction of the practice trial sequences.

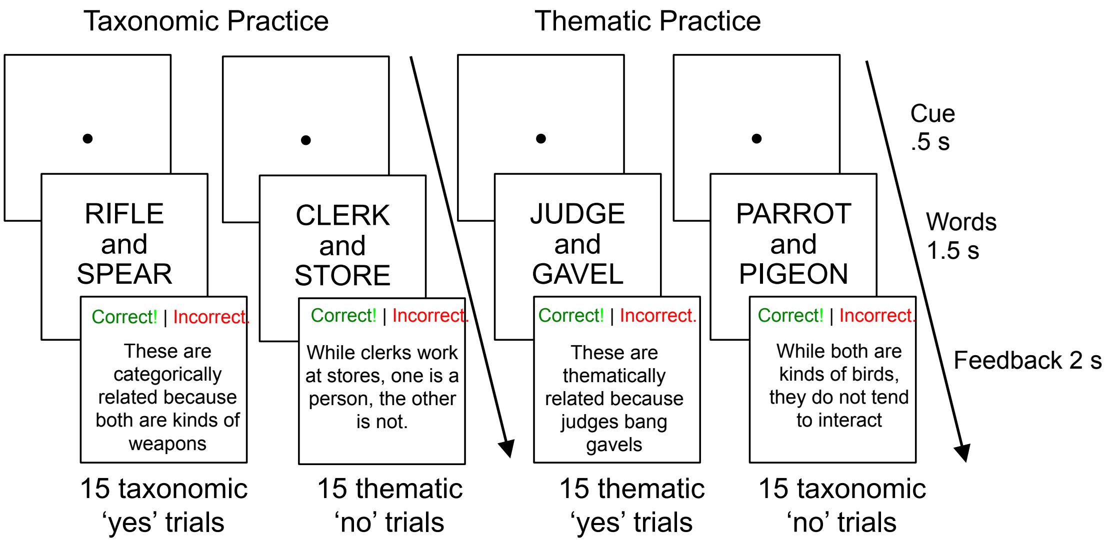

Figure 1: Trial sequences in the practice sessions prior to each task.

Procedure. The main experiment presented a total of 320 word pairs across four 7-min scans (80 pairs in each). Task order was counterbalanced such that half of the participants completed the taxonomic task in runs 1-2 and the thematic task in runs 3-4 and vice versa. None of the items were repeated within or across tasks. Stimuli were presented using MATLAB (MathWorks) Psychtoolbox and MGL scripts (available at http://justingardner.net/mgl) via a MacMini computer and projected on a screen in the magnet. The trial sequence in the main experiment was similar to the practice, only trials pseudo-randomly ordered and were jittered by a variable ITI (2 - 6 s, with optimal fMRI schedules programmed in Optseq2). Feedback (i.e., "Incorrect") was only provided on incorrect trials. Figure 2 shows the trial sequence.

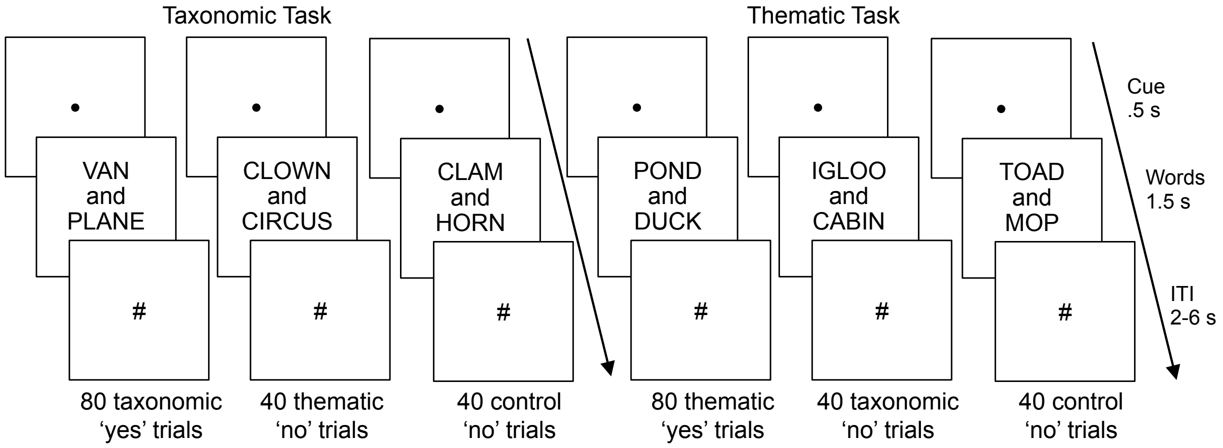

Figure 2: Trial sequence in the fMRI experiment.

Data acquisition. Scans were collected with a 3-T Siemens Allegra system with a Nova Medical NM-011 head coil at New York University's Center for Brain Imaging (CBI) lab. 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 x 240 mm], slices = 34 [slice thickness = 3 mm, voxel resolution = 3 x 3 x 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 angular gyrus. Magnetic field estimates were acquired with a field map sequence for later use in image distortion correction. Standard sagittal, high-resolution structural images were acquired with a magnetization-prepared rapid-acquisition gradient echo (MPRAGE) (voxel resolution: 1 x 1 x 1 mm).

Data analysis. For each subject, time-series data from each of the four runs were individually modeled using general linear modelling (GLM; FILM; FMRIB Improved Linear Model). The GLM for each run included the following main event-related regressors: correct responses from taxonomic, thematic, and filler trials. Incorrect trials (~18% of the data) were modeled as a variable of no interest. Event files for each variable were read into FEAT, with the duration of the main event-related variables set to the mean response time (RT) across all events types in the run, and the weight of each event set to 1. The model additionally included RT and relatedness score as parametric regressors. Rest periods, which included fixation and time between trials when the stimulus was not present, were defined implicitly by exclusion from the model. This resulted in three contrast parameter estimate images (COPEs) for each run. Estimates from runs in the same task of the same type were combined within-subject in fixed effects models that forced random effects variance to zero, and higher-level contrasts for each subject were computed using fixed effects models in FLAME (FMRIB's Local Analysis of Mixed Effects; Beckmann, Jenkinson, & Smith, 2003; Woolrich, 2008; Woolrich et al., 2004). The higher-level COPEs for each subject were as follows: (1) TaxYes vs. Fill, (2) ThemNo vs. Fill, (3) TaxNo vs. Fill, (4) ThemYes vs. Fill, (5) ThemNo vs. ThemYes, (6) ThemNo vs. TaxNo, (7) ThemNo vs. TaxYes, (8) ThemYes vs. TaxYes, and (9) TaxNo vs. TaxYes. Group-level analyses employed COPEs of the same type in one-sample *t*-tests in FLAME. Resulting Z-statistic images were thresholded to define contiguous clusters, with voxel inclusion set at $Z > 2.3$ (corrected cluster significance of $p = .05$; Worsley, 2001).

Regions-of-Interest analysis. We used anatomical labels from the Harvard Oxford Cortical (available in FSLview) to form two main ROIs for each hemisphere. The ATL ROI comprised a merger of the temporal pole, anterior inferior temporal gyrus, and middle temporal gyrus (corresponding to BA 38, anterior BA 20, and anterior BA 21, respectively). The TPJ ROI comprised the angular gyrus (BA 39) and the anterior and posterior supramarginal gyrus. Figure 3 shows the regions of interest.

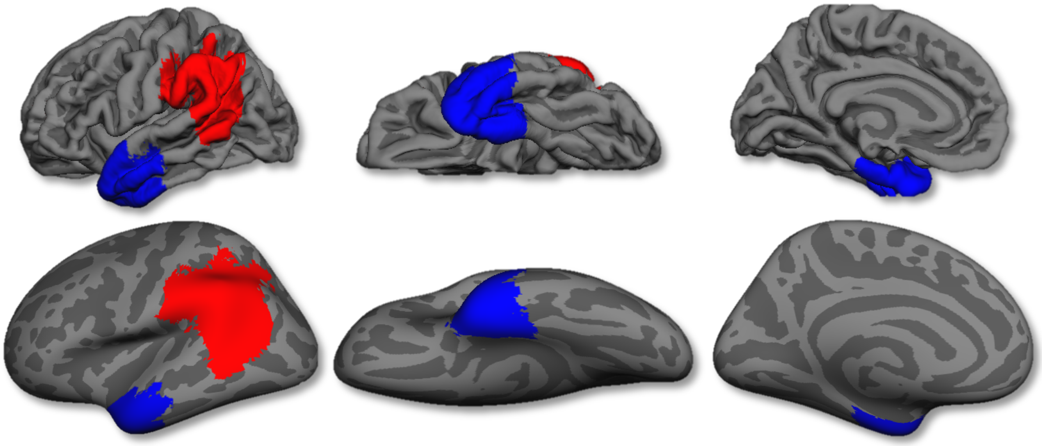

Figure 3: The ATL (blue) and TPJ (red) ROIs shown on the standard FreeSurfer brain. Pial surface (top); inflated surface (bottom). From left to right: sagittal, inferior, and medial views.

Whole-Brain analysis. A second analysis examined significant activity in each individual contrast (threshold for cluster significance at $p < .001$). We additionally employed contrasts weighted by relatedness strength to determine whether certain areas responded differentially to relatedness score in a given condition.

3.4 Results

Behavioral results. Average accuracy across trials was lowest for both ThemNo (75%) and TaxNo (77%). The same average accuracy score (84%) was found for TaxYes and ThemYes. For fillers, accuracy was considerably higher in the thematic (88%) than in the task (80%). When excluding responses past the 1.5 s response window, lowest accuracy was found for TaxNo (81%) and ThemNo (84%). Accuracy for fillers in the thematic task (95%) was slightly higher than in the taxonomic task (91%). Accuracy for ThemYes (89%) was slightly lower than TaxYes (92%).

Average response time (RT) on accurate trials was slowest for fillers (1101-1105 ms), which were always negative, and fastest for positive trials (982-999 ms). While RT differences within filler and positive conditions were negligible, negative trials showed a 40 ms difference between ThemNo (1091 ms) and TaxNo (1051 ms), suggesting greater effort required in rejecting thematic pairs. Single-trial correlations with relatedness score showed significant correlations with relatedness score: positive thematic pairs, in which higher relatedness was associated with faster responses ($r = -.205$, $p < .0001$), as well as control pairs in the thematic task, in which higher relatedness was associated with slower responses ($r = .104$, $p < .001$). None of the other correlations between relatedness and RT were significant ($p > .05$). Figure 4 plots RT and accuracy by condition and task and Table 3 shows the means and standard deviations.

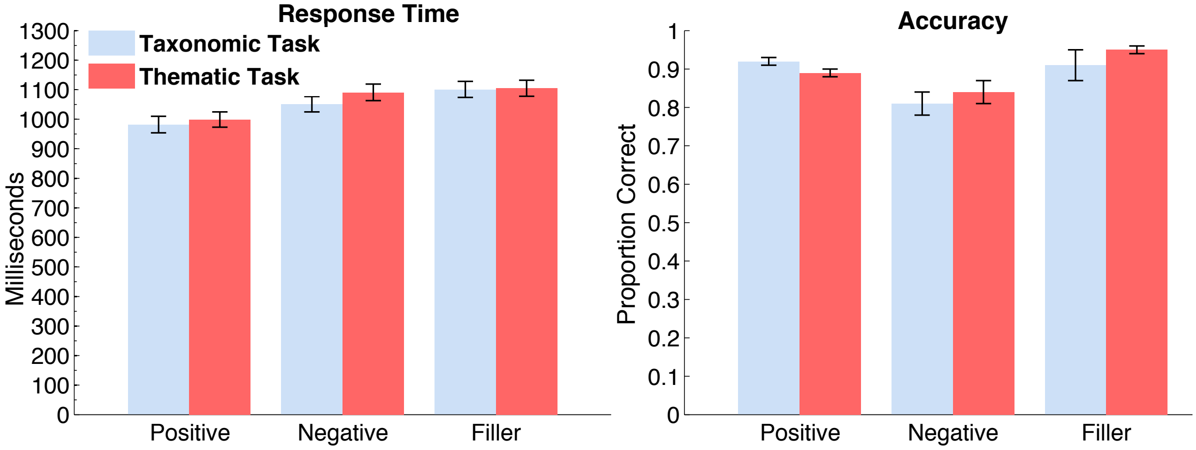

Figure 4: Average RT and proportion correct for accurate trials.

Condition	All Trials				Accurate Trials			
	RT (ms)		Accuracy		RT (ms)		Accuracy	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
TaxYes	1028	125	0.84	0.1	982	113	0.92	0.05
ThemYes	1039	116	0.84	0.09	999	107	0.89	0.05
TaxNo	1061	118	0.77	0.12	1051	107	0.81	0.13
ThemNo	1127	132	0.75	0.16	1091	112	0.84	0.12
Fill1	1147	118	0.8	0.18	1101	108	0.91	0.15
Fill2	1128	119	0.88	0.08	1105	111	0.95	0.05

Table 3: Response Time and Accuracy for Analyses Including All Trials and Correct Trials. Fill1 and Fill2 correspond to control (unrelated) pairs in the taxonomic and thematic task, respectively.

Regions-of-Interest results. The ROI contrasts focused on differences in activation within the TPJ and ATL ROIs. While activation within the ATL showed significant increase over rest in all conditions, none of the higher-level contrasts between conditions were significant. As for the TPJ, within-task contrasts between related and filler conditions showed significantly greater TPJ activation for ThemNo, ThemYes, and TaxNo over unrelated pairs. The contrast between TaxYes and filler did not yield a significant difference. Contrasts with related pairs showed greater TPJ activation for ThemNo relative to TaxNo. None of the other contrasts between related conditions yielded significant differences following cluster correction. In sum, the ROI analysis demonstrated more activity for: rejecting taxonomic and thematic pairs relative to rejecting unrelated pairs, accepting thematic (but not accepting taxonomic) pairs relative to rejecting unrelated pairs, as well as rejecting thematic relative to rejecting taxonomic pairs.

Figures 5 and 6 show the activation maps.

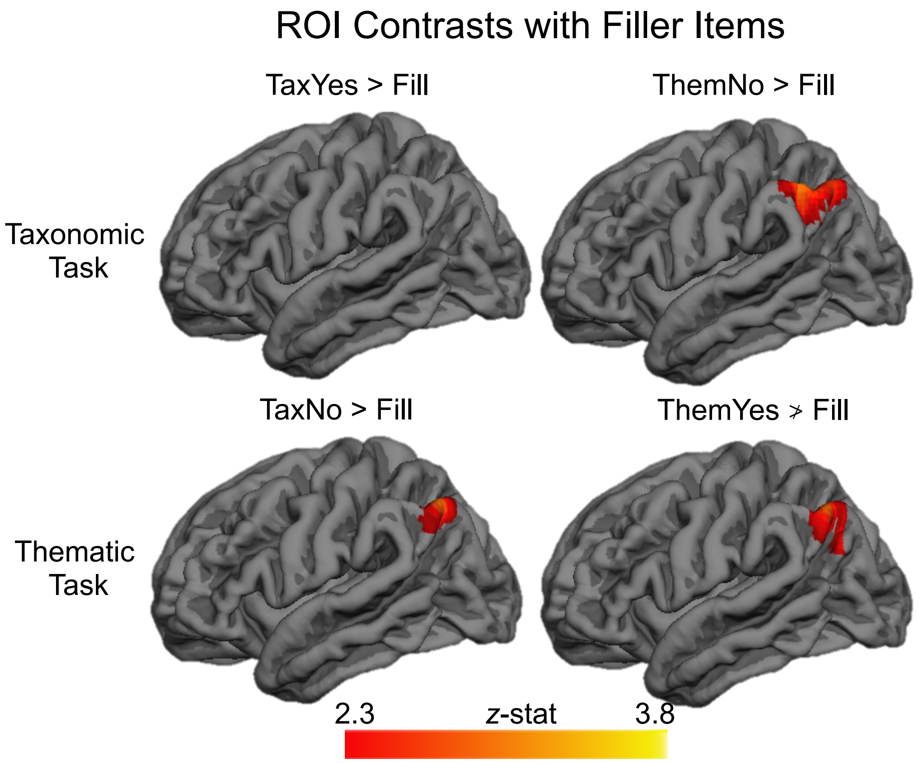

Figure 5: ROI activation for contrasts with filler pairs shown on the FreeSurfer brain with cluster correction applied (voxel inclusion threshold $z = 2.3$, cluster significance threshold $p < .05$).

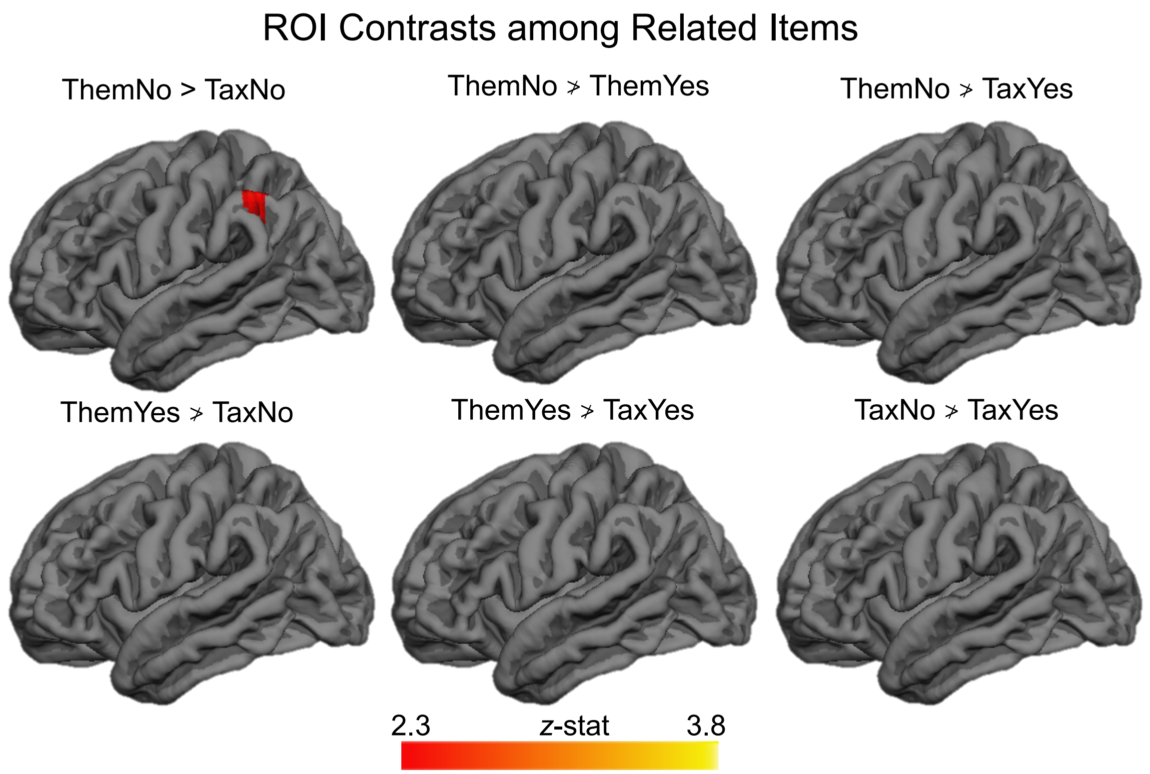

Figure 6: ROI activation for contrasts among related conditions shown on the FreeSurfer brain with cluster correction applied (voxel inclusion threshold $z = 2.3$, cluster significance threshold $p < .05$).

Whole-Brain analysis results. We supplemented our ROI analysis by contrasting activation between task conditions over the entire brain. To visualize the whole-brain contrasts between related and unrelated items, we initially used an uncorrected significance threshold of $p < .001$. Contrasts revealed clusters within the TPJ region that generally reflected greater activation for related relative to unrelated pairs, negative (related) pairs relative to positive pairs, and for thematic pairs relative to taxonomic pairs. The whole-brain analysis confirmed each result of our ROI analysis in that significantly greater activation within the TPJ area was found for negative taxonomic, negative thematic, and positive thematic pairs (but not positive taxonomic pairs) relative to unrelated pairs. The analysis additionally confirmed the TPJ ROI result of greater activity for negative thematic pairs relative to negative taxonomic pairs. Figures 7 and 8 show activation maps from the whole-brain analysis.

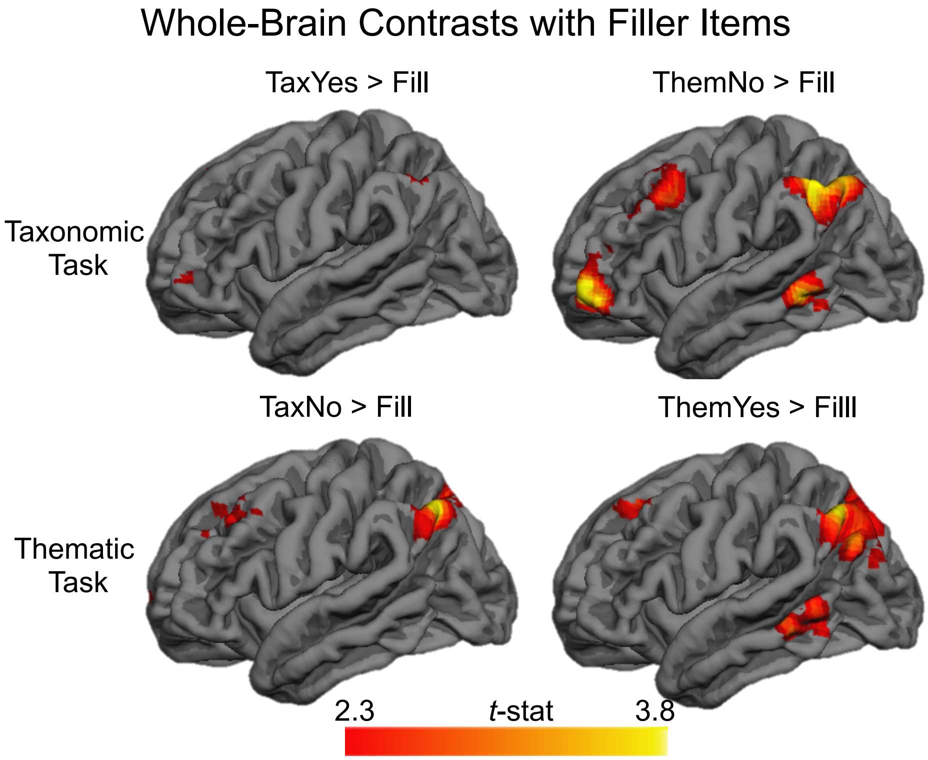

Figure 7: Whole-brain activation for contrasts with fillers shown on the FreeSurfer brain. Activation maps are shown at uncorrected $p < .001$.

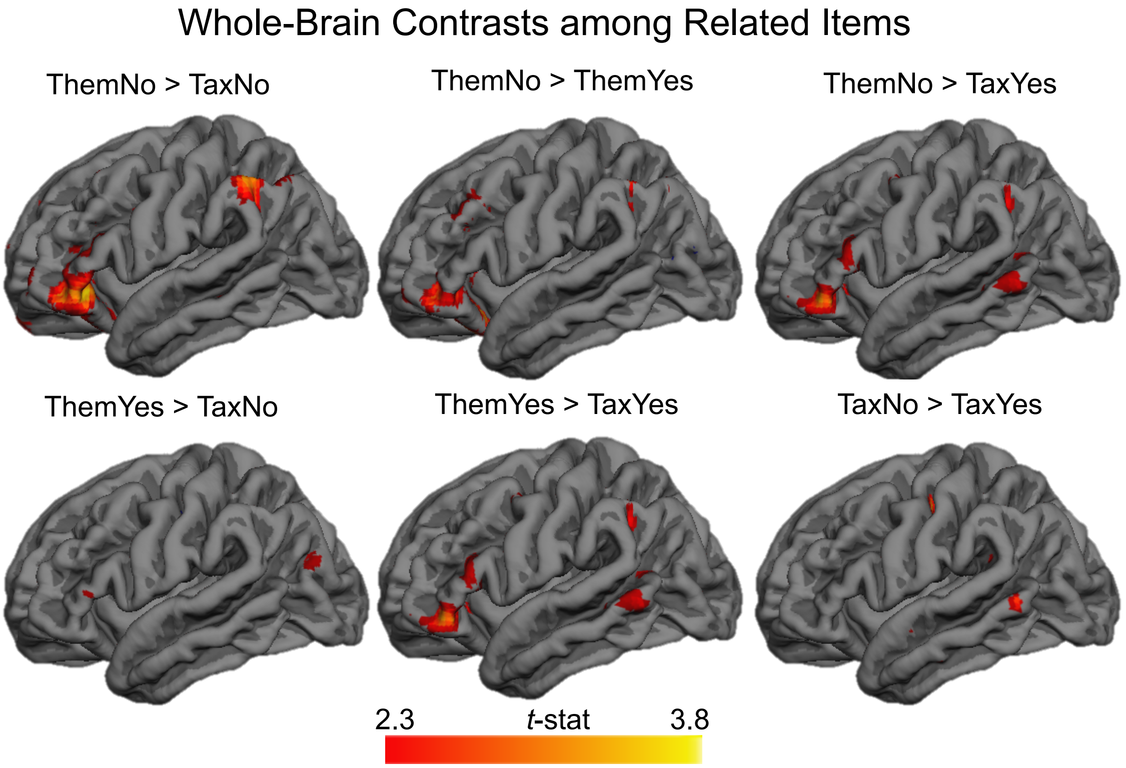

Figure 8: Whole-brain activation maps for related contrasts shown on the FreeSurfer brain. Activation maps are shown at uncorrected $p < .001$.

3.5 Discussion

This experiment was designed to test predictions for inhibition vs. semantic processing of thematic relations. While our primary goal was to determine whether previously reported associations between TPJ responses and thematically related word pairs reflect specialization for thematic concepts or inhibition of irrelevant activity, we were additionally interested in replicating taxonomic effects on ATL responses. Higher-level contrasts with ATL activation did not survive correction for multiple comparisons. Indeed, our fMRI acquisition procedure was not optimized for the ATL, which is notoriously difficult to image. According to Patterson, Nestor, and Rogers (2007), “the anterior temporal lobe is ‘shy’ to fMRI” because “the signal-to-noise ratio diminishes substantially near the temporal poles, owing to their proximity to air filled sinuses (the so-called ‘susceptibility artefact’)”, (p. 981). The remainder of this discussion thus focuses on the TPJ results.

First, the contrasts with filler conditions in the ROI analysis showed significantly greater TPJ activity for positive thematic, negative thematic, and negative taxonomic conditions (ThemYes > Fill, ThemNo > Fill, TaxNo > Fill). These results do not strongly speak to either account *per se*. Note, however, that the contrast between the positive taxonomic and filler condition did not reveal a significant difference at the TPJ. That is, the TPJ was very active for negative trials and thematic trials, but not for positive taxonomic trials responses. Thus, greater activation for negative taxonomic pairs seems to be related to inhibitory but not semantic processes. These effects were also evident in the whole-brain analysis, which revealed significant clusters for these contrasts in regions overlapping with the TPJ label. The crucial differences were revealed by the contrasts between the related conditions. The whole-brain analyses showed, in general, significant clusters in the TPJ, with greater activity for (1) thematic over taxonomic conditions and (2) negative over positive conditions. When these contrasts were made in the ROI analyses, it was found that the negative thematic condition (e.g., DOG-SEAL) yielded greater activity than the negative taxonomic condition (e.g., DOG-LEASH) (ThemNo > TaxNo). Although it could be argued that the thematic items were simply more difficult to reject, as suggested by our behavioral results, recall that our MRI model included both response time and relatedness score as parametric regressors. Moreover, our conditions were closely matched on a number of linguistic and psycholinguistic variables, such as word frequency and concreteness.

We propose that the difference in activity for ThemNo > TaxNo likely reflects both thematic semantics and inhibition in the TPJ.

Recall that Lewis et al. (2015; Chapter 2) previously found TPJ involvement in both taxonomic and thematic relations during processing of sequentially presented word-pairs. As mentioned in the introduction, it was speculated that the apparent involvement of the TPJ in processing taxonomic pairs (e.g., DOG-SEAL) could have reflected inhibition of thematic relations (e.g., *leash, bone, collar*) that were initially activated by the first word yet required inhibition in order to successfully process the taxonomic relation. Although the results of this fMRI experiment are less discriminating than originally hoped, the findings suggest a role of the TPJ both in thematic processing and inhibition. Given the large size of our TPJ ROI, it could be useful to focus on isolating which particular TPJ areas are uniquely involved in thematic but not inhibitory processing and vice-versa.

GENERAL DISCUSSION

The primary goal of this dissertation was to illuminate the neural mechanisms of conceptual relations. We approached this goal by focusing on ATL and TPJ responses to taxonomic and thematic stimuli. Unlike previous studies, we simultaneously controlled for multiple factors that could cloud interpretation of apparent thematic vs. taxonomic differences. First, we used word stimuli (rather than pictures) to limit incidental visual activation. Second, we matched conditions on linguistic and psycholinguistic variables to ensure that brain responses reflected conceptual semantics as opposed to lexical properties. We also employed continuous values of relatedness strength as a more sensitive index of differences in conceptual processing and to equate conditions on association strength.

Our MEG experiment (Chapter 2) shed light not just on temporal aspects of conceptual processing as it relates to lexical access, but also the degree of involvement of the ATL and TPJ in taxonomic and thematic relations. To recap, the two-systems view predicted ATL selectivity for taxonomic relations, which we confirmed, as the ATL showed sensitivity to taxonomic but not thematic variables. The two-systems view also predicted TPJ specificity for thematic relations, which we did not confirm, as the TPJ showed sensitivity to both thematic and taxonomic variables. We speculated that the TPJ's apparent involvement in taxonomic relations could reflect inhibition of pre-activated thematic concepts. This could work as follows: In our semantic priming task, the first word (e.g., DOG) activated taxonomic links to features (e.g., *barks, wet nose*) and thematic links to associates (e.g., *leash, collar*). Once the second word appeared (e.g., SEAL), the thematic associates for the first word would then be inhibited to successfully compute the taxonomic relation with the second word. Given that we chose taxonomic pairs that did not additionally share a thematic relationship (e.g., CAT → MOUSE), it is especially unlikely that any pre-activated thematic relations would be relevant to processing the taxonomic relation.

Results of our fMRI experiment (Chapter 3) suggest that the TPJ plays a multifaceted role in conceptual relations. As in Chapter 2, we carried on the tradition of using word stimuli and controlling for linguistic and psycholinguistic differences between conditions. We additionally equated the two conditions on relatedness score, which was of particular importance for our study given that thematic pairs tend to be more highly associated. Building on results from our MEG experiment, we tested whether the TPJ reflects thematic semantic processing or semantic inhibition. We approached this by teaching participants how to inhibit taxonomic and thematic relations prior to making taxonomic and thematic judgments in the scanner. In doing so we created a context in which participants had to inhibit or process the opposite relation, which our MEG experiment did not explicitly require. Recall that the inhibition account predicted stronger TPJ involvement for negative trials than positive trials, while the thematic account predicted more involvement for thematic trials than taxonomic trials. Results of our whole-brain analyses support both these predictions, as contrasts negative > positive and thematic > taxonomic revealed significant clusters in overlap with the TPJ region. Our primary focus was on results from our ROI analysis, which also yielded results in support of both accounts. Relative to unrelated pairs, the TPJ was more active for both negative taxonomic and thematic trials and for positive thematic trials, but not for positive taxonomic trials. Crucially, our ROI analysis found greater TPJ activity for negative thematic relative to negative taxonomic pairs, even though we equated these conditions on linguistic and psycholinguistic variables, on relatedness score, and included response time and relatedness score as parametric regressors.

We propose that the TPJ plays a role both in inhibition and thematic activation, depending on the computation required. The TPJ effects in our experiment likely correspond to the following: (1) thematic activation on positive thematic trials, (2) inhibition on negative taxonomic trials, and (3) thematic activation + thematic inhibition on negative thematic trials. We propose that (3) could reflect initial activation of thematic relations coupled with (or followed by) thematic inhibition in attempts to process the word-pair based on taxonomic criteria. In the Schwartz study, it is likely that patients with TPJ damage offered thematic substitutions (even though the task required taxonomic labels) due to loss of the ability to inhibit contextually relevant relations with the target concept.

A number of unanswered questions remain, such as whether thematic concepts exist in the TPJ *per se,* or whether the TPJ reflects relational computations specific to thematic information (i.e., it is uncertain whether the process of inhibiting features is just similar to the process of activating features). Specific to our results, it is unclear whether the negative taxonomic > filler effect in our fMRI experiment reflects inhibition of taxonomic relations or of thematic representations related to the taxonomic word pair. As yet, both are equally probable. One might ask why taxonomic but not unrelated word pairs would entail inhibition of thematic relations. One possibility is that taxonomic pairs required computing the taxonomic relation, which could entail inhibition of pre-activated thematic relations. In contrast, relational processing of the fillers might be attempted but ultimately fail, thus any pre-activated thematic information would not require inhibition.

Another uncertainty concerns why the contrast between negative thematic and negative taxonomic trials was greater than for the contrast between negative thematic and positive taxonomic trials. Earlier, we argued that the TPJ responses reflect activation + inhibition for negative thematic trials, inhibition for negative taxonomic trials, and no involvement for positive taxonomic trials. It is unclear, therefore, why the TPJ ROI analysis did not reveal a significant ThemNo > TaxYes result. Despite these and other remaining questions, the present studies make an important contribution to our understanding of concepts in the brain. Our MEG experiment helped shed light on where and when thematic and taxonomic relations are computed, while our fMRI experiment helped characterize the neural mechanisms involved in conceptual selection.

APPENDIX A: SUPPORTING INFORMATION FOR CHAPTER 1

Stimuli in descending order of imageability

goose
pill
bulb
chimp
sled
jeep
fist
gym
lid
shelf
wasp
cheese
dice
pearl
wrist
car
girl
boat
beard
tub
kite
tooth
yacht
kilt
tongue
blood
chef
clerk
fern
nut
yarn
child
lip
cat
trout
church
golf
lake
mouse
bird
clock
mouth
peach
queen
shirt
crate
skull
broom
lawn
hill
house
keg
mink
roof
web
bun
dorm
smog
spine
thorn
door
pond
fork
hoof
hen
sheep
tomb
blouse
sheet
juice
noose
tent
stove
grape
rug
dime
jug
monk
rod
barn
clown
morgue
rat
fruit
wig
porch
rib
wool
cage
glass
goat
king
stool
globe
mud
flute
vest
crutch
cub
grain
prom
gown
song
moth
snail
wheat
lung
shawl
disc
dusk
lamp
wife
beak
brick
desk
mug
scab
brain
gem
van
hemp
moss
blade
flesh
slush
horn
scar
tube
vine
vase
hat
birch
cheek
throat
hoop
hut
pub
rum
snack
cone
fang
beast
mast
slime
dome
chest
dot
goal
shrub
bench
hive
knob
gift
town
ghost
fuzz
path
sleeve
sperm
cloth
dirt
sleet
nymph
ridge
cob
den
mound
niece
tweed
wart
wick
fin
mutt
quill
yeast
thief
womb
hog
brass
pork
dude
gauze
reed
wench
badge
ranch
veal
tang
filth
groin
tribe
lint
pouch
wand
wreath
food
tool
tusk
brat
couch
graph
lad
gas
lice
crotch
mile
fig
ledge
plaque
shrine
silk
rim
rice
mace
mob
snout
spud
stag
swine
valve
wealth
death
crumb
font
belt
chrome
gasp
gulf
lard
clan
gap
marsh
thug
bib
pint
booth
nerve
news
life
brute
truce
breath
duke
grief
hunk
scope
smudge
twine
yam
height
crime
chive
grove
hearth
rump
self
latch
mosque
dean
grid
hick
pal
gust
lair
broth
chunk
cove
dune
loin
malt
smock
snob
barb
month
thong
speech
clove
pest
knoll
lust
germ
folk
jazz
josh
lobe
loft
slate
theft
health
zone
balm
cod
dill
finch
punk
slaw
gloom
year
growth
smut
blotch
chore
greed
math
rink
slab
steed
trough
haste
volt
slang
grub
jab
sloth
stooge
loss
threat
spite
chic
hag
hutch
noon
sheath
luck
dell
length
theme
deed
noise
skit
tab
cult
debt
lymph
fraud
batch
bunt
clique
creed
drought
gene
grime
runt
scum
wrath
truth
drake
husk
musk
wad
noun
bile
kin
nook
myth
thing
verb
crock
depth
faith
fame
zeal
curd
fright
rift
romp
realm
glade
zest
bliss
grail
fact
douche
dud
guild
sham
siege
thirst
whiz
wisp
welt
niche
shank
fad
pox
prude
choice
mead
pun
quirk
crude
norm
vogue
farce
stein
beck
sheen
bout
pence
whim

APPENDIX B: SUPPORTING INFORMATION FOR CHAPTER 2

Test stimuli (taxonomic prime/thematic prime -->target)

Animate	**Clothing**	**Food**	**Tools/Miscellaneous**	**Dishes/Appliances**
psychic/play→actress	bib/chef→apron	mango/peeler→apple	crutch/wound→bandage	bell/fire→alarm
teacher/mask→bandit	medal/cop→badge	cracker/bakery→bread	clipboard/leaflet→binder	toaster/smoothie→blender
seal/dam→beaver	mascara/cheek→blush	pork/ketchup→burger	fog/drizzle→cloud	squeegee/debris→broom
gymnast/gloves→boxer	vein/fracture→bone	yogurt/bun→butter	dollar/meter→coin	vial/water→canteen
centipede/net→butterfly	vault/clothing→closet	scone/candles→cake	hammock/toddler→crib	furnace/smoke→chimney
wolf/lap→cat	barrette/hair→comb	muffin/bowl→cereal	pamphlet/padlock→diary	spatula/pasta→colander
lamb/milk→cow	broach/gown→corsage	crouton/dip→chips	locker/socks→drawer	speakers/program→computer
professor/patient→doctor	cot/infant→cradle	raisin/box→chocolate	package/letter→envelope	pan/stew→crock
mule/leash→dog	camera/eyes→glasses	brandy/thermos→cocoa	emblem/pole→flag	rag/chalk→eraser
hamster/stable→horse	dishes/hamper→laundry	grape/palm→coconut	yarn/teeth→floss	teapot/whiskey→flask
gardener/mop→janitor	detergent/skin→lotion	nectar/mug→coffee	sword/bullet→gun	pot/coaster→glass
spider/scalp→lice	knot/neck→noose	ale/bottle→cola	drill/nail→hammer	rod/bait→hook
idol/costume→mascot	gold/oyster→pearl	candy/wrapper→gum	lighter/candle→match	bin/jam→jar
bear/banana→monkey	baton/cheerleader→pompom	liquor/straw→juice	marker/note→pencil	canister/beer→keg
mulch/tree→moss	sombrero/storm→poncho	soda/pitcher→lemonade	quilt/head→pillow	fryer/whistle→kettle
sparrow/nest→owl	string/braid→ribbon	clam/trap→lobster	spade/leaves→rake	flashlight/book→lamp
donkey/sty→pig	bracelet/finger→ring	pea/tears→onion	scissors/beard→razor	latch/key→lock
rat/kennel→puppy	slip/throat→scarf	tomato/brine→pickle	magazine/actor→script	recipe/diner→menu
skunk/burrow→rabbit	blouse/lady→skirt	casserole/plate→sandwich	brick/tar→shingle	wood/rust→metal
dove/worm→robin	robe/artist→smock	herring/can→sardines	hoe/earth→shovel	microwave/food→refrigerator
belch/nose→sneeze	skates/laces→sneakers	porridge/spoon→soup	shampoo/cloth→soap	carafe/cup→saucer
dentist/scalpel→surgeon	mittens/legs→stockings	rice/fork→spaghetti	smudge/postcard→stamp	blemish/bleach→stain
cactus/trellis→vine	uniform/groom→tuxedo	ham/grill→steak	painting/chisel→statue	jug/ashes→urn
hawk/carcass→vulture	rainboots/rain→umbrella	dumpling/soysauce→sushi	staple/paper→tape	goblet/bouquet→vase
hostess/tip→waiter	turban/bride→veil	vodka/sugar→tea	lane/map→route	toilet/bath→tub

Association strengths of the stimuli

We checked the association strengths of our prime-target pairs against the list of South Florida free association normative data from Nelson et al., (2004). The corpus listed values for 101 of the taxonomic, 123 of the thematic, and 137 of our filler primes. Nineteen taxonomic targets, 49 thematic targets, and no filler targets were listed as responses to these primes. The "Forward Cue to Target Strength" of the pairs was low for both taxonomic pairs (*M* = .01, *SD* = .02) and thematic pairs (*M* = .06, *SD* = .13). We have included the association data in the Appendix of the Revision.

-

Measure	Taxonomic	Thematic
Forward Cue to Target Strength	.01(.02)	.06(.13)
Backward Target to Cue Strength	.02(.05)	.15(.21)
Mediated Strength	.04(.07)	.01(.02)
Overlapping Associated Strength	.05(.06)	.03(.06)

Means(SDs) of association strengths by condition

Instructions for Mechanical Turk relatedness task (Study 1)

People can make various predictions about objects, people, or animals after reading a short story about them. Today we would like you to make judgments about objects, people, or animals outside the context of a story. You will be shown a pair of words. We would like you to rate the relatedness of the two words on a scale of 1–7. 1 means the words are “not at all related” and 7 means the words are “highly related.” Make your decision by pressing the corresponding number key on your keyboard. Please use all the numbers in the 1–7 scale, not just one or two. There is no right or wrong answer to these questions – we're just interested in your opinion. You have as much time as you want, but it is usually best to just go with your first reaction about how related the words are.

Instructions for Mechanical Turk relatedness task (Study 2)

In this experiment, you will be shown pairs of words. We would like you to rate the relatedness of the two words on a scale of 1 to 7. We consider words to be "related" if they are the same general kind of thing (like velcro and zipper-both are kinds of fasteners) or if they are related to one another (like pants and zipper-zippers open pants). Examples of unrelated items would include words like desk and harp, or cafe and harp. Of course, words can be more or less related, and we would like you to decide just how related (if at all) each pair of words is. 1 means the words are “not at all related” and 7 means the words are “highly related.” Make your decision by pressing the corresponding number key on your keyboard. Please use all the numbers in the 1-7 scale, not just one or two. There is no right or wrong answer to these questions--we're just interested in your opinion. You have as much time as you want, but it is usually best to just go with your first reaction about how related the words are.

Instructions for MEG relatedness task

In this experiment, you will view pairs of related and unrelated word pairs. We consider words to be "related" if they are the same general kind of thing (like velcro and zipper-both are kinds of fasteners) or if they are related to one another (like pants and zipper-zippers open pants). Examples of unrelated items would include words like desk and harp, or cafe and harp. Each trial will commence with a fixation cross. Look at that cross. Next, a word will appear, followed by a second word. Respond “yes” with your index finger if they are related. Respond “no” with your middle finger if they are not related. Please respond as quickly and as accurately as possible.

APPENDIX C: SUPPORTING INFORMATION FOR CHAPTER 3

Taxonomic pairs

shampoo-soap
dollar-coin
tornado-cyclone
bazaar-market
bell-alarm
medal-badge
lighter-match
crypt-grave
candy-gum
aisle-path
furnace-boiler
soccer-golf
whiskey-gin
drum-organ
sea-lake
van-jet
fog-cloud
ruby-emerald
horn-whistle
recipe-menu
trout-bass
chair-sofa
stairs-elevator
guitar-tuba
marker-pencil
blouse-skirt
stool-bench
toaster-blender
vodka-bourbon
plate-saucer
basement-hall
porridge-soup
pear-cherry
bib-apron
latch-lock
blemish-stain
igloo-cabin
ceiling-wall
aquarium-zoo
nose-ear
string-ribbon
lane-route
daisy-violet
mango-apple
sombrero-cap
clam-lobster
hotel-mansion
bracelet-ring
rag-towel
toilet-tub
oak-elm
scone-cake
hut-cottage
scissors-razor
dentist-surgeon
sparrow-owl
pot-skillet
waltz-tango
violin-flute
goat-zebra
tuna-cod
tape-glue
vial-carafe
rose-tulip
crow-canary
pie-cookie
pliers-drill
pastor-monk
square-circle
plateau-cliff
lemon-kiwi
gnat-moth
wool-silk
butler-waiter
cart-stroller
piano-harp
pepper-nutmeg
turnip-spinach
clarinet-banjo
beets-celery
article-essay
polka-mambo
falcon-parakeet
cello-saxophone
flamingo-stork
cardinal-eagle
sapphire-quartz
scooter-tricycle
turban-beret
general-colonel
cannon-pistol
nylon-satin
ostrich-pelican
oboe-accordion
wasp-ladybug
mosquito-beetle
peach-lime
locker-safe
tennis-lacrosse
raven-rooster
hurricane-typhoon
lichen-moss
bronze-platinum
pen-paintbrush
orchid-sunflower
canyon-gully
asparagus-squash
hummingbird-woodpecker
hacksaw-sledgehammer
tangerine-watermelon
centipede-butterfly
paprika-oregano
package-envelope
lettuce-cauliflower
vanilla-cinnamon
speakers-computer
tugboat-steamship
mattress-cushion
rainboots-umbrella
flashlight-lamp

Thematic pairs

reef-ocean
crayon-color
hose-garden
tutu-ballet
rug-floor
straw-juice
ink-printer
bride-bouquet
shutter-window
rabbi-temple
relish-hotdogs
cabinet-kitchen
snorkel-swim
pitchfork-hay
moat-bridge
witch-broom
tourist-bus
hinge-door
staple-paper
chisel-statue
carton-milk
swatter-bug
garbage-odor
flea-dog
carcass-vulture
faucet-leak
sheets-bed
throne-prince
clog-drain
jar-jelly
hanger-coat
anchor-boat
kettle-tea
scarf-neck
pond-duck
donor-blood
tank-soldier
telescope-star
chime-wind
bait-fish
sleeve-arm
roof-ladder
axe-tree
chlorine-pool
list-groceries
safari-jungle
trap-mouse
porch-swing
shirt-iron
robe-bath
crutch-leg
burrow-rabbit
bag-lunch
pony-carriage
stew-crock
lap-cat
bun-butter
zipper-pants
raft-river
vent-air
leaflet-binder
vase-flower
coaster-glass
tray-table
poodle-vet
cast-limb
tar-shingle
storm-poncho
pasta-colander
artist-smock
cobweb-spider
trellis-vine
grove-orange
thermos-cocoa
bluebird-worm
syrup-pancake
tepee-Indian
ashtray-cigarette
shovel-dirt
grill-charcoal
cork-wine
dam-beaver
outlet-plug
balloon-helium
thread-needle
arrow-target
pasture-cow
bandage-cut
acorn-squirrel
thimble-finger
hiker-mountain
helmet-motorcycle
seashore-shell
beaker-chemistry
ticket-concert
mower-lawn
cocoon-caterpillar
eyelash-mascara
plaza-shopping
astronaut-moon
pier-fishing
aspirin-headache
tobacco-pipe
seagull-beach
protractor-geometry
headband-sweat
hostess-restaurant
dice-gamble
bar-cocktail
frosting-cupcake
newsstand-magazine
blackboard-teacher
contractor-building
handkerchief-nose
decoration-party
ambulance-hospital
scarecrow-corn
airport-baggage
iceberg-penguin
banana-monkey

Taxonomic task: Instructions, practice pairs, and feedback

Practice Instructions

You will decide whether words relate categorically by being the same general kind of thing. Such items belong to the same category and share features. Examples of categorically related items include:

MELON and BERRY: both are kinds of fruits
SALT and SUGAR: both are kinds of condiments
BIRD and LAMB: both are kinds are kinds of animals
BARN and CONDO: both are kinds of buildings.

Not ALL related items are categorically related Examples include:

HEN and CAGE: while hens live in cages, only one is an animal.
PRAM and BABY: while babies ride in prams, only one is a person.
SAW and WOOD: while saws cut wood, only one is a tool.
CAPTAIN and SHIP: while captains steer ships, only one is a boat.

Lets practice making categorical judgments about word pairs. Respond YES only for categorically related items. Respond NO for all other word pairs. Each trial will begin with a dot, followed by the word pair. Decide if the words relate categorically. Press with your left index finger for YES Press with your left middle finger for NO.

Practice Items: Feedback for (-) Response to (+) Pair

COCKROACH-TERMITE. These two are categorically related because both are kinds of insects.
CLOVER-DANDELION. These two are categorically related because both are kinds of plants.
COKE-MILKSHAKE. These two are categorically related because both are kinds of beverages.
PARKA-CAPE. These two are categorically related because both are kinds of clothing.
FERRY-CANOO. These two are categorically related because both are kinds of boats.
SLACKS-SHORTS. These two are categorically related because both are kinds of clothing.
VEST-GOWN. These two are categorically related because both are kinds of clothing.
VELVET-DENIM. These two are categorically related because both are kinds of fabric.
SWORDFISH-GUPPY. These two are categorically related because both are kinds of fish.
BLUEBERRY-KUMQUAT. These two are categorically related because both are kinds of fruit.
RULER-STAPLER. These two are categorically related because both are kinds of office supplies.
PANSY-LILY. These two are categorically related because both are kinds of flowers.
BICYCLE-TAXI. These two are categorically related because both are kinds of vehicles.
RIFLE-SPEAR. These two are categorically related because both are kinds of weapons.
MARTINI-SHERRY. These two are categorically related because both are kinds of beverages.

Practice Items: Feedback for (+) Response to (-) Pair

CALCULATOR-MATH. While calculators do math, one is a machine, the other is not.
POLICE-HANDCUFFS. While police use handcuffs, one is a person, the other is not.
EGYPT-MUMMY. While Egypt has mummies, one is a country, the other is not.
WARDROBE-CLOTHES. While wardrobes store clothes, one is furniture, the other is not.
SHUTTLE-SPACE. While shuttles fly in space, one is a vehicle, the other is not.
CASHIER-MONEY. While cashiers handle money, one is a person, the other is not.
LETTER-MAILBOX. While letters go in mailboxes, one is paper, the other is not..
MINISTER-CHURCH. While ministers work at churches, one is a person, the other is not.
PLUMBER-PIPES. While plumbers work with pipes, one is a person, the other is not..
FARM-CATTLE. While farms have cattle, one is a place, the other is not.
HALO-ANGEL. While halos go above angels, one is a spirit, the other is not.
CLERK-STORE. While clerks work at stores, one is a person, the other is not.
PILOT-PLANE. While pilots fly planes, one is a person, the other is not.
THERMOMETER-FEVER. While thermometers test for fevers, one is an instrument, the other is not.
STOCKING-CHRISTMAS. While stockings are hung at Christmas, one is a holiday, the other is not.

Thematic Task: Instructions, practice pairs, and feedback

Practice Instructions

You will decide if words relate thematically by tending to occur or interact together. Such items share a spatial or functional relationship. Examples of thematically related items include:
JAM and TOAST: jam is spread on toast.
BEE and HIVE: bees live in hives.
QUEEN and CROWN: queens wear crowns.
BRACES and TEETH: braces straighten teeth.
Not ALL related items are thematically related. Examples include:
LIMO and TRACTOR: both are vehicles but do not interact.
KOI and EEL: both are fish but do not interact.
STEER and RHINO: both are mammals but do not interact.
BOOK and MAP: both are reading material but do not interact.
Lets practice making thematic judgments. Respond YES only for thematically related items. Respond NO for all other word pairs. Each trial will begin with a dot, followed by the word pair. Decide if the words relate thematically. Press with your left index finger for YES. Press with your left middle finger for NO.

Practice Items: Feedback for (-) Response to (+) Pairs
JUDGE-GAVEL. These are related because judges bang gavels.
SANTA-REEINDEER. These are related because Santa has reindeer.
DESERT-CAMEL. These are related because camels live in deserts.
KNIGHT-ARMOR. These are related because knights wear armor.
PITCHER-BASEBALL. These are related because pitchers throw baseballs.
BAND-MUSIC. These are related because bands play music.
CLOWN-CIRCUS. These are related because clowns work at circuses.
HAND-GLOVE. These are related because gloves go on hands.
HORSE-SADDLE. These are related because horses wear saddles.
BANK-LOAN. These are related because banks make loans.
WAND-MAGIC. These are related because wands are used for magic.
COWBOY-RODEO. These are related because cowboys ride at rodeos.
BUTCHER-MEAT. These are related because butchers cut up meat.
SHEPHERD-SHEEP. These are related because shepherds watch over sheep.
TURKEY-THANKSGIVING. These are related because turkey is eaten at Thanksgiving.

Practice Items: Feedback for (+) Response to (-) Pairs
PARROT-PIGEON. While both are kinds of birds, parrots and pigeons do not often interact.
TAFFY-LICORICE. While both are kinds of candy, taffy and licorice do not often interact.
LIBRARY-MUSEUM. While both are kinds of places, libraries and museums do not often interact.
SALMON-GOLDFISH. While both are kinds of fish, salmon and goldfish do not often interact.
SANDALS-SKIS. While both are kinds of foot-wear, sandals and skis do not often interact.
PLUM-PINEAPPLE. While both are kinds of fruit, plums and pineapples do not often interact.
VALLEY-CAVERN. While both are kinds of land forms, valleys and caverns do not often interact.
RESORT-HOSPICE. While both are kinds of places, resorts and hospices do not often interact.
GOVERNOR-SHERIFF. While both are kinds of professions, governors and sheriffs do not often interact.
PYRAMID-TOWER. While both are kinds of structures, pyramids and towers do not often interact.
KNIFE-TORPEDO. While both are kinds of weapons, knives and revolvers do not often interact.
JACKS-PUZZLE. While both are kinds of games, jacks and puzzles do not often interact.
DRAGONFLY-MAGGOT. While both are kinds of insects, dragonflies and maggots do not often interact.
SCIENTIST-MERCHANT. While both are kinds of professions, scientists and merchants do not often interact.
GARNET-TURQUOISE. While both are kinds of stones, garnets and turquoise do not often interact.

REFERENCES

Adachi, Y., M., Shimogawara, M., Higuchi, M., Haruta, Y., & Ochiai, M. (2001). Reduction of non-periodic environmental magnetic noise in MEG measurement by continuously adjusted least squares method. *IEEE Transactions on Applied Superconductivity, 11*, 669–672.

Andersson, J. L., Jenkinson, M., & Smith, S. (2007a). Non-linear optimisation. FMRIB technical report TR07JA1. *University of Oxford FMRIB Centre: Oxford, UK.*

Andersson, J. L., Jenkinson, M., & Smith, S. (2007b). Non-linear registration, aka Spatial normalisation FMRIB technical report TR07JA2. *FMRIB Analysis Group of the University of Oxford*, *2*.

Baayen, R., Piepenbrock, R., & Gulikers, L. (1995). The celex lexical database (release 2) [CD ROM]. Philadelphia, PA: Linguistic Data Consortium, University of Pennsylvania.

Balota, D., Yap, M., Cortese, M., Hutchison, K., Kessler, B., Loftis, B., Neely, J., Nelson, D., Simpson, G., & Treiman, R. (2007). The English lexicon project. *Behavior Research Methods, 39,* 445.

Barrós-Loscertales, A., González, J., Pulvermüller, F., Ventura-Campos, N., Bustamante, J., Costumero, V., Parcet, M., & Ávila, C. (2011). Reading salt activates gustatory brain regions: fMRI evidence for semantic grounding in a novel sensory modality. *Cerebral Cortex,* doi:10.1093/cercor/bhr324.

Battig, W. F., & Montague, W. E. (1969). Category norms of verbal items in 56 categories A replication and extension of the Connecticut category norms. *Journal of Experimental Psychology*, *80*, 1.

Beckmann, C. F., Jenkinson, M., & Smith, S. M. (2003). General multilevel linear modeling for group analysis in FMRI. *Neuroimage*, *20*, 1052-1063.

Bemis, D. K., & Pylkkänen, L. (2011). Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience, 31,* 2801.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society*, *57*, 289-300.

Benjamini, Y., & Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, 1165-1188.

Bickhard, M. (2008). Is embodiment necessary? In P. Calvo & T. Gomila (Eds.), *Handbook of cognitive science: An embodied approach* (pp. 29–40). Amsterdam: Elsevier.

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,* 2767-2796.

Carlson, T. A., Simmons, R. A., Kriegeskorte, N., & Slevc, L. R. (2013). The emergence of semantic meaning in the ventral temporal pathway. *Journal of Cognitive Neuroscience, 26,* 120-131.

Coltheart, M. (1981). The MRC Psycholinguistic Database, *Quarterly Journal of Experimental Psychology, 33A,* 497-505.

Cortese, M., & Fugett, A. (2004). Imagery ratings for 3,000 monosyllabic words. *Behavior Research Methods, 3,* 384-387.

Cortese, M., & Khanna, M. (2008). Age of acquisition ratings for 3,000 monosyllabic words. *Behavior Research Methods, 40, 791-794*.

Dale, A. M., Liu, A.K., Fischl, B. R., Buckner, R. L., Belliveau, J. W., Lewine, J. D., & Halgren, E. (2000). Dynamic statistical parametric mapping: Combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron, 26*, 55-67.

Davey, J., Cornelissen, P. L., Thompson, H. E., Sonkusare, S., Hallam, G., Smallwood, J., & Jefferies, E. (2015). Automatic and controlled semantic retrieval: TMS reveals distinct contributions of posterior middle temporal gyrus and angular gyrus. *The Journal of Neuroscience*, *35*, 15230-15239.

Davis, M., & Johnsrude, I. Hierarchical processing in spoken language comprehension. (2003). *Journal of Neuroscience, 23,* 2423-3431.

de Zubicaray, G. I., Hansen, S., & McMahon, K. L. (2013). Differential processing of thematic and categorical conceptual relations in spoken word production. *Journal of Experimental Psychology: General*, *142*, 131.

Dronkers N. Wilkins D., Van Valin R., Redfern B., & Jaeger J. (2004). Lesion analysis of the brain areas involved in language comprehension. *Cognition, 92*, 145–177.

Estes Z., Golonka, S., & Jones L. L. (2011). Thematic thinking: The apprehension and consequences of thematic relations. *Psychology of Learning and Motivation: Advances in Research and Theory, 54,* 249–294.

Ettinger, A., Linzen, T., & Marantz, A. (2014). The role of morphology in phoneme prediction: Evidence from MEG. *Brain and Language, 129,* 14-23.

Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition, 28*, 3–71.

Fruchter, J., Stockall, L., & Marantz, A. (2013). MEG masked priming evidence for form-based decomposition of irregular verbs. *Frontiers in human neuroscience*, *7*, 798.

Gallese, V. & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in reason and language. *Cognitive Neuropsychology, 22*, 455–479.

Ganis, G., Thompson, W., & Kosslyn, S. (2004). Brain areas underlying visual mental imagery and visual perceptual. *Cognitive Brain Research, 20*, 226-241.

Genovese, C. R., Lazar, N. E., & Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage, 15,* 870-878.

González J., Barrós-Loscertales, A., Pulvermüller, F., Meseguer, V., Sanju´an, A., et al. (2006). Reading cinnamon activates olfactory brain regions. *Neuroimage, 32,* 906–12.

Helenius P., Salmelin, R., Service, E., Connolly, J., Leinonen, S., & Lyytinen, H. (2002). Cortical activation during spoken-word segmentation in nonreading impaired and dyslexic adults. *Journal of Neuroscience, 22,* 2936-2944.

Hickok, G. & Poeppel, D. (2007). The cortical organization of speech processing. *Nature Reviews Neuroscience, 8,* 393–402.

Holmes, V., & Langford, J. (1976). Comprehension and recall of abstract and concrete sentences. *Journal of Verbal Learning and Verbal Behavior, 15,* 559-566.

Inhelder, B., & Piaget, J. (1964). *The early growth of logic in the child: Classification and seriation.* London: Routledge and Kegan Paul.

Jackson, R. L., Hoffman, P., Pobric, G., & Ralph, M. A. L. (2015). The nature and neural correlates of semantic association versus conceptual similarity. *Cerebral Cortex*, *25*, 4319-4333.

James, T., Stevenson, R., Kim, S., VanDerKlok, R., & James, K. (2011). Shape from sound: Evidence for a shape operator in the lateral occipital cortex. *Neuropsychologia, 49,* 1807-1815.

Jefferies, E., & Lambon Ralph, M. A. (2006). Semantic impairment in stroke aphasia versus semantic dementia: A case-series comparison. *Brain, 129*, 2132-2147.

Jefferies, E., & Ralph, M. A. L. (2006). Semantic impairment in stroke aphasia versus semantic dementia: a case-series comparison. *Brain*, *129*, 2132-2147.

Jenkinson, M., & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical image analysis*, *5*, 143-156.

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, C. F., Behrens, T. E., Woolrich, M. W., & Smith, S. M. (2012). Fsl. *Neuroimage*, *62*, 782-790.

Kalénine, S., Peyrin, C., Pichat, C., Segebarth, C., Bonthoux, F., & Baciu, M. (2009). The sensory-motor specificity of taxonomic and thematic conceptual relations: A behavioral and fMRI study. *NeuroImage, 44,* 1152-1162.

Kircher, T., Sass, K., Sachs, O., & Krach, S. (2009). Priming words with pictures: Neural correlates of semantic associations in a cross‐modal priming task using fMRI. *Human Brain Mapping*, *30*, 4116-4128.

Klimesch, W., Schimke, H., & Schwaiger, J. (1994). Episodic and semantic memory: an analysis in the EEG theta and alpha band. *Electroencephalography and clinical neurophysiology*, *91*, 428-441.

Kotz, S. A., Cappa, S. F., von Cramon, D. Y., & Friederici, A. D. (2002). Modulation of the lexical–semantic network by auditory semantic priming: An event-related functional MRI study. *Neuroimage*, *17*, 1761-1772.

Lambon Ralph, M. A., McClelland, J. L., Patterson, K. Galton, C. J., & Hodges, J. R. (2001). No right to speak? The relationship between object naming and semantic impairment: Neurophysiological evidence and a computational model. *Journal of Cognitive Neuroscience, 13,* 341-356.

Lau, E., Phillips, C., & Poeppel, D. (2008). A cortical network for semantics: (de)constructing the N400. *Nature Reviews Neuroscience, 9,* 920-933.

Lewis, G., & Poeppel, D. (2014). The role of visual representations during the lexical access of spoken words. *Brain and Language, 134,* 1-10.

Lewis, G., Poeppel, D., & Murphy, G. L. (2015). The neural bases of taxonomic and thematic conceptual relations: An MEG study. *Neuropsychologia*, *68*, 176-189.

Lewis, G., Solomyak, O., & Marantz, A. (2011). The neural basis of obligatory decomposition of suffixed words. *Brain & Language, 118,* 118-127.

Lin, E. L., & Murphy, G. L. (2001). Thematic relations in adults' concepts. *Journal of experimental psychology: General*, *130*, 3.

Luce, P., & Pisoni, B. (1998). Recognizing spoken words: The neighborhood activation model. *Ear and Hearing, 19,* 1-36.

Lund, K., & Burgess, C. (1996) Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers, 28***,** 203-208.

Maguire, M. J., Brier, M. R., & Ferree, T. C. (2010). EEG theta and alpha responses reveal qualitative differences in processing taxonomic versus thematic semantic relationships. *Brain and language*, *114*, 16-25.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods, 164*, 177-190.

Markman, E. M. (1989). *Categorization and naming in children: Problems of induction.* Cambridge, MA: MIT Press.

Marslen-Wilson, W. (1987) Functional parallelism in spoken word recognition. *Cognition, 25,* 71-102.

Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology, 58,* 25-45.

Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: Structure and processes. *Current Opinion in Neurobiology*, *11*, 194–201.

McClellend, J., & Elman, J. (1986). The TRACE model of speech perception.*Cognitive Psychology, 18*, 1-86.

Mesulam M., Rogalski, E., Wieneke, C., Cobia, D., Rademaker, A., Thompson, C., & Weintraub, S. (2009). Neurology of anomia in the semantic variant of primary progressive aphasia. *Brain, 132,* 2553–2565.

Meteyard, L., & Vigliocco, G. (2008). The role of sensory and motor information in semantic representation: A review. In P. Calvo & A. Gomila (Eds.), *Handbook of cognitive science: An embodied approach* (pp. 293–312). London, United Kingdom: Academic Press, Elsevier.

Mirman, D., & Graziano, K. M. (2012). Damage to temporo-parietal cortex decreases incidental activation of thematic relations during spoken word comprehension. *Neuropsychologia*, *50*, 1990-1997.

Murphy, G. L. (2001). Causes of taxonomic sorting by adults: A test of the thematic-to-taxonomic shift. *Psychonomic Bulletin & Review, 8*, 834-839.

Murphy, G. L. (2010). What are categories and concepts? In D. Mareschal, P. C. Quinn, & S. E. G. Lea (Eds.), *The making of human concepts* (pp. 11-28). Oxford: Oxford University Press.

Noonan, K. A., Jeffries, E., Corbett, F., & Lambon Ralph, M. A. (2009). Elucidating the nature of deregulated semantic cognition in semantic aphasia: Evidence for the roles of prefrontal and temporo-parietal cortices. *Journal of Cognitive Neuroscience, 22,* 1597-1613.

Paivo, A., Yuille, J., & Madigan, S. (1968). Concreteness, imagery, and meaningfulness for 925 nouns. *Journal of Experimental Psychology, 76,* 1-25.

Patterson, K., Lambon Ralph, M. A., Jefferies, E., Woollams, A., Jones, R., Hodges, J. R., et al. (2006). “Presemantic” cognition in semantic dementia: Six deficits in search of an explanation. *Journal of Cognitive Neuroscience, 18*, 169–183.

Patterson, K., Nestoer, P., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience, 8,* 976-987.

Pobric, G., Ralph, M., & Jefferies, E. (2009). The role of the anterior temporal lobes in the comprehension of concrete and abstract words: rTMS evidence. *Cortex, 45,* 1104-1110.

Poeppel, D., Idsardi, W., & Wassenhove, V. (2008). Speech perception at the interface of neurobiology and linguistics. *Philisophical Transactions of the Royal Society of London Biological Sciences, 363,* 1071-1086.

Price, C. (2012). A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading*. Neuroimage,* 62, 816-847.

Pulvermüller, F. (1999). Words in the brain’s language. *Behavioral and Brain Sciences, 22,* 253-336.

Pulvermüller, F. (2005). Brain mechanisms linking language and action. *Nature Reviews Neuroscience*, 6, 576-582.

Pulvermüller, F., & Hauk, O. (2006). Category-specific conceptual processing of color and form in left fronto-temporal cortex. *Cerebral Cortex, 8,* 1193-1201.

Pulvermüller, F., Preissl, H., Lutzenberger, W. & Birbaumer, N. (1996). Brain rhythms of language: Nouns versus verbs. *European Journal of Neuroscience, 8,* 937–941.

Putnam, H. (1973). Meaning and reference. *Journal of Philosophy, LXX*, 699-711.

Pylkkänen, L., Llinás, R., & Murphy, G. L. (2006). The representation of polysemy: MEG evidence. *Journal of Cognitive Neuroscience, 18*, 97-109.

Pylkkänen, L., & Marantz, A. (2003). Tracking the time course of word recognition with MEG. *Trends in Cognitive Science, 7,* 187-189.

Rastle, K., Harrington, J., & Coltheart, M. (2002). 358,534 nonwords: The ARC Nonword Database. *Quarterly Journal of Experimental Psychology, 55A,* 1339-1362.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59-108.

Rodd, J., Davis, H., & Johnsrude, I. (2005). The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. *Cerebral Cortex, 15,* 1261–1269.

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., & Patterson, K. (2004). Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review, 11,* 205-235.

Sachs, O., Weis, S., Krings, T., Huber, W., & Kircher, T. (2008a). Categorical and thematic knowledge representation in the brain: Neural correlates of taxonomic and thematic conceptual relations. *Neuropsychologia*, *46*, 409-418.

Sachs, O., Weis, S., Zellagui, N., Huber, W., Zvyagintsev, M., Mathiak, K., & Kircher, T. (2008b). Automatic processing of semantic relations in fMRI: neural activation during semantic priming of taxonomic and thematic categories. *Brain research*, *1218*, 194-205.

Sachs, O., Weis, S., Zellagui, N., Sass, K., Huber, W., Zvyagintsev, M., ... & Kircher, T. (2011). How different types of conceptual relations modulate brain activation during semantic priming. *Journal of Cognitive Neuroscience*, *23*, 1263-1273.

Schirmer, A., Fox, P., & Grandjean, D. (2012). On the spatial organization of sound processing in the human temporal lobe: a meta-analysis. *Neuroimage, 63,* 137-147.

Schwartz, M. F., Kimberg, D. Y., Walker G. M., Brecher A., Faseyitan O. K., Dell, G. S., Mirman, D., & Coslett, H. B. (2011). Neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. *Proceedings of the National Academy of Sciences of the United States of America, 108,* 8520-8524.

Simon, D. A., Lewis, G., & Marantz, A. (2012). Disambiguating form and lexical frequency of MEG responses using homonyms. *Language & Cognitive Processes, 27*, 275-287.

Smiley, S. S., & Brown, A. L. (1979). Conceptual preferences for thematic or taxonomic relations: A nonmonotonic age trend from preschool to old age. *Journal of Experimental Child Psychology, 28*, 249-257.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human brain mapping*, *17*, 143-155.

Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., ... & Niazy, R. K. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, *23*, S208-S219.

Solomyak, O., & Marantz, A. (2009). Lexical access in early stages of visual word processing: A single-trial correlational MEG study of heteronym recognition. *Brain and Language, 108,* 191–196.

Solomyak, O., & Marantz, A. (2010). MEG evidence for early morphological decomposition in visual word recognition: A single-trial correlational MEG study. *Journal of Cognitive Neuroscience, 22,* 2042–2057.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*, 1632.

Turkeltaub, P., & Coslett, H. (2010). Localization of sublexical speech perception components. *Brain and Language, 114*, 1-15.

Turken, U., & Dronkers, N. (2011). The neural architecture of the language comprehension network: converging evidence from lesion and connectivity analyses. *Frontiers in Systems Neuroscience, 5,* doi:10.3389/fnsys.2011.00001

Tyler, L., Voice, J., & Moss, H. (2000). The interaction of meaning and sound in spoken word recognition. *Psychonomic Bulletin & Review, 7,* 320–326.

Vitevitch, M., & Luce, P. (1998). When words compete: Levels of processing in spoken word perception. *Psychological Science, 9*, 325–329.

Vivas, L., García García, R., Perea Bartolomé, M. V., Leite D'almeida, A., & Ladera Fernández, V. (2016). Recognition of thematic and taxonomic conceptual relations in patients with aphasia. *Aphasiology*, *30*, 1-21.

West, W., & Holcomb, P. (2000). Imaginal, semantic, and surface-level processing of concrete and abstract words: An electrophysiological investigation. *Journal of Cognitive Neuroscience, 12,* 1024-1037.

Westerlund, M., & Pylkkänen, L. (2014). The role of the left anterior temporal lobe in semantic composition vs. semantic memory. *Neuropsychologia, 57,* 59-70.

Whitney, C., Kirk, M., O'Sullivan, J., Ralph, M. A. L., & 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. *Journal of Cognitive Neuroscience*, *24*, 133-147.

Woolrich, M. (2008). Robust group analysis using outlier inference. *Neuroimage*, *41*, 286-301.

Woolrich, M. W., Behrens, T. E., Beckmann, C. F., Jenkinson, M., & Smith, S. M. (2004). Multilevel linear modelling for FMRI group analysis using Bayesian inference. *Neuroimage*, *21*, 1732-1747.

Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., ... & Smith, S. M. (2009). Bayesian analysis of neuroimaging data in FSL. *Neuroimage*, *45*, S173-S186.

Worsley, K. J. (2001). Statistical analysis of activation images. *Functional MRI: An introduction to methods*, *14*, 251-270.

Yuan, J., & Liberman, M. (2008). Speaker identification on the SCOTUS corpus. *Journal of the Acoustical Society of America*, *123*, 3878.

Zhuang, J., Randall, B., Stamatakis, E., Marslen-Wilson, W., & Tyler, L. (2011). The interaction of lexical semantics and cohort competition in spoken word recognition: an fMRI study. *Cognitive Neuroscience, 12,* 778-3790**.**