Restrictive vs. non-restrictive composition: a magnetoencephalography study

Timothy Leffel^{a,*}, Miriam Lauter^{a,b,c}, Masha Westerlund^{b,c} and Liina Pykkänen^{a,b,c}

^aDepartment of Linguistics, New York University, 10 Washington Place, New York, NY 10003, USA; ^bDepartment of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA; ^cNYUAD Institute, New York University Abu Dhabi, PO Box 129188, Abu Dhabi, United Arab Emirates

(Received 24 November 2013; accepted 6 August 2014)

Recent research on the brain mechanisms underlying language processing has implicated the left anterior temporal lobe (LATL) as a central region for the composition of simple phrases. Because these studies typically present their critical stimuli without contextual information, the sensitivity of LATL responses to contextual factors is unknown. In this magnetoencephalography (MEG) study, we employed a simple question-answer paradigm to manipulate whether a prenominal adjective or determiner is interpreted restrictively, i.e., as limiting the set of entities under discussion. Our results show that the LATL is sensitive to restriction, with restrictive composition eliciting higher responses than non-restrictive composition. However, this effect was only observed when the restricting element was a determiner, adjectival stimuli showing the opposite pattern, which we hypothesise to be driven by the special pragmatic properties of non-restrictive adjectives. Overall, our results demonstrate a robust sensitivity of the LATL to high level contextual and potentially also pragmatic factors.

Keywords: left anterior temporal lobe; composition; MEG; pragmatics

1. Introduction

1.1. The interplay between semantic composition and referent identification

Research on language processing has shown that the meanings of complex phrases such as modified nouns (e.g. *red dog*, *dog over there*) are constructed incrementally during comprehension (e.g. Grodner & Sedivy, 2011; Noveck & Reboul, 2008; Sedivy, Tanenhaus, Chambers, & Carlson, 1999; Wolter, Gorman, & Tanenhaus, 2011). Incremental semantic processing allows language users to integrate the referential contribution of individual words as they are encountered, resulting in rapid association between phrases and the objects in the world that they refer to. This kind of association is sometimes called ‘reference resolution’ (or ‘referent identification’), a phenomenon that has been studied extensively in a variety of experimental paradigms (Arnold & Griffin, 2007; Arts, Maes, Noordman, & Jansen, 2011; Barr, 2008; Engelhardt et al., 2006; Grodner, Gibson, & Watson, 2005; Grodner, Klein, Carbary, & Tanenhaus, 2010; Hanna & Tanenhaus, 2004; Hanna, Tanenhaus, & Trueswell, 2003; Wolter et al., 2011).

Numerous studies on reference resolution have shown that the default interpretation of noun modifiers is ‘contrastive’, in the sense that language users typically assume that modifiers narrow the reference set of the nouns they modify (Heller, Grodner, & Tanenhaus, 2008; Sedivy, 2003; Sedivy et al., 1999; Wolter et al., 2011). For

example, subjects tend to interpret an expression like *the blue cup* as picking out the unique blue member from a set of cups. However, this tendency can be overridden if a ‘non-contrastive’ modifier is supported by the pragmatic context or the specific communicative goal (Arnold & Griffin, 2007; Arts et al., 2011; Engelhardt, Bailey, & Ferreira, 2006; Rohde, Levy, & Kehler, 2011; Sedivy, 2003). For example, if one’s goal is to collect all the blue objects, then *pick up the blue cup* is a felicitous instruction even in a context in which there is only one cup – in this case, the function of the modifier is to provide a reason for why the cup should be picked up, rather than to distinguish it from other cups.

In linguistic theory, contrastive modification is called ‘restrictive modification’, and the process of composing contrastive modifiers with their noun phrases is known as ‘predicate modification’ (Heim & Kratzer, 1998). ‘Non-contrastive’ modification is known in the linguistic literature as ‘non-restrictive modification’. The difference between restrictive and non-restrictive modification is most easily illustrated with relative clauses: the difference between *the person who I met* and *John, who I met* is that the relative clause *who I met* is restrictive in the former example and non-restrictive in the latter (non-restrictive relative clauses are also distinguished by an intonational break between head noun and modifier).

The goal of the current study was to investigate whether brain responses hypothesised to reflect composition,

*Corresponding author. Email: tim.leffel@nyu.edu

specifically within the left anterior temporal lobe (LATL), are sensitive to the restrictiveness or non-restrictiveness of composition. A combinatory operation that simply composes the meanings of elements A and B together would not show such sensitivity. In contrast, an operation that more specifically performs set restriction, i.e., narrowing down the set of individuals under discussion, would only show a combinatory response for restrictive composition.

1.2. Restrictive vs. non-restrictive nominal modification

In English, prenominal adjectives can be interpreted either restrictively or non-restrictively (Bolinger, 1967; Cinque, 2010; Larson, 1998, 1999; Morzycki, 2008; Vendler, 1968). Typically, however, context, lexical semantics and world knowledge are jointly sufficient to eliminate this ambiguity. For example, in *I visited my Italian friend*, *Italian* is used to pick out a specific individual from a set of friends (restrictive), whereas in *I visited my sick mother*, *sick* would typically not serve to pick out a mother from a set (non-restrictive).

This study investigated the processing of restrictive and non-restrictive composition in simple noun phrases consisting of a determiner (e.g. *the*, *his*), a noun (e.g. *cup*, *chicken*) and, in half of the critical trials, a prenominal attributive adjective (e.g. *blue*, as in *the blue cup*). To achieve a minimal contrast in restrictiveness, we exploited the disambiguating function of context. Specifically, we used a simple question-answer paradigm where by varying the type of the question, we induced either restrictive or non-restrictive interpretations of adjectives appearing within the answers. In addition, our design included a non-adjectival version of this contrast that simply varied the restrictiveness of prenominal determiners. This contrast was included as the interpretation of non-restrictive adjectives is, in fact, in many ways special (as we elaborate below), and thus they cannot by themselves constitute a ‘baseline’ condition for the examination of restrictiveness. However, the unique properties of non-restrictive adjectives also allowed us to test the sensitivity of combinatory brain responses to a certain kind of pragmatic enrichment that takes place in the interpretation of non-restrictive modification structures.

Non-restrictive interpretations of attributive adjectives intuitively require a certain kind of pragmatic licensing that is not required in restrictive composition. In the case of *my sick mother*, for example, the mother’s sickness would usually be understood to explain the need for a visit. Surprisingly, although the restrictive/non-restrictive ambiguity in English prenominal adjectives is a well-known phenomenon, there has not yet been any systematic research on the pragmatics of non-restrictive modification. Here, we propose that the function of *sick* in this case – and of non-restrictive modifiers in many other contexts – is to establish an explanation-type

discourse coherence relation (Asher & Lascarides, 1993; Kehler, 2002; Koornneef & Sanders, 2013) that links the truth-conditional meaning of the adjective to the truth-conditional meaning of the rest of the sentence/text (or to a subpart of it).¹ This conforms to the intuition that non-restrictive modification can very naturally be paraphrased with overt markers of explanation, as in *I visited my mother because she is sick* or *My mother is sick, and so I visited her*. However, the core prediction of this hypothesis is that a non-restrictive adjective should be anomalous if it cannot naturally be interpreted as providing an explanation for what is asserted in the rest of the sentence. Nouns that typically have only a single referent are good test nouns for the felicity of non-restrictive modifiers, as restrictive modification should in such cases be impossible. Thus, compare *I visited my sick mother* with the minimally different *I visited my tall mother*. The latter sentence is considerably less natural than the former, which is precisely what our coherence-based hypothesis predicts: there is no natural causal or explanatory association between being tall and needing to be visited. We conclude that non-restrictive modifiers are pragmatically licensed only if they establish a discourse coherence relation, which is often manifest as an explanation of some other event described in the sentence. A related finding within restrictive modifiers has been reported by Rohde et al. (2011), who showed that processing is facilitated when restrictive modification is also explanatory. Non-restrictive modification has not, however, yet figured in the psycho- or neurolinguistic research on syntax and semantics.

1.3. The LATL as a locus of composition

A large body of neurolinguistic work has focused on understanding the internal architecture of composition. Many of these studies have implicated the LATL as playing an important role in combinatory processes (Dronkers & Wilkins, 2004). Structured sentences elicit greater LATL activity than meaningless sentences or word lists (Friederici, Meyer, & von Cramon, 2000; Humphries, Binder, Medler, & Liebenthal, 2006, 2007; Mazoyer et al., 1993; Pallier, Devauchelle, & Dehaene, 2011; Rogalsky & Hickok, 2009; Stowe et al., 1998; Xu, Kemeny, Park, Frattali, & Braun, 2005). Recent work has focused directly on simple composition – the procedure of combining two linearly adjacent words and computing the meaning of the resulting phrase (Bemis & Pyllkänen, 2011, 2012, 2013). The most consistent finding in these studies has been that the LATL exhibits increased activity during the processing of simple adjective-noun phrases (e.g. *red boat*) compared to various non-compositional control trials (e.g. *xkh boat*). Relatedly, a separate set of experiments has implicated the LATL in ‘conceptual combination’, a term used in the psychology literature for the operation required to form

complex concepts from simpler ones, e.g., the concepts denoted by *male child* or *boy* from the individual words *male* and *child* (Baron and Osherson, 2011; Baron et al., 2010). These results show that the LATL is sensitive to very basic combinatory operations. However, although this work solidifies the role of the LATL in composition, none of these results can speak to the question of restrictiveness; the adjective-noun and conceptual combination experiments most likely only involved restrictive interpretations (due to a lack of contextual information to support a non-restrictive interpretation) and the sentence vs. word list contrast is too gross to speak to detailed functional hypotheses.

The literature on the LATL also extends well beyond online language processing, one of the prominent hypotheses being that it is a ‘semantic hub’ in which many kinds of meaning-related tasks are performed. Damage to the anterior temporal lobes causes severe amodal semantic memory deficits (Gainotti, 2006, 2007, 2011; Garrard & Carroll, 2006; Garrard & Hodges, 2000; Hodges, Graham, & Patterson, 1995; Hodges, Patterson, Oxbury, & Funnell, 1992; Mummery et al., 1999, 2000; Patterson et al., 2006; Rogers et al., 2004; Snowden, Goulding, & Neary, 1989). Several haemodynamic experiments have shown that increased ATL activity is observed when subjects access ‘specific level’ concepts, as opposed to ‘basic concepts’, a finding that is broadly compatible with the hypothesis that the LATL performs set restriction (Bright, Moss, & Tyler, 2004; Clarke, Taylor, Devereux, Randall, & Tyler, 2013; Clarke, Taylor, & Tyler, 2011; Gauthier, Anderson, Tarr, Skudlarski, & Gore, 1997; Rogers et al., 2006; Tyler et al., 2004). The LATL has additionally been implicated in processing ‘unique’ stimuli (Damasio, Grabowski, Tranel, Hichwa, & Damasio, 1996; Gainotti, 2007; Gainotti, Barbier, & Marra, 2003; Grabowski et al., 2001; Ross & Olson, 2012), as well as some functions surely required for pragmatic reasoning about language – social conceptual knowledge (Zahn et al., 2007), social cognition (Olson et al., 2007; Ross & Olson, 2010; Simmons & Martin, 2009; Simmons, Reddish, Bellgowan, & Martin, 2010) and theory of mind (Gallagher & Frith, 2003), for example.

One open question both within the composition and the semantic memory literature has been to what extent the LATL effects discussed earlier are uniquely left lateral, as opposed to bilateral. Previous results in our lab have primarily been located in the LATL (Bemis & Pyllkänen, 2011, 2013; Brennan et al., 2010), although some studies have also implicated the right anterior temporal lobe (RATL) as subserving semantic memory (Lambon Ralph et al., 2009; Lambon Ralph, Cipolotti, Manes, & Patterson, 2010). Given this profile, we treated both the LATL and RATLs as regions of interest (ROIs) in this study.

1.4. Current study

In the context of the broad literature cited above, it is plausible that the LATL could be simultaneously sensitive to both restriction and the pragmatic processes associated with non-restrictive adjectives (discussed in Section 1.2). To enable our design to detect both types of effects, we manipulated the restrictiveness of two different prenominal categories: adjectives, which in their non-restrictive readings elicit special pragmatic inferencing, and determiners, which do not. We thus had four critical conditions in this experiment: restrictive adjective (Restr-Adj), non-restrictive adjective (NonRestr-Adj), restrictive determiner (Restr-Det) and non-restrictive determiner (NonRestr-Det). The critical stimuli were presented as answers to questions, which were formulated to elicit either restrictive or non-restrictive interpretations of particular words in the answers. The lexical material in the critical stimuli was kept constant across the restrictiveness manipulation, thus allowing us to isolate different interpretations of the same phrase (the answer) by varying the background information (the question). To illustrate the paradigm, consider the noun phrase *his fat chicken*. This phrase could be used to answer either of the questions in (1)–(2). However, each question imposes a certain interpretation on *fat*: in (1) *fat* must be understood restrictively (since the question establishes a multitude of chickens), while in (2) *fat* must be non-restrictively (since the question tells us there is only a single chicken under discussion).

- (1) Restrictive, adjective (target word **bolded** throughout):
Q: Which chicken should the farmer slaughter next? A: His fat **chicken**.
(2) Non-restrictive, adjective:
Q: Will the farmer slaughter his chicken or his lamb? A: His fat **chicken**.

We used a parallel strategy to manipulate the restrictiveness/contrastiveness of prenominal determiners. It should be noted that within linguistic theory, the terms ‘restrictive’ and ‘non-restrictive’ are typically applied only to modifiers (e.g. adjectives, adverbs, relative clauses) and not to other categories such as determiners or nouns. However, in this paper, we use the term ‘(non-)restrictive’ in a slightly extended but nevertheless intuitive sense: the possessive determiner *his* in (3) is used contrastively – it specifies a particular chicken from a multitude of chickens – and so by analogy with (1) we call *his* restrictive in this context. Similarly, since *his* in (4) is not used contrastively, we call it non-restrictive, as in (2).

- (3) Restrictive, determiner:
Q: Will the farmer slaughter his chicken or Mary’s chicken? A: His **chicken**.
(4) Non-restrictive, determiner:
Q: Will the farmer slaughter his chicken or his lamb? A: His **chicken**.

The stimuli were organised into sets of seven trials each: in each set was one question-answer pair for each of the types illustrated in (1)–(4), and also three anomalous question-answer pairs as fillers, whose answers were in various ways infelicitous/incoherent relative to the question presented before them. Some examples of anomalous question-answer pairs are shown in (5).

(5) Anomalous fillers:

Q: Will the farmer slaughter his chicken or Mary's chicken? A: His lamb.

Q: Will the farmer slaughter his chicken or his lamb? A: His fat pig.

Q: Which chicken will the farmer slaughter next? A: His chicken.

Lexical items were held constant within each set, so that, e.g. the seven question-answer pairs shown in (1)–(5) constitute a complete set. The distance between the target word in the question and successive answer varied across question-answer pairs, but the wording of questions was quite varied, and there was no systematic difference in the distance between the target word in the question and answer across conditions.

Magnetoencephalography (MEG) was used to measure neural activity time-locked to the target noun *chicken* across these four conditions. In the restrictive conditions, the word immediately preceding the noun serves to narrow down reference within the set of individuals introduced in the question, while in the non-restrictive conditions, this is not the case.

Our aim was to determine whether the restrictiveness of composition in the answer phrases, as determined by the preceding question, has an influence on the composition effect in the LATL. In prior studies on adjective-noun combinations, these composition effects have been observed at around 200–250 ms post-target onset (Bemis & Pykkänen, 2011, 2012, 2013; Westerlund & Pykkänen, 2014). However, given that the lexical material in the answer phrases was rather constrained, we defined our interval of interest, 0–300 ms, to also capture potential earlier effects, given that effects of various linguistic variables have been attested as early as 100 ms (Dikker, Rabagliati, & Pykkänen, 2009; Pulvermüller, Shtyrov, & Hauk, 2009) or even 80 ms (Keuper et al., 2013, 2014). We hypothesised that if previous composition effects in the LATL (Bemis & Pykkänen, 2011, 2013) in fact reflected set restriction at the noun – i.e., for *yellow boat*, a narrowing down of the set to only those entities that are both yellow and a boat – then the restrictive conditions should elicit larger combinatory responses in the LATL than the non-restrictive conditions. Intuitive motivation for this hypothesis comes from considering together results from three different bodies of research previously discussed. First, there is strong evidence that the default interpretation of an attributive adjective is

restrictive (or ‘contrastive’), and that non-restrictive interpretations require stricter conditions of use (see e.g. Grodner & Sedivy, 2011, pp. 267–268). Therefore, if a subject is required to compose an adjective with a noun in the absence of context, we assume that the adjective receives its default restrictive interpretation in this situation. In previous experiments investigating the role of the LATL in simple composition, the critical stimuli were adjective-noun phrases presented in isolation or ‘out-of-the-blue’. Therefore, we assume that subjects interpreted the adjectives in these experiments restrictively.

Additionally, the LATL has been implicated in ‘conceptual combination’, a concept from cognitive psychology that corresponds fairly closely with the linguistic notion of restrictive modification (Baron & Osherson, 2011; Baron et al., 2010). Finally, studies on the role of the LATL in semantic memory focus on its increasing activation as a concept gets more specific, and arguably the reference set gets smaller (e.g. calling a concept a ‘bird’ as opposed to an ‘animal’). This type of process analogises to the linguist’s ‘restrictive modification’. The hypothesis that restriction increases LATL amplitude is therefore natural in light of such results.

Note that the predictabilities of the target noun *chicken* across the four critical conditions (1–4) vary: the target nouns of the restrictive conditions are more predictable than those of the non-restrictive conditions (see Section 2 for details). Previous work on prediction in language processing suggests that the more predictable a stimulus is, the lower the amplitude of the response it produces, and this result has been shown in a multitude of brain regions, including the LATL (e.g., Halgren et al., 2002; Lau, Gramfort, Hamalainen, & Kuperberg, 2013). Because prediction tends to decrease the amplitude of language processing responses, the differences in predictability between our conditions would tend to bias us against finding increases for the restrictive conditions, and thus the opposite finding would in fact constitute rather strong evidence for a robust effect of restriction.

Alternatively, if the LATL performs composition in a very general sense, i.e., any instance of composing A and B together engages the LATL, there should be no difference between the restrictive and non-restrictive conditions.

Finally, given the multifaceted effect profile observed for the LATL in prior research, it could plausibly show sensitivity both to restriction and the pragmatic processes required when interpreting non-restrictive adjectives. Sensitivity to restriction is a prediction that naturally arises from the specificity and uniqueness effects observed in the semantic memory literature (Grabowski et al., 2001; Rogers et al., 2006) – in a sense these could be interpreted as restriction effects within single words/concepts. Sensitivity to pragmatic processes would be consonant with the LATL’s proposed role in social cognition and theory of mind (Olson et al., 2007). If this combination of effects

were present, we would expect an increase for restriction when special pragmatics is not involved, i.e., in the Determiner conditions. However, in the adjectival cases, there should be less contrast between the restrictive and non-restrictive cases, as in a sense both amplitudes would be elevated: the restrictive adjectives eliciting increased activation due to their restrictiveness and the non-restrictive adjectives due to their special pragmatic properties. The special pragmatics of the non-restrictive adjective could of course elicit even higher amplitudes than the restrictive adjectives, which would constitute stronger evidence for the LATL's sensitivity to pragmatic inferencing. The predictions of all three hypotheses are sketched in Figure 1.

2. Materials and methods

2.1. Participants

Thirty-three participants participated in the experiment. All participants were right-handed native English speakers with normal or corrected-to-normal vision. All procedures were approved by New York University's Committee on Activities Involving Human participants and informed written consent was obtained from each participant. Over the course of recording this experiment, a significant noise issue arose, interfering with a number of recording sessions. The issue was subsequently resolved, but we were unable to use data from nine participants, and one further participant was excluded due to excess head movement. Two participants failed to understand the task and were excluded from future analysis, and a further three were excluded for having task accuracy below two standard deviations from the mean. Three subjects were determined to have responded with the opposite button on the button box for the duration of the experiment because their accuracy was within two standard deviations of 0 (meaning perfectly incorrect) for all conditions. These

participants were included. This left us with a total of 18 participants for final analysis (8 female, mean age = 23 years, SD = 3.9).

2.2. Materials and design

The notion of 'restriction' that we investigated corresponds quite closely to the notion of 'contrastive interpretation' explored in a number of psycholinguistic studies (e.g. Sedivy et al., 1999; Sedivy, 2003), as discussed in Section 1.1. Since contrast is a contextually determined property, we used the kind of question-answer pairs illustrated in 1.4 to form the stimuli for this experiment. The idea is that the question provides background information in the context of which the answer is to be interpreted. This allowed us to use the same noun phrase for both restrictive and non-restrictive trials, keeping lexical information matched across conditions. Recall that in (1) above, for example, the question establishes the existence of a multitude of chickens, and thus *fat* in the answer is interpreted restrictively. By contrast, the question in (2) establishes that there is just a single chicken in the discourse context, and hence the modifier *fat* provides an extra piece of information about that chicken, and is interpreted non-restrictively (note also the inferred information that the fatness of the chicken provides a reason for the farmer to slaughter it). All critical stimuli were noun phrases containing either the definite article *the* or a possessive determiner, e.g. *his*, *Mary's*.

Importantly, there is a predictability confound inherent in this design, namely that any felicitous answer to the question in (1) will contain the noun *chicken*, whereas in (2), *his (fat) lamb* is also a possible answer. Thus, it is not possible to rule out the hypothesis that differences between Restr-Adj and NonRestr-Adj are due to effects of predictability (see e.g. Dikker & Pykkänen, 2012). However, the Det conditions illustrated in (3)–(4) above

Figure 1. Hypothesised LATL activity profiles. If the LATL only computes restrictive composition, both restrictive conditions should elicit increased amplitudes (Hypothesis 1). In contrast, an increase for both adjectivally modified conditions would suggest a more general role for the LATL in composition (Hypothesis 2). Finally, sensitivity to both restriction and pragmatic inferencing predicts an increase for restriction for the Determiner cases (with no specially interpreted adjectives) and a potential increase for the non-restrictive adjectives over the restrictive ones, driven by pragmatics (Hypothesis 3).

serve essentially as predictability controls: these conditions (Restr-Det and NonRestr-Det) differ in a way that roughly parallels the difference between Restr-Adj and NonRestr-Adj, as discussed in Section 1.4. In (3) (Restr-Det), when the determiner *his* is presented, the identity of the noun *chicken* is already guaranteed. But in (4) (NonRestr-Det), no such inference can be made. Similarly, the noun *chicken* can be inferred immediately after the question in (Restr-Adj) but not in (NonRestr-Adj). In this way, Restr-Det and NonRestr-Det mimic the predictability difference between Restr-Adj and NonRestr-Adj.²

In sum, we employed a 2×2 design with Restriction (Restr vs. NonRestr) and Category (Adjective vs. Determiner) as factors. Before the experiment, we conducted a well-formedness norming study on our experimental stimuli using Amazon Mechanical Turk (www.mturk.com). One hundred and five respondents were asked to ‘indicate on a scale from 1 to 7, where 1 = Very Unnatural and 7 = Very Natural, how natural or unnatural the answers sound given the question’. The naturalness measure was explained to respondents in the following way: ‘If a sentence is a natural sentence, it should be easy to imagine using such a sentence if an appropriate situation arises. If a sentence is unnatural, it should be hard or impossible to imagine saying it’. Stimuli were divided into seven blocks, and each respondent saw only one block, such that each block was seen by 15 respondents.

Based on these naturalness ratings, we selected 46 sets of stimuli (each set consisting of seven items as described above) from our original inventory of 53 sets of manually constructed question-answer pairs. On average, participants rated the determiner conditions, Restr-Det ($6.48 \pm .45$) and NonRestr-Det ($6.48 \pm .79$), somewhat higher than the adjective conditions, Restr-Adj ($5.62 \pm .65$) and NonRestr-Adj ($5.80 \pm .48$). Though all natural conditions received high ratings and the difference between the Adj and Det conditions was small, a two-way analysis of variance (ANOVA) revealed a significant main effect of Category (Adj or Det) ($F_{(1,45)} = 69.56, p < 0.001$). There

was no significant interaction of Restriction and Category. All anomalous filler conditions were, on average, rated far below the mid-point of the scale: Anom1 ($1.92 \pm .71$), Anom2 ($1.91 \pm .81$) and Anom3 ($2.05 \pm .87$).

During MEG recordings, the task proceeded as follows. For each trial, subjects were presented with a question, followed by a word-by-word presentation of a two- or three-word answer to the question. To ensure that participants were paying attention, participants were instructed to decide whether the answer presented on the screen constituted a ‘natural’ answer to the question that preceded it. Specifically, they were told to imagine that the answers were provided by a robot designed to mimic natural human communication, and that their input was intended to help the robot learn. Since the task was designed to make sure participants were paying attention, participants whose task accuracy was below two standard deviations from the mean were excluded from further analysis.

All participants saw all stimuli. All trials contained a question, fixation cross, a determiner and a noun, but only the Adj trials included an adjective between the determiner and the noun. At the beginning of each trial, the question appeared on the screen and remained until participants pressed a button to advance. After the button press, the fixation cross, determiner and adjective (if present) were presented sequentially for 300 ms each, with a 300 ms blank screen in between. The noun remained on the screen until participants pressed a button to indicate their response to the task. No feedback was given during experimental blocks. The inter-trial interval was normally distributed with a mean of 400 ms (SD = 100 ms). See Figure 2 for a visual representation of the stimuli and task.

Each word or group of words was displayed in white Courier type on a grey background. Stimuli were presented in pseudorandom order, using PsychToolBox (Brainard, 1997; Pelli, 1997) and projected onto a 7.3×5.5 -inch screen with a resolution of 1024×768 pixels, placed approximately 45 cm from the participant’s eye. Words subtended between 1.5° and 6° .

Figure 2. Trial structure.

2.3. Procedure

Before the MEG measurement, participants were given written and verbal instructions about the experimental task and practiced a set of 10 trial items outside of the MEG room. A Polhemus Fastrak 3D digitiser (Polhemus, VT, USA) was used to digitise participants' head shapes and record the position of five marker coils placed around the participants' faces. The marker coils allowed us to measure the position of the participants' heads with respect to the MEG sensors at the beginning and end of the experiment. We were thus able to constrain source localisation by co-registering the participants' head shapes to an average of the two marker position measurements.

During the experiment, participants lay in a dimly lit, magnetically shielded room. Blinks were recorded using an SR Research Eyelink 1000 Arm-Mounted Eyetracker sampling at 1000 Hz (SR-Research, Osgoode, ON, Canada). Participants were given three rest periods between blocks. MEG data were collected using a whole-head 157-channel axial gradiometer system (Kanazawa Institute of Technology, Tokyo, Japan) sampling at 1000 Hz with a low-pass filter at 200 Hz and a notch filter at 60 Hz. The entire recording session, including preparation time and practice, lasted approximately an hour and a half.

2.4. Behavioural analysis

Behavioural data for the three subjects who reversed the response buttons were included with their responses corrected. Participants' responses were analysed for accuracy and reaction time (RT). We performed a 2×2 ANOVA with Restriction and Category as factors on both accuracy and RT data.

2.5. MEG analysis

Continuously recorded MEG data were noise-reduced using the Continuously Adjusted Least Squares Method (CALM; Adachi et al., 2001), high-pass filtered at 1 Hz and epoched from 1300 ms prior to the onset of the noun to 400 ms post-onset. The long pre-noun window was used in order to allow analysis on the determiner and adjective (in trials that included one), both of which preceded the noun. An activity baseline (channel noise covariance matrix) was taken from the 100 ms before the determiner in each condition, which corresponded to 1300–1200 ms pre-stimulus in the three-word trials, and to 700–600 ms pre-stimulus in the two-word trials. We chose this early baseline rather than one immediately prior to the target noun in order to avoid introducing differences in the conditions based on the amount and type of lexical material presented before the target stimulus. We cleaned the raw data of artefacts by rejecting trials for which either the maximum amplitude exceeded 3000 ft, or the participant blinked within the critical time window. Blinks

were defined as a loss of pupil signal to the eyetracker. For participants for whom no eyetracker data was collected, eyeblinks were manually removed from the epoch by visual inspection. We then averaged the remaining data for each participant for each condition and low-pass filtered the waveforms at 40 Hz.

Next, we performed a ROI analysis on activity time-locked to the target noun, focusing on the LATL and RATL. This was followed by a whole-brain source analysis primarily aimed at verifying that effects seen in the ROI analysis in fact reflected activity in the temporal poles, rather than spillover from adjacent areas.

2.5.1. Minimum norm estimates

We created separate distributed L2 minimum norm source estimates for each condition average for each subject, using BESA 5.1 (MEGIS Software GmbH, Gräfelfing, Germany). In order to construct these source estimates, two shells, containing 713 evenly distributed regional sources each, were placed at 10% and 30% below a standard smoothed brain surface. Each regional source contained two orthogonally oriented dipoles, and the total activity at the source was calculated by taking the root mean square of the two dipoles. The largest source, either on the 10% or the 30% shell, was selected by BESA at each location. Minimum norm images were depth and spatiotemporally weighted, using a signal subspace correlation measure (Mosher & Leahy, 1998) and ranged from 1.27 to 2.27 in signal-to-noise ratio (SNR) across subjects.

2.5.2. ROI analysis

We performed ROI analyses time-locked to presentation of the target noun and focusing on the left and right anterior temporal poles (BA 38). We first used the Talarach daemon (Lancaster et al., 1997, 2000) to automatically assign Brodmann area labels to the 713 sources on the smooth BESA cortex. Then, we performed 2×2 cluster permutation ANOVAs (Maris & Oostenveld, 2007), with Restriction and Category as factors, on the sources corresponding to both left and right BA 38. We used the F values from the repeated measures ANOVA as our test statistic. Throughout, we used 10,000 permutations and the same cluster selection criteria as Bemis and Pylkkänen (2011); that is, we only considered clusters that maintained a significance of $p = 0.3$ for at least 10 consecutive time points. Previous experiments have shown that composition-related responses reliably occur around 200–250 ms (Bemis & Pylkkänen, 2011, 2013); therefore, we constrained our time window of interest to 0–300 ms post-presentation of the noun. Lastly, we performed follow-up two-tailed cluster permutation t -tests in the same time window to investigate the effect of Restriction within each Category. All p -values reported

were false discovery rate (FDR) corrected (Genovese, Lazar, & Nichols, 2002) to control for multiple comparisons across the ROIs.

2.5.3. Whole-brain analysis

Whole-brain analyses were conducted for each comparison by performing a paired *t*-test over each source at each time point, and plotting activity on the standard BESA brain that was significant at $p < 0.1$, uncorrected for multiple comparisons, for at least five adjacent sources and at least five consecutive time points. Because this test is very liberal, we used the whole-brain analysis only as illustration to confirm that our ROI analysis reflected activity localised to the temporal poles, rather than spillover activity from adjacent regions outside the temporal lobe.

3. Results

3.1. Behavioural results

Participants' response accuracy was in general quite high; across all subjects and all conditions, mean accuracy was 94.1% (SD = 6.5%). On average, participants responded within 886 ms (SD = 289 ms).

3.1.1. Reaction times

Overall, participants were slower to reject the anomalous filler items (mean = 1011 ms, SD = 285 ms) than to accept the natural ones (mean = 762 ms, SD = 240 ms). A paired two-tailed *t*-test confirmed that this difference was significant ($t_{(8)} = 19.28$, $p < 0.0001$). We assume this difference is simply a consequence of a general communicative principle according to which language users attempt to interpret others' utterances as felicitous whenever possible.

A 2×2 ANOVA with Restriction and Category as factors yielded a main effect of both Restriction ($F_{(1,17)} = 38.88$, $p < 0.0001$) and Category ($F_{(1,17)} = 20.85$, $p < 0.001$). On average, participants were faster to respond in the Restr conditions (Restr-Adj mean = 684 ms, SD = 252 ms; Restr-Det mean = 763 ms, SD = 247 ms) than in the NonRestr conditions (NonRestr-Adj mean = 744 ms, SD = 216 ms; NonRestr-Det mean = 857 ms, SD = 278 ms). Presumably, the increased speed for Restr trials was due to noun predictability (see Section 1.3). Participants were faster to respond in the Adj conditions than in the Det conditions. There was no significant interaction.

3.1.2. Accuracy

On average, participants were also less accurate in the anomalous filler trials ($t_{(8)} = 8.01$, $p < 0.0001$; 90.1%, SD = 7.0%) than in the natural trials (98.3%, SD = 1.0%), although we note that accuracy was almost at ceiling across the board.

A 2×2 ANOVA yielded a significant effect of Category ($F_{(1,17)} = 18.5$, $p < 0.001$) but not of Restriction.

Overall, participants were more accurate in the Adj conditions than the Det conditions. Condition means were as follows: Restr-Adj (99.5%, SD = .9%), Non-Restr-Adj (98.4%, SD = 1.8%), Restr-Det (97.7%, SD = 1.7%) and NonRestr-Det (97.7%, SD = 2.0%). Once again, there was no significant interaction. An important final observation about the accuracy results should be noted: recall from the discussion in Section 1.2 that non-restrictive adjectives are semantically/pragmatically anomalous if their extra pragmatic inference is not computed. We therefore take the high accuracies of the NonRestr-Adj condition as evidence that our participants were indeed correctly computing these inferences.

3.2. MEG data

3.2.1. ROI results: LATL and RATL

A cluster permutation ANOVA (Restriction $\times$ Category) yielded no main effects of Restriction or Category in either ROI. However, there were significant early interaction clusters between Restriction and Category, both in the LATL (72–197 ms, $p = 0.011$) and in the RATL (119–262 ms, $p = 0.009$, see Figure 4 for mean activity by condition in the LATL and RATL; see Figure 3 for root mean square sensor data across all subjects and conditions). Follow-up cluster permutation *t*-tests showed that these interactions were driven by opposite effects of Restriction in the unmodified (Det) and modified (Adj) conditions. In the LATL, pairwise comparisons of Restr-Det [average $\pm$ SD: $7.57 \pm .93$ nanoampere meters (nAm)] and NonRestr-Det ($5.10 \pm .58$ nAm) revealed significant clusters of increased activity in the Restr-Det condition in the LATL between 117 ms and 194 ms ($p = .047$). Similarly, in the RATL, pairwise comparisons revealed a significant increase for Restr-Det ($8.96 \pm .94$ nAm) over NonRestr-Det ($6.44 \pm .76$ nAm) between 108 ms and 261 ms ($p = .005$). For the adjectivally modified conditions, on the other hand, pairwise comparisons indicated the opposite pattern of activity. In both the LATL and the RATL, NonRestr-Adj (LATL: $7.25 \pm .77$ nAm, RATL: $7.74 \pm .89$ nAm) elicited larger amplitudes in both hemispheres than Restr-Adj (LATL: $5.33 \pm .58$ nAm, RATL: $5.76 \pm .66$ nAm), but not significantly so (although a potentially revealing near-significant cluster was observed in the LATL from 69 ms to 135 ms, $p = .15$).

These findings fit the interaction profile of Hypothesis 3: restriction increases LATL (and RATL) amplitudes as long as additional pragmatics are not involved (see Figure 1). This pattern could suggest that the ATLs are implicated in the additional pragmatic inferencing associated with non-restrictive *adjectival* modification.

Crucially, our effect pattern does not appear to be confounded with the predictabilities of the target nouns: a predictability-driven effect would have shown similar

Figure 3. Root mean square over all sensors across all subjects and conditions with box depicting the analysed time window. Below the plot, the trial structure is given by Category (Adj or Det) and shaded regions indicate a time window used for baseline correction.

responses for the restrictive cases throughout. However, given that within the Adj conditions, restrictive stimuli did elicit decreased amplitudes, our conclusions about any amplitude enhancement due to pragmatics in the LATL must remain more tentative than our conclusions about the restriction-driven increases within the Det stimuli.

3.2.2. Whole-brain results

LATL activity seen in the whole-brain comparisons conformed well to the results of the ROI analysis (Figure 4). Within the Det conditions, Restr-Det showed increased activity over NonRestr-Det in both temporal poles in the time window of the ROI interactions. In contrast, within the adjectival conditions, NonRestr-Adj elicited larger amplitudes than Restr-Adj in LATL and to some extent also in the RATL. These findings confirm that the interactions found in our ROI analysis does indeed localise to the temporal poles.

The whole-brain comparisons also revealed a cluster of sources localising to the left inferior frontal gyrus (IFG) that showed increased activity for Restr-Det over Non-Restr-Det, a similar pattern to the temporal poles. Though these effects represent an uncorrected comparison and would require further research to confirm, we draw attention to it because the LIFG is central to a prominent model of composition, or ‘unification’ (Hagoort, Baggio, & Willems, 2009), as well as being sensitive to manipulations of syntactic complexity (Friederici, Rüschemeyer, Hahne, & Fiebach, 2003; Grodzinsky & Santi, 2008), cognitive control (Badre & Wagner, 2007) and semantic retrieval (Thompson-Schill, D’Esposito, Aguirre, & Farah, 1997), amongst other factors. Therefore, it is intriguing that it may also show certain sensitivity to restrictiveness in the absence of pragmatic constraints.

4. Discussion

In this work, we exploited the semantic contrast between restrictive and non-restrictive composition to shed light on the nature of the computations carried out by the LATL and, more tentatively, its right hemisphere homologue. Restrictiveness was manipulated both in prenominal determiner and adjective positions, the former allowing for a pure contrast in restrictiveness while in the latter case, non-restrictiveness correlated with additional pragmatic inferencing – namely establishing something like a cause–effect coherence relation (see Section 1.2). We found that restrictive composition recruits the LATL and RATL but only when added pragmatic processing is not involved; that is, in the unmodified conditions. These findings are consistent with our third hypothesis, namely that the ATLs would show sensitivity both to restriction and to the pragmatic process required by the interpretation of non-restrictive adjectives (see Figure 1). Importantly, this hypothesis was agnostic as to whether pragmatic processes would elicit greater activity in the ATLs than restriction processes, thus it is not falsified by the fact that we do not see a significant increase for the non-restrictive adjective over the restrictive adjective condition. We do note, however, that non-restrictive but pragmatically rich adjectives trended towards larger amplitudes at a subsequent noun than their restrictive counterparts. In sum, these effects demonstrate the sensitivity of the LATL to factors at a ‘higher’ representational level than what has been demonstrated in prior research. As our restriction manipulation simply varied the context, the critical stimuli across the restrictiveness manipulation were lexically, syntactically and, in fact, truth-conditionally matched.³ Thus, they only differed in the way the composition interacted with reference resolution and the way it fed pragmatic processing.

Figure 4. ROI and whole-brain results for localised activity during the comprehension of the critical nouns. Cluster-based permutation tests (Maris and Oostenveld, 2007) revealed significant interactions between Restriction and Category both in the left and right temporal poles (BA 38), such that for determiners (right column) restrictive composition elicited increased LATL amplitudes as compared to non-restrictive composition, whereas the reverse pattern was observed for adjectives (left column). The histograms plot mean activations per condition for the interval during which the interaction was reliable (left BA 38: 72–197 ms, $p = 0.0108$; right BA 38: 119–262 ms, $p = .0087$). The rectangles in the waveform graphs show the time window of the interaction, and the shaded regions give the FDR significant cluster for the pairwise comparison between Restr-Det and NonRestr-Det (pairwise comparisons within adjectives did not reach significance). The bottom row plots uncorrected whole-brain contrasts of the effect of Restriction for determiners (right) and adjectives (left). Activity corresponding to the ROI analyses is boxed, showing an increase in left BA 38 for Restriction in determiners and a decrease in adjectives. For determiners, a parallel increase is also seen in right BA 38, whereas adjectives show no obvious right temporal pole effect, which is consistent with our ROI analysis.

A central question within the composition work on the LATL has been whether its contribution is syntactic or semantic (though see Rogalsky and Hickok (2009) for discussion of the possibility that spatially distinct regions within the LATL may be involved). Core evidence for a syntactic role has comprised of studies demonstrating increased LATL activity for ‘jabberwocky’ sentences, i.e., grammatical sentences made up of pseudoword stems and functional morphology, over unstructured lists of pseudowords (Friederici et al., 2000; Humphries et al., 2006; Mazoyer, et al. 1993). This research has typically assumed that pseudoword sentences do not elicit semantic processing, but although such sentences may elicit less semantic processing than regular grammatical sentences, they still convey a rich array of information about events, individuals and relations among them. For example, Humphries et al.’s (2006) pseudoword sentence *the solims on a sonting grilloted a yome and a sovir* clearly communicates that individuals belonging to the solim-category participated in a grilloting-activity that affected individuals belonging to the yome and sovir-categories. In other words, we quite intuitively semantically compose such sentences and thus data such as these do not rule out a semantic role for the LATL. Further, the conceptual specificity effects within the semantic memory literature are obviously not syntactic effects, and thus any hypothesis aiming for a unified explanation of the composition and specificity effects would need to be of a semantic nature. Our study shows that the LATL is sensitive to interpretative factors extending even beyond truth-conditional semantics. These findings strongly suggest that the LATL’s role is not syntactic structure building but rather some function contributing to meaning composition (Pallier et al., 2011; Vandenberghe, Nobre, & Price, 2002).

Compared to prior MEG studies on combinatory effects in the LATL, the effects observed here are very early. The simple adjective-noun composition tasks conducted by Bemis and Pykkänen (2011, 2012, 2013) systematically elicited combinatorial LATL effects at around ~180–250 ms after the noun onset. Our LATL interaction took place considerably earlier, at 72–197 ms. Early effects of context integration and semantic processing in the range of 80–100 ms are not without precedent (Pulvermüller et al., 2009) or even 80 ms (Keuper et al., 2013, 2014), particularly when lexical material is tightly controlled across conditions, as it was in our design. Additionally, in the current design, participants were presented with background information – via a context-setting question – prior to the composition part of each trial. This background information introduced the target noun into the discourse. Further, the restrictiveness manipulation took place on the word preceding the noun, and thus already at the onset of the (rather predictable) noun, participants could have anticipated whether to compose restrictively or non-restrictively.

A literature focusing on the top-down effects of context has shown modulations of responses ranging from very early sensory responses (Dikker et al., 2009; Dikker & Pykkänen, 2011; Kim & Lai, 2012) to later components (e.g. the N400, Halgren et al., 2002; Lau, Gramfort, Hamalainen, & Kuperberg, 2013). We believe that in our experiment it is plausible for contextual information to have constrained composition such that the word could start to be composed as soon as it was encountered. In the context of these factors, the earlier onset of our effects is unsurprising.

While discussing anticipation, it is also important to note that while the target noun is more predictable in the restrictive conditions than the non-restrictive conditions, we see an interaction of Restriction $\times$ Category, not a main effect of Restriction, thus making it unlikely that our effects are solely due to predictability (see Section 1.4 of the Introduction).

Although our hypotheses were mainly focused on the LATL, the RATL revealed a pattern of effects similar to the LATL: in both hemispheres, Restriction and Category interacted such that Restriction increased ATL amplitude only for the Determiner contrast, where special pragmatics were not at play. Overall, these results add to the body of work suggesting that semantic processing in the ATLs is essentially bilateral (e.g., Visser, Jefferies, & Ralph, 2010).

5. Conclusion

In this study, we aimed to contribute to the functional understanding of the LATL by testing whether its activity profile might fit a relatively narrow computational profile reflecting restriction, i.e., the narrowing down of the set of entities under discussion, an interpretation consistent with both the composition and specificity literatures on this region. We manipulated the restrictiveness of composition in two environments, one of which allowed for a relatively pure manipulation of restrictiveness (Det) and another in which non-restrictiveness induced special pragmatic processes (Adj). The LATL as well as the RATL showed increased amplitudes for Restriction when special pragmatics were not involved but an opposite pattern was observed when non-restrictiveness brought about a pragmatic inference. Thus, although we found that the LATL is sensitive to restriction, restriction alone does not underlie its functional profile: added pragmatic demands seem to also affect LATL amplitude. Overall, these findings are consistent with a central role of the LATL in semantic combinatory operations and demonstrate its sensitivity to subtle and previously uninvestigated semantic and pragmatic factors.

Acknowledgement

We thank Rebecca Egbert and Paul Del Prato for their help in data analysis and collection.

Funding

This research was supported by the National Science Foundation [grant number BCS-1221723] (Liina Pykkänen) and [grant number G1001] from the NYUAD Institute, New York University Abu Dhabi (Liina Pykkänen).

Notes

1. We simplify matters here by discussing only explanatory-type uses of non-restrictive modifiers in single sentences. There are other uses as well, as in *My sick mother is still relatively independent*, in which the relation between ‘my mother is sick’ and ‘she is still relatively independent’ appears to be that of ‘Violated Expectation’ (a discourse coherence relation introduced by Kehler, 2002). Determining the range of coherence relations that non-restrictive adjectives can establish, and what implications of a discourse they can establish coherence with, are both topics of ongoing research.
2. It should be noted that there are also more subtle ‘category predictability’ differences between conditions. For example, once an adjective is displayed (in Adj trials), the syntactic category (but not always the identity) of the next word is guaranteed to be Noun, whereas in the Det trials, there is no point in the presentation of the stimulus at which the category of a subsequent word can be deduced.
3. It should be noted that the restrictive vs. non-restrictive contrast has been hypothesised to correlate with distinct syntactic representations (Cinque, 2010; Larson, 1998, 1999) but this possibility does not pertain to our Det stimuli.

References

- Adachi, Y., Shimogawara, M., Higuchi, M., Haruta, Y., & Ochiai, M. (2001). Reduction of non-periodic environmental magnetic noise in MEG measurement by continuously adjusted least squares method. *IEEE Transactions on Applied Superconductivity*, 11(1), 669–672.
- Arnold, J. E., & Griffin, Z. M. (2007). The effect of additional characters on choice of referring expression: Everyone counts. *Journal of Memory and Language*, 56, 521–536. doi:10.1016/j.jml.2006.09.007
- Arts, A., Maes, A., Noordman, L., & Jansen, C. (2011). Over specification facilitates object identification. *Journal of Pragmatics*, 43, 361–374. doi:10.1016/j.pragma.2010.07.013
- Asher, N., & Lascarides, A. (1993). Temporal interpretation, discourse relations and common sense entailment. *Linguistics and Philosophy*, 16, 437–493. doi:10.1007/BF00986208
- Badre, D., & Wagner, A. D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia*, 45, 2883–2901. doi:10.1016/j.neuropsychologia.2007.06.015
- Baron, S. G., & Osherson, D. (2011). Evidence for conceptual combination in the left anterior temporal lobe. *Neuroimage*, 55(4), 1847–1852.
- Baron, S. G., Thompson-Schill, S. L., Weber, M., & Osherson, D. (2010). An early stage of conceptual combination: Superimposition of constituent concepts in left anterolateral temporal lobe. *Cognitive Neuroscience*, 1(1), 44–51.
- Barr, D. J. (2008). Pragmatic expectations and linguistic evidence: Listeners anticipate but do not integrate common ground. *Cognition*, 109(1), 18–40. doi:10.1016/j.cognition.2008.07.005
- Bemis, D. K., & Pykkänen, L. (2011). Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience*, 31(8), 2801–2814. doi:10.1523/JNEUROSCI.5003-10.2011
- Bemis, D. K., & Pykkänen, L. (2012). Basic linguistic composition recruits the left anterior temporal lobe and left angular gyrus during both listening and reading. *Cerebral Cortex*, 23(8), 1859–1873. doi:10.1093/cercor/bhs179
- Bemis, D. K., & Pykkänen, L. (2013). Flexible composition: MEG evidence for the deployment of basic combinatorial linguistic mechanisms in response to task demands. *PLOS ONE*, 8, e73949. doi:10.1371/journal.pone.0073949.g004
- Bolinger, D. (1967). Adjectives in English: Attribution and predication. *Lingua*, 18, 1–34. doi:10.1016/0024-3841(67)90018-6
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 443–446.
- Brennan, J., Nir, Y., Hasson, U., Malach, R., Heeger, D. J., & Pykkänen, L. (2010). Syntactic structure building in the anterior temporal lobe during natural story listening. *Brain Lang*, 120(2), 163–173.
- Bright, P., Moss, H., & Tyler, L. (2004). Unitary vs multiple semantics: PET studies of word and picture processing. *Brain and Language*, 89, 417–432. doi:10.1016/j.bandl.2004.01.010
- Cinque, G. (2010). *The syntax of adjectives: A comparative study* (vol. 57). Cambridge, MA: MIT Press.
- Clarke, A., Taylor, K. I., Devereux, B., Randall, B., & Tyler, L. K. (2013). From perception to conception: How meaningful objects are processed over time. *Cerebral Cortex*, 23, 187–197. doi:10.1093/cercor/bhs002
- Clarke, A., Taylor, K. I., & Tyler, L. K. (2011). The evolution of meaning: Spatio-temporal dynamics of visual object recognition. *Journal of Cognitive Neuroscience*, 23, 1887–1899. doi:10.1016/S0042-6989(02)00298-5
- Damasio, H., Grabowski, T. J., Tranel, D., Hichwa, R. D., & Damasio, A. R. (1996). A neural basis for lexical retrieval. *Nature*, 380, 499–505. doi:10.1038/380499a0
- Dikker, S., & Pykkänen, L. (2011). Before the N400: Effects of lexical-semantic violations in visual cortex. *Brain and Language*, 118(1–2), 23–28. doi:10.1016/j.bandl.2011.02.006
- Dikker, S., & Pykkänen, L. (2012). Predicting language: MEG evidence for lexical preactivation. *Brain and Language*, 127(1), 55–64.
- Dikker, S., Rabagliati, H., & Pykkänen, L. (2009). Sensitivity to Syntax in Visual Cortex. *Cognition*, 110(3), 293–321. doi:10.1016/j.cognition.2008.09.008
- Dronkers, N. F., & Wilkins, D. P. (2004). Lesion analysis of the brain areas involved in language comprehension. *Cognition*, 92(1), 145–177. doi:10.1016/j.cognition.2003.11.002
- Engelhardt, P. E., Bailey, K. G. D., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean Maxim of quantity? *Journal of Memory and Language*, 54, 554–573. doi:10.1016/j.jml.2005.12.009
- Friederici, A. D., Meyer, M., & von Cramon, D. Y. (2000). Auditory language comprehension: An event-related fMRI study on the processing of syntactic and lexical information. *Brain and Language*, 74, 289–300. doi:10.1006/brln.2000.2313
- Friederici, A. D., Rüschemeyer, S. A., Hahne, A., & Fiebach, C. J. (2003). The role of left inferior frontal and superior temporal cortex in sentence comprehension: Localizing syntactic and semantic processes. *Cerebral Cortex*, 13, 170–177. doi:10.1093/cercor/13.2.170
- Gainotti, G. (2006). Anatomical functional and cognitive determinants of semantic memory disorders. *Neuroscience & Biobehavioral Reviews*, 30, 577–594. doi:10.1016/j.neubiorev.2005.11.001

- Gainotti, G. (2007). Different patterns of famous people recognition disorders in patients with right and left anterior temporal lesions: A systematic review. *Neuropsychologia*, 45, 1591–1607. doi:10.1016/j.neuropsychologia.2006.12.013
- Gainotti, G. (2011). The format of conceptual representations disrupted in semantic dementia: A position paper. *Cortex*, 48, 521–529. doi:10.1016/j.cortex.2011.06.019
- Gainotti, G., Barbier, A., & Marra, C. (2003). Slowly progressive defect in recognition of familiar people in a patient with right anterior temporal atrophy. *Brain*, 126, 792–803. doi:10.1093/brain/awg092
- Gallagher, H. L., & Frith, C. D. (2003). Functional imaging of 'theory of mind'. *Trends in Cognitive Sciences*, 7(2), 77–83. doi:10.1016/S1364-6613(02)00025-6
- Garrard, P., & Carroll, E. (2006). Lost in semantic space: A multi-modal, non-verbal assessment of feature knowledge in semantic dementia. *Brain*, 129, 1152–1163. doi:10.1093/brain/awl069
- Garrard, P., & Hodges, J. (2000). Semantic dementia: Clinical, radiological and pathological perspectives. *Journal of Neurology*, 247, 409–422. doi:10.1007/s004150070169
- Gauthier, I., Anderson, A. W., Tarr, M. J., Skudlarski, P., & Gore, J. C. (1997). Levels of categorization in visual recognition studied using functional magnetic resonance imaging. *Current Biology*, 7, 645–651. doi:10.1016/S0960-9822(06)00291-0
- Genovese, C. R., Lazar, N. A., & Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, 15, 870–878. doi:10.1006/nimg.2001.1037
- Grabowski, T. J., Damasio, H., Tranel, D., BolesPonto, L. L., Hichwa, R. D., & Damasio, A. R. (2001). A role for left temporal pole in the retrieval of words for unique entities. *Human Brain Mapping*, 13(4), 199–212.
- Grodner, D., Gibson, E., & Watson, D. (2005). The influence of contextual contrast on syntactic processing: Evidence for strong-interaction in sentence comprehension. *Cognition*, 95, 275–296. doi:10.1016/j.cognition.2004.01.007
- Grodner, D. J., Klein, N. M., Carbary, K. M., & Tanenhaus, M. K. (2010). "Some," and possibly all, scalar inferences are not delayed: Evidence for immediate pragmatic enrichment. *Cognition*, 116(1), 42–55. doi:10.1016/j.cognition.2010.03.014
- Grodner, D. J., & Sedivy, J. C. (2011). The effect of speaker-specific information on pragmatic inferences (239–272). In N. J. Pearlmuter & E. Gibson (Eds.), *The Processing and Acquisition of Reference*. MIT Press.
- Grodzinsky, Y., & Santi, A. (2008). The battle for Broca's region. *Trends in Cognitive Sciences*, 12, 474–480. doi:10.1016/j.tics.2008.09.001
- Hagoort, P., Baggio, G., & Willems, R. M. (2009). Semantic unification. *The Cognitive Neurosciences*, 4, 819–836.
- Halgren, E., Dhond, R. P., Christensen, N., Van Petten, C., Marinkovic, K., Lewine, J. D., & Dale, A. M. (2002). N400-like magnetoencephalography responses modulated by semantic context, word frequency, and lexical class in sentences. *Neuroimage*, 17, 1101–1116. doi:10.1006/nimg.2002.1268
- Hanna, J. E., & Tanenhaus, M. K. (2004). Pragmatic effects on reference resolution in a collaborative task: Evidence from eye movements. *Cognitive Science*, 28(1), 105–115. doi:10.1207/s15516709cog2801_5
- Hanna, J. E., Tanenhaus, M. K., & Trueswell, J. C. (2003). The effects of common ground and perspective on domains of referential interpretation. *Journal of Memory and Language*, 49(1), 43–61. doi:10.1016/S0749-596X(03)00022-6
- Heim, I., & Kratzer, A. (1988). *Semantics in generative grammar*. Malden: Blackwell.
- Heller, D., Grodner, D., & Tanenhaus, M. K. (2008). The role of perspective in identifying domains of reference. *Cognition*, 108, 831–836. doi:10.1016/j.cognition.2008.04.008
- Hodges, J. R., Graham, N., & Patterson, K. (1995). Charting the progression in semantic dementia: Implications for the organization of semantic memory. *Semantic Knowledge and Semantic Representations*, 3(3–4), 463–495.
- Hodges, J. R., Patterson, K., Oxbury, S., & Funnell, E. (1992). Progressive fluent aphasia with temporal lobe atrophy. *Brain*, 115, 1783–1806. doi:10.1093/brain/115.6.1783
- Humphries, C., Binder, J. R., Medler, D. A., & Liebenthal, E. (2006). Syntactic and semantic modulation of neural activity during auditory sentence comprehension. *Journal of Cognitive Neuroscience*, 18, 665–679. doi:10.1006/nimg.2001.0933
- Humphries, C., Binder, J. R., Medler, D. A., & Liebenthal, E. (2007). Time course of semantic processes during sentence comprehension: An fMRI study. *NeuroImage*, 36, 924–932. doi:10.1016/j.neuroimage.2007.03.059
- Kehler, A. (2002). *Coherence, reference, and the theory of grammar*. Stanford, CA: CSLI Press.
- Keuper, K., Zwanzger, P., Nordt, M., Eden, A., Laeger, I., Zwitserlood, P., & Dobel, C. (2014). How "love" and "hate" differ from "sleep": Using combined electro/magnetoencephalographic data to reveal the sources of early cortical responses to emotional words. *Human Brain Mapping*, 35, 875–888. doi:10.1002/hbm.22220
- Keuper, K., Zwitserlood, P., Rehbein, M. A., Eden, A. S., Laeger, I., Junghöfer, M., ... Dobel, C. (2013). Early prefrontal brain responses to the hedonic quality of emotional words – A simultaneous EEG and MEG study. *PLoS One*, 8, e70788. doi:10.1371/journal.pone.0070788.s001
- Kim, A., & Lai, V. (2012). Rapid interactions between lexical semantic and word form analysis during word recognition in context: Evidence from ERPs. *Journal of Cognitive Neuroscience*, 24(5), 1104–1112.
- Koornneef, A. W., & Sanders, T. J. M. (2013). Establishing coherence relations in discourse: The influence of implicit causality and connectives on pronoun resolution. *Language and Cognitive Processes*, 28, 1169–1206. doi:10.1080/01690965.2012.699076
- Lambon Ralph, M. A., Cipolotti, L., Manes, F., & Patterson, K. (2010). Taking both sides: Do unilateral anterior temporal lobe lesions disrupt semantic memory? *Brain*, 133(11), 3243–3255.
- Lambon Ralph, M. A., Pobric, G., & Jefferies, E. (2009). Conceptual knowledge is underpinned by the temporal pole bilaterally: Convergent evidence from rTMS. *Cerebral Cortex*, 19(4), 832.
- Lancaster, J. L., Summerlin, J. L., Rainey, L., Freitas, C. S., & Fox, P. T. (1997). The Talairach Daemon, a database server for Talairach atlas labels. *Neuroimage*, 5(4), 5634.
- Lancaster, J. L., Woldorff, M. G., Parsons, L. M., Liotti, M., Freitas, C. S., Rainey, L., Kochunov, P. V., Nickerson, D., Mikiten, S. A., & Fox, P. T. (2000). Automated Talairach atlas labels for functional brain mapping. *Human Brain Mapping*, 10(3), 120–131.
- Larson, R. K. (1998). Events and modification in nominals. *Proceedings of SALT*, 8, 145–168.
- Larson, R. K. (1999). Semantics of adjectival modification. In *Lectures presented at the Dutch National Graduate School (LOT)*, Amsterdam, The Netherlands.
- Lau, E. F., Gramfort, A., Hamalainen, M. S., & Kuperberg, G. R. (2013). Automatic semantic facilitation in anterior temporal cortex revealed through multimodal neuroimaging. *The Journal of Neuroscience*, 33, 17174–17181. doi:10.1523/JNEUROSCI.1018-13.2013

- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190.
- Mazoyer, B. M., Tzourio, N., Frak, V., Syrota, A., Murayama, N., Levrier, O., ... Mehler, J. (1993). The cortical representation of speech. *Journal of Cognitive Neuroscience*, 5, 467–479. doi:10.1126/science.1589767
- Morzycki, M. (2008). Nonrestrictive modifiers in nonparenthetical positions. In L. McNally & C. Kennedy (Eds.), *Adjectives and adverbs: Syntax, semantics, and discourse* (pp. 101–122). Oxford: Oxford University Press.
- Mosher, J. C., & Leahy, R. M. (1998). Recursive MUSIC: A framework for EEG and MEG source localization. *IEEE Transactions on Biomedical Engineering*, 45(11), 1342–1354.
- Mummery, C., Patterson, K., Price, C., Ashburner, J., Frackowiak, R., & Hodges, J. R. (2000). A voxel-based morphometry study of semantic dementia: Relationship between temporal lobe atrophy and semantic memory. *Annals of Neurology*, 47(1), 36–45. doi:10.1002/1531-8249(200001)47:1<36::AID-ANA8>3.0.CO;2-L
- Mummery, C., Patterson, K., Wise, R., Vandenberg, R., Price, C., & Hodges, J. (1999). Disrupted temporal lobe connections in semantic dementia. *Brain*, 122(1), 61–73. doi:10.1093/brain/122.1.61
- Noveck, I. A., & Reboul, A. (2008). Experimental pragmatics: A Gricean turn in the study of language. *Trends in Cognitive Sciences*, 12, 425–431. doi:10.1016/j.tics.2008.07.009
- Olson, I. R., Plotzker, A., & Ezzyat, Y. (2007). The Enigmatic temporal pole: A review of findings on social and emotional processing. *Brain*, 130(7), 1718–1731.
- Pallier, C., Devauchelle, A.-D., & Dehaene, S. (2011). Cortical representation of the constituent structure of sentences. *Proceedings of the National Academy of Sciences*, 108, 2522–2527. doi:10.1073/pnas.1018711108
- Patterson, K., Ralph, M. A. L., Jefferies, E., Woollams, A., Jones, R., Hodges, J. R., & Rogers, T. T. (2006). 'Pre-semantic' cognition in semantic dementia: Six deficits in search of an explanation. *Journal of Cognitive Neuroscience*, 18, 169–183. doi:10.1162/089892904322755494
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Pulvermüller, F., Shtyrov, Y., & Hauk, O. (2009). Understanding in an instant: Neurophysiological evidence for mechanistic language circuits in the brain. *Brain and Language*, 110(2), 81–94. doi:10.1016/j.bandl.2008.12.001
- Rogalsky, C., & Hickok, G. (2009). Selective attention to semantic and syntactic features modulates sentence processing networks in anterior temporal cortex. *Cerebral Cortex*, 19, 786–796. doi:10.1093/cercor/bhn126
- Rogers, T. T., Hocking, J., Noppeney, U., Mechelli, A., Gorno-Tempini, M. L., Patterson, K., & Price, C. J. (2006). Anterior temporal cortex and semantic memory: Reconciling findings from neuropsychology and functional imaging. *Cognitive, Affective, & Behavioral Neuroscience*, 6, 201–213.
- 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*, 111, 205–235. doi:10.1037/0033-295X.111.1.205
- Rohde, H., Levy, R., & Kehler, A. (2011). Anticipating explanations in relative clause processing. *Cognition*, 118, 339–358. doi:10.1016/j.cognition.2010.10.016
- Ross, L. A., & Olson, I. R. (2010). Social cognition and the anterior temporal lobes. *NeuroImage*, 49, 3452–3462. doi:10.1016/j.neuroimage.2009.11.012
- Ross, L. A., & Olson, I. R. (2012). What's unique about unique entities? An fMRI investigation of the semantics of famous faces and landmarks. *Cerebral Cortex*, 22, 2005–2015. doi:10.1093/cercor/bhr274
- Sedivy, J. C. (2003). Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of Psycholinguistic Research*, 32(1), 3–23. doi:10.1023/A:1021928914454
- Sedivy, J. C., Tanenhaus, M. K., Chambers, C. G., & Carlson, G. N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, 71(2), 109–147. doi:10.1016/S0010-0277(99)00025-6
- Simmons, W. K., & Martin, A. (2009). The anterior temporal lobes and the functional architecture of semantic memory. *Journal of the International Neuropsychological Society*, 15, 645–649. doi:10.1017/S1355617709990348
- Simmons, W. K., Reddish, M., Bellgowan, P. S. F., & Martin, A. (2010). The selectivity and functional connectivity of the anterior temporal lobes. *Cerebral Cortex*, 20, 813–825. doi:10.1093/cercor/bhp149
- Snowden, J. S., Goulding, P., & Neary, D. (1989). Semantic dementia: A form of circumscribed cerebral atrophy. *Behavioural Neurology*, 2(3), 167–182.
- Stowe, L. A., Broere, C. A. J., Paans, A. M. J., Wijers, A. A., Mulder, G., Vaalburg, W., & Zwarts, F. (1998). Localizing components of a complex task: Sentence processing and working memory. *NeuroReport*, 9, 2995–2999. doi:10.1097/00001756-199809140-00014
- Thompson-Schill, S. L., D'Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: A reevaluation. *Proceedings of the National Academy of Sciences*, 94, 14792–14797. doi:10.1073/pnas.94.26.14792
- Tyler, L. K., Stamatakis, E. A., Bright, P., Acres, K., Abdallah, S., Rodd, J., & Moss, H. (2004). Processing objects at different levels of specificity. *Journal of Cognitive Neuroscience*, 16, 351–362. doi:10.1093/brain/107.3.829
- Vandenberghe, R., Nobre, A. C., & Price, C. J. (2002). The response of left temporal cortex to sentences. *Journal of Cognitive Neuroscience*, 14, 550–560. doi:10.1006/brln.1993.1054
- Vendler, Z. 1968. *Adjectives and nominalizations*. Paris, France: Walter De Gruyter.
- Visser, M., Jefferies, E., & Ralph, M. L. (2010). Semantic processing in the anterior temporal lobes: A meta-analysis of the functional neuroimaging literature. *Journal of Cognitive Neuroscience*, 22, 1083–1094. doi:10.1093/bmb/65.1.95
- Westerlund, M., & Pykkänen, L. (2014). The role of the left anterior temporal lobe in semantic composition vs. semantic memory. *Neuropsychologia*, 57, 59–70. doi:10.1016/j.neuropsychologia.2014.03.001
- Wolter, L., Gorman, K. S., & Tanenhaus, M. K. (2011). Scalar reference, contrast and discourse: Separating effects of linguistic discourse from availability of the referent. *Journal of Memory and Language*, 65, 299–317. doi:10.1016/j.jml.2011.04.010
- Xu, J., Kemeny, S., Park, G., Frattali, C., & Braun, A. (2005). Language in context: Emergent features of word, sentence, and narrative comprehension. *NeuroImage*, 25, 1002–1015. doi:10.1016/j.neuroimage.2004.12.013
- Zahn, R., Moll, J., Krueger, F., Huey, E. D., Garrido, G., & Grafman, J. (2007). Social concepts are represented in the superior anterior temporal cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 104(15), 6430–6435.