*Research Article*

# How the Brain Processes Causal Inferences in Text

## A Theoretical Account of Generation and Integration Component Processes Utilizing Both Cerebral Hemispheres

**Robert A. Mason and Marcel Adam Just**

*Center for Cognitive Brain Imaging, Carnegie Mellon University*

ABSTRACT—*Theoretical models of text processing, such as the construction-integration framework, pose fundamental questions about causal inference making that are not easily addressed by behavioral studies. In particular, a common result is that causal relatedness has a different effect on text reading times than on memory for the text: Whereas reading times increase linearly as causal relatedness decreases, memory for the text is best for events that are related by a moderate degree of causal relatedness and is poorer for events with low and high relatedness. Our functional magnetic resonance imaging study of the processing of two-sentence passages that varied in their degree of causal relatedness suggests that the inference process can be analyzed into two components, generation and integration, that are subserved by two large-scale cortical networks (a reasoning system in dorsolateral prefrontal cortex and the right-hemisphere language areas). These two cortical networks, which are distinguishable from the classical left-hemisphere language areas, approximately correspond to the two functional relations observed in the behavioral results.*

In order to fully understand a narrative text, readers must be able to mentally link together successive events to form a coherent representation of the story. Often, the events in the story are not explicitly related to each other, and the reader must connect them by generating and integrating inferences. Through the history of discourse-processing research, cognitive psychologists have struggled with the questions of under what circumstances and how such connecting causal inferences are generated. With the advent of functional magnetic resonance imaging (fMRI) techniques, it is now possible to integrate cognitive behavioral findings with brain-imaging research to inform answers to these questions by investigating the neural bases of the component processes involved in the generation and integration of inferences.

Address correspondence to Robert Mason, Center for Cognitive Brain Imaging, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213; e-mail: rmason@andrew.cmu.edu.

The inference process requires that a reader first generate a possible inference and then attempt to integrate it into the internal representation of the text. The construction-integration (CI) model of text comprehension (Kintsch, 1988) is consistent with this general description of inferencing. According to the CI model, an initial process in which the reader will liberally generate many possible inferences is followed by a second process of integrating those inferences that have a high degree of connection with the text base into the representation of the text. A successful integration of an inference will then result in a text representation that involves both the specific propositions contained in the text and those inferred propositions that were generated by the reader to connect information in the text.

Although various types of inferences have been studied and categorized (e.g., Singer, 1994; van den Broek, 1994), our study focused on causal inferences. Keenan, Baillet, and Brown (1984) and Myers, Shinjo, and Duffy (1987) investigated the relationship between memory for text and the degree of causal relatedness between the sentences within the text. They created sentence pairs that varied across four levels of intersentence causal relatedness. An "outcome" sentence, such as *The next day his body was covered with bruises*, was preceded by one of four different types of sentences describing antecedent conditions:

| | |
|---|---|
| highly related: | *Joey's big brother punched him again and again.* |
| moderately related 1: | *Racing down the hill, Joey fell off his bike.* |
| moderately related 2: | *Joey's crazy mother became furiously angry with him.* |
| distantly related: | *Joey went to a neighbor's house to play.* |

Myers et al. (1987) and Keenan et al. (1984) found that reading times increased as the relatedness of the sentences decreased. Somewhat counterintuitively, memory for the two-sentence passages (measured by various recall and recognition tests) followed an inverted-U-shaped function: The moderately related sentences were recalled (and recognized) better than the highly related or distantly related pairs. This is counterintuitive because the pattern is not a monotonic function of either reading time or causal relatedness.

 Copyright © 2004 American Psychological Society 

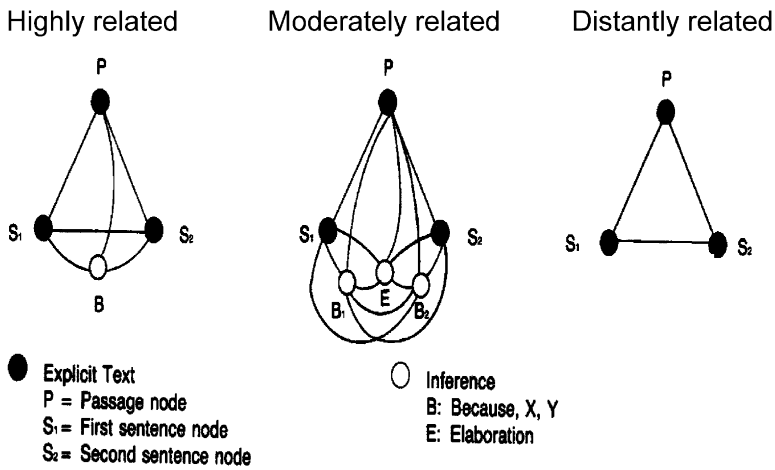


**Fig. 1.** Possible representation of the reader's internal network corresponding to each of the three types of causally related sentence pairs (reprinted from Myers & Duffy, 1990, with permission).

Myers et al. (1987) and Myers and Duffy (1990) suggested that the increased recall for the intermediate causally related sentences was a result of the readers having generated a causal link between the two sentences. The generation of a causal link is presumed to produce a larger and perhaps richer interconnected network of nodes in the reader's text-base representation, as shown in Figure 1. This larger, richer network can provide additional retrieval cues for recall, thus resulting in higher recall.

We hypothesized that the reading of the moderately related sentence pairs was accompanied by both the generation and the integration of causal inferences. In the highly related sentences, this inferencing process was unnecessary, resulting in faster reading times and more sparse text representations. In contrast, the distantly related sentences had slower reading times, presumably as a result of a liberal generation of possible inferences to connect them, but lower recall, as a result of lack of success in integrating any inference.

Using functional brain imaging, we attempted to find specific areas in the brain that respond differently to the two component processes of inference making in discourse comprehension. Studies have shown that fMRI is an excellent measure of the intensity of cognitive processing, specifically during reading (Just, Carpenter, Keller, Eddy, & Thulborn, 1996); this is true even when additional processing is not accompanied by an increase in reading times (Mason, Just, Keller, & Carpenter, in press). Consider the component process of generating possible inferences. If the linear increase in reading times with increasing causal distance is due to the generation of possible inferences, then there may exist a set of brain areas that shows a similar linear increase in brain activation. Similarly, if the inverted-U-shaped function relating recall to degree of causal relatedness is due to integration processes, there may exist a set of cortical areas involved in inference integration that shows a similar inverted-U function. Finally, it is quite likely that a set of brain areas involved with the basic levels of sentence processing (e.g., lexical access, syntactic parsing) will show equivalent activation across the levels of causal relatedness.

It has been proposed that the right-hemisphere homologues of the left-hemisphere language areas (superior, middle, and inferior temporal gyri; inferior frontal gyrus, including pars opercularis and pars triangularis; and the inferior parietal area) are extensively utilized in discourse processing, particularly inference generation (for an extensive review, see Beeman, 1998). Evidence for the role of the right hemisphere in discourse processing has come from two lines of research. The first is neuropsychological investigations of patients with lesions who have trouble with aspects of discourse processing. Patients with lesions to the right hemisphere generally have trouble drawing inferences in order to integrate sentences and maintain coherence (Beeman, 1993; Brownell, Potter, Bihrle, & Gardner, 1986); they do not mistakenly recall inferences, presumably because they never generated them (Grafman, Salazar, Vance, Weingartner, & Amin, 1987); and they make elaborative inferences more easily than bridging inferences (Tompkins & Mateer, 1985). The second research area that has provided evidence for the right hemisphere's involvement in discourse processing is experimental studies in which critical words from a text are presented separately to one of the two visual hemifields under the assumption that they will be processed first by the contralateral hemisphere. Beeman et al. (1994) have shown that when probes are inference related, they are primed in the left visual field (right hemisphere) immediately, and subsequently are primed in both the left visual field (right hemisphere) and the right visual field (left hemisphere).

Additionally, there is a small set of brain-imaging studies that is beginning to illuminate brain function in discourse processing. Several positron emission tomography (PET) studies have examined discourse processing at a broad level; by comparing the comprehension of stories with the comprehension of unrelated sentences, Fletcher et al. (1995) and Mazoyer et al. (1993) were able to attribute activation in the left frontal gyrus to story processing. Bilateral inferior frontal and bilateral middle temporal gyri were activated when a moral judgment was required after reading a set of Aesop's fables (Nichelli et al., 1995). St. George, Kutas, Martinez, and Sereno (1999), using fMRI, found greater right-hemisphere activation when the stories were not preceded by a title than when they were, but Maguire, Frith, and Morris (1999) did not find an increase in right-hemisphere activation in a similar task. Bottini et al. (1994) found several areas of right-hemisphere activation during the processing of figurative language.

Finally, two recent fMRI studies examined cortical activation across multiple sentences that varied in their coherence. Robertson et al. (2000) found greater right-hemisphere activation for lists of sentences that used definite articles rather than indefinite articles to anaphorically relate the nouns in a text. Ferstl and von Cramon (2001) compared sentences that were either coherent or incoherent, as well as cohesive or incohesive. The coherence manipulation was similar to the contrast between stories and unrelated sentences in Mazoyer et al. and Fletcher et al.; similar areas were activated. The cohesiveness manipulation involved adding lexical connectives to the pairs of sentences to make them easier to understand as a single unit; unlike St. George et al. and Robertson et al., Ferstl and von Cramon did not find any additional right-hemisphere activation. In general, there are some indications that the right hemisphere is involved in discourse comprehension, but the results from imaging have been inconsistent, and thus the mapping of cognitive processes onto the right hemisphere is not well defined.

Our aim in this experiment was to investigate the large-scale cortical networks underlying various component processes of inference making. We expected that two-sentence passages that varied in their level of causal relatedness would yield differential orderings of brain activation across the conditions, thereby indicating which areas are involved in three facets of discourse comprehension: generation of inferences, integration of inferences, and basic sentence processing.

## METHOD

### Participants

The participants were 13 right-handed paid volunteer students (6 females). Each participant gave signed informed consent (approved by the University of Pittsburgh and the Carnegie Mellon Institutional Review Boards). Participants were familiarized with the scanner, the fMRI procedure, and the sentence comprehension task before the study started.

### Materials and Procedure

The stimulus items were pairs of sentences, taken from Myers et al. (1987), that varied in their degree of causal relatedness. In order to reduce the amount of time that participants spent in the scanner, we used only one of the two moderately related conditions in the current study, resulting in a total of three levels of causal relatedness: high, moderate, and distant. (The relatively small differences between the two moderately related conditions in the studies by Keenan et al., 1984, and Myers et al., 1987, provide an additional justification for collapsing the two moderately related conditions.) Participants read a total of 40 two-sentence passages, 10 passages in each of the three conditions and an additional 10 filler passages. The order of the passages was the same for all participants, with the experimental condition varying according to a Latin square design.

Four 30-s fixation epochs, consisting of an *X* at the center of the screen, provided a baseline activation measure. They were presented at the beginning, end, and approximate trisections of the study. In addition, the remaining interpassage intervals were filled with 12-s rest periods, also consisting of a centered *X*, to allow the hemodynamic response to approach baseline between test epochs.

In each trial, the first sentence of a sentence pair was presented for 5 s, and then was replaced by the second sentence for 5 s. An *X* then appeared on the screen for the rest period. The passage presentation and the 12-s rest that followed constituted 22 s of data acquisition. The filler items followed the same procedure as the experimental items with the exception that comprehension probes followed the filler items, to encourage participants to fully process the two sentences as a text. At the end of the presentation of the second sentence in each filler item, a yes/no comprehension probe was presented, with up to 5 s allocated for response. Comprehension probes were not presented in the experimental trials so as not to contaminate the hemodynamic response with the processing associated with reading and answering a probe question.

### Scanning Procedures

A 16-slice oblique axial prescription (approximately 10° angle relative to a straight axial) was set. This prescription covered the inferior, middle, and superior portions of the temporal lobe (including Wernicke's area); the inferior parietal lobe; and the inferior frontal gyrus (including Broca's area). The temporal pole and the medial orbital frontal regions, which are highly susceptible to artifact, were not included in our slice prescription. Figure 2 shows the location of the slices for one of the participants. The onset of each sentence was synchronized with the beginning of the acquisition of the most superior slice.

Cerebral activation was measured using blood-oxygenation-level-dependent (BOLD) contrast (Kwong, 1992; Ogawa, Lee, Kay, & Tank, 1990). Images were acquired on a 3.0-T scanner at the MR Research Center at the University of Pittsburgh Medical Center using a spiral pulse sequence in which slices were not interleaved. The acquisition parameters for the spiral scan pulse sequence with 16 oblique axial slices were as follows: TR = 1 s, TE = 18 ms, flip angle = 70°, acquisition matrix = 64 × 64, slice thickness = 3.2 mm, gap = 1 mm, RF head coil. The structural images with which the functionals were co-registered were 124-slice axial $T_1$-weighted 3D SPGR volume scans that were acquired in the same session for each participant with TR = 25 ms, TE = 4 ms, flip angle = 40°, and field of view = 24 cm; the matrix size was 256 × 192.

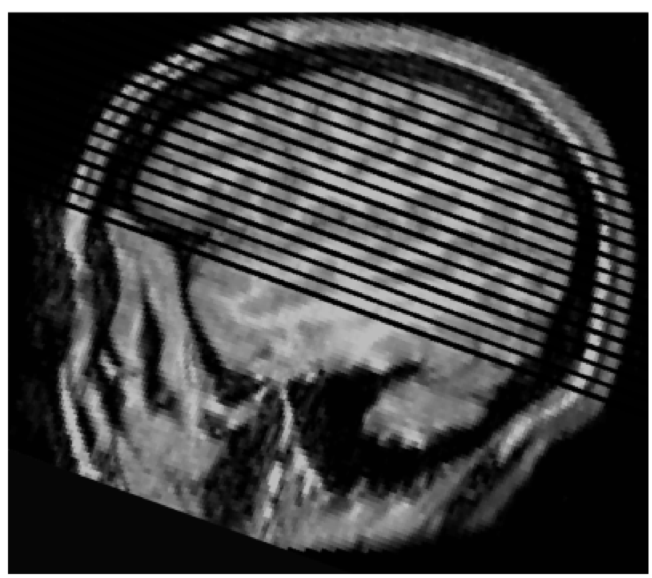

**Fig. 2.** The slice prescription for a typical participant.

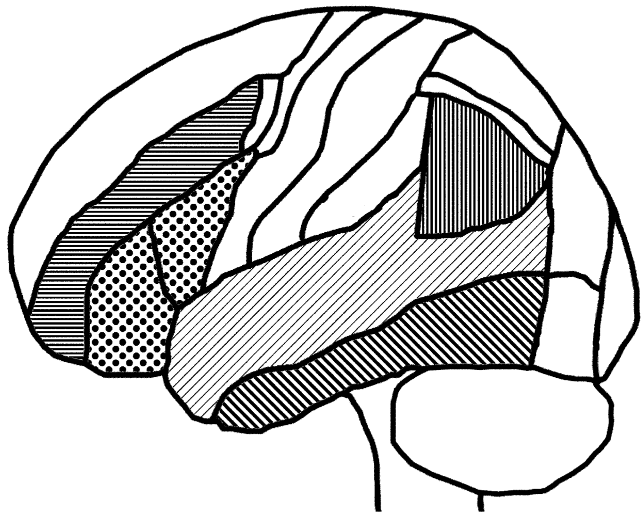

**Fig. 3.** **The anatomical areas included in the regions of interest: inferior frontal (pars opercularis and pars triangularis; black dots), medial and superior temporal (thin diagonal lines), inferior temporal (thick diagonal lines), inferior parietal (vertical lines), and dorsolateral prefrontal cortex (horizontal lines).**

### Data Analysis

To compare the amount of activation in a given area across experimental conditions, we drew anatomically defined regions of interest (ROIs) for each participant using the parcelation scheme described by Rademacher, Galaburda, Kennedy, Filipek, and Caviness (1992) and further refined by Caviness, Meyer, Makris, and Kennedy (1996). In order to examine how the degree of causal relatedness affected the volume of the activation in each of these regions, it was important to use an a priori, independent method of defining the ROIs. The ROIs in the functional images, as shown schematically in Figure 3, were defined for each participant with respect to co-registered structural images.[1]

The image preprocessing corrected for in-plane head motion and signal drift using procedures and software developed by Eddy, Fitzgerald, Genovese, Mockus, and Noll (1996). Data sets with large amounts of in-plane or out-of-plane motion were discarded without further analysis. The voxels of interest within the ROIs were identified by computing separate voxel-wise *t* statistics (using a high threshold of $t \geq 6.0$ to account for multiple comparisons) that compared the activation for the baseline fixation condition and the experimental conditions for all voxels within the ROIs (for a discussion of issues concerning baseline selection, see Binder et al., 1999, and Newman, Twieg, & Carpenter, 2001).

The main focus of the data analysis was on three groupings of ROIs: the left-hemisphere language areas, the right-hemisphere language areas, and the dorsolateral prefrontal cortex (DLPFC):

| | |
|---|---|
| left-hemisphere language areas: | inferior frontal gyrus (pars opercularis and pars triangularis) |
| | inferior temporal gyrus |
| | middle temporal gyrus |
| | superior temporal gyrus |
| | inferior parietal area (angular gyrus) |
| right-hemisphere language areas: | inferior frontal gyrus (pars opercularis and pars triangularis) |
| | inferior temporal gyrus |
| | middle temporal gyrus |
| | superior temporal gyrus |
| | inferior parietal area (angular gyrus) |
| DLPFC: | left DLPFC |
| | right DLPFC |

## RESULTS

As predicted, the three groups of ROIs were affected differently by the degree of causal relatedness. Activation in the right-hemisphere language areas showed the same inverted-U-shaped function as the recall results from Keenan et al. (1984) and Myers et al. (1987), with the activation being the highest for the moderately related sentences. This pattern of results can be seen in Figure 4. Moreover, the activation for the moderately related sentences was as high in the right hemisphere as in the left hemisphere, which is unusual in a sentence reading comprehension task (e.g., Just et al., 1996; Mason et al., in press). As expected, activation in the left-hemisphere language areas did not differ reliably across levels of causal relatedness. However, activation in DLPFC showed a nonsignificant trend toward increasing linearly as the level of relatedness decreased.

In the right-hemisphere language areas, the number of activated voxels was greatest in the moderately related condition, but in the left-hemisphere language areas, the activation did not vary as a function of the experimental manipulation; this hemisphere-by-causal-relatedness interaction was significant, $F(2, 24) = 3.70$, $MSE = 2.333$, $p < .05$. Planned comparisons showed that the number of activated voxels was greater in the left hemisphere than the right for the combination of the highly related and distantly related conditions; this effect was significant, $F(1, 12) = 6.08$, $MSE = 20.268$, $p < .05$. Only for the moderately related condition were the right-hemisphere language areas as active as the left-hemisphere language areas.

To show that the laterality effects do not apply to all of the activated cortical areas, we compared the activation across the hemispheres in the DLPFC. This region was important for the task and even showed a trend toward a linear increase across conditions; activation in the area increased as the sentences became more distantly related, although this trend did not reach significance. In contrast to the language areas, DLPFC did not show differential activation across the two hemispheres.

For completeness, the average number of activated voxels within each ROI is presented in Figure 5. The pattern of activation for the right-hemisphere areas was fairly stable across all of the individual anatomical ROIs, with the exception of pars opercularis, which was relatively inactive. Although the linear increase for DLPFC was not significant, the trend was consistent with the order of the means for both hemispheres. Finally, the three inference conditions had no systematic impact on the left-hemisphere ROIs.

[1]The interrater reliability of this ROI-defining procedure was evaluated for four ROIs in 2 participants in another study in this laboratory. The reliability measure was obtained by dividing the size of the set of voxels that overlapped between two raters by the mean of their two set sizes. The resulting eight reliability measures were in the range from 78 to 91%, with a mean of 84%, as high as the reliability reported by the developers of the parcelation scheme.

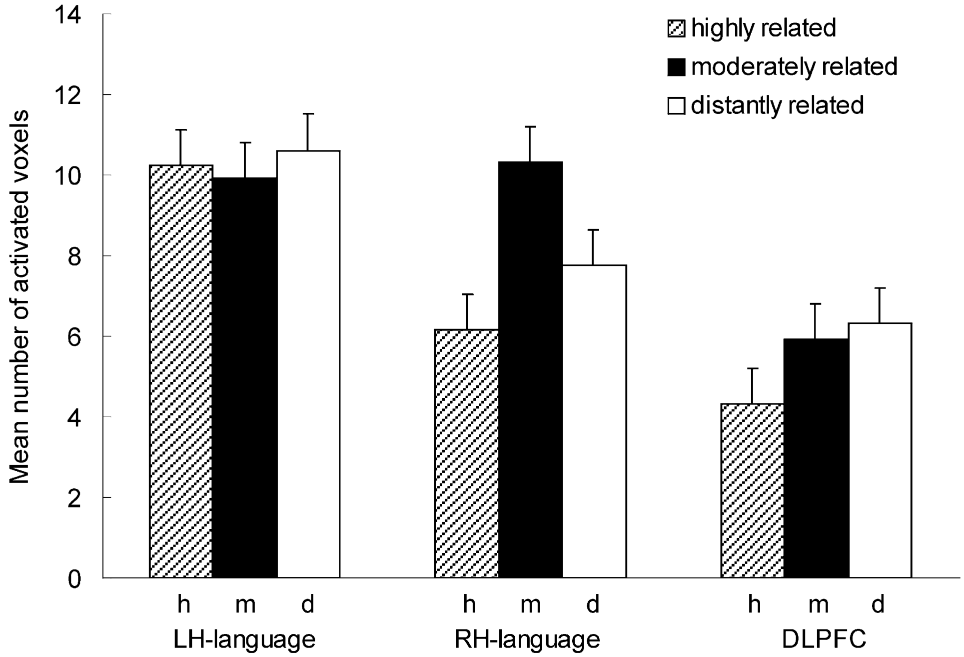


**Fig. 4.** The average number of activated voxels for the three experimental conditions averaged across the regions of interest within each of the three critical networks, the left-hemisphere (LH) language areas, right-hemisphere (RH) language areas, and dorsolateral prefrontal cortex (DLPFC).

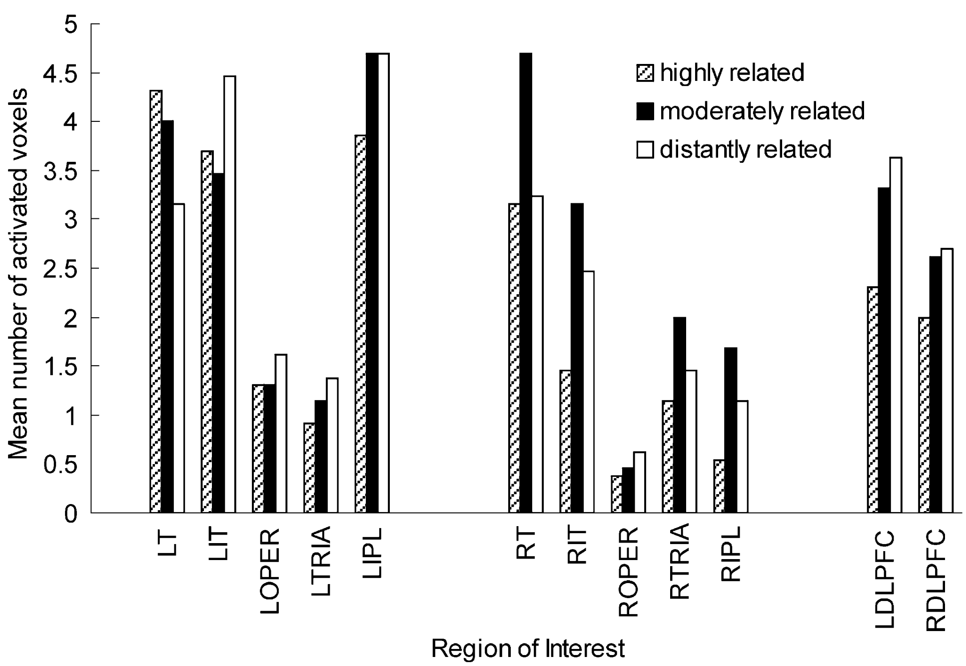


**Fig. 5.** The average number of activated voxels for the three experimental conditions for the left-hemisphere (L) language-network regions of interest (ROIs), right-hemisphere (R) language-network ROIs, and bilateral dorsolateral prefrontal cortex (DLPFC). T = middle and superior temporal; IT = inferior temporal; OPER = opercularis; TRIA = triangularis; IPL = inferior parietal lobe.

## DISCUSSION

These results are consistent with the hypothesized two-stage process of inferencing. The right-hemisphere language areas may be involved in the integration of inferences once those inferences have been generated with DLPFC involvement. The pattern of cortical activation across the two hemispheres enables us to make a distinction between the generation and integration components of inference processing. The right-hemisphere results mirror the memory performance of participants in the Myers et al. study (1987); thus, the difference in processing due to causal relatedness may be attributed to an integration of inferences that clearly involves the right hemisphere. Conversely, the lack of an effect in the left hemisphere can be attributed to the construction of the text-base representation of propositions from the text, which would not have varied across conditions.

The nonsignificant trend in the DLPFC was for activation to increase linearly across levels of causal relatedness. As the relatedness of the sentences decreased, the attempts at generating possible inferences increased, particularly in the distant condition, in which it was difficult to find an inference that could be successfully integrated into the text representation. Thus, the activation pattern in DLPFC is consistent with the generation component of inferencing. This generation process may be automatic, proceed in parallel, and operate on easily available information. The pattern of the activation in the DLPFC was similar to the pattern of reading times reported by Keenan et al. (1984) and Myers et al. (1987). It could be that the increase in reading time across the conditions is attributable to the increase in the need to generate an inference. That the longest reading times occurred for the most distantly related sentences could be due to the need to continually generate additional possible inferences as a result of the failure to successfully integrate any one inference.

The inverted-U-shaped function of activation that was found in the right-hemisphere language areas is consistent with the integration component of successful inferencing. According to Myers et al. (1987), the moderately related sentence pairs are remembered better than the other pairs because of the successful integration of a causal inference linking the two sentences. This reasoning led us to hypothesize that the brain activation in some cortical network should correspond to the integration of an inference into the internal representation of the text. The pattern of activation for the right-hemisphere language areas resembles the pattern of recognition and recall results reported by Keenan et al. (1984) and Myers et al. (1987). As a result of this integration process, the right-hemisphere activation in the moderately related condition was comparable in volume to the activation in the left hemisphere. In contrast, the integration of inferences did not occur in the distantly related and the highly related conditions, and hence the activation in the right-hemisphere language areas was equal for those two conditions and was lower in these conditions than in the moderately related condition.

Although previous brain-imaging results have shown that the right-hemisphere homologues of the left-hemisphere language areas can be differentially affected by an experimental manipulation in a language task (Just et al., 1996), to our knowledge, this is one of the few imaging studies that has shown activation in the right-hemisphere homologues that is equal to the activation in the left-hemisphere areas in a "pure" language task (i.e., one that does not involve manipulations of spatial information). Previously, additional right frontal activation was observed for reading sentences that contained definite as opposed to indefinite articles (Robertson et al., 2000). The additional activation was believed to be a result of mapping a representation of current information onto previous information. This mapping process may be analogous to the integration process described here. The new results converge with those of Robertson et al. Right-hemisphere involvement was demonstrated in both studies. In the current study, equal activation was found in the right and left hemispheres in the moderately related condition, as compared with a fixation baseline; in contrast, their results showed the right hemisphere was more active than the left hemisphere, but without a significant increase from baseline.[2] These results are also consistent with Ferstl and von Cramon's (2001) lack of differential right-hemisphere activation if we assume that their cohesive and incohesive conditions are similar to our highly related condition and distantly related condition, respectively, in which the right-hemisphere activation was less than the left-hemisphere activation.

The activation in the left hemisphere is consistent with the expectation that basic language processing (e.g., lexical access, syntactic parsing) would not be affected by the experimental manipulation. Considering that the task involved sentence reading, we expected to find activation in the left-hemisphere language areas. Furthermore, because the experimental manipulation was at the discourse level, we expected that brain areas that are primarily utilized in lower levels of sentence processing would be active, but that their activation would not necessarily vary as a function of the experimental manipulation. The theories of Beeman (1998) and the data obtained from patients who have damage to their right hemisphere (e.g., Beeman, 1993; Brownell et al., 1986; Grafman et al., 1987) led to the hypothesis that the left hemisphere would not be affected by the experimental manipulation, and this is what we found.

In summary, the results show how different parts of the brain may be differentially involved in the component processes of making inferences during the reading of narrative text. The relevant parts of the brain appear to be dynamically recruited as their involvement in text comprehension becomes needed. Bilateral DLPFC is involved whenever the generation of an inference is necessary to maintain coherence in the text. If an inference is successfully generated, the right-hemisphere language areas play an active role in integrating that inference into the reader's internal representation of the text.

***Acknowledgments***—This research was supported by National Institute of Mental Health Grant MH-29617 and National Institute of Mental Health Senior Scientist Award MH-00662. Portions of this work were reported at the 42nd Annual Meeting of the Psychonomic Society, Orlando, Florida, November 2001. We appreciate the assistance of members of the Center for Cognitive Brain Imaging in conducting the experiments and defining ROIs. We thank Victor A. Stenger for assistance with the spiral pulse sequence. We also thank

[2]It should be noted that the fMRI signal associated with the integration process appeared over several different ROIs in the right hemisphere. It is possible that the various cortical regions are utilized for even more specialized component processes that make up the integration process. This would be consistent with what is known about left-hemisphere language processing. Another possibility is that the integration process is accomplished as a whole across the various right-hemisphere cortical regions. Although this seems unlikely in view of the distinguishable functions within the left hemisphere, some evidence suggests that cognitive processing occurs in a more diffuse manner in the right hemisphere (Beeman, 1998; Semmes, 1968). The current experiments do not allow us to differentiate between these two hypotheses.

Tim Keller, Zohar Eviatar, Shasank Varma, and Franz Schmalhofer for helpful comments on earlier versions of this manuscript.

## REFERENCES

Beeman, M. (1993). Semantic processing in the right hemisphere may contribute to drawing inferences during comprehension. *Brain and Language*, *44*, 80–120.

Beeman, M. (1998). Coarse semantic coding and discourse comprehension. In M. Beeman & C. Chiarello (Eds.), *Getting it right: The cognitive neuroscience of right hemisphere language comprehension* (pp. 225–284). Mahwah, NJ: Erlbaum.

Beeman, M., Friedman, R.B., Grafman, J., Perez, E., Diamond, S., & Lindsay, M.B. (1994). Summation priming and coarse coding in the right hemisphere. *Journal of Cognitive Neuroscience*, *6*, 26–45.

Binder, J.R., Frost, J.A., Hammeke, T.A., Bellgowan, P.S., Rao, S.M., & Cox, R.W. (1999). Conceptual processing during conscious resting state: A functional MRI study. *Journal of Cognitive Neuroscience*, *11*, 80–95.

Bottini, G., Corcoran, R., Sterzi, R., Paulesu, E., Schenone, P., Scarpa, P., Frackowiak, R.S.J., & Frith, C.D. (1994). The role of the right hemisphere in the interpretation of figurative aspects of language: A positron emission tomography activation study. *Brain*, *117*, 1241–1253.

Brownell, H.H., Potter, H.H., Bihrle, A.M., & Gardner, H. (1986). Inference deficits in right brain-damaged patients. *Brain and Language*, *29*, 310–321.

Caviness, V.A., Meyer, J., Makris, N., & Kennedy, D.N. (1996). MRI-based topographic parcellation of human neocortex: An anatomically specified method with estimate of reliability. *Journal of Cognitive Neuroscience*, *8*, 566–587.

Eddy, W.F., Fitzgerald, M., Genovese, C.R., Mockus, A., & Noll, D.C. (1996). Functional imaging analysis—computational olio. In A. Prat (Ed.), *Proceedings in computational statistics* (pp. 39–49). Heidelberg, Germany: Physica-Verlag.

Ferstl, E.C., & von Cramon, D.Y. (2001). The role of coherence and cohesion in text comprehension: An event-related fMRI study. *Cognitive Brain Research*, *11*, 325–340.

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

Grafman, J., Salazar, A.M., Vance, S.C., Weingartner, H., & Amin, D. (1987). Immediate memory for story discourse in Vietnam veterans with penetrating brain wounds. *Journal of Clinical and Experimental Neuropsychology*, *9*, 23–34.

Just, M.A., Carpenter, P.A., Keller, T.A., Eddy, W.F., & Thulborn, K.R. (1996). Brain activation modulated by sentence comprehension. *Science*, *274*, 114–116.

Keenan, J.M., Baillet, S.D., & Brown, P. (1984). The effects of causal cohesion on comprehension and memory. *Journal of Verbal Learning and Verbal Behavior*, *23*, 115–126.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, *95*, 163–182.

Kwong, K.K., Belliveau, J.W., Chester, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., & Turner, R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences, USA*, *89*, 5675–5679.

Maguire, E.A., Frith, C.D., & Morris, R.G. (1999). The functional neuroanatomy of comprehension and memory: The importance of prior knowledge. *Brain*, *122*, 1839–1850.

Mason, R.A., Just, M.A., Keller, T.A., & Carpenter, P.A. (in press). What brain imaging reveals about the processing of syntactically ambiguous sentences. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Mazoyer, B.M., Tzourio, N., Frak, V., Syrota, A., Murayama, N., Levrier, O., Salamon, G., Dehaene, S., Cohen, L., & Mehler, J. (1993). The cortical representation of speech. *Journal of Cognitive Neuroscience*, *5*, 467–479.

Myers, J.L., & Duffy, S.A. (1990). Causal inferences and text memory. *The Psychology of Learning and Motivation*, *25*, 159–173.

Myers, J.L., Shinjo, M., & Duffy, S.A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language*, *26*, 453–465.

Newman, S.D., Twieg, D.B., & Carpenter, P.A. (2001). Baseline conditions and subtractive logic in neuroimaging. *Human Brain Mapping*, *14*, 228–235.

Nichelli, P., Grafman, J., Pietrini, P., Clark, K., Lee, K.Y., & Miletich, R. (1995). Where the brain appreciates the moral of a story. *NeuroReport*, *6*, 2309–2313.

Ogawa, S., Lee, T.M., Kay, A.R., & Tank, D.W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences, USA*, *87*, 9868–9872.

Rademacher, J., Galaburda, A.M., Kennedy, D.N., Filipek, P.A., & Caviness, V.S., Jr. (1992). Human cerebral cortex: Localization, parcellation, and morphometry with magnetic resonance imaging. *Journal of Cognitive Neuroscience*, *4*, 352–374.

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

Semmes, J. (1968). Hemispheric specialization: A possible clue to mechanism. *Neuropsychologia*, *6*, 11–26.

Singer, M. (1994). Discourse inference processes. In M.A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 479–515). San Diego, CA: Academic Press.

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.

Tompkins, C.A., & Mateer, C.A. (1985). Right hemisphere appreciation of prosodic and linguistic indications of implicit attitude. *Brain and Language*, *24*, 185–203.

van den Broek, P.W. (1994). Comprehension and memory of narrative texts inferences and coherence. In M.A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 539–588). San Diego, CA: Academic Press.

(Received 6/27/02; Revision accepted 12/20/02)