NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 33 (2006) 343 – 361

Making sense of discourse: An fMRI study of causal inferencing across sentences

Gina R. Kuperberg,[a,b,c,*] Balaji M. Lakshmanan,[b] David N. Caplan,[d] and Phillip J. Holcomb[a]

[a]*Department of Psychology, Tufts University, 490 Boston Avenue, Medford, MA 02155, USA*
[b]*MGH/MIT/HMS Athinoula A. Martinos Center for Biomedical Imaging, Bldg 149, 13th Street, Charlestown, MA 02129, USA*
[c]*Department of Psychiatry, Massachusetts General Hospital, Bldg 149, 13th Street, Charlestown, MA 02129, USA*
[d]*Department of Neurology, Massachusetts General Hospital, Boston, MA 02114, USA*

Received 19 February 2006; revised 31 May 2006; accepted 4 June 2006
Available online 28 July 2006

To build up coherence between sentences (comprehend discourse), we must draw inferences, i.e. activate and integrate information that is not actually stated. We used event-related fMRI to determine the localization and extent of brain activity mediating causal inferencing across short, three-sentence scenarios. Participants read and made causal coherence judgments to sentences that were highly causally related, intermediately related or unrelated to their preceding two-sentence contexts. The highly related and intermediately related scenarios were matched in terms of semantic similarities between their individual component words. A pre-rating study established that causal inferences were generated to the intermediately related but not to the highly related or unrelated scenarios. In the scanner, sentences that were intermediately related (relative to highly related or unrelated) to their preceding contexts were associated with longer judgment reaction times and sustained increases in hemodynamic activity within left lateral temporal/inferior parietal/prefrontal cortices, the right inferior prefrontal gyrus and bilateral superior medial prefrontal cortices. In contrast, sentences that were unrelated (relative to highly related) to their preceding contexts were associated with only transient increases in activity (at, but not after, the peak of the hemodynamic response) within the right lateral temporal cortex and the right inferior prefrontal gyrus. These data suggest that, to make sense of discourse, we activate a large bilateral cortical network in response to what is not explicitly stated. We suggest that this network reflects the activation, retrieval and integration of information from long-term semantic memory into incoming discourse structure during causal inferencing.

Keywords: Language; Discourse; Sentences; fMRI; Inferences; Causal relationships; Coherence; Semantics; Temporal cortex; Inferior prefrontal cortex; Dorsolateral prefrontal cortex; Frontomedian cortex

Introduction

As we read text or listen to spoken language, we access the meaning of individual words, combine words with syntactic structure to build up propositional meaning and combine propositions to determine the meaning of discourse as a whole (Gernsbacher, 1990; Kintsch, 1988; McKoon and Ratcliff, 1992; Van Dijk and Kintsch, 1983). Yet, much of what we read or hear is not explicitly stated. Our representation of discourse as a coherent structure, rather than as a series of unrelated propositions, is dependent on our ability to draw inferences that establish consistency and coherence between individual events (Graesser et al., 1994; McKoon and Ratcliff, 1992; van den Broek, 1994).

The focus of the current study is on causal inferences—information that is activated about causal relationships between sentences but that is not actually stated (Singer, 1994; van den Broek, 1994). Causal relationships and inferences are essential for constructing and remembering the content of narrative texts. In a series of important studies, Keenan et al. (1984), Myers et al. (1987) and Myers and Duffy (1990) examined reading times and recall of sentences such as "The next day his body was covered in bruises" when preceded by (1) highly causally related sentences (e.g. "Joey's brother punched him again and again"), (2) intermediately causally related sentences (e.g. "Joey's brother became furiously angry with him") and (3) unrelated sentences (e.g. "Joey went to a neighbor's house to play"). They demonstrated that reading times were longer and cued recall was poorer to the unrelated sentences than to the highly related sentences. Of particular interest was the pattern of findings to the intermediately related sentences. Although the reading times to these sentences were in between those to the unrelated and highly related sentences, their cued recall was even better than that of the highly related sentences. Myers et al. (1987) explained this non-linear pattern of recall by suggesting that readers generated an elaborative causal inference to the intermediately related sentence pairs, consolidating them in episodic memory and facilitating their cued retrieval. In contrast, less elaboration was necessary to establish

* Corresponding author. CNY-2, Bldg 149, Massachusetts General Hospital (East), Charlestown, MA 02129, USA.
E-mail address: kuperber@nmr.mgh.harvard.edu (G.R. Kuperberg).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2006.06.001

coherence between the highly related sentence pairs[1] and, although readers might have made an attempt to search for a connection between the unrelated pairs, no consistent inference was generated to connect them.

Neuropsychological and fMRI studies have provided some insights into the brain regions mediating the generation and integration of causal inferences. Lesion studies have classically implicated the right hemisphere in inferencing, as well as in other higher order language processes such as the interpretation of jokes (Brownell and Gardner, 1988; Brownell et al., 1983; Shammi and Stuss, 1999) and metaphor (Winner and Gardner, 1977). Patients with right hemisphere lesions sometimes produce speech that is socially inappropriate with tangential relationships between sentences, even when other aspects of their language comprehension and production are normal (Joanette et al., 1990). This contrasts with the classic aphasia syndromes that are associated with lesions around the left perisylvian cortex and in which clinical problems are generally at the level of generating and/or comprehending individual words and sentences (Caplan, 1992). Controlled experimental studies on patients with right hemisphere lesions report abnormalities in comprehending discourse that requires the generation of causal inferences for coherence (Beeman, 1993; Brownell et al., 1986). For example, Beeman (1993) reported that right hemisphere patients failed to generate appropriate inferences when asked explicit comprehension questions and were slower than healthy controls to make lexical decisions on inference-related (relative to unrelated) probe words.

The specificity of discourse-level processes to the right hemisphere, however, is debated. Zaidel et al. (2002) demonstrated that patients with left hemisphere lesions can sometimes perform just as badly as patients with right hemisphere lesions on so-called 'right hemisphere' tasks, and Ferstl et al. (2002) demonstrated that patients with left or bilateral frontal lobe lesions, compared with patients with left temporal or right frontal lesions, made more errors and took longer to judge causal relationships between related (relative to unrelated) sentence pairs. Evidence from divided visual field studies suggests that both the right and left hemispheres are able to extract the gist of a discourse message (its thematic and semantic information) (Long and Baynes, 2002; Long et al., 2005), although only the left hemisphere appears to extract structural information from propositions within sentences, perhaps because of its role in syntactic processing or in combining syntactic and semantic information (Faust et al., 2003; Long and Baynes, 2002; Long et al., 2005).

Neuroimaging studies of discourse processing have implicated multiple regions across both right and left hemispheres in establishing discourse coherence. Studies that have compared whole stories with a resting baseline (Mazoyer et al., 1993), unrelated words (Xu et al., 2005) and unrelated sentences (Fletcher et al., 1995; Xu et al., 2005) have highlighted the roles of regions outside the left perisylvian cortex in higher-level language processing, including bilateral anterior temporal cortices (Fletcher et al., 1995; Mazoyer et al., 1993; Xu et al., 2005), the medial prefrontal cortices (Fletcher et al., 1995; Gallagher et al., 2000; Xu et al., 2005) and temporal and prefrontal regions within the right hemisphere (Xu et al., 2005). Other fMRI studies that have directly compared coherent and incoherent text have also implicated the right hemisphere in establishing discourse coherence, supporting some of the neuropsychological literature. For example, the right temporal cortex was engaged to a greater degree to stories presented without an explanatory title (less coherent) than to stories presented with such a title context (more coherent) (St George et al., 1999), and the right inferior prefrontal cortex showed more activity to sentences without definite articles (less coherent) than to stories with definite articles (more coherent) (Robertson et al., 2000). The interpretation of many of these studies, however, is limited because of their blocked designs and their focus on overall global coherence rather than on the inferential processes engaged as connections are established between individual sentences. The development of event-related fMRI has allowed a more precise examination of the brain regions mediating causal inferencing. Studies by two research groups are particularly relevant to the current study.

First, Ferstl and von Cramon contrasted brain activity to coherent (related in meaning) and incoherent (unrelated in meaning) sentence pairs, presented both visually (Ferstl and von Cramon, 2001)[2] and aurally (Ferstl and von Cramon, 2002), as subjects judged whether the sentences were pragmatically or logically linked. In both studies, the main region activated to the coherent (relative to the incoherent) sentence pairs was the left medial prefrontal cortex. This was interpreted as mediating controlled, elaborative inferential processing. Of particular note, this contrast failed to activate the right or left lateral temporal or inferior prefrontal cortices.

Another study by Mason and Just (2004), however, did implicate right-lateralized brain regions in causal inferencing. Like the studies of Keenan et al. (1984), Myers et al. (1987) and Myers and Duffy (1990) described above, Mason and Just (2004) used scenarios with three levels of casual relatedness. The right temporal cortex showed more activity to the intermediately related sentences than to both the highly related and the unrelated scenarios, mirroring the non-linear pattern of recall across these types of scenarios reported in the studies by Keenan, Myers and colleagues. The authors argued that this right-sided activation reflected participants' tendencies to generate elaborative causal inferences to the intermediately related sentences. There were, however, problems with this study that limit the interpretation and generalizability of these findings. First, large portions of brain (including medial prefrontal orbitofrontal cortices) were not scanned and/or analyzed at all. Second, because activity was only averaged across voxels within large regions of interest (ROIs), localized activity within these ROIs may not have been detected. Third, there were only ten stimuli per condition, again limiting power to detect BOLD modulation. Fourth, although the authors did not give details about how they analyzed the fMRI data, the reported effects appear to have captured BOLD activity to all the sentences within the discourse scenarios rather than specifically to the final sentence when inferences were most likely to have been generated. This might have further reduced power to detect inference-related BOLD activity.

The current study examined the brain regions mediating the generation and integration of causal inferences in short, three-sentence scenarios (a two-sentence context and a critical final sentence). Like

[1] We do not imply that highly related scenarios are not associated with any type of inference generation. Studies by Singer and others suggest that elaborative inferences, drawing upon world knowledge, are generated even to highly related sentence pairs (Singer, 1994).

[2] In Ferstl and von Cramon (2001), coherence was crossed with cohesion (the inclusion of words that explicitly indicated a connection between the two sentences).

Ferstl and von Cramon (2001, 2002), we included both causally related and unrelated scenarios and focused on the hemodynamic response to critical final sentences on which participants generated inferences to link them to their preceding discourse context. However, following Keenan et al. (1984) and Myers et al. (1987) and as in the fMRI study by Mason and Just (2004), we included two types of related scenarios—those with final sentences that were highly causally related to their preceding contexts and that did not require readers to generate causal inferences to establish coherence, and those with final sentences that were intermediately related to their preceding context and in which readers were more likely to generate causal inferences to establish coherence.

We approached the analysis of fMRI data in two complementary ways. First, we separately modeled the hemodynamic response to the context (the first two sentences) and the final sentence (that distinguished between the three conditions). The advantage of this approach is that it allowed us to test our *a priori* hypotheses regarding differences between the three conditions at the point of the final sentence when inferential processes were most likely to have taken place. However, this approach makes specific assumptions about the shape of the hemodynamic response function as well as the time course of inferential processes. We therefore also adopted a more flexible finite impulse response (FIR) model to construct hemodynamic time courses without assumptions about its shape (Burock et al., 1998; Dale, 1999), focusing on ROIs that were defined on purely anatomical grounds (Caviness et al., 1996; Rademacher et al., 1992). This approach allowed a finer grained analysis of the time course of the hemodynamic response to each of the three different types of scenarios.

On the basis of the studies reviewed above, including the neuroimaging studies by Ferstl and von Cramon (2001, 2002) and Mason and Just (2004), we hypothesized that the generation and integration of causal inferences would be associated with increased hemodynamic activity within a network that included left and right lateral temporal, inferior prefrontal and superior medial prefrontal cortices. We predicted that these regions would show increased activity to the intermediately related sentences relative to both the highly related and the unrelated scenarios.

Methods

Construction of stimuli

Two hundred and forty sets of three-sentence scenarios, each with three conditions – highly related, intermediately related and unrelated – were constructed as described in Table 1. The scenarios were divided into three lists (each with 240 sentences, 80 in each condition) that were counterbalanced between subjects such that no participant encountered the same final sentence more than once and such that, across all subjects, all final sentences were seen in all conditions. Discourse scenarios were then randomized within lists.

In initially constructing these scenarios and assigning them to the highly related and intermediately related conditions, two experimenters employed two checks. To be assigned to the highly related condition, sentence 3 followed by “because” followed by sentence 2 had to make intuitive sense (“The boys were having an argument. The next morning they had many bruises *because* they hit each other”). In addition, when presented with sentence 1 and sentence 3 and then asked why sentence 3 occurred, the answer had to be what was described in sentence 2 (“The boys were having an argument. The next morning they had many bruises.” Question: “Why did they have many bruises?” Answer: “Because they hit each other”). To be assigned to the intermediately related condition, sentence 3 followed by “because” followed by sentence 2 had to appear somewhat anomalous (“The boys were having an argument. The next morning they had many bruises *because* they got more and more angry”); and, when presented with sentence 1 and sentence 3 and then asked why sentence 3 occurred, the answer could not be what was described in sentence 2 of the intermediately related condition but, rather, had to be what was described in sentence 2 of the highly related condition. For example, in the scenario, “The boys were having an argument. The next morning they had

Table 1

Scenario type	Construction and explanation[a]	Example
1. Highly related	The first sentence sets up a fairly non-constraining context. The second and third sentences are explicitly causally linked in meaning.	*“The boys were having an argument. They began hitting each other. The next day they had bruises.”*
2. Intermediately related	The same first and third sentences as in scenario type 1. The second sentence is constructed such that the reader is required to make an inference to connect the second and third sentences.	*“The boys were having an argument. They became more and more angry. The next day they had bruises.”* [In this example, the reader would have to infer that the boys hit each other.]
3. Unrelated	The same third sentence as in scenario type 1. The first and second sentences are taken from scenario type 1 from another list. (counterbalancing).	*“The boys were unsure about the weather. At noon they started to hike. The next day they had bruises.”* [In this example, there is no clear inference connecting the final sentence with its preceding context.]

In all scenarios, the first two sentences each contained between 4 and 10 words. Each final sentence contained 5 or 6 words. Pronouns were made consistent across all scenarios.

[a] See Methods for more detailed explanation of how these scenarios were constructed and pre-rated.

many bruises." Question: "Why did they have many bruises?", the expected answer would be "because they hit each other", and not "because they got more and more angry".

These checks were also carried out for the unrelated scenarios that were created after counterbalancing across three lists: first, sentence 3 followed by "because" followed by sentence 2 had to be clearly anomalous ("The boys were unsure about the weather. The next morning they had bruises *because* at noon they started to hike"); second, when presented with sentence 1 and sentence 3 and then asked why sentence 3 occurred, the answer had to be "I don't know" or nothing related to the previous two sentences ("The boys were unsure about the weather. The next morning they had many bruises." Question: "Why did they have many bruises?" Answer: "?").

The content words within the highly related and intermediately related conditions were matched in terms of numbers of word repetitions. In addition, highly related and intermediately related scenarios were also matched in terms of their semantic similarity values (SSVs). These were calculated using a Latent Semantic Analysis (LSA) (Landauer and Dumais, 1997; Landauer et al., 1998); available on the Internet at http://lsa.colorado.edu). In this LSA, pairwise comparisons using tasaALL space (corresponding to a 1st year college student reading level) were carried out such that all content words in each scenario were compared on a term-by-term basis, yielding SSVs for each scenario. There was no significant difference in SSVs between the highly related scenarios (mean: 0.18, SD: 0.086) and the intermediately related scenarios (mean: 0.17, SD: 0.00), $t(478)=0.87$, $p=0.38$. The SSVs of the intermediately related scenarios were significantly greater than those of the unrelated scenarios (mean: 0.13, SD: 0.07), $t(478)=5.9$, $p<0.001$.

Pretests of stimuli

In order to corroborate our previous assessments of causal relationship and inference generation, two pretests were carried out on two separate groups of participants who had no knowledge of the stimuli and who did not take part in the fMRI study.

Pretest 1: verification of inference generation. The aim of this pretest was to verify that participants did not generate consistent inferences to the highly related or the unrelated scenarios but that they did generate inferences to the intermediately related scenarios. The discourse scenarios were presented in random order to twelve Tufts undergraduate students (four for each of the three list). At the end of each discourse scenario, a "Why" question was presented, in italics. This question was constructed from the final sentence in each scenario (e.g., for the example given in Table 1, "Why did they have bruises?"). Subjects were given the following specific instructions followed by some examples:

> *Try to answer the question in italics on the basis of what you have just read. You should write a one-sentence answer that should be as clear and simple as possible. It may or may not be similar to one of the sentences that you have read. You may not be able to answer the question at all, in which case you should indicate 'don't know'.*

Inspection of subjects' answers indicated that, for the highly related scenarios, subjects wrote responses that were very similar to the second sentence for that scenario: that is, they repeated what they had just read. For the intermediately related scenarios, subjects wrote responses that were very similar to the second sentence of the highly related condition for that scenario, even though they had not seen that sentence: they made the expected inference. For the causally unrelated scenarios, subjects either indicated 'don't know' or, rarely, they wrote responses that were very different from the second sentence of the highly related condition for that scenario: they either failed to make an inference or any inferences generated were inconsistent across subjects.

Pretest 2: ratings of causal relatedness. A rating study was conducted (a) to obtain ratings of how related the final sentence was to the previous two sentences and (b) to objectively determine the word on which subjects made their rating decision – henceforth termed the 'critical word' – and to ensure that it was the same across the three experimental conditions. Scenarios were presented, in random order, to twelve undergraduate students from Tufts University who did not participate in Pretest 1 (four for each of the three list). The third sentence in each scenario was underlined and called the resulting sentence. Participants were given the following instructions:

> *Read each group of sentences once. (1) Give a rating or 1, 2 or 3 according to how strongly the third underlined sentence – the resulting sentence – is causally connected to the previous two sentences. (2) Select one word in the third underlined sentence – the resulting sentence – that you judge to be the first indication that the final sentence would or would not be causally connected to the previous two sentences.*

As expected, subjects and items analyses revealed significant differences in subjects' ratings across the three scenario types, $F1(2,18)=510.9$, $p<0.0001$, $F2(2,719)=2273.9$, $p<0.0001$. Subjects rated the highly related scenarios (mean: 1.08, SD: 0.07) as being significantly more related than the intermediately related scenarios (mean: 1.58, SD: 0.27), $t(11)=7.7$, $p<0.0001$, that were, in turn, rated as significantly more related than the unrelated scenarios (mean: 2.82, SD: 0.14), $t(11)=23.4$, $p<0.0001$.

In 90% ($N=648$) of the scenarios, the majority of raters selected the same word to be the critical word both (a) within and (b) across experimental conditions. In 10% ($N=72$) scenarios, there was not a majority opinion within and/or between scenarios as to which word was critical. These scenarios were further edited and subsequently presented to a further nine subjects (three per list) for rating until there was a majority opinion as to which was the critical word across all three conditions for each scenario.

The critical word was the third, fourth or fifth word in the final sentence. In half the sentences, the critical verb was the final word of the sentence, and in the other half, the sentence continued from one to three additional words.

fMRI study

Participants

Fifteen (12 male and 3 female; mean age: 40, SD: 11) subjects were recruited by advertisement. All participants were right-handed as assessed using the modified Edinburgh Handedness Inventory (Oldfield, 1971; White and Ashton, 1976). Selection criteria required all participants to have normal or corrected-to-normal vision and to be native speakers of English. In addition, volunteers were not taking any medication and were screened to exclude the presence of psychiatric and neurological disorders and to exclude the contraindications for MRI. Written consent was obtained from all subjects before participation according to the established guidelines of the Massachusetts General Hospital Institutional Review Board.

Stimulus presentation and task

During scanning, each discourse trial (three sentences) began with the presentation of a yellow fixation point at the center of the screen for 500 ms followed by a 100 ms interstimulus interval (ISI). The first two sentences were presented successively, each for 3.4 s with an interstimulus interval of 100 ms. The third sentence was then presented word by word; each word appeared on the screen for 500 ms with an ISI of 100 ms separating words. Given that the third sentence constituted 5 or 6 words, it lasted between 3 s and 3.6 s. The final word of the last sentence was followed by a "?" that cued subjects to make their responses and, depending on the number of words in the third sentence, lasted either 2.8 s or 3.4 s before the next trial started. Each trial thus lasted for 14 s. The subjects' task was to press one of three buttons on a response box (using the index, middle and third fingers of their left hand), depending on the difficulty of connecting the last sentence with the previous two sentences. Subjects' ratings and judgment reaction times (RTs) were recorded. The fingers corresponding to responses were counterbalanced across subjects. Subjects were instructed to wait until the "?" cue before responding. Sentence trials were pseudorandomly presented among fixation trials (16% of the time) in which subjects were asked to fixate on a white "+" symbol for variable durations, ranging from 2 to 14 s. The random interleaving of these 'fixation' or 'null-events' among the discourse scenarios is critical for the efficient estimation of the entire hemodynamic response in rapid event-related fMRI experimental designs, enabling the deconvolution of the recorded fMRI time courses (Burock et al., 1998).

Behavioral data analysis

RTs to each scenario type (collapsed across individual scenario items) were entered into repeated-measures ANOVAs with subjects as a random effect and scenario type as a within-subject factor (subjects analyses). In addition, we conducted corresponding items analyses ANOVAs in which the dependent variable was the RT for each individual scenario, collapsed over individual subjects. Here, scenario type was a between-subject factor. Significant effects of scenario type in all these ANOVAs were followed up using planned *t* tests comparing the different scenario types. Alpha was set to 0.05.

We also examined participants' ratings in relation to the three scenario types and repeated all RT analyses (a) including only trials in which participants' ratings were consistent with the three scenario types (i.e. highly related and rated 1 *vs.* intermediately related and rated 2 *vs.* unrelated and rated 3) and (b) comparing participants' ratings, regardless of the *a priori* conditions (i.e. rating 1 *vs.* rating 2 *vs.* rating 3).

For all these analyses, results on raw data are reported but all analyses were repeated after logarithmically transforming the data to reduce skew and revealed the same pattern of findings.

MRI data acquisition

Each participant was scanned on a separate structural and functional session. On both sessions, head motion was minimized using pillows and cushions around the head and a forehead strap.

In the structural session, two sets of high-resolution anatomical images were acquired in a 1.5 T whole-body Siemens Sonata scanner (Siemens Medical Systems, Iselin, NJ) using a T1-weighted MP-RAGE sequence (TR = 7.25 ms, TE = 3.0 ms, and flip angle = 7°). Volumes consisted of 128 sagittal slices with an effective thickness of 1.33 mm. The in-plane resolution was 1.33 mm × 1.0 mm (192 × 256 matrix, 256 mm FOV).

In the functional session, volume acquisitions were acquired using a T2*-weighted gradient-echo pulse sequence (TR = 2 s, TE = 25 ms, and flip angle = 90°) in a 3.0 T head-only Siemens Allegra scanner (Siemens Medical Systems, Iselin, NJ). The volume was comprised of 33 transverse slices aligned along the AC–PC plane. Slices were 3 mm thick with a distance of 0.9 mm between slices. The in-plane resolution was 3.13 × 3.13 mm (64 × 64 matrix, 200 mm FOV), and each functional run consisted of 120 such volume acquisitions for a total of 3960 images.

MRI data analysis

Reconstruction of cortical surfaces from structural MRI data. Following motion correction, the two high-resolution structural scans for each participant were averaged to increase contrast between gray and white matter, and the resulting volume was used to reconstruct a model of each individual's cortical surface (Dale et al., 1999; Dale and Sereno, 1993; Fischl et al., 2001). The surface representing the gray/white border was inflated (Dale and Sereno, 1993; Fischl et al., 1999a), differences between individuals in the depth of gyri/sulci were normalized and, for the purposes of averaging functional data across subjects (see below), each subject's reconstructed cortical surface was morphed/registered to an average spherical surface representation that optimally aligned sulcal and gyral features across subjects while minimizing metric distortion (Fischl et al., 1999a,b). This cortical reconstruction was carried out using FreeSurfer, developed at the Martinos Center for Biomedical Imaging, Charlestown, MA (http://surfer.nmr.mgh.harvard.edu/).

Analysis of individual functional MRI data by modeling with a hemodynamic response function (HRF). For each participant, the acquired native functional volumes were first corrected for potential motion using the AFNI algorithm (Cox, 1996). Next, the functional volumes were spatially smoothed using a 3-D Gaussian filter with a full-width half-max (FWHM) of 6 mm. Global intensity variations across runs and participants were removed by rescaling all voxels and time points of each run such that the mean in-brain intensity was fixed at an arbitrary value of 1000.

The functional images were then analyzed with a General Linear Model (GLM) using the FreeSurfer Functional Analysis Stream (FS-FAST). The HRF for each condition was modeled using two components, each constituting a canonical HRF (Friston et al., 1998), convolved with a box car of an appropriate length. The first component was modeled as a single regressor and lasted for approximately the first two-thirds (9.33 s) of the trial. It corresponded to the first two sentences and the beginning of the third sentence up to the onset of the critical word that did not differ between experimental conditions. The second component lasted for the final one-third of the trial (4.66 s) – from the onset of the critical word until the onset of the next trial – and was modeled separately for each condition. In addition, mean offset and linear trend regressors were included to remove low-frequency drift.

Construction of group cortical statistical maps. The GLM parameter estimates and residual error variances of each participant's functional data were resampled onto his or her inflated cortical surface and then onto the average cortical spherical representation (see anatomical reconstruction above). Each participant's data was then smoothed on the surface tessellation using an iterative nearest-neighbor averaging procedure equivalent to applying a two-dimensional Gaussian smoothing kernel with an FWHM of approximately 8.5 mm. Because this smoothing procedure was restricted to the cortical surface, averaging data across sulci or outside gray matter was avoided.

Three contrasts of interest – intermediately related *vs*. unrelated, intermediately related *vs*. highly related, and unrelated *vs*. highly related – were generated using the regression weights of the second canonical HRF component. These contrasts were tested using a t statistic at each voxel on the spherical surface using a random effects model, thus generating three cortical statistical maps. To correct for multiple comparisons, we identified significant clusters of activated voxels on the basis of a Monte Carlo simulation (Doherty et al., 2004) using a cluster size threshold of 300 mm^2 and threshold for rejection of the null hypothesis at $p<0.05$. The accompanying Talairach coordinates reported in the tables correspond to the vertices within each cluster with the minimum local p value.

Cortical statistical maps were also generated (a) including only trials in which participants' ratings were consistent with the three scenario types (i.e. highly related and rated 1 *vs*. intermediately related and rated 2 *vs*. unrelated and rated 3) and (b) comparing participants' ratings, regardless of the *a priori* conditions (i.e. rating 1 *vs*. rating 2 *vs*. rating 3).

Analysis of functional MRI data by modeling with an FIR. The model described above assumes a particular shape to the hemodynamic response function and also assumes that inference generation began at the point of presentation of the critical word. In order to examine the hemodynamic time course to each scenario type more closely and to determine if these assumptions were justified, we also analyzed each individual's functional data using a finite impulse response model (FIR) which gives an estimate of the hemodynamic response average at each TR within a peristimulus window, without any assumption about the overall shape of the hemodynamic response (Burock et al., 1998; Dale, 1999).

We first examined the hemodynamic response of all voxels averaged within various anatomical ROIs. These regions were selected on the left and right hemispheres on the basis of the fMRI literature on discourse processing (see Introduction) and are indicated in Fig. 1. They were the anterior inferior prefrontal gyri (pars triangularis and orbitalis), middle frontal gyri, the lateral temporal cortex (superior temporal sulci and middle temporal gyri), angular gyri and, extending onto the medial surface, the superior prefrontal gyri (see Fig. 1). All these ROIs were delineated purely on anatomical grounds and were defined on each individual cortical surface in accordance with the MGH Center for Morphometric Analysis (CMA) parcellation system (Caviness et al., 1996; Rademacher et al., 1992), using automated cortical parcellation methods described by Fischl et al. (2002, 2004). The FIR parameter estimates were mapped onto each ROI and, for each scenario type at each post-stimulus delay, the mean percent signal change relative to a pre-stimulus baseline of the first three volume acquisitions (−4 to 0 s relative to sentence onset) was averaged across all vertices within the ROI, averaged across subjects and then plotted to generate hemodynamic time courses (see Figs. 5

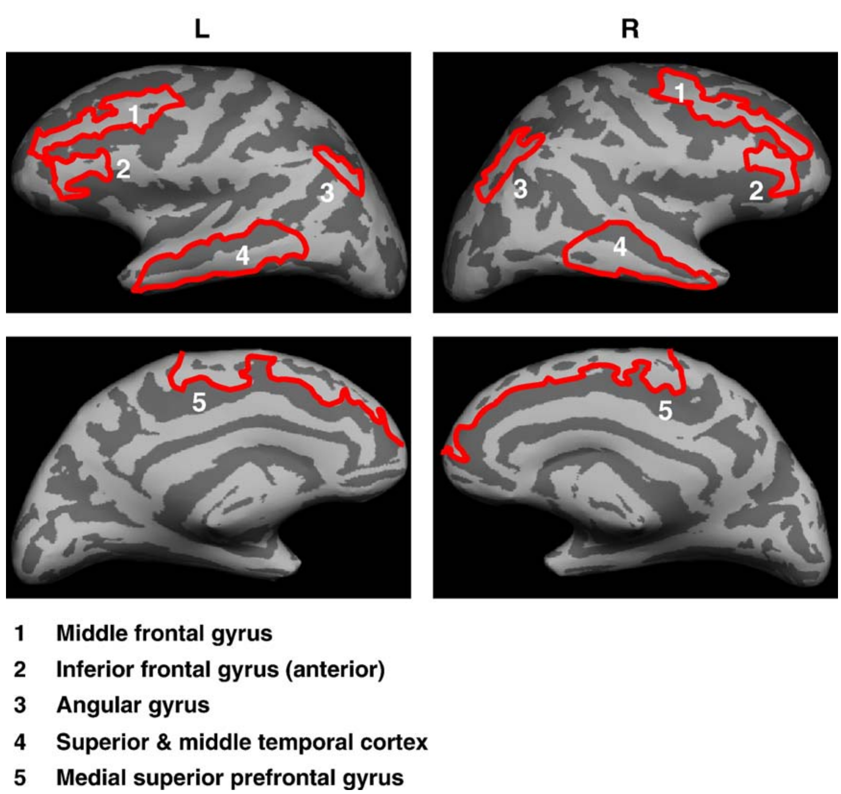

Fig. 1. For consistency and ease of visualization, anatomical regions of interest are indicated on the same average cortical surface (generated by averaging the sulcal and gyral folding patterns of all participants) as that used to display activation in Figs. 3 and 4. However, in the ROI analyses themselves, these ROIs were generated on each individual participant's reconstructed cortical surface (see Methods). Light gray: gyri. Dark gray: sulci.

and 6). Planned repeated-measures 2 (scenario type) × 2 (TR interval: 16, 18 s) ANOVAs were then conducted to examine how the percent BOLD signal change was modulated by each scenario type in relation to one another at the peak (16 s) and just after the peak (18 s) of the hemodynamic response.

The FIR model was also used to construct cortical statistical maps at the peak (16 s) and straight after the peak (18 s) of the hemodynamic response, serving as an independent check on findings revealed by the ROI analyses.

Results

Behavioral data

RTs to the three scenario types are shown in Fig. 2. An overall ANOVA revealed a significant main effect of scenario type, $F1(2,28)=20.2$, $p<0.00001$; $F2(2,717)=58$, $p<0.000001$. This arose because RTs were significantly longer to the intermediately related scenarios than to the unrelated scenarios, $t1(14)=4.86$, $p<0.00001$, $t2(478)=9.45$, $p<0.00001$, and highly related scenarios, $t1(14)=6.65$, $p<0.00001$, $t2(478)=8.94$, $p<0.00001$. There was no significant difference in RTs between the unrelated and highly related scenarios, $t1(14)=0.165$, $p=0.87$, $t2(478)=0.62$, $p=0.54$.

The percentages of trials given ratings 1, 2 and 3 to each scenario type are shown in Table 2. The pattern of RTs was the same (a) when we included only trials in which participants' ratings were consistent with the three scenario types (i.e. highly related and rated 1 *vs.* intermediately related and rated 2 *vs.* unrelated and rated 3) and (b) when we compared participants' ratings, regardless of their *a priori* conditions (i.e. rating 1 *vs.* rating 2 *vs.* rating 3).

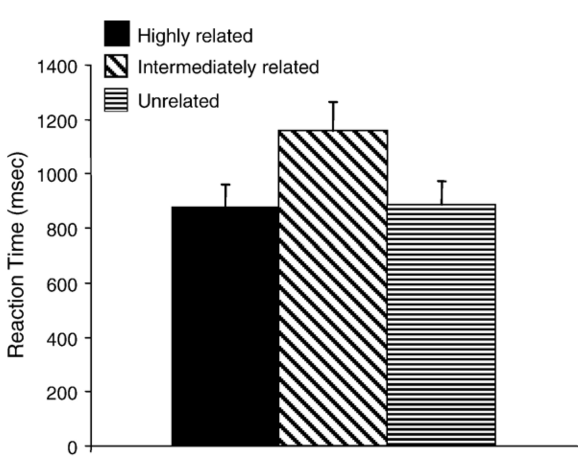

Fig. 2. Reaction times to the three types of discourse scenarios. Means and standard errors are shown.

Table 2

Ratings and scenario type (%)	Highly related	Intermediately related	Unrelated
Rating 1	85 (8.6)	49.8 (22.6)	2 (2.4)
Rating 2	10.3 (7.6)	37.4 (19.6)	13.4 (19.4)
Rating 3	2.3 (2.6)	9.1 (5.6)	82.4 (22.1)

Mean percentage of trials in each of the three *a priori* scenario types that were rated by participants, inside the scanner, as 1 (highly related), 2 (intermediately related) and 3 (unrelated).

fMRI data

Cortical statistical map analyses (modeled using the canonical HRF) using all responses

Intermediately related vs. highly related scenarios. There was more activity to the intermediately related scenarios than to the highly related scenarios within a widespread network on the lateral surfaces including bilateral anterior inferior prefrontal cortices, bilateral inferior parietal lobules, the left middle frontal gyrus and the left middle temporal gyrus and also within bilateral superior frontal gyri on the medial surface (yellow clusters in Fig. 3, Table 3A all responses).

Several other regions showed less activity to the intermediately related than to the highly related scenarios (blue clusters in Fig. 3, Table 3B all responses). These included sensory and motor/premotor cortices including the left occipital cortex and bilateral superior and inferior regions around the central sulcus.

Intermediately related vs. unrelated scenarios. With the exception of the right inferior parietal lobule, many of the same regions that showed more activity to the intermediately related scenarios than to the highly related scenarios (above) also showed more activity to the intermediately related scenarios than to the unrelated scenarios (yellow clusters in Fig. 4, Table 4A all responses; differences within the right inferior frontal gyrus reached voxel-level but not cluster-level significance). In addition, some of the same sensory and motor/premotor regions that showed less activity to the intermediately related than to the highly related scenarios (above) also showed less activity to intermediately related relative to unrelated scenarios (blue clusters in Fig. 4, Table 4B all responses).

Regions of overlap between the contrasts of (a) intermediately related vs. highly related scenarios (b) intermediately related vs. unrelated scenarios are illustrated in Fig. 5.

Unrelated vs. highly related scenarios. There were no regions that showed significantly more activity to unrelated than to highly related scenarios. Only three small clusters – the right superior precentral gyrus on the lateral surface (310 mm^2; tal: 34 −18 55, $p_{max}<0.004$) and bilateral motor cortices on the superior medial surface (left: 338 mm^2, tal: −16 0 70, $p_{max}<0.0008$; right: 322 mm^2, tal: 11 3 48, $p_{max}<0.0036$) – showed less activity to unrelated than to highly related scenarios.

Cortical statistical map analyses (modeled using canonical HRF) in which conditions were defined by subjects' ratings

As shown in the final two columns of Tables 3 and 4, many (although not all) of the regions that showed more or less activity to the intermediately related than to the highly related and unrelated scenarios were also modulated (a) in an analysis that included only trials in which participants' ratings were consistent with the three scenario types (i.e. highly related and rated 1 *vs.* intermediately related and rated 2 vs. unrelated and rated 3) and (b) in an analysis that compared participants' ratings, regardless of their *a priori* conditions (i.e. rating 1 *vs.* rating 2 *vs.* rating 3).

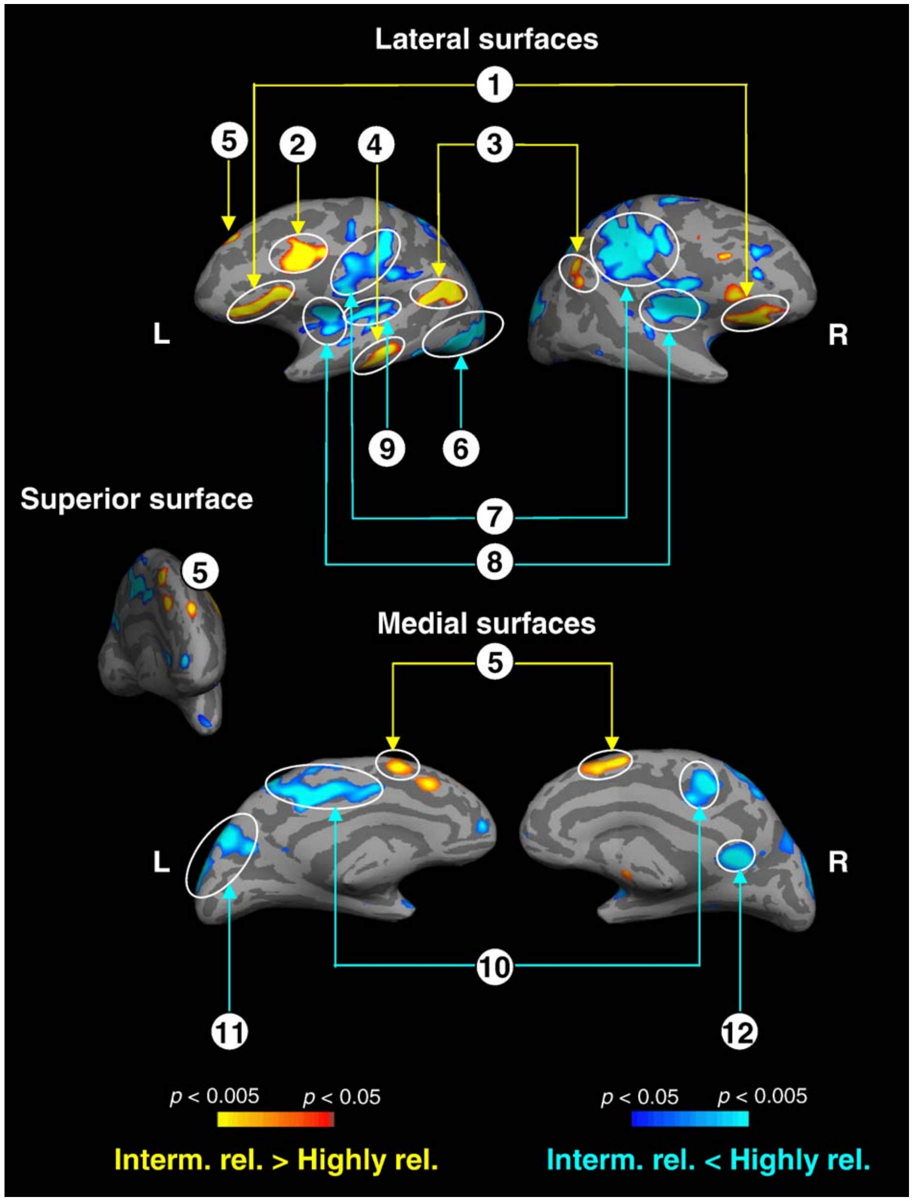

Fig. 3. Cortical statistical maps comparing intermediately related and highly related discourse scenarios (all responses). Yellow–red: more activity to intermediately related than to highly related scenarios. Blue: less activity to intermediately related than to highly related scenarios. All clusters circled are significant at a cluster-level $p<0.05$. Cluster numbers correspond directly to those regions reported in Table 3.

Region of interest analyses

The hemodynamic time courses in various regions of interest are shown in Figs. 5 and 6. The peak of the hemodynamic response fell at 16 s, and divergence between experimental conditions appeared to occur at and after the peak. This generally supported the assumptions made in modeling the data with the canonical HRF as described above. It suggested that inference generation occurred at or after the presentation of the final sentence rather than earlier in the discourse scenarios.

A series of 3 (Scenario Type) × 2 (Time: 16 s, 18 s) repeated-measures ANOVAs were carried out to examine the modulation of activity to the three scenario types at and just after the peak of the hemodynamic response in each ROI. As expected, in most of these regions (bilateral anterior inferior prefrontal gyri, bilateral medial superior prefrontal gyri, bilateral superior/middle temporal cortices and the left middle frontal gyrus), there were significant main effects of time (all Fs > 8.91, all ps < 0.001), reflecting the peak of the hemodynamic response at 16 s. Of most interest, these analyses revealed main effects of scenario type that reached significance in bilateral anterior inferior prefrontal gyri (left: $F(2,27)=13.35$, $p<0.0001$; right: $F(2,27)=5.02$, $p<0.014$), the left angular gyrus ($F(2,24)=5.35$, $p<0.015$), the left medial superior prefrontal gyrus ($F(2,28)=6.14$, $p<0.006$), the left middle frontal gyrus ($F(2,27)=7.23$, $p<0.003$), and that approached significance in the right middle frontal gyrus ($F(2,23)=2.7$, $p<0.09$) and in the left superior/middle temporal cortex ($F(1,20)=3.13$, $p<0.08$). All these main effects arose because of significant quadratic effects across the highly related, intermediately related and unrelated scenarios (in all these regions, Fs > 5.06, ps < 0.041); there were no linear increases in

Table 3
Intermediately related *vs.* highly related

All responses						Conditions defined by scenario type *and* by participants' ratings	Conditions defined by participants' ratings
Region	L/R	BA	Area mm^2	*p* value	Talairach (*x*, *y*, *z*)	Intermediately related and Rating 2 *vs.* Highly related and Rating 1	Rating 2 *vs.* Rating 1
A. Intermediately related > highly related							
1. Inferior prefrontal gyrus	L	47/(45)	1129	0.0005	−43, 38, −6	*	*
	R		966	0.0014	32, 19, 3	*	*
2. Middle frontal gyrus	L	8/(9)	1231	0.0003	−43, 12, 35	*	*
3. Inferior parietal lobule	L	40	993	0.00001	−45, −58, 37	*	*
(angular gyrus)	R	7	378	0.0018	38, −62, 49	^	^
4. Middle temporal gyrus	L	21	432	0.0003	−56, −32, −5	^	^
5. Superior and medial	L	6/(8, 9)	324	0.002	−5, 13, 47	*	*
frontal gyrus	R		516	0.001	6, 20, 51	*	*
B. Intermediately related < highly related							
6. Lateral occipital cortex (fusiform)	L	19	5929	0.00001	−38, −67, 7	*	^
7. Superior pericentral	L	4	15081	0.00056	−57, −8, 17	^	–
cortex	R		4868	0.0001	46, −16, 60	*	^
8. Insula and inferior	L	–/4	689	0.000001	−37, −3, 17	^	^
pericentral cortex	R		1525	0.00004	45, −4, 18	*	–
9. Insula/transverse temporal gyrus	L	–/41	1257	0.00023	−29, −28, 15	^	^
10. Paracentral lobule	L	5/7	1856	0.00058	−9, −39, 68	*	*
(medial)	R	5	781	0.00068	9, −24, 48	*	^
11. Medial occipital cortex (cuneus)	L	18/30	512	0.00059	−7, −72, 46	*	*
12. Retrosplenial cortex	R	29	452	0.0013	10, −47, 13	^	^

Talairach coordinates and approximate Brodmann area (BA) correspond to the local minimum *p* values for each cluster of activated vertices on the cortical surface; when a region appears to extend into another BA, this area is indicated in brackets. Cluster # corresponds directly to cluster labels in Fig. 3. All the regions shown for analyses with all responses reached cluster-level significance ($p<0.05$). In the analyses that included only trials in which participants' ratings were consistent with these scenario types and that compared participants' ratings, regardless of their a priori conditions, * indicates that some of the indicated region reached cluster-level significance ($p<0.05$) and ^ indicates that some of the indicated region reached voxel-level significance ($p<0.05$).

BOLD activity across these scenarios (in all regions, Fs<1.6, ps>0.23).

To follow up these main effects of scenario type, a series of planned 2×2 repeated-measures ANOVAs, comparing each scenario type to one another, were carried out, mirroring the statistical voxel-based cortical map contrasts described above (Table 5).

Intermediately related vs. highly related scenarios. The hemodynamic response at and immediately after the peak (16–18 s) was significantly greater to the intermediately related scenarios (blue dashed lines) than to the highly related scenarios (black solid lines) in bilateral anterior inferior prefrontal gyri, the left middle frontal gyrus and in bilateral medial prefrontal gyri. This difference also reached significance in the left angular gyrus and approached significance in the left superior/middle temporal cortex (see Figs. 5 and 6 and Table 5 first column). There were no scenario type by time interactions in any of these ROIs (all Fs<3.07, all ps>0.1).

Intermediately related vs. unrelated scenarios. ANOVAs revealed significantly greater hemodynamic responses to the intermediately related scenarios (blue dashed lines) than to the unrelated scenarios (red dotted lines) within the same set of ROIs as described above as well as in the right middle frontal gyrus (Figs. 5 and 6, Table 5 second column). Again, there were no scenario type by time interactions in any of these ROIs (all Fs<2.68, all ps>0.124).

Unrelated vs. highly related scenarios. ANOVAs examining this contrast did not show significant main effects of scenario type in any of the ROIs (all Fs<1.6, all ps>0.23). However, there were significant scenario type by time interactions within two right-sided regions: the right inferior prefrontal gyrus and the right lateral temporal cortex (Table 5, 3rd column). There were no scenario type by time interactions in any of the other ROIs (all Fs<2.53, all ps>0.14). Examination of these hemodynamic time courses within these two right-sided temporal and inferior prefrontal ROIs (shown in Fig. 7A right) suggested that these scenario type by time interactions reflected a transient increase at the peak of the hemodynamic response to the unrelated scenarios that, straight after the peak, dipped down such that activity was the same as to the highly related scenarios.

An FIR analysis at all voxels across the cortex confirmed this pattern: as shown in Fig. 7B, at the peak of the response at time point 16 s, unrelated scenarios were associated with more activity

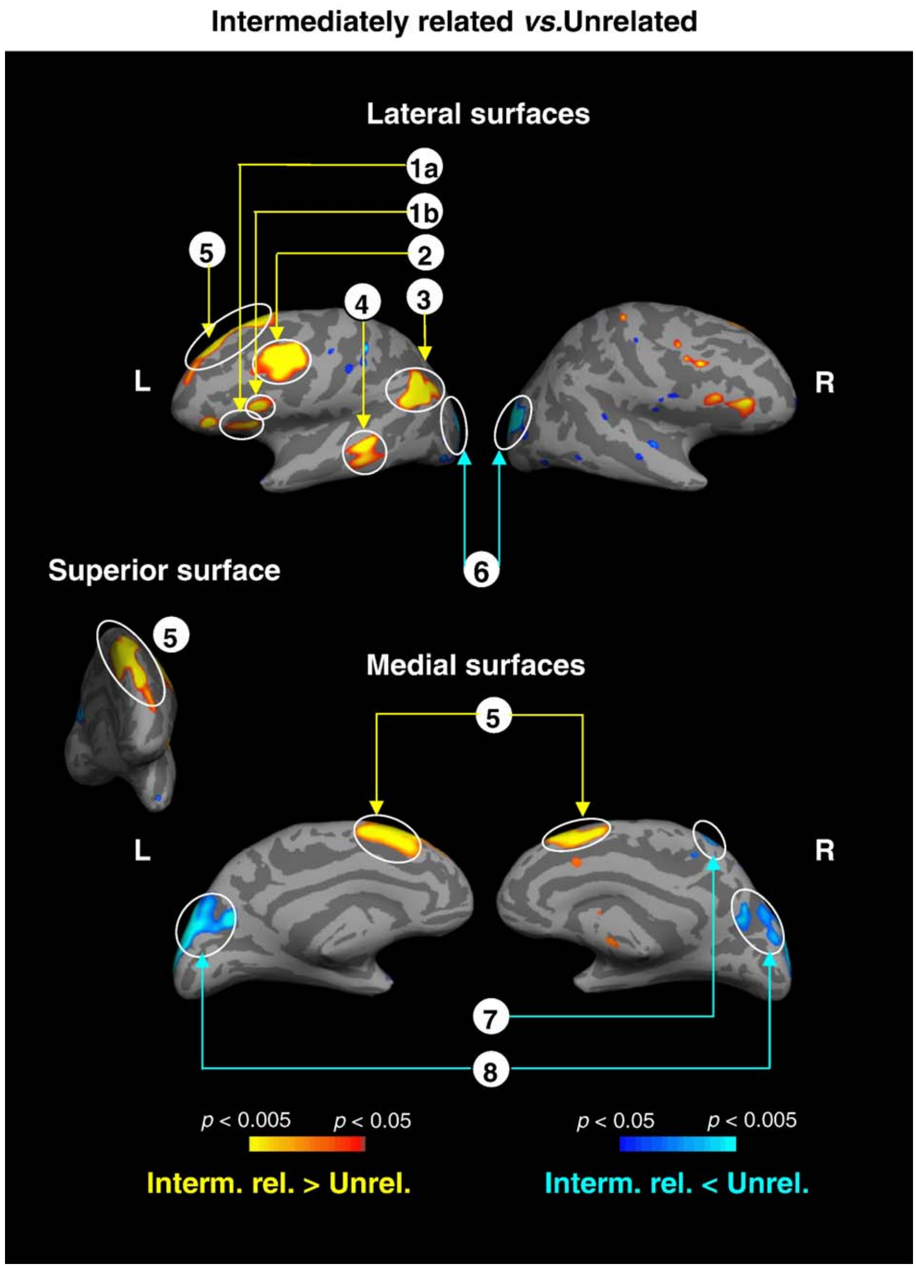

Fig. 4. Intermediately related vs. unrelated. Cortical statistical maps comparing intermediately related and unrelated discourse scenarios (all responses). Yellow–red: more activity to intermediately related than to unrelated scenarios. Blue: less activity to intermediately related than to unrelated scenarios. All clusters circled are significant at a cluster-level $p<0.05$. Cluster numbers correspond directly to those regions reported in Table 4.

than highly related scenarios in (1) the right inferior prefrontal cortex (387 mm^2, tal: 36 34 1, $p_{\max}<0.0003$) and (2) the right lateral temporal cortex (173 mm^2, tal: 46 −15 −12, $p_{\max}<0.003$), but there were no such increases in activity straight after the peak.

Summary

The cortical statistical map and the ROI analyses supported each other in demonstrating increased activity of a left lateral temporal/inferior parietal/prefrontal network to intermediately related scenarios relative to both highly related and unrelated scenarios. The right inferior prefrontal gyrus also showed more activity to the intermediately related scenarios relative to both the highly related scenarios (reaching significance on both ROI and cortical statistical map analyses) and to the unrelated scenarios (reaching significance on the ROI analysis and approaching significance on the cortical statistical map analyses). There were, however, a few discrepancies between the findings of the ROI and cortical statistical map analyses. For example, the right middle frontal gyrus showed significantly more activity to the intermediately related scenarios than to the unrelated scenarios in the ROI analysis but not in the cortical statistical map analysis, and the right inferior parietal lobule showed significantly more activity to the intermediately related scenarios than to the highly related scenarios in the cortical statistical map analysis, but not in the ROI analysis. These discrepancies probably reflect the

Table 4
Intermediately related *vs*. unrelated

All responses						Conditions defined by scenario type and by participants' ratings	Conditions defined by participants' ratings
Region	R/L	BA	Area mm^2	p value	Talairach (x, y, z)	Intermediately related and Rating 2 *vs*. Unrelated and Rating 3	Rating 2 *vs*. Rating 3
A. Intermediately related > unrelated							
1a. Inferior frontal gyrus (pars orbitalis)	L	47	316	0.0015	−27 20 6	*	*
1b. Inferior frontal gyrus (pars triangularis)	L	45	376	0.0014	−47 23 15	*	*
2. Middle frontal gyrus	L	6/(9)	1797	0.0001	−35 12 48	*	*
3. Inferior parietal lobule	L	40	1735	0.0002	−42 −53 37	*	*
4. Middle temporal gyrus	L	21	635	0.0012	−54 −31 4	^	*
5. Superior frontal gyrus	L	9	2907	0.0003	−14 48 30	*	*
	R	6/8 (9)	1095	0.0004	9 22 48	*	*
B. Intermediately related < unrelated							
6. Middle occipital gyrus	L	18	3668	0.00001	−22 −95 15	*	*
	R		2634	0.0001	33 −88 19	*	*
7. Paracentral lobule	R	4	2634	0.0001	7 −30 70	^	*
8. Precuneus	L	7	461	0.0004	−12 −58 27	*	*
	R		846	0.0032	5 −69 33	*	*

Legend: as in Table 3.

different sensitivities of these two methods in detecting activity at different spatial resolutions: a cortical statistical mapping approach has more power to detect changes in localized regions while an ROI approach has more power to detect modulation that occurs across many voxels within relatively large anatomical boundaries.

The contrast between the unrelated and highly related scenarios was particularly interesting. The cortical statistical map analysis that modeled the hemodynamic response to a known hemodynamic response function failed to demonstrate any regions that showed more activity to the unrelated than to the highly related scenarios. However, a more fine-grained FIR analysis, both within ROIs and at all voxels across the cortex, demonstrated a transient increase in activity within right inferior prefrontal and right temporal cortices at the peak of the hemodynamic response, but not after the peak.

In addition, the cortical statistical maps revealed less activity to intermediately related sentences than to both highly related and unrelated sentences within several sensory and motor/premotor regions.

Discussion

We have demonstrated increased RTs and increased temporal/inferior parietal/prefrontal hemodynamic activity in association with reading sentences that were intermediately related, relative to highly related or unrelated, to their preceding two-sentence contexts. Below, we first consider these behavioral and fMRI findings in relation to the cognitive processes engaged as participants comprehended and judged the coherence of these three-sentence scenarios. We then speculate on the possible roles of different regions within this widespread network in establishing coherence to the intermediately related scenarios. Finally, we suggest ways in which this neurocognitive model can be tested and refined in future research.

Participants took longer to make coherence judgments to the intermediately related scenarios than to the other two types of scenarios. The hemodynamic data showed the same non-linear pattern of response. There were no regions that showed incremental increases in activity across the three scenario types.[3] Moreover, the increased hemodynamic response to the intermediately related scenarios was maximal at a time point corresponding to the presentation of the final sentence and participants' coherence judgments.[4] We interpret these longer RTs and the increased hemodynamic response at the point of the third sentence as reflecting the generation of inferences as participants encountered the third sentences and attempted to causally link them to their preceding contexts. Supporting this interpretation, our pre-rating study established that the intermediately related scenarios, but not the other two types of scenarios, were associated with the generation of causal inferences. This interpretation is also in line with the conclusions of Keenan et al. (1984), Myers et al. (1987) and Myers and Duffy (1990) who explained the superior recall of intermediately related scenarios, relative to highly related and unrelated scenarios, as resulting from the generation of elaborative causal inferences.

Unlike the studies by Keenan et al. (1984), Myers et al. (1987) and Myers and Duffy (1990) in which participants simply

[3] This contrasts with the findings of Mason and Just (2004) who reported a linear increase in activity across these three types of scenario within the dorsolateral prefrontal cortex (DLPFC). The reason for this discrepancy between the two studies is unclear, particularly as Mason and Just (2004) also used a coherence judgment task and reported the same non-linear pattern of reaction times across the three scenario types as observed in the current study. Of note, however, the increase in DLPFC activity across the three scenario types in Mason and Just (2004) did not reach significance and should therefore be interpreted with caution.

[4] This does not rule out the possibility that, in some scenarios, participants may have also engaged in some predictive forward inferencing at the point of encountering the second sentence. Indeed, in some regions, the hemodynamic time courses did reveal some intriguing hints of a divergence in the BOLD response between scenario types before the onset of the final sentence (see Virtue et al., 2006 for similar observations). This will be followed up systematically in future studies by examining the relationship between this early hemodynamic divergence and measures of predictability at the point of the second sentence within each scenario.

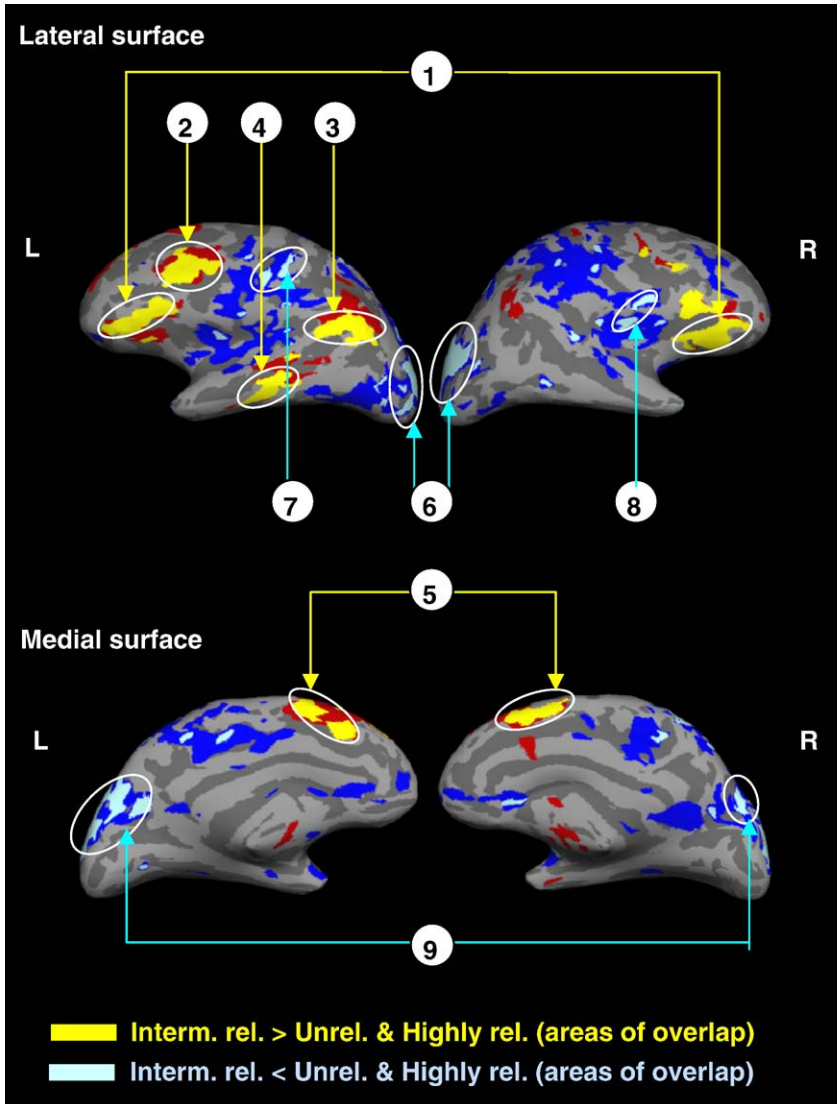

Fig. 5. Areas of overlap showing modulation of activity to intermediately related scenarios (relative to both other scenario types) Interm. rel. > Unrel. and Highly rel. (areas of overlap). Overlap statistical maps showing regions with *more* activity to intermediately related scenarios than to *both* unrelated *and* highly related scenarios (yellow) and regions that showed *less* activity to intermediately related than to both unrelated *and* highly related scenarios (pale blue). Red: regions that showed more activity to intermediately related scenarios than to unrelated scenarios (Fig. 3, red–yellow) *or* highly related scenarios (Fig. 4, red–yellow). Blue: regions that showed less activity to intermediately related scenarios than to unrelated scenarios (Fig. 3, blue) *or* highly related scenarios (Fig. 4, blue). All clusters circled are significant at a cluster-level $p < 0.05$. Cluster numbers correspond directly to those regions reported in Tables 3 and 4.

read text for comprehension, participants in the current study were required to make explicit coherence judgments about the relationship of the final sentence in each scenario to its preceding two-sentence context. This additional task requirement is particularly likely to have affected participants' processing of the causally unrelated scenarios that were not matched with the other two scenario types in terms of the semantic similarity values (SSVs) of their component content words, as determined by an LSA (Landauer and Dumais, 1997; Landauer et al., 1998). Semantically unrelated words may have provided an obvious strategic clue to participants that these scenarios were unrelated, leading to the relatively short RTs and the relatively transient BOLD response (in right temporal–prefrontal cortices) to these scenarios. In normal comprehension, where completely unrelated scenarios are less likely to be encountered, readers might make more of an attempt to seek coherence, even if this attempt is ultimately unsuccessful. Indeed, in the studies by Keenan et al. (1984), Myers et al. (1987) and Myers and Duffy (1990), even though the causally unrelated scenarios were associated with worse recognition at recall than the intermediately related scenarios, they were associated with the longest reading times.

Given these caveats in interpreting how the unrelated scenarios were processed in the current study, the increases in RT and BOLD activity to the intermediately related relative to the highly related scenarios are of particular interest. The intermediately related and highly related scenarios were closely matched in terms of the semantic relationships (SSVs) between their component content words (Landauer and Dumais, 1997; Landauer et al., 1998). Thus, these increases in RT and BOLD must have been driven by discourse-level inferential processes rather than by differences in

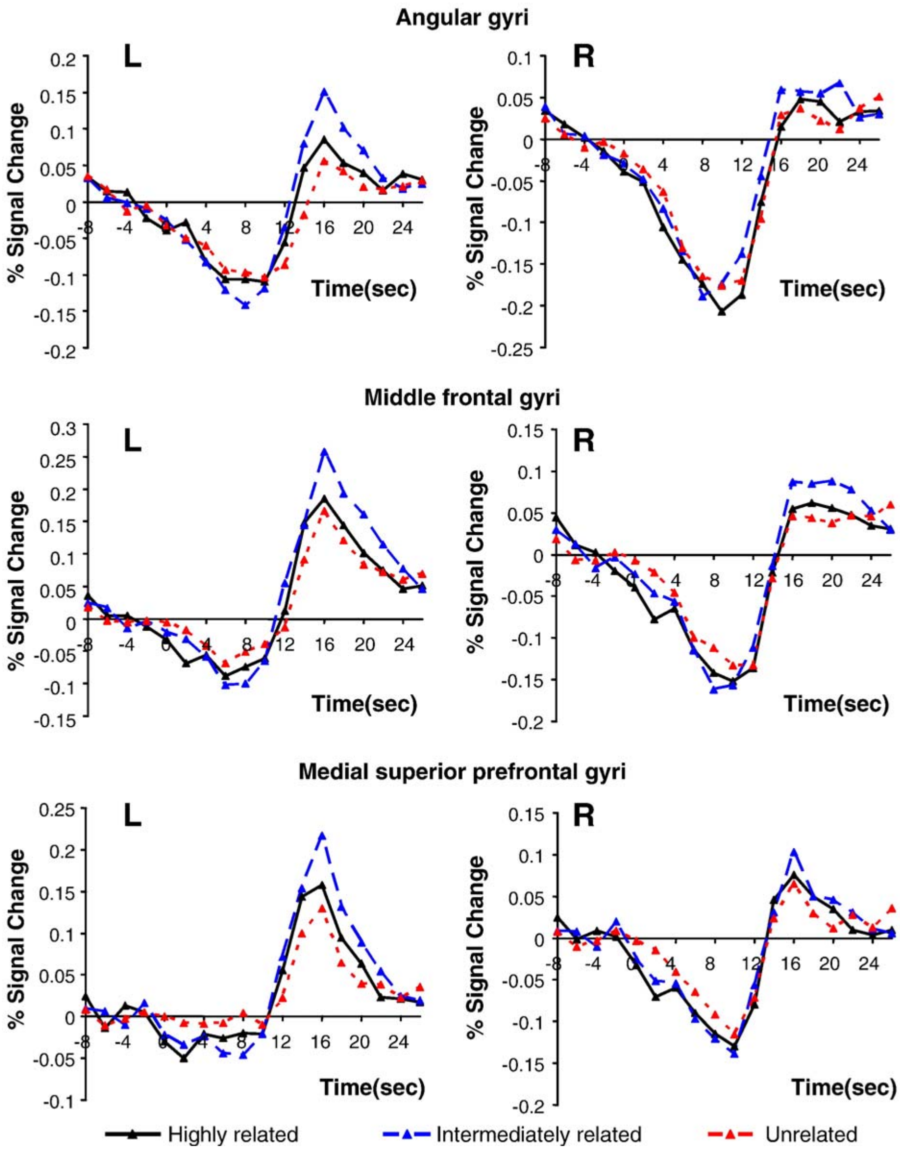

Fig. 6. Regions of interest: hemodynamic responses.

lexico-semantic associations between the content words of these two scenario types. Once again, however, we note that these increases in RTs and BOLD reflected not only normal comprehension processes, but also additional (although probably overlapping) neurocognitive processes involved in making the coherence judgments to the intermediately related scenarios. Moreover, as discussed further below, it is likely that the requirement to make the coherence judgment interacted with normal comprehension processes, leading to a relatively increased influence of top–down strategic processes on both RT and BOLD activity to these scenarios.

Neurocognitive networks mediating causal inferencing

Many different cognitive processes have been implicated as contributing to causal inferencing. These include the automatic activation of semantic information within long-term semantic memory, a more explicit retrieval process, the strategic search for coherence, the appropriate selection of relevant semantic information that has been activated or retrieved, its short-term retention within working memory, its integration into incoming discourse structure and its encoding into long-term memory (van den Broek, 1994). This study did not explicitly address the question of how each component of the widespread temporal/inferior parietal/prefrontal network that was activated to the intermediately related scenarios mediated each of these cognitive processes, but below we speculate on their possible functional roles, based on what we know from other neuroimaging and lesion studies.

There is strong evidence that inference generation can involve the automatic activation of semantic associations. This can occur either through priming by individual component words in a constrained discourse context (Kintsch and Mross, 1985; McKoon and Ratcliff, 1989; Till et al., 1988), or through the construction of higher-order discourse representations (or macropropositions) (Guindon and Kintsch, 1984). In the current investigation, such an implicit activation of semantic information may have been reflected by the increased hemodynamic activity within the temporal cortices that have been implicated in the storage of semantic information by

Table 5
ANOVAs in regions of interest

Contrast		Intermediately related ≥ Highly related			Intermediately related ≥ Unrelated			Unrelated ≥ Highly related		
Region	R/L	Effect and DOF	F	p	Effect and DOF	F	p	Effect and DOF	F	p
Inferior frontal gyrus (anterior)	L	Main effect scenario: 1,14	28.2	0.0001	Main effect scenario: 1,14	14.5	0.002	NS		
	R	Main effect scenario: 1,14	9.7	0.008	Main effect scenario: 1,14	5.7	0.032	Time × Scenario interaction 1,14	5.4	0.035
Middle frontal gyrus	L	Main effect scenario: 1,14	6.3	0.025	Main effect scenario: 1,14	16.7	0.001	NS		
	R	NS			Main effect scenario: 1,14	8.8	0.01	NS		
Superior and middle temporal cortex	L	Main effect scenario: 1,14	4.1	0.061	Main effect scenario: 1,14	6.9	0.02	NS		
	R	Time × Scenario interaction 1,14	12.6	0.0032	NS			Time × Scenario interaction 1,14	9.9	0.007
Angular gyrus	L	Main effect scenario: 1,14	7.4	0.017	Main effect Scenario: 1,14	11.3	0.005	NS		
	R	NS			NS			NS		
Medial superior prefrontal gyrus	L	Main effect scenario: 1,14	4.47	0.053	Main effect scenario: 1,14	13.54	0.002	NS		
	R	NS			NS			NS		

Results of a series of 2 (Scenario Type) × 2 (Time: 16, 18) ANOVAs in cortical regions of interest at and straight after the peak of the hemodynamic response. NS: non-significant, that is, main effects and interactions did not reach or approach significance (all $ps > 0.05$).

studies in non-human primates (Miyashita, 1993; Tanaka, 1997), lesion studies in humans (e.g. Alexander et al., 1989; de Renzi et al., 1987; Hodges et al., 1992) and neuroimaging studies (e.g. Chao et al., 1999; Martin and Chao, 2001). While the implicit activation of semantic concepts (such as during lexical decision tasks) is often associated with *decreases* in temporal and/or inferior prefrontal hemodynamic activity to subsequently presented semantically related targets in priming paradigms (Copland et al., 2003; Giesbrecht et al., 2004; Kotz et al., 2002; Matsumoto et al., 2005; Mummery et al., 1999; Rissman et al., 2003; Rossell et al., 2003; Wheatley et al., 2005), there is some evidence from masked repetition priming studies (Schnyer et al., 2002) and MEG studies (Dhond et al., 2001; Marinkovic et al., 2003) that *increases* in hemodynamic activity within the temporal cortex may reflect the early activation of semantic concepts.

A transient, automatic activation of semantic information within temporal cortices, however, does not, in its own right, constitute an inference. It is likely that at least some of the recruitment of lateral temporal cortices to the intermediately related scenarios in the current study reflected the activation of semantic information through controlled processes (Badre et al., 2005; Bokde et al., 2001; Wagner et al., 2001). Activity within the left inferior prefrontal cortex may have mediated such controlled retrieval (top–down activation) of semantic information (Wagner et al., 2001) and/or the selection of the most appropriate activated concepts (Fletcher et al., 2000; Moss et al., 2005; Thompson-Schill et al., 1997, 1999) for integration. Indeed, it is possible that different parts of the inferior prefrontal cortex may have played differential roles, with more anterior inferior prefrontal regions mediating controlled retrieval and more posterior regions mediating semantic selection (Badre et al., 2005).

As inferences are generated during normal comprehension, activated semantic information is thought to be held temporarily within working memory for integration into the incoming discourse structure. There are at least two sources of evidence for this. First, when activated semantic information contradicts the meaning of subsequently presented information, reading times are slowed (Klin, 1995, Experiment 3). Second, individuals with high verbal working memory spans perform better during text inferencing than individuals with low working memory spans (Singer and Ritchot, 1996). In the current study, working memory demands in association with processing the intermediately related scenarios may have been particularly high as activated semantic information had to be held online not only for integration with incoming discourse context, but also to perform the coherence judgment task. Both posterior inferior prefrontal cortices as well as dorsolateral prefrontal regions have been implicated in holding verbal semantic information within working memory (Barde and Thompson-Schill, 2002);[5] it is therefore possible that, in the current study, the recruitment of more posterior parts of the inferior prefrontal cortex (that included BA 45) and of the middle frontal gyrus (that included the posterior dorsolateral prefrontal cortex, BA 9) to the intermediately related scenarios was driven, in part, by the increased working memory load associated with establishing coherence to these scenarios. This would be consistent with recent fMRI findings demonstrating increased inferior prefrontal activity in association with inference generation specifically in individuals with high working memory spans (Virtue et al., 2006).

Finally, as the studies by Keenan et al. (1984), Myers et al. (1987) and Myers and Duffy (1990) clearly demonstrate, readers appear to actually encode elaborative inferences within long-term memory, leading to the superior recall of intermediately

[5] Some research suggests that ventral regions of the prefrontal cortex are involved in simply holding information online in working memory, while dorsolateral regions are involved in manipulating such information to perform a given cognitive task (D'Esposito et al., 1999), but the findings by Barde and Thompson-Schill (2002) fail to support this neuroanatomical distinction (see also, Veltman et al., 2003).

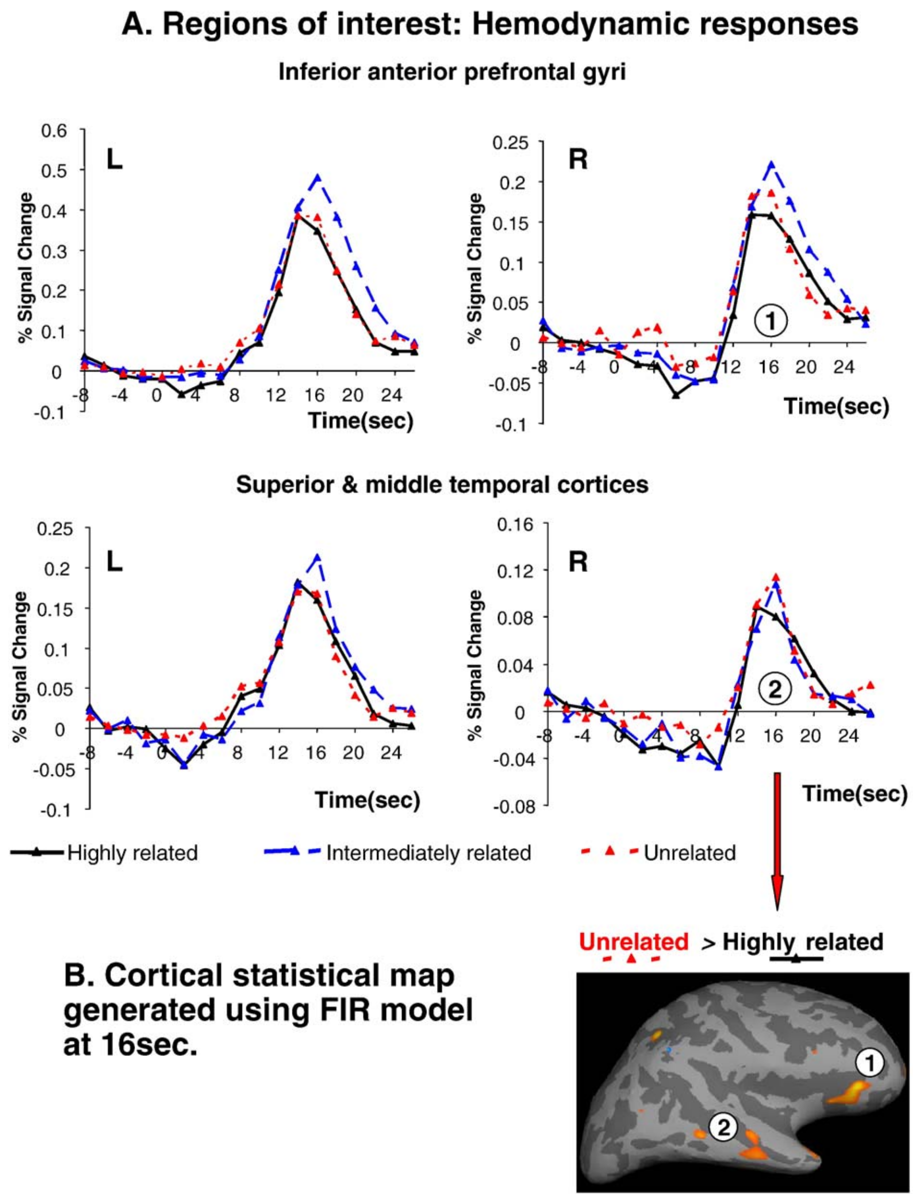

Fig. 7. (A) Regions of interest: Hemodynamic responses. (B) Cortical statistical map generated using FIR model at 16 s.

related scenarios (see also Klin, 1995, Experiment 3). Just as for semantic retrieval, semantic encoding is also thought to be mediated through interactions between temporal and inferior prefrontal cortices (Demb et al., 1995; Kirchhoff et al., 2000; Wagner et al., 1998).[6] Indeed, in single words studies, the degree of activity in both inferior prefrontal cortices and temporal cortices predicts the success of later recall (Wagner et al., 1998). Future studies will determine if this is also true of recall for intermediately related discourse scenarios that, as Keenan et al. (1984), Myers et al. (1987) and Myers and Duffy (1990) established, show superior recall in comparison with highly related and unrelated scenarios.

In addition to engaging the middle prefrontal gyrus (including the posterior dorsolateral prefrontal cortex, approximately BA 9), the intermediately related scenarios also engaged superior/medial prefrontal cortices relative to the other scenario types. Activation of superior and medial prefrontal cortices has been reported in association with establishing coherence in several studies of discourse processing (Ferstl and von Cramon, 2001; Ferstl and von Cramon, 2002; Xu et al., 2005). For example, Ferstl and von Cramon (2001, 2002) report activity within the medial prefrontal

[6] Within the inferior prefrontal cortex, some have claimed a functional segregation according to the type of information that is being encoded, with the left anterior inferior prefrontal regions encoding semantic information and more posterior regions encoding phonological information (Poldrack et al., 1999). Other studies, however, suggest that the anterior left inferior prefrontal gyrus mediates controlled processing regardless of stimulus domain (Gold and Buckner, 2002; Gold et al., 2005) or that, as discussed above, functional specialization within prefrontal regions is driven primarily by processing demands, with posterior inferior prefrontal regions being engaged in tasks with increased working memory requirements.

cortex (encompassing posterior to anterior regions: BA 6, 8, 9, 10 respectively) to coherent relative to unrelated scenarios. In the current study, activity within the superior medial prefrontal cortex was observed bilaterally and was slightly more posterior than that reported by Ferstl et al. (encompassing BAs 6, 8 and, particularly in the contrast between intermediately related and unrelated scenarios, also BA 9 more anteriorly on the left superior surface). Ferstl et al. have suggested that the activity within the superior medial prefrontal cortex reflects controlled, strategic, higher-order inferential processes—an idea that is consistent with its activation in association with making explicit coherence judgments about the intermediately related scenarios in the current study. A related idea is that activity within this overall region reflects participants' efforts to make sense of the sequential order of the events described in the intermediately related scenarios (Crozier et al., 1999; Ruby et al., 2002).[7]

Interestingly, in the studies by Ferstl and von Cramon (2001, 2002), activity within superior/medial prefrontal cortices was not accompanied by modulation of activity within temporal and inferior prefrontal cortices. Unlike the current study that explicitly included scenarios with a range of causal relationships, Ferstl and von Cramon (2001, 2002) only included coherent and unrelated scenarios.[8] One possibility is that the range of causal relationships in the current study encouraged more implicit inferencing than in the studies by Ferstl and von Cramon, leading to more engagement of temporal and inferior prefrontal cortices. On this account, the role of superior medial prefrontal regions in inferencing is distinct from that of temporal and lateral inferior prefrontal cortices that may mediate more implicit semantic retrieval and integration processes. Consistent with this idea, modulation of the superior medial prefrontal cortices has not been not reported in studies in which participants do not make explicit judgments but simply read text for comprehension (Robertson et al., 2000; Virtue et al., 2006). More generally, this view would be consistent with models of inferencing that suggest a neurocognitive distinction between a coherence-based processing mechanism that involves a strategic, backward elaborative search and a more implicit process of semantic activation (Gernsbacher, 1990, but see Golden and Rumelhart, 1993; St John, 1992).

Roles of the left and right hemispheres

As reviewed in the Introduction section, there has been extensive debate about the role of the right hemisphere in inferencing. Although the right hemisphere has been classically associated with higher-level language processes, both lesion studies (Zaidel et al., 2002) and visual field studies (Long and Baynes, 2002; Long et al., 2005) have implicated both hemispheres in these processes. The findings in the current study are consistent with this general view: activity in association with the intermediately related, relative to the highly related scenarios, was bilaterally distributed. Indeed, the most extensive activity was observed in the left hemisphere and, on the right, only inferior prefrontal and superior medial prefrontal regions showed more activity to the intermediately related than the highly related scenarios.

Given that the right hemisphere does not play an exclusive role in inferencing, this raises the related question of whether its role is distinct from that of the left hemisphere. For example, it has been suggested that the right hemisphere plays a specific role in the activation, retrieval and integration of *diffuse* semantic associations that lead to causal inference generation (Beeman et al., 1994; Faust and Lavidor, 2003). The current study does not directly address this hypothesis. There was, however, one piece of evidence suggesting a functional distinction between the right and left hemispheres in this paradigm: only right-sided inferior prefrontal and temporal cortices showed the short-lived increased activity at the peak of the hemodynamic response to the unrelated scenarios relative to the highly related scenarios. This observation should be interpreted with caution because, as discussed above, the unrelated and highly related scenarios were not matched in terms of the number of lexico-semantic associations between their individual component words. Moreover, transient increases in hemodynamic activity do not necessarily reflect the time course of underlying neural activity (Logothetis and Pfeuffer, 2004). Despite these caveats, however, this observation is interesting as it raises the possibility that the right inferior prefrontal gyrus may play an initial role in the detection of incoherence within discourse. On this account, inference generation may have been triggered by the detection of incoherence by the right hemisphere that, in the presence of semantic associations, led to more sustained right-sided activity and to the recruitment of the large semantic processing network in the left hemisphere.

Deactivation

Finally, it is worth considering the network of regions that showed reduced activity to intermediately related relative to unrelated and highly related scenarios. Some of these regions (e.g. the precuneus and retrosplenial cortex) constituted part of a resting-state 'default' network that is deactivated with respect to low-level resting conditions and whose activity is differentially modulated (shows less or more deactivation) across a variety of tasks and stimuli (Gusnard and Raichle, 2001). Other regions showing reduced activity to intermediately related scenarios, however, lay outside this resting state network and encompassed primary sensory, motor and premotor regions. Differential deactivation of the resting state network during attention-requiring, goal-directed tasks is thought to reflect a redirection of processing resources from default resting activity (Gusnard and Raichle, 2001). It is possible that, when tasks are particularly complex (such as making coherence judgments to intermediately related scenarios), involving multiple cognitive processes and engaging large, widespread brain networks (in

[7] In the current study, activity on the superior medial prefrontal surface to the intermediately related scenarios encompassed some of BA 9 and extended backwards to include BAs 6 and 8. It may be that a putative role of posterior medial superior prefrontal cortices in examining temporal relationship between events to make inferences during discourse comprehension is related to the more general roles of BAs 6 and 8 in action planning and sequencing of goal-directed activity (Rushworth et al., 2004). However, we cannot exclude the possibility that activation of posterior medial prefrontal regions (BAs 6 and 8) in the current study reflected participants' uncertainties in making coherence judgments to the intermediately related scenarios relative to the other types of scenarios (Ridderinkhof et al., 2004).

[8] The coherent scenarios used by Ferstl and von Cramon (2001, 2002) probably included many sentences that were intermediately related to their context, but these scenarios were not explicitly designed or counterbalanced such that, across subjects, the same final sentence was also preceded by a highly related context.

this case, lateral temporal/inferior parietal/prefrontal and medial prefrontal cortices), processing resources are diverted not only from the default resting network, but also from these sensory–motor regions. Future studies will explore whether reciprocal modulation within sensory–motor regions is also seen when participants perform complex tasks in domains other than language.

Conclusions and future directions

In summary, the current study documents the recruitment of a widespread temporal/inferior parietal/prefrontal cortical network in association with sentences that are intermediately causally related relative to highly related or unrelated to their preceding context. We suggest that this network mediates the generation and integration of causal inferences as subjects establish coherence during discourse processing. It seems unlikely that any particular area or hemisphere plays an exclusive role in inferencing, but rather that multiple regions act in consort. Building upon previous studies, we have hypothesized distinct roles for different components of this network in causal inferencing. We have suggested that the modulation of temporal cortices reflects the activation of stored semantic information, that inferior prefrontal regions may mediate the retrieval and/or selection of such information and that posterior inferior prefrontal regions and posterior dorsolateral prefrontal regions may mediate the maintenance and manipulation of such information within working memory as it is integrated into incoming structure and encoded in long-term memory, once again through inferior prefrontal and temporal interactions. Based on previous functional neuroimaging work (Crozier et al., 1999; Ferstl et al., 2002; Ferstl and von Cramon, 2001; Ferstl and von Cramon, 2002; Crozier et al., 1999; Ruby et al., 2002) and psycholinguistic models of inference processing (Gernsbacher, 1990; Sanford, 1990; van den Broek, 1994), we have also suggested that superior medial prefrontal regions may be involved in a directed search for meaning and possibly, more specifically, in examining temporal, sequential relationships between events to generate inferences. Finally, we have tentatively suggested that right temporal and inferior prefrontal regions may play a role in the initial detection of incoherence within intermediately related scenarios and that, in the presence of semantic associations, this triggers the subsequent recruitment of the rest of the network to generate and integrate causal inferences to achieve coherence.

These hypotheses are clearly speculative at this stage. One way of testing them is to manipulate the task performed by participants as well as other experimental parameters to determine whether different components of this network can be modulated independently of one another. As discussed above, some evidence that this may be the case comes from previous studies of discourse processing that have highlighted different components of this network, depending on the task, the type of relationship between sentences and the working memory span of individual participants. Another way to explore the functional role of different components of this network during inferencing is to examine their precise timing of activation. For example, if inferior prefrontal regions are involved in the top–down retrieval of semantic information during inferencing, this would predict early prefrontal activity that precedes activity within temporal cortices. A role of prefrontal cortices in selection, working memory, integration and encoding, however, would predict prefrontal activity that follows temporal activity. Finally, if, as we have speculated, right-sided temporal and inferior prefrontal cortices detect initial incoherence of intermediately related scenarios, this would predict that they are activated before the left hemisphere to these types of scenarios. fMRI, of course, is not suitable for examining the timing of activation across brain regions, but, in conjunction with techniques such as EEG and MEG (Dale et al., 2000), it may be possible to test such hypotheses directly.

Acknowledgments

This research was supported by NIMH (R01 MH071635) and the Institute for Mental Illness and Neuroscience Discovery (MIND). Gina R. Kuperberg was also supported by a NARSAD (with the Sidney Baer Trust) and a Claflin Distinguished Scholars Award from Massachusetts General Hospital. We thank Marianna Eddy, Jordana Cotton and Karin Blais for their help in data acquisition and Daya Gulabani for her help in analyzing the data. We are also grateful to Tatiana Sitnikova and Debra Long for their very insightful comments about the interpretation of these findings.

References

Alexander, M.P., Hiltbrunner, B., Fischer, R.S., 1989. Distributed anatomy of transcortical sensory aphasia. Arch. Neurol. 46, 885–892.

Badre, D., Poldrack, R.A., Pare-Blagoev, E.J., Insler, R.Z., Wagner, A.D., 2005. Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. Neuron 47, 907–918.

Barde, L.H., Thompson-Schill, S.L., 2002. Models of functional organization of the lateral prefrontal cortex in verbal working memory: evidence in favor of the process model. J. Cogn. Neurosci. 14, 1054–1063.

Beeman, M., 1993. Semantic processing in the right hemisphere may contribute to drawing inferences from discourse. Brain Lang. 44, 80–120.

Beeman, M., Friedman, R.B., Grafman, J., Perez, E., Diamond, S., Lindsay, M.B., 1994. Summation priming and coarse coding in the right hemisphere. J. Cogn. Neurosci. 6, 26–45.

Bokde, A.L., Tagamets, M.A., Friedman, R.B., Horwitz, B., 2001. Functional interactions of the inferior frontal cortex during the processing of words and word-like stimuli. Neuron 30, 609–617.

Brownell, H.H., Gardner, H., 1988. Neuropsychological insights into humour. In: Durant, J., Miller, J. (Eds.), Laughing Matters. Longman Scientific, Essex, pp. 17–34.

Brownell, H.H., Michel, D., Powelson, J.A., Gardner, H., 1983. Surprise but not coherence: sensitivity to verbal humor in right hemisphere patients. Brain Lang. 18, 20–27.

Brownell, H.H., Potter, H.H., Bihrle, A.M., Gardner, H., 1986. Inference deficits in right brain-damaged patients. Brain Lang. 29, 310–321.

Burock, M.A., Buckner, R.L., Woldorff, M.G., Rosen, B.R., Dale, A.M., 1998. Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. NeuroReport 9, 3735–3739.

Caplan, D., 1992. Language Structure, Processing and Disorders. MIT Press, Cambridge, MA.

Caviness, V.S., Meyer, J., Makris, N., Kennedy, D.N., 1996. MRI-based topographic parcellation of human neocortex: an anatomically specified method with estimate of reliability. J. Cogn. Neurosci. 8, 566–587.

Chao, L.L., Haxby, J.V., Martin, A., 1999. Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. Nat. Neurosci. 2, 913–919.

Copland, D.A., de Zubicaray, G.I., McMahon, K., Wilson, S.J., Eastburn, M., Chenery, H.J., 2003. Brain activity during automatic semantic priming revealed by event-related functional magnetic resonance imaging. NeuroImage 20, 302–310.

Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput. Biomed. Res. 29, 162–173.

Crozier, S., Sirigu, A., Lehericy, S., van de Moortele, P.F., Pillon, B., Grafman, J., Agid, Y., Dubois, B., LeBihan, D., 1999. Distinct prefrontal activations in processing sequence at the sentence and script level: an fMRI study. Neuropsychologia 37, 1469–1476.

Dale, A.M., 1999. Optimal experimental design for event-related fMRI. Hum. Brain Mapp. 8, 109–114.

Dale, A.M., Sereno, M.I., 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. J. Cogn. Neurosci. 5, 162–176.

Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis: I. Segmentation and surface reconstruction. NeuroImage 9, 179–194.

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.

de Renzi, E., Zambolin, A., Crisi, G., 1987. The pattern of neuropsychological impairment associated with left posterior cerebral artery infarcts. Brain 110, 1099–1116.

Demb, J.B., Desmond, J.E., Wagner, A.D., Vaidya, C.J., Glover, G.H., Gabrieli, J.D., 1995. Semantic encoding and retrieval in the left inferior prefrontal cortex: a functional MRI study of task difficulty and process specificity. J. Neurosci. 15, 5870–5878.

D'Esposito, M., Postle, B.R., Ballard, D., Lease, J., 1999. Maintenance versus manipulation of information held in working memory: an event-related fMRI study. Brain Cogn. 41, 66–86.

Dhond, R.P., Buckner, R.L., Dale, A.M., Marinkovic, K., Halgren, E., 2001. Spatiotemporal maps of brain activity underlying word generation and their modification during repetition priming. J. Neurosci. 21, 3564–3571.

Doherty, C.P., West, W.C., Dilley, L.C., Shattuck-Hufnagel, S., Caplan, D., 2004. Question/statement judgments: an fMRI study of intonation processing. Hum. Brain Mapp. 23, 85–98.

Faust, M., Lavidor, M., 2003. Semantically convergent and semantically divergent priming in the cerebral hemispheres: lexical decision and semantic judgment. Cogn. Brain Res. 17, 585–597.

Faust, M., Bar-lev, A., Chiarello, C., 2003. Sentence priming effects in the two cerebral hemispheres: influences of lexical relatedness, word order, and sentence anomaly. Neuropsychologia 41, 480–492.

Ferstl, E.C., von Cramon, D.Y., 2001. The role of coherence and cohesion in text comprehension: an event-related fMRI study. Cogn. Brain Res. 11, 325–340.

Ferstl, E.C., von Cramon, D.Y., 2002. What does the frontomedian cortex contribute to language processing: coherence or theory of mind? NeuroImage 17, 1599–1612.

Ferstl, E.C., Guthke, T., von Cramon, D.Y., 2002. Text comprehension after brain injury: left prefrontal lesions affect inference processes. Neuropsychology 16, 292–308.

Fischl, B., Sereno, M.I., Dale, A.M., 1999a. Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. NeuroImage 9, 195–207.

Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M., 1999b. High-resolution intersubject averaging and a coordinate system for the cortical surface. Hum. Brain Mapp. 8, 272–284.

Fischl, B., Liu, A., Dale, A.M., 2001. Automated manifold surgery: constructing geometrically accurate and topologically correct models of the human cerebral cortex. IEEE Trans. Med. Imag. 20 (1), 70–80.

Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, R., Dale, A.M., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron 33, 341–355.

Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Segonne, F., Salat, D.H., Busa, E., Seidman, L.J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., Dale, A.M., 2004. Automatically parcellating the human cerebral cortex. Cereb. Cortex 14, 11–22.

Fletcher, P.C., Happe, F., Frith, U., Baker, S.C., Dolan, R.J., Frackowiak, R. S., Frith, C.D., 1995. Other minds in the brain: a functional imaging study of "theory of mind" in story comprehension. Cognition 57, 109–128.

Fletcher, P.C., Shallice, T., Dolan, R.J., 2000. "Sculpting the response space"—An account of left prefrontal activation at encoding. NeuroImage 12, 404–417.

Friston, K.J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M.D., Turner, R., 1998. Event-related fMRI: characterizing differential responses. NeuroImage 7, 30–40.

Gallagher, H.L., Happe, F., Brunswick, N., Fletcher, P.C., Frith, U., Frith, C. D., 2000. Reading the mind in cartoons and stories: an fMRI study of 'theory of mind' in verbal and nonverbal tasks. Neuropsychologia 38, 11–21.

Gernsbacher, M.A., 1990. Language Comprehension as Structure Building. Erlbaum, Hillsdale, NJ.

Giesbrecht, B., Camblin, C.C., Swaab, T.Y., 2004. Separable effects of semantic priming and imageability on word processing in human cortex. Cereb. Cortex 14, 521–529.

Gold, B.T., Buckner, R.L., 2002. Common prefrontal regions coactivate with dissociable posterior regions during controlled semantic and phonological tasks. Neuron 35, 803–812.

Gold, B.T., Balota, D.A., Kirchhoff, B.A., Buckner, R.L., 2005. Common and dissociable activation patterns associated with controlled semantic and phonological processing: evidence from FMRI adaptation. Cereb. Cortex 15, 1438–1450.

Golden, R.M., Rumelhart, D.E., 1993. A parallel distributed processing model of story comprehension and recall. Discourse Process. 16, 203–237.

Graesser, A.C., Singer, M., Trabasso, T., 1994. Constructing inferences during narrative text comprehension. Psychol. Rev. 101, 371–395.

Guindon, R., Kintsch, W., 1984. Priming macropropositions: evidence for the primacy of macropropositions in the memory for text. J. Verbal Learn. Verbal Behav. 23, 508–518.

Gusnard, D.A., Raichle, M.E., 2001. Searching for a baseline: functional imaging and the resting human brain. Nat. Rev., Neurosci. 2, 685–694.

Hodges, J.R., Patterson, K., Oxbury, S., Funnell, E., 1992. Semantic dementia. Brain 115, 1783–1806.

Joanette, Y., Goulet, P., Hannequin, D., 1990. Right Hemisphere and Verbal Communication. Springer-Verlag, New York.

Keenan, J.M., Baillet, S.D., Brown, P., 1984. The effect of causal cohesion on comprehension and memory. J. Verbal Learn. Verbal Behav. 23, 115–126.

Kintsch, W., 1988. The role of knowledge in discourse comprehension: a construction–integration model. Psychol. Rev. 95, 163–182.

Kintsch, W., Mross, E.F., 1985. Context effects in word identification. J. Mem. Lang. 24, 336–349.

Kirchhoff, B.A., Wagner, A.D., Maril, A., Stern, C.E., 2000. Prefrontal–temporal circuitry for episodic encoding and subsequent memory. J. Neurosci. 20, 6173–6180.

Klin, C.M., 1995. Causal inferences in reading: from immediate activation to long-term memory. J. Exp. Psychol., Learn. Mem. Cogn. 21, 1483–1494.

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.

Landauer, T.K., Dumais, S.T., 1997. A solution to Plato's problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. Psychol. Rev. 104, 211–240.

Landauer, T.K., Foltz, P.W., Dumais, S.T., 1998. Introduction to latent semantic analysis. Discourse Process. 25, 259–284.

Logothetis, N.K., Pfeuffer, J., 2004. On the nature of the BOLD fMRI contrast mechanism. Magn. Reson. Imaging 22, 1517–1531.

Long, D.L., Baynes, K., 2002. Discourse representation in the two cerebral hemispheres. J. Cogn. Neurosci. 14, 228–242.

Long, D.L., Baynes, K., Prat, C.S., 2005. The propositional structure of discourse in the two cerebral hemispheres. Brain Lang. 95, 383–394.

Marinkovic, K., Dhond, R.P., Dale, A.M., Glessner, M., Carr, V., Halgren, E., 2003. Spatiotemporal dynamics of modality-specific and supramodal word processing. Neuron 38, 487–497.

Martin, A., Chao, L.L., 2001. Semantic memory and the brain: structure and processes. Curr. Opin. Neurobiol. 11, 194–201.

Mason, R.A., Just, M.A., 2004. How the brain processes causal inferences in text. Psychol. Sci. 15, 1–7.

Matsumoto, A., Iidaka, T., Haneda, K., Okada, T., Sadato, N., 2005. Linking semantic priming effect in functional MRI and event-related potentials. NeuroImage 24, 624–634.

Mazoyer, B.M., Tzourio, N., Frak, V., Syrota, A., Murayama, O., Levrier, O., Salamon, G., Dehaene, S., Cohen, L., Mehler, J., 1993. The cortical representation of speech. J. Cogn. Neurosci. 5, 467–479.

McKoon, G., Ratcliff, R., 1989. Semantic associations and elaborative inference. J. Exp. Psychol., Learn. Mem. Cogn. 15, 326–338.

McKoon, G., Ratcliff, R., 1992. Inference during reading. Psychol. Rev. 99, 440–466.

Miyashita, Y., 1993. Inferior temporal cortex: where visual perception meets memory. Annu. Rev. Neurosci. 16, 245–263.

Moss, H.E., Abdallah, S., Fletcher, P., Bright, P., Pilgrim, L., Acres, K., Tyler, L.K., 2005. Selecting among competing alternatives: selection and retrieval in the left inferior frontal gyrus. Cereb. Cortex. 15, 1723–1735.

Mummery, C.J., Shallice, T., Price, C.J., 1999. Dual-process model in semantic priming: a functional imaging perspective. NeuroImage 9, 516–525.

Myers, J.K., Duffy, S.A., 1990. Causal inferences and text memory. In: Graesser, A.C., Bower, G.H. (Eds.), The Psychology of Learning and Motivation: Inferences and Text Comprehension. Academic Press, San Diego, pp. 159–173.

Myers, J.L., Shinjo, M., Duffy, S.A., 1987. Degree of causal relatedness and memory. J. Mem. Lang. 26, 453–465.

Oldfield, R.C., 1971. The assessment and analysis of handedness: the Edinburgh inventory. Neuropsychologia 9, 97–113.

Poldrack, R.A., Wagner, A.D., Prull, M.W., Desmond, J.E., Glover, G.H., Gabrieli, J.D., 1999. Functional specialization for semantic and phonological processing in the left inferior prefrontal cortex. NeuroImage 10, 15–35.

Rademacher, J., Galaburda, A.M., Kennedy, D.N., Filipek, P.A., Caviness, V.S., 1992. Human cerebral cortex localization, parcellation, and morphometry with magnetic resonance imaging. J. Cogn. Neurosci. 4, 352–373.

Ridderinkhof, K.R., Ullsperger, M., Crone, E.A., Nieuwenhuis, S., 2004. The role of the medial frontal cortex in cognitive control. Science 306, 443–447.

Rissman, J., Eliassen, J.C., Blumstein, S.E., 2003. An event-related FMRI investigation of implicit semantic priming. J. Cogn. Neurosci. 15, 1160–1175.

Robertson, D.A., Gernsbacher, M.A., Guidotti, S.J., Robertson, R.R., Irwin, W., Mock, B.J., Campana, M.E., 2000. Functional neuroanatomy of the cognitive process of mapping during discourse comprehension. Psychol. Sci. 11, 255–260.

Rossell, S.L., Price, C.J., Nobre, A.C., 2003. The anatomy and time course of semantic priming investigated by fMRI and ERPs. Neuropsychologia 41, 550–564.

Ruby, P., Sirigu, A., Decety, J., 2002. Distinct areas in parietal cortex involved in long-term and short-term action planning: a PET investigation. Cortex 38, 321–339.

Rushworth, M.F., Walton, M.E., Kennerley, S.W., Bannerman, D.M., 2004. Action sets and decisions in the medial frontal cortex. Trends Cogn. Sci. 8, 410–417.

Sanford, A., 1990. On the nature of text-driven inference. In: Balota, D.A., d'Arcais, F., Rayner, K. (Eds.), Comprehension Processes in Reading. Erlbaum, Hillsdale, NJ.

Schnyer, D.M., Ryan, L., Trouard, T., Forster, K., 2002. Masked word repetition results in increased fMRI signal: a framework for understanding signal changes in priming. NeuroReport 13, 281–284.

Shammi, P., Stuss, D.T., 1999. Humour appreciation: a role of the right frontal lobe. Brain 122 (Pt. 4), 657–666.

Singer, M., 1994. Discourse inference processes. In: Gernsbacher, M.A. (Ed.), Handbook of Psycholinguistics. Academic Press, San Diego, pp. 479–510.

Singer, M., Ritchot, K.F., 1996. The role of working memory capacity and knowledge access in text inference processing. Mem. Cogn. 24, 733–743.

St George, M., Kutas, M., Martinez, A., Sereno, M.I., 1999. Semantic integration in reading: engagement of the right hemisphere during discourse processing. Brain 122, 1317–1325.

St John, M.F., 1992. The story gestalt: a model of knowledge intensive processes in text comprehension. Cogn. Sci. 16, 271–306.

Tanaka, K., 1997. Mechanisms of visual object recognition: monkey and human studies. Curr. Opin. Neurobiol. 7, 523–529.

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. Proc. Natl. Acad. Sci. U. S. A. 94, 14792–14797.

Thompson-Schill, S.L., D'Esposito, M., Kan, I.P., 1999. Effects of repetition and competition on activity in left prefrontal cortex during word generation. Neuron 23, 513–522.

Till, R.E., Mross, E.F., Kintsch, W., 1988. Time course of priming for associate and inference words in a discourse context. Mem. Cogn. 16, 283–298.

van den Broek, P., 1994. Comprehension and memory of narrative texts: inferences and coherence. In: Gernsbacher, M.A. (Ed.), Handbook of Psycholinguistics. Academic Press, San Diego, pp. 539–583.

Van Dijk, T.A., Kintsch, W., 1983. Strategies of Discourse Comprehension. Academic Press, New York.

Veltman, D.J., Rombouts, S.A., Dolan, R.J., 2003. Maintenance versus manipulation in verbal working memory revisited: an fMRI study. NeuroImage 18, 247–256.

Virtue, S., Haberman, J., Clancy, Z., Parrish, T., Jung Beeman, M., 2006. Neural activity of inferences during story comprehension. Brain Res.

Wagner, A.D., Schacter, D.L., Rotte, M., Koutstaal, W., Maril, A., Dale, A., Rosen, B., Buckner, R.L., 1998. Building memories: remembering and forgetting verbal experiences as predicted by brain activity. Science 281, 1188–1191.

Wagner, A.D., Pare-Blagoev, E.J., Clark, J., Poldrack, R.A., 2001. Recovering meaning: left prefrontal cortex guides controlled semantic retrieval. Neuron 31, 329–338.

Wheatley, T., Weisberg, J., Beauchamp, M.S., Martin, A., 2005. Automatic priming of semantically related words reduces activity in the fusiform gyrus. J. Cogn. Neurosci. 17, 1871–1885.

White, K., Ashton, R., 1976. Handedness assessment inventory. Neuropsychologia 14, 261–264.

Winner, E., Gardner, H., 1977. The comprehension of metaphor in brain damaged patients. Brain 100, 719–727.

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.

Zaidel, E., Kasher, A., Soroker, N., Batori, G., 2002. Effects of right and left hemisphere damage on performance of the "Right Hemisphere Communication Battery". Brain Lang. 80, 510–535.