Interpretation-mediated changes in neural activity during language comprehension

Emily A. Cooper [a],*, Uri Hasson [b], Steven L. Small [c]

[a] *Helen Wills Neuroscience Institute, UC Berkeley, Berkeley, CA, USA*
[b] *Center for Mind/Brain Sciences and Faculty of Cognitive Science, University of Trento, Trento, Italy*
[c] *Departments of Psychology and Neurology, University of Chicago, Chicago, IL, USA*

ARTICLE INFO

ABSTRACT

Using functional magnetic resonance imaging (fMRI), we identified cortical regions mediating interpretive processes that take place during language comprehension. We manipulated participants' interpretation of texts by asking them to focus on action-, space-, or time-related features while listening to identical short stories. We identify several cortical regions where activity varied significantly in response to this attention manipulation, even though the content being processed was exactly the same. Activity in the posterior and anterior sections of the left inferior frontal gyrus (IFG), which are thought to have different sensitivities to high-level language processing, was modulated by the listeners' attentional focus, but in ways that were quite different. The posterior left IFG (Pars Opercularis) showed different activity levels for the three conditions. However, a population coding analysis demonstrated similar distributions of activity across conditions. This suggests that while the gain of the response in the Pars Opercularis was modulated, its core organization was relatively invariant across the experimental conditions. In the anterior left IFG (Pars Triangularis), the analysis of population codes revealed different activity patterns between conditions: there was little similarity between activity during time–attention and action- and space–attention, however there were similar activity patterns while attending to space and action information. In addition, both the left superior temporal gyrus and sulcus showed greater activity in the space and action attention conditions when contrasted with time attention. We discuss these findings in light of work on the role of left IFG in processing semantic information in language, and in light of theories suggesting that temporal information in language is processed in the brain using similar mechanisms as spatial information. Our findings suggest that a substantial source of variance in neural activity during language comprehension emerges from the internally-driven, information-seeking preferences of listeners rather than the syntactic or semantic properties of a text.

© 2011 Elsevier Inc. All rights reserved.

Introduction

The experience of getting lost in a story is a familiar one. Reading even the most detailed text engages mental processes in which readers make inferences and develop interpretations based on their prior knowledge and intrinsic motivation (Buckner and Carroll, 2006; Friese et al., 2008). Our goal in the present study was to use functional magnetic resonance imaging (fMRI) to identify the neural circuitry mediating these interpretive processes of language.

Previous work has shown that neural activity evoked by listening to language depends to some extent on the nature of the actors and events that are described. Much of this work has focused on the finding that language describing actions engages low-level sensory and motor systems more than language that is unrelated to action (e.g., Tettamanti et al., 2005; Beilock et al., 2008). Deen and McCarthy (2010) demonstrated that listening to stories that specifically emphasize biological motion activates brain regions involved in perception of biological motion. Differences have also been reported between regions involved in processing narratives about more abstract concepts, such as emotions and the passage of time (Ferstl and von Cramon, 2007). Other aspects of a narrative, like changes in spatial location, interactions with objects, or the introduction of character goals, may also be associated with different types of neural activity (Speer et al., 2009).

Several prior works have tried to dissociate neural activity driven by properties of a text (i.e., its specific lexical, syntactic, or semantic features) from that driven by interpretive processes. These studies have focused largely on examining the integration of new information gained from the text with prior knowledge, by manipulating the coherence of a passage or sentence. For example, St. George et al. (1999) had participants listen to identical paragraphs either with or without a title that framed the text with a specific schema (eg. "Horseback riding"). The untitled paragraphs were perceived as incoherent, but participants had no trouble understanding the titled versions. St. George et al. reported greater activation overall to the incoherent untitled paragraphs, mostly in the right hemisphere. However, other

* Corresponding author. Banks Lab University of California, Berkeley, 360 Minor Hall, Berkeley, CA 94720-2020, USA. Fax: +1 510 643 5109.
E-mail address: emilycooper@berkeley.edu (E.A. Cooper).

studies that manipulated text coherence by varying coherence relations between sentences within a text (e.g. Xu et al., 2005; Hasson et al., 2007) have shown that fronto-temporal regions are increasingly engaged as language becomes more informative.

The premise underlying our current work is that even when a story is coherent and unambiguous, it can have many possible interpretations. These interpretations depend on the background knowledge of the reader or his/her particular goals when reading the text. In the present study, we therefore examined differences in neural activity arising when participants interpreted the same stories with different information-seeking goals. To accomplish this, we framed each story with a short introduction, which instructed participants to focus on a single feature of the story: the specific events that occur, the locations of these events, or their time course (See Table 1 for an example core story framed using three different introductions). Thus, different participants listened to the exact same words and sentences but focused on different aspects of their meaning. We collected whole-brain blood oxygen level dependent (BOLD) fMRI data in order to determine which regions of the brain were sensitive to this manipulation.

Using this paradigm, we tested several hypotheses about the neural circuitry mediating top-down interpretation. The first hypothesis was based on research showing that the posterior and anterior sectors of the left inferior frontal gyrus (IFG) may be differentially specialized for processing meaning. Several studies have indicated that the left posterior IFG (Pars Opercularis) is sensitive to relatively non-semantic aspects of language, such as syntax and phonology (Gough et al., 2005; Pulvermuller et al., 1999; Poldrack et al., 1999). Yet activity in this region is also sometimes linked to semantic processing: Thompson-Schill et al. (1999) showed increased left posterior IFG activity during word generation, suggesting that this region is involved in semantic retrieval. Tettamanti et al. (2005) suggested that this same region plays an important role in processing action-related information during language comprehension, since it was more sensitive to action sentences than to neutral ones. Adding to this literature are studies showing interactions between posterior and anterior IFG during letter fluency and category fluency. Paulesu et al. (1997) found that Pars Opercularis was selectively active during the former (phonemic

generation), whereas Pars Triangularis was selectively active during the latter (semantic generation). However, Heim et al. (2008) used the same tasks and found that Pars Opercularis and Triangularis were active during both, but that Pars Opercularis showed reliably greater activation during phonological generation than during other conditions.

These findings have left open the issue of the relative role of these sectors of the left IFG in mediating higher-level comprehension processes of the sort examined in the current study. At the level of discourse, several studies suggest that anterior aspects of IFG (i.e., Pars Triangularis and Pars Orbitalis) are particularly involved in higher-level language comprehension functions such as integration of sentences (e.g., Lei et al., 2006; Hasson et al., 2007). These studies used narratives to set up contextual expectations during auditory comprehension. Lei et al. (2006) found the left anterior IFG to be particularly sensitive to sentences that were inconsistent with prior context. Hasson et al. (2007) similarly found greater activation in the anterior left IFG to more surprising story segments than to less surprising ones. The authors theorize that the left anterior IFG is sensitive to high-level information integration.

Our first hypothesis addressed the roles of anterior and posterior left IFG in interpretation-dependent language processes. First, because both portions of left IFG are known to mediate aspects of language comprehension, we expected that all three conditions should be associated with above-baseline activity in both. However, to the extent that left Pars Opercularis is involved in low-level semantic processing (i.e., independent of top-down influences), it would differentiate less strongly among the conditions than would Pars Triangularis. This would lead to an interaction between regions and conditions, with reliable differences between action, space, and time attention conditions in Pars Triangularis, but weaker effects (if any) in Pars Opercularis.

Our second hypothesis was more exploratory and dealt with the neural processing of space- and time-related information in language more generally. A prominent position in linguistics and philosophy is that abstract linguistic expressions (e.g., "to move forward in time") are grounded in actual sensory-motor experience (Lakoff and Johnson, 1980; Sweetser, 1992). This position has inspired a large body of work in cognitive psychology (e.g., Boroditsky and Ramscar, 2002) leading to a view that temporal information may be interpreted via spatial analogy (see Kemmerer, 2005, for a thorough exposition and critique of the time-as-space framework). Our study aimed to investigate the biological basis of this view by examining neural activity in the two experimental conditions where participants focused on time-related vs. space-related information while hearing the exact same stories. In this way we could determine if top-down attention processes are associated with different activity patterns, while controlling for effects that might be caused by differences in lower level linguistic phenomena.

If spatial and temporal information in language are processed by similar neural systems, then the patterns of BOLD responses in those two attention conditions would be reliably different from the third attention condition (focusing on actions), but not from each other. Beyond examining activity in lateral temporal regions that are often associated with semantic processes, we were also interested in several frontal regions that are known to be sensitive to language context: the inferior and middle frontal gyri and the precentral gyrus (Hasson et al., 2009; See Ferstl et al., 2008 for meta-analysis and review). If reliable differences were found between the time and space conditions, this would suggest they are not mediated via a shared neural system. We proposed the same hypothesis for comparisons between the time and action conditions, based on the idea that time conceptualized as "motion through space" in language may in fact be more related to action language processing than spatial language.

Table 1

Focus manipulation for a "core" story text. Participants listened to twelve short stories. Different conditions were created using a framing device, such that each of twelve "core" stories was cued to all three concepts of interest between subjects: action, space, and time. Participants were reminded of the dimension they should attend to by two sentences that introduced each of the stories and also by general instructions given before each group of four stories (see Methods).

Action introduction	Space introduction	Time introduction
Aaron had done many things while in search of his lost pet lion. I'm going to tell you a story about all the things that he did.	Aaron had gone many places in search of his lost pet lion. I'm going to tell you a story about all the places that he went.	Aaron had spent lots of time searching for his lost pet lion. I'm going to tell you a story about the timeline of his search.

Core story

When the lion ran away one winter night, Aaron went looking in the Lincoln Park Zoo. He jumped the fence into the lion den, and spent hours calling it by name. But when security dragged him away, he decided to search in Africa. So in the spring, he stowed away on an ocean liner, crushed underneath a large pile of life jackets. And as the boat neared the Angolan coast, Aaron jumped into the ocean and swam his way to land. He spent months hiking through the African planes but his lion was definitely not there. And in Nigeria, he clung to the top of a tree for three days when he found himself in the middle of a fight between two prides. Aaron went home to Chicago that fall and he found his pet lion sitting waiting on his doorstep. A few months later when the lion ran away again, he followed stealthily behind. And together they ended up in Mexico, laying on the warm beach to wait out the Chicago winter.

Methods

Participants

Twelve participants (four men and eight women, mean age 22.4 years) were recruited from the student and employee populations at The University of Chicago. All were right-handed, native English speakers, and had normal hearing and normal (corrected) vision. Participants provided written informed consent and the Institutional Review Board of The University of Chicago Biological Sciences Division approved the study.

Materials and procedure

The stimuli consisted of twelve short auditory stories, each twelve sentences long, which were presented to participants via headphones during an fMRI scan. Participants listened to stories in three different conditions, differentiated by short introductions that focused their attention on the action-, space-, or time-related information. Several studies have shown that directing participants' attention to specific features of linguistic input in this way can facilitate feature-specific brain activity (Mitchell et al., 2003; Hugdahl et al., 2003). Further, behavioral data suggest that readers readily separate out and focus attention on one dimension of a narrative at a time (Therriault et al., 2006). Each experimental run consisted of four stories within one of the three conditions (action–attention, space–attention, and time–attention), and the experiment consisted of three runs, one for each condition. Each run contained four 90 second stories separated by 30-second breaks to allow the hemodynamic response to return to baseline level. Thus, each functional run lasted about eight minutes in total. The 90 second period for each story consisted of (i) an eight second introduction; (ii) a two second interval; and (iii) the main content (80 s). Stories were preceded by a two second orientation tone.

Before each experimental run, participants were also given explicit verbal instructions to focus on a single dimension of the stories: the actions, locations, or time-course of events. Between participants, each of the twelve stories was cued by all three dimensions. The serial order of the stories remained static, whereas the framing conditions were rotated for each participant. Participants were told to pay attention because they would be quizzed following one of the three story runs. This instruction was included to motivate the participants to attend carefully. Due to time constraints, the participants were never actually quizzed.

Stimulus construction

As explained, the stories were designed to direct the participants' attention maximally to one dimension at a time. To do this, three interchangeable introductions were written for each of the twelve core stories, arranged to focus attention on the dimension of interest. These introductions were two sentences long and spoken in the voice of the narrator directing the listener to the subject matter of the story they were about to hear (e.g., "I'm going to tell you a story about all the places that someone went." See Table 1 for a complete example).

Our goal with the core stories was to create narratives for which it was equally easy for listeners to focus in and attend to all three dimensions. Behavioral work has shown that increased processing of a given dimension in a story is prompted when there is a change along that dimension (ie. the characters move to a new location, or the story jumps ahead in time) (Zwaan et al., 1995; Therriault et al., 2006). Thus, we designed our stories such that each sentence included a change in the time, location, and actions in the story. All core stories were ten sentences long. We used longer narratives than is typical in fMRI studies of language in order to maintain participant interest, allow for the development of more engaging and complex plots, and

to be able to collect sufficient data points per experimental condition. Although each participant heard four stories in each of the three conditions, we were able to collect 160 whole-brain images (40 per story) per condition due to their duration. This amount of data per condition is typical of other studies of language using fMRI (eg. Hasson et al., 2007 [120 volumes per condition]; Saygin et al., 2010 [126 volumes per condition]).

While each sentence was written so that it contained changes in the space-, time-, and action-related information, there was no further control for the amount of information related to each condition, and no post-hoc behavioral assessments of the participants processing of the different types of information. There is currently no agreement in the literature as to how to objectively quantify "equal amounts" of temporal, spatial and action information (Therriault et al., 2006). However, our strategy of including changes in each dimension in each sentence tended to lead to similar amounts of words and syllables (in the example story from Table 1, there are 21 verbs referring to actions (26 syllables), 20 nouns referring to specific places (37 syllables), and 19 words or phrases referring to temporal relationships (21 syllables)).

Stimulus presentation

All stimuli were spoken by a female speaker in a quiet recording room, recorded to digital tape, and converted to computer files (16-bit mono, 48 khz sampling rate). Stimulus volume was mean-normalized to 21 dbfs. Each core story was recorded separately from the three types of introduction and then concatenated with them, such that every story had three different versions.

fMRI data acquisition

Scans were acquired on a 3 Tesla scanner (GE Signa) using spiral acquisition with a standard head coil (Noll et al., 1995). Two volumetric T1-weighted scans (120 axial slices, $1.5 \times 0.938 \times 0.938$ mm resolution) were acquired and averaged to provide a high-quality structural image on which to identify anatomical landmarks and onto which functional activation maps could be superimposed. For the functional scans, thirty-two spiral T2* gradient echo images covering the entire brain were collected every two seconds in the axial plane (32 slices, slice thickness of 3.8 mm no-gap, interleaved acquisition, matrix size $= 64 \times 64$, TR $= 2$ s; TE $= 30$; flip angle $= 77$). On the individual participant level, functional runs were spatially registered in three-dimensional space to a single reference time point in order to correct for head movement, using AFNI (http://www.afni.nimh.nih.gov/afni/). For each participant, the raw signal in each voxel was scaled to its mean signal during the study. Time points associated with extreme movement were removed from the regression models (about 1% of the data).

Data analysis

Anatomical images were aligned to the functional volumes automatically (Saad et al., 2009) and alignment was manually verified and adjusted when needed. Time series were despiked, mean-normalized, and spatially smoothed with a 6 mm kernel to increase the signal-to-noise ratio. Data were analyzed using multiple linear regression with AFNI's 3dDeconvolve utility (Cox, 1996). Regressors were waveforms with similarity to the hemodynamic response, generated by convolving a gamma variant function with the onset time and duration of the stories of interest. There were three such regressors of interest for each participant (time-, space-, and action–attention manipulations). The remaining regressors were the mean, linear and quadratic trend components of the time series, and the six motion parameters for each of the functional runs. Time points from the introductory sentences were excluded from the regression model

due to their statistical colinearity with the immediately following attention conditions.

To define anatomical regions of interest (ROIs) on each individual's brain, the anatomical volumes were inflated to a surface representation using the FreeSurfer software package (version 4.3; Fischl et al., 1999). ROIs were defined separately for each participant in his or her original brain space. The regions were delineated on the basis of anatomical features alone, using the automatic parcellation procedure implemented in FreeSurfer. These parcellation methods have been shown to be comparable in accuracy to manual parcellation (Fischl et al., 2002, 2004) and their statistical knowledge base derives from a training set incorporating the anatomical conventions of Duvernoy (1991). All ROIs were delineated for each participant in their original space (See Fig. 1 for an example). The resulting surface representations, as well as their parcellation into regions, were imported into SUMA for functional analysis (Fischl et al., 2004; Saad et al., 2004). The functional data were then projected from the 3-dimensional volumes onto the 2-dimensional surfaces in SUMA.

Statistical analyses for ROIs were conducted using the "R" statistical software package (http://www.r-project.org/). The regions analyzed included each of the three anatomical parts of the IFG (Pars Opercularis, Pars Triangularis, and Pars Orbitalis), and two frontal regions associated with complex language processing and language context (middle frontal and precentral gyrus). The temporal poles, which are often active during language comprehension (Ferstl et al., 2008), were not included in the ROI analysis because the signal drop out in our fMRI scan precluded a meaningful analysis. Initially, we had also planned to include lateral temporal regions (MTG, STG) within the set of anatomically defined ROIs. However, we excluded these regions from the ROI analysis after finding between-condition differences in these regions in the whole-brain analysis. An ROI analysis of a brain region identified in the whole-brain analysis would constitute a non-independent test. All ROI ANOVAs that tested for between-region differences were evaluated for violations of sphericity to take into account the possibility of inter-subject correlations across regions. No such violations were found in any of the analyses, and all results remained reliable at the reported levels of significance after Greenhouse–Geisser corrections.

Whole brain analyses were conducted to identify regions showing differential activity vs. baseline for each of the three attention conditions, as well as regions showing differential activity across the three conditions. All analyses were controlled for family-wise error rate of $p<.05$ using cluster thresholds. For the whole brain analysis the single-voxel threshold was set at an alpha level of $p<.001$, and for the between-condition contrasts the single voxel level threshold was set

at $p<.01$. We chose a relatively strict threshold for the tests against baseline based on the fact that these tests are intrinsically more powerful than between-conditions contrasts. Matching cluster extent thresholds for these two voxel-level alpha levels were determined using Monte-Carlo simulations. These simulations used an estimate of the spatial smoothing of the experimental dataset (from the residuals of the regression) to determine the sampling distribution of the maximum cluster size likely to be found by chance (Forman et al., 1995). Cluster extent thresholds were 300 mm^2 for between-condition comparisons, and 120 mm^2 for comparisons against baseline.

Results

Whole brain analysis

To verify our basic results against prior work, we conducted a whole-brain analysis to identify brain regions that showed above or below baseline activity at the group level in any condition. We compared these with previous findings for language comprehension. The results are presented in Fig. 2 and Table 2. Across story conditions, brain areas showing reliable activity ($p<.05$, family-wise error rate corrected for multiple comparisons) were those commonly found for the comprehension of spoken language, and included bilateral temporal regions, and left frontal regions (Fig. 2a). Regions showing deactivation patterns included both anterior and posterior midline regions, as well as the intraparietal sulcus (Fig. 2b). The distribution of deactivation was highly similar to that found in our prior work on discourse comprehension (Hasson et al., 2007).

To determine the effects of the attention manipulation, we conducted a whole-brain analysis testing for differences across conditions. Two of the contrasts – Action vs. Time and Space vs. Time – showed reliable differences in activity. In the left hemisphere, both action-attention and space-attention revealed greater activity overall than time-attention. In both cases, this increased activity was found in lateral temporal regions. The Action vs. Time contrast revealed reliably greater activity during action–attention in the left anterior superior temporal gyrus (STG) and the left posterior superior temporal sulcus (STS). The Space vs. Time contrast revealed greater activity for space–attention in the left anterior STG just posterior to the temporal pole and in the left posterior STS. In the right hemisphere, greater activity was found for this contrast in the anterior STG (see Fig. 3 and Table 3). No other contrasts were reliable at the whole-brain level when corrected for family-wise error.

Region of interest analysis

For this analysis, regions of interest were defined anatomically. Of the ROIs we examined, the MFG, precentral gyrus, and IFG Pars Orbitalis showed no reliable impact of condition, and we do not discuss them further. To understand patterns of activation in the Pars Opercularis (POp) and Pars Triangularis (PTr) of the left IFG, we conducted several analyses examining activity in these two areas. These consisted of: (1) a group-level random effects analysis of Beta values (representing percent signal change), (2) a group-level random effects analysis summarizing reliable activity relative to baseline and between-condition differences propagated from the individual level analyses, and (3) a group-level random effects analysis examining population coding separately within these two regions.

Beta values for the three conditions reflect the degree of BOLD signal intensity in each condition. A 2 (Region) × 3 (Condition) ANOVA for POp and PTr revealed a marginal interaction $F(2, 22) = 2.79$, $p = .08$. Follow up contrasts revealed that in the PTr activity during space-attention was stronger than that during time-attention: $T(11) = 2.73$, $p = .02$. No other contrast was reliable in either PTr or
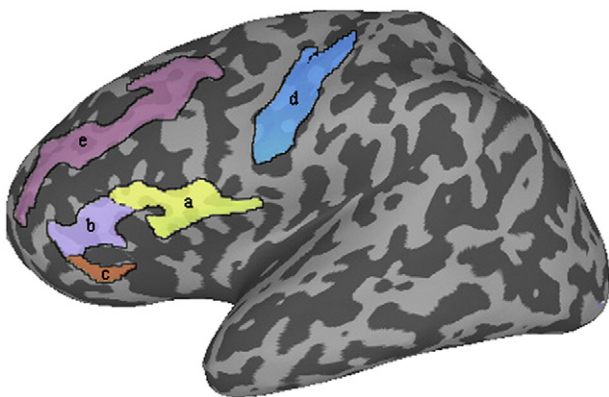

Fig. 1. Example of anatomical region of interest (ROI) delineation on an individual participant's brain. ROI's were defined in each participant's original brain space and delineated by anatomical features alone, using the automatic parcellation procedure implemented in FreeSurfer (Fischl et al., 2002, 2004). All ROI's were in the left hemisphere. (a) IFG Pars Opercularis; (b) IFG Pars Triangularis; (c) IFG Pars Orbitalis; (d) precentral gyrus; (e) middle frontal gyrus.
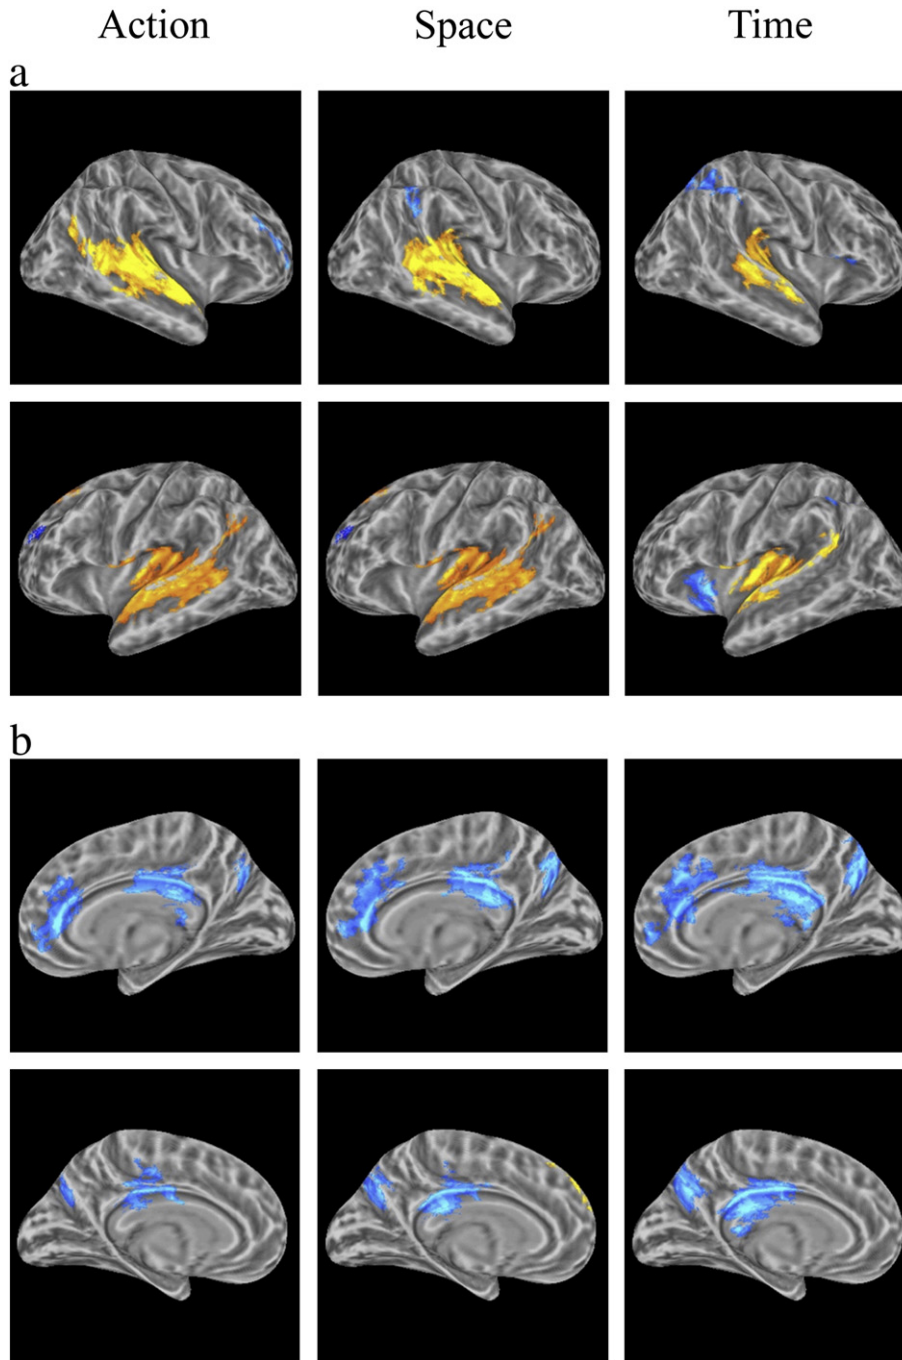
Fig. 2. Group level activation maps (vs. baseline) for the action, space, and time attention conditions in (a) lateral and (b) midline cortex. Regions showing positive activation relative to baseline are presented in warm colors and those showing relative deactivation are presented in cold colors. Individual voxel threshold is $p<.001$ (Family-wise error rate of $p<.05$ using cluster threshold).

POp. The analysis of Beta values, while typical in neuroimaging studies, can suffer from limitations of power. This particularly holds given our sample size ($N=12$) and the fact that Beta estimates do not carry information about whether there was a reliable difference between conditions at the single participant level (see Heller et al., 2007). To address this limitation, we conducted several other analyses at the group level.

The first analysis examined the proportion of voxels that showed reliable activity for each of the experimental conditions in POp and PTr. For each participant, we determined activity in these two regions by conducting a General Linear Test (vs. baseline) on the voxel's time series. We then quantified the percentage of voxels that showed reliable activity at the single participant level ($p<.001$), and

conducted a 2 (Region)×3 (Condition) repeated measures ANOVA with participants as a random factor to determine if the percentage of active voxels differed as a function of condition and region (see Sridharan et al., 2007 for similar analyses). The ANOVA revealed a main effect of Region ($F(1, 11)=18.78$, $p<.01$) and Condition ($F(2,22)=4.29$, $p<.05$), but no interaction. The main effect of Region reflected the fact that POp had a larger percentage of active voxels than PTr (POp $M=20\%$, $SE=4\%$; PTr $M=7\%$, $SE=2\%$). See Fig. 4a. The main effect of condition reflected a decrease in activity from action–attention to space–attention and space–attention to time–attention. Contrasts against chance probability (0.1% nominal value for $p<.001$) indicated that all three conditions had above-chance activity in POp (all T values>4.1, all p values<.001). In PTr,

Table 2

Regions showing reliable departures from baseline for each condition. Individual voxel threshold is $p<.001$ (Family-wise error rate of $p<.05$). Center of mass is defined by Talairach and Tournoux coordinates in the volume space. Area = cluster area in mm^2. [T] = maximum (activation) or minimum (deactivation) T value in the cluster. BA = Brodmann Area. L = left. R = right. G = gyrus. IPL = inferior parietal lobule. TTG = transverse temporal gyrus. MTG = middle temporal gyrus. STG = superior temporal gyrus. MFG = middle frontal gyrus. SFG = superior frontal gyrus. dMPFC = dorsomedial prefrontal cortex. IFG = inferior frontal gyrus. Post. = posterior. Cent. = central. Ant. = anterior.

Region (BA)	Action attention					Space attention					Time attention				
	Talairach			Area [T]		Talairach			Area [T]		Talairach			Area [T]	
	x	y	z			x	y	z			x	y	z		
Activation															
Temporal and parietal															
L. Insula (BA13)	−39	−36	18	3546	[28.6]	−39	−36	19	3764	[18.1]	–	–	–	–	
L. IPL (BA40)	−47	−51	41	224	[7.0]	–	–	–	–		–	–	–	–	
L. TTG (BA41)	–	–	–	–		–	–	–	–		−35	−27	13	2197	[9.2]
R. MTG (BA20)	53	−34	−8	3251	[16.2]	54	−33	−8	3265	[14.9]	–	–	–	–	
R. STG (BA21)	–	–	–	–		–	–	–	–		56	−19	−2	1816	[8.2]
Frontal															
L. Precentral Gyrus (BA6)	–	–	–	–		−53	−3	7	252	[6.5]	−52	−1	6	229	[6.7]
L. MFG (BA6)	–	–	–	–		−6	36	33	237	[5.0]	–	–	–	–	
Deactivation															
Temporal and parietal															
L. Insula (BA13)	–	–	–	–		–	–	–	–		−39	0	3	441	[−8.0]
R. Insula (BA13)	–	–	–	–		–	–	–	–		33	−7	10	273	[−7.0]
L. IPL (BA40)	–	–	–	–		–	–	–	–		−29	−35	36	158	[−4.6]
L. Intraparietal Sulcus	−27	−36	41	147	[−4.4]	–	–	–	–		–	–	–	–	
R. Angular G (BA39)	–	–	–	–		–	–	–	–		34	−56	34	1015	[−8.6]
R. Supramarginal G (BA40)	–	–	–	–		51	−55	33	260	[−6.6]	–	–	–	–	
Medial regions															
L. Post. Cingulate (BA23)	0	−29	20	722	[−6.4]	−1	−28	17	724	[−7.0]	−9	−30	19	1248	[−9.9]
R. Cent. Cingulate (BA24)	7	−14	40	924	[−8.6]	6	−15	42	1015	[−8.8]	7	−13	38	1520	[−8.6]
R. Ant. Cingulate (BA32)	–	–	–	–		14	35	20	928	[−7.4]	13	33	22	1097	[−5.5]
L. Precuneus (BA31)	−10	−64	27	172	[−4.7]	–	–	–	–		–	–	–	–	
L. Precuneus (BA7)	–	–	–	–		–	–	–	–		−10	−66	30	588	[−8.3]
											−19	−65	42	124	[−4.3]
R. Precuneus (BA7)	10	−60	34	187	[−6.2]	8	−59	34	386	[−7.1]	11	−54	40	655	[−8.5]
Frontal															
L. SFG (BA9)	–	–	–	–		−10	58	25	163	[−5.8]		–	–	–	
R. dMPFC (BA32)	14	37	14	819	[−6.4]	–	–	–	–		–	–	–	–	
R. IFG (BA46)	51	28	11	276	[−6.1]	–	–	–	–		–	–	–	–	
Occipital															
L. Cuneus (BA7)	–	–	–	–		−10	−66	31	433	[−6.0]	–	–	–	–	

activity was reliably above chance for the action ($T(11) = 3.01$, $p<.01$) and space ($T(11) = 4.9$, $p<.001$), but not for the time–attention condition, with 5 of the 12 participants showing no active voxels at this threshold for this condition.

We conducted a similar analysis to examine the proportion of voxels that differentiated between conditions in POp and PTr. We summarized between-condition differences at the single participant level by determining the percentage of voxels showing a reliable between-condition difference at a level of $p<.001$. This percentage of 'discriminating' voxels was established separately for each participant and then analyzed at the group level (see Fig. 4b). For POp on average 12% of voxels showed a reliable difference for the Action vs. Time contrast ($SE=4\%$), 12% for the Space vs. Action contrast ($SE=4\%$),

and 9% for the Time vs. Space contrast ($SE=2.5\%$). These percentages reliably exceeded that expected by chance (i.e., 0.1%) for all three contrasts (all $ps<.02$). The same analysis conducted for PTr revealed slightly lower percentages: Action vs. Time: 8% ($SE=3\%$), Space vs. Action: 4% ($SE=1\%$), Time vs. Space: 4% ($SE=1.5\%$). These percentages also reliably exceeded that expected by chance for all three contrasts (all $ps<.05$). This analysis indicates that participants dissociated between all three conditions at rates that largely exceed chance by one or two orders of magnitude, in both POp and PTr. This analysis was highly sensitive because it utilized the strong power of the between-condition test on the single voxel level, and propagated this information to the group-level analysis. Simulations conducted to evaluate the sensitivity of these general linear tests showed they had
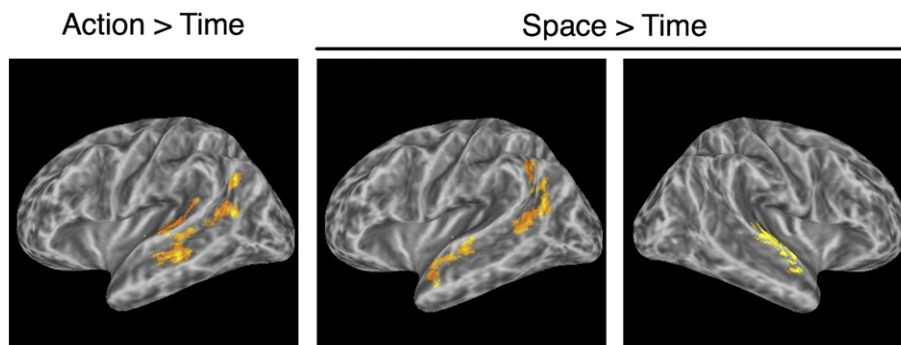

Fig. 3. Group level contrast maps for differences between action, space, and time attention conditions. Differences in the left hemisphere were found for both space and action vs. time, and in the right hemisphere for space vs. time. Individual voxel threshold is $p<.01$ (Family-wise error rate of $p<.05$ using cluster threshold).

Table 3
Regions showing differences across conditions. Individual voxel threshold $p<.01$ (Family-wise error rate of $p<.05$). Center of mass is defined by Talairach and Tournoux coordinates in the volume space. Area = cluster area in mm². [T] = maximum (activation) or minimum (deactivation) T value in the cluster. BA = Brodmann Area. L = left. R = right. Ant. = anterior. Post. = posterior. STG = superior temporal gyrus. STS = superior temporal sulcus/\.

Region (BA)	Talairach			Area [T]
	x	y	z	
Action>time attention				
L. Ant. STG (BA42)	−55	−35	15	824 [5.3]
L. Post. STS (BA39)	−50	−57	22	352 [5.7]
Space>time attention				
L. Post. STS (BA40)	−44	−48	32	522 [4.7]
L. Ant. STG (BA41)	−59	−20	6	312 [4.4]
R. Ant. STG (BA22)	45	−14	−8	423 [5.9]

high power: on the single voxel level, the probability of finding signal difference of 0.2% (1/5th of a percent) exceeded 80% given our design and given the mean tSNR in left IFG (~150; see Supplementary materials for full simulation details).

The analyses of baseline-relative and between-condition differences reported above indicate that the three conditions prompted different levels of activation in POp and PTr, and that (with the exception of the time condition in PTr), all induced above chance activity in these regions. This leaves open the question of whether

these different levels were accompanied by qualitatively different activity distributions within the region, or alternatively, whether they indexed similar patterns of activity that simply differed in response magnitude. If a brain region codes similarly for the three conditions of interest, the distribution of activity values (T-values) in these regions' voxels should be similar across the three conditions. However, if two conditions prompt a different organization of activity, the correlation between the activity maps in these conditions should be low. To assess similarity in population codes, we quantified the pairwise degree of similarity between each two conditions (Action vs. Time, Space vs. Action, Time vs. Space). These three correlations were determined separately for each participant and a group analysis was conducted on these correlation values (the analysis was conducted on the surface representation of each participant and thus avoided inclusion of white-matter areas). As shown in Fig. 5, the analysis indicated that POp coded all three conditions in a highly similar manner. In contrast, in PTr, the space and action conditions were coded similarly, but both of these conditions differed from the time condition.

A 2 (Region) × 3 (Correlation: Action vs. Time, Space vs. Action, Time vs. Space) ANOVA of these pair-wise correlation values revealed a main effect of region, since correlations were higher on average in POp than in PTr (POp $M = .64$, SE = .07; PTr $M = .39$, SE = .1; $F(1,11) = 26.5$, $p<.001$). The ANOVA also revealed a reliable Condition × Region interaction ($F(2,22) = 3.9$, $p<.05$). ANOVAs conducted for each region separately revealed reliable differentiation between correlation structures in PTr ($F(2, 22) = 4.67$, $p<.02$), but no differences in POp ($F(2, 22)<1$, $p>.47$). Follow up t-tests showed that in PTr, the correlation between population codes in the space- and action–attention conditions (mean $r = 0.57$) showed greater similarity than between both the space and time conditions (mean $r = 0.28$; $T(11) = 2.4$, $p<.05$), and between the action and time conditions (mean $r = 0.38$, $T(11) = 2.52$, $p<.05$). For POp, none of the contrasts approached significance (all $ps>.25$). To verify the construct validity of the procedure we conducted the same analysis against data collected in the transverse temporal gyrus, an anatomical landmark subsuming the location of the human primary auditory cortex (Morosan et al., 2001). Activity patterns in this region were expected to be highly similar across orientation conditions, and this is what we found. All three pair wise correlations were high (Action vs. Time: mean $r = 0.72$, median = 0.82; Space vs. Action: mean $r = 0.79$, median = 0.90, Time vs. Space: mean $r = 0.72$, median = 0.80;). A repeated measures ANOVA revealed no difference between these correlation values ($F<1$; $p>0.55$).
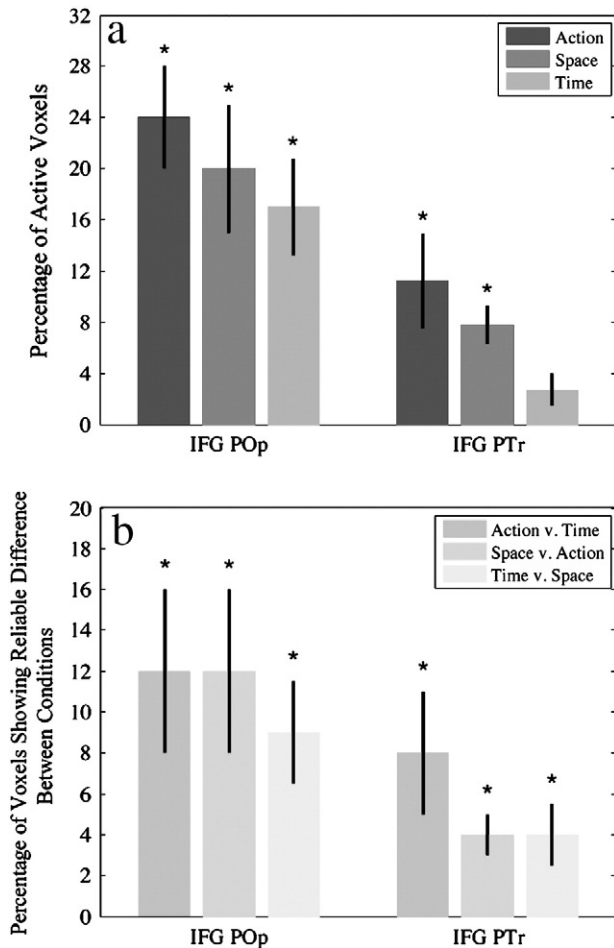

Fig. 4. Percentage of voxels showing (a) reliable activity relative to baseline and (b) reliable differentiation between conditions in left IFG Pars Opercularis (POp) and Pars Triangularis (PTr). Percentages that are reliably different from zero are marked with an asterisk. Error bars represent standard error of the mean across participants.
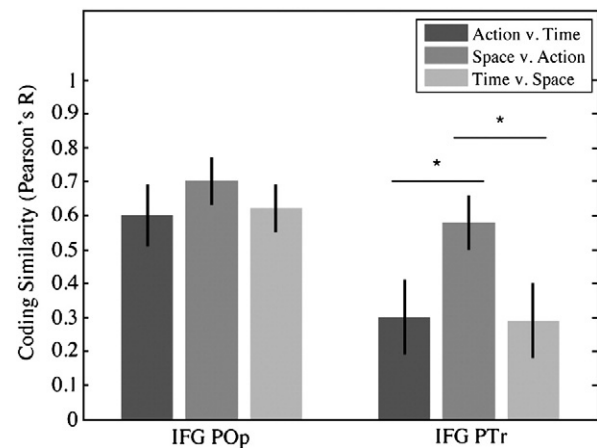

Fig. 5. Mean between-condition correlation between T values in the different conditions in left IFG Pars Opercularis (Pop) and Pars Triangularis (PTr). Reliable differences in correlation are marked with an asterisk. Error bars represent standard error of mean across participants.

Discussion

Our goal was to examine differences in the neural basis of natural language processing arising solely from the interpretations and information-seeking preferences of listeners. We found that engaging in different types of top-down attentional focus can greatly alter the brain activation associated with comprehension of identical auditory narratives.

Interpretation of language and the left inferior frontal gyrus

Our first hypothesis dealt with the roles of the posterior (POp) and anterior (PTr) portions of the left IFG on semantic language processing. Dissociations between the POp and PTr have been repeatedly documented in studies involving lexical and discourse-level semantics, with the emerging view that the PTr is involved in higher-level semantic processing, while the POp may be involved in both semantic and phonological processing (Bookheimer, 2002). Our finding that these two regions showed above baseline activity in almost all conditions (with the exception of time-attention in PTr) supports their general involvement in language processing. We also found differences in the mean percentage of active voxels across participants for each condition, suggesting that activity in both POp and PTr varies as function of top-down involvement. Interestingly, our analysis of the population code differences across conditions revealed highly similar activity patterns in POp. In contrast, activity patterns in PTr were highly similar for action–attention and space–attention, but both were different from time–attention. These novel findings suggest that POp operated similarly during the three conditions, whereas PTr had a different mode of organization during attention to time than attention to space or to actions. One possible interpretation is that different information seeking states of the listener induce a modulation of sensitivity in the POp, but do not affect its internal organization (akin to gain modulation). Our finding of between-condition differentiation in PTr, and in particular those identified by the population code analysis, may be consistent with prior work showing that the PTr has greater activation when story segments are more informative in context (Hasson et al., 2007). It is possible that action- and space-related information are more informative to listeners than time-related information (reflected in the above baseline activity), and also rely on similar interpretive processes (reflected in the higher correlation in population coding).

Another implication of these results relates to research showing that the posterior aspect of the left IFG is sensitive to action information in language (Tettamanti et al., 2005). One account explains the involvement of posterior IFG in action language processing by linking it to a low-level system that mediates non-linguistic functions and that is sensitive to both the observation and execution of action in humans (Molnar-Szakacs et al., 2005). In our study, all of the stimuli contained action information, and POp showed above-baseline activity for all participants. Our detailed analysis showed that a large proportion of voxels in the region demonstrated maximal activity in the action condition, suggesting that monitoring action information indeed modulates overall activity in POp. However, the fact that the activity organization in POp was highly similar across the three conditions suggests that there was not a unique mode of processing while monitoring action information. Understanding how attention processes during language can induce activity changes while at the same time maintaining the topology of activation patterns is an important topic for future work.

Spatial and temporal language and the lateral temporal cortex

We also investigated whether top-down interpretive processing could dissociate between the monitoring of spatial and temporal information (or action and temporal information). In the whole brain analysis, we found greater activation in the left STS during action and space attention than during time attention. Previous studies have suggested a role for the posterior STS in interpretation of language about motion. Dick et al. (2009) identified this region as being active during the processing of meaningful gestures, and Deen and McCarthy (2010) found increases when processing short stories that contained references to biological motion. The effect of our top-down attention manipulation in this region suggests that this content-specific activity can be modulated by interpretive processes, such that activity is increased when focusing on aspects of a story that involve movement (like actions and spaces). Focusing on time-related information does not appear to load on this region in the same way.

We also found greater activity in the left STG during attention to actions and space than to time, and greater activity in the right STG during attention to space than to time. This pattern of activity is consistent with findings associating these two regions with high level discourse comprehension. Hasson et al. (2007) found that increased activity in these regions is associated with correct subsequent memory for story content. This finding is consistent with our interpretation of the differences in PTr activity, suggesting that focusing on space- and action-related aspects of narrative was more informative as to the overall story.

There are several possibilities for why attention to space and actions resulted in greater activity in the temporal lobe than did attention to time. First, this could suggest a dissociation between space- and time-related neural processing. While time-as-space is a useful conceptual metaphor, there is no *a priori* reason to suppose that such high-level analogies between domains are realized in neural organization. At the same time, although activity in inferior frontal and temporal regions does differ for these two conditions, this does not exclude the possibility that parietal representations of space and time are the same (see Walsh, 2003). While this was not the focus of the current study, future work may reveal that certain systems code spatial and temporal information in a similar way, whereas frontal and temporal regions do not.

Another possibility is that maintaining temporal and spatial information does rely on a single system, but that the uni-dimensionality of a time line (vs. the two- or three-dimensionality of a spatial map) results in simpler updating and contextualizing of temporal "locations" and requires fewer resources (du Boisgueheneuc et al., 2006). Behavioral research suggests that when people attend to spatial information during narrative comprehension, they engage their spatial working memory. Specifically, they construct mental representations that are bound in space. For example, objects that are described as being occluded from protagonists' viewpoint are less accessible during narrative comprehension (Horton and Rapp, 2003), and objects in the same room as the protagonist are more accessible (Bower and Morrow, 1990). This sort of working memory may not be engaged by maintaining temporal information. This question can be explored by future work in which the relative degree of temporal and spatial change is parametrically manipulated to examine whether there are neural systems that respond similarly to increases in complexity across both domains. It may also be fruitful to compare patterns of activation when spatial language is restricted to terms that refer to spatial relationships and general locations, rather than specific places.

Based on the finding that both space- and action-related attention resulted in several regions of greater activity than time-related attention, but no regions showed the reverse relationship, it is interesting to consider the possibility that attention to the temporal relationships in a narrative was occurring in all conditions. Attention to temporal relationships in a narrative could be considered a default condition, onto which the space- and action-attention conditions added additional demands.

Deactivation in language processing

All three experimental conditions evoked patterns of deactivation that were highly similar to those reported previously in studies of narrative processing. Deactivation patterns are thought to reflect a suspension of 'default mode' operations that take place during rest, and have been documented in several studies (see Buckner et al., 2008 for meta-analysis). Prior work has shown that variations in the degree of deactivation appear to be related to depth of processing: the degree of deactivation predicts subsequent memory (Hasson et al., 2007; Otten and Rugg, 2001) and is modulated by task difficulty (McKiernan et al., 2003). In the current study, between-condition differences, on both the whole-brain and ROI level, were manifested as modulations of above-baseline activation rather than below-baseline deactivation. While it is difficult to interpret null effects, this pattern is consistent with the idea that between-condition differences did not originate in differential difficulty, but in modulation of interpretive, information-seeking processes.

Issues and future directions

It is widely believed that textual representations exist at multiple levels. These levels include a representation of the surface level features, as well as a high-level representation of the events being described, called the *situation model* (van Dijk and Kintsch, 1983). In this study, we attempted to measure the effects of changing the high-level interpretation of text by keeping the surface level features the same. We inferred that differences in neural activity when attending to action, space, and time were caused by differences in this top-down interpretation. However, we cannot exclude the possibility that some amount of the variability between conditions was related to differences in the surface-level features of the action-, space-, and time-related information in our stories. This relationship between textual properties and cognitive processes is itself a large and fundamental question in discourse processing, to which there is no obvious answer. Future work could attempt to develop a metric whereby surface-level features related to different types of information can be equated with the level of cognitive processing or type of neural activity involved in comprehending them. Such a metric may need to consider things such as number of syllables and words, behavioral measurements of memory and reading times, as well as introspective reports. One implication of the present work is that the internal motivations and focus of the listener may also play an important role in determining the relationship between textual features and neural activity.

Supplementary materials related to this article can be found online at doi:10.1016/j.neuroimage.2011.01.003.

References

Beilock, S.L., Lyons, I.M., Mattarella-Mincke, A., Nusbaum, H.C., Small, S.L., 2008. Sports experience changes the neural processing of action language. Proc. Natl Acad. Sci. 105, 13269–13273.

Bookheimer, S.Y., 2002. Functional MRI of Language: new approaches to understanding the cortical organization of semantic processing. Annu. Rev. Neurosci. 25, 151–188.

Boroditsky, L., Ramscar, M., 2002. The roles of body and mind in abstract thought. Psychol. Sci. 13, 185–189.

Bower, G.H., Morrow, D.G., 1990. Mental models in narrative comprehension. Science 247, 44.

Buckner, R.L., Carroll, D.C., 2006. Self-projection and the brain. Trends Cogn. Sci. 11, 49–57.

Buckner, R.L., Andrews-Hanna, J.R., Schacter, D.L., 2008. The brain's default network. Ann. NY Acad. Sci. 1124, 1–38.

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

Dick, A.S., Goldin-Meadow, S., Hasson, U., Skipper, J.I., Small, S.L., 2009. Co-speech gestures influence neural activity in brain regions associated with processing semantic information. Hum. Brain Mapp. 30, 3509–3526.

Deen, B., McCarthy, G., 2010. Reading about the actions of others: biological motion imagery and action congruency influence brain activity. Neuropsychologia 48, 1607–1615.

du Boisgueheneuc, F., Levy, R., Volle, E., Seassau, M., Duffau, H., Kinkingnehun, S., Samson, Y., Zhang, S., Dubois, B., 2006. Functions of the left superior frontal gyrus in humans: a lesion study. Brain 129, 3315–3328.

Duvernoy, H., 1991. The Human Brain: Surface, Three-Dimensional Sectional Anatomy and MRI. Springer-Verlag, Vienna, Austria.

Ferstl, E.C., Neumann, J., Bogler, C., von Cramon, D.Y., 2008. The extended language network: a meta-analysis of neuroimaging studies on text comprehension. Hum. Brain Mapp. 29, 581–593.

Ferstl, E.C., von Cramon, D.Y., 2007. Time, space and emotion: fMRI reveals content specific activation during text comprehension. Neurosci. Lett. 427, 159–164.

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, B., Dale, A.M., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron 33, 341–355.

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

Fischl, B., van der Kouwe, A., Destrieux, C., Halgreen, E., Segonne, F., Salat, D.H., Busa, E., Seidman, L.J., Goldstein, J., Kennedt, D., et al., 2004. Automatically parcellating the human cerebral cortex. Cereb. Cortex 14, 11–22.

Forman, S.D., Cohen, J.D., Fitzgerald, M., Eddy, W.F., Mintun, M.A., Noll, D.C., 1995. Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. Magn. Reson. Med. 33, 636–647.

Friese, U., Rutschmann, R., Raabe, M., Schmalhofer, F., 2008. Neural indicators of inference processes in text comprehension: an event-related functional magnetic resonance imaging study. J. Cogn. Neurosci. 20, 2110–2124.

Gough, P.M., Nobre, A.C., Devlin, J.T., 2005. Dissociating linguistic processes in the left inferior frontal cortex with transcranial magnetic stimulation. J. Neurosci. 25, 8010–8016.

Hasson, U., Nusbaum, H., Small, S.L., 2007. Brain networks subserving the extraction of sentence information and its encoding in memory. Cereb. Cortex 17, 2899–2913.

Hasson, U., Nusbaum, H., Small, S.L., 2009. Task-dependent organization of brain regions active during rest. Proc. Natl Acad. Sci. USA 206, 10841–10846.

Heim, S., Eickhoff, S.B., Amunts, K., 2008. Specialisation in Broca's region for semantic, phonological, and syntactic fluency? Neuroimage 40, 1362–1368.

Heller, R., Golland, Y., Malach, R., Benjamini, Y., 2007. Conjunction group analysis: an alternative to mixed/random effect analysis. Neuroimage 37, 1178–1185.

Horton, W.S., Rapp, D.N., 2003. Occlusion and the accessibility of information in narrative comprehension. Psychon. Bull. Rev. 10, 104–109.

Hugdahl, K., Thomsen, T., Ersland, L., Rimol, L.M., Niemi, J., 2003. The effects of attention on speech perception: an fMRI study. Brain Lang. 85, 37–48.

Kemmerer, D., 2005. The spatial and temporal meanings of English prepositions can be independently impaired. Neuropsychologia 43, 797–806.

Lakoff, G., Johnson, M., 1980. Conceptual metaphor in everyday language. J. Philos. 77, 453–486.

Lei, M., Ho-Ling, L., Hua, J., Yen-Bee, N., Chongde, L., 2006. Passive reactivation of background information from long-term memory during reading. NeuroReport 17, 1887–1891.

McKiernan, K.A., Kaufman, J.N., Kucera-Thompson, J., Binder, J.R., 2003. A parametric manipulation of factors affecting task-induced deactivation in functional neuroimaging. J. Cogn. Neurosci. 15, 394–408.

Mitchell, R.L.C., Elliott, R., Barry, M., Cruttenden, A., Woodruff, P.W.R., 2003. The neural response to emotional prosody, as revealed by functional magnetic resonance imaging. Neuropsychologia 41, 1410–1421.

Molnar-Szakacs, I., Iacoboni, M., Koski, L., Mazziotta, J.C., 2005. Functional segregation within pars opercularis of the inferior frontal gyrus: evidence from fMRI studies of imitation and action observation. Cereb. Cortex 15, 986–994.

Morosan, P., Rademacher, J., Schleicher, A., Amunts, K., Schormann, T., Zilles, T., 2001. Human primary auditory cortex: cytoarchitectonic subdivisions and mapping into a spatial reference system. Neuroimage 13, 684–701.

Noll, D.C., Cohen, J.D., Meyer, C.H., Schneider, W., 1995. Spiral k-space MRI imaging of cortical activation. J. Magn. Reson. Imaging 5, 49–56.

Otten, L.J., Rugg, M.D., 2001. Task-dependency of the neural correlates of episodic encoding as measured by fMRI. Cereb. Cortex 11, 1150–1160.

Paulesu, E., Goldacre, B., Scifo, P., Cappa, S.F., Gilardi, M.C., Castiglioni, I., Perani, D., Fazio, F., 1997. Functional heterogeneity of left inferior frontal cortex as revealed by fMRI. NeuroReport 8, 2011–2016.

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.

Pulvermuller, F., Lutzenberger, W., Preissl, H., 1999. Nouns and verbs in the intact brain: evidence from event-related potentials and high-frequency cortical responses. Cereb. Cortex 9, 497–506.

Saad, Z.S., Glen, D.R., Chen, G., Beauchamp, M.S., Desai, R., Cox, R.W., 2009. A new method for improving functional-to-structural alignment using local Pearson correlation. Neuroimage 44, 839–848.

Saad, Z.S., Reynolds, R.C., Argall, B., Japee, S., Cox, R.W., 2004. SUMA: an Interface for Surface-based Intra- and Inter-Subjects Analysis with AFNI. IEEE International Symposium on Biomedical Imaging, Arlington, VA, pp. 1510–1513.

Saygin, A.P., McCullough, S., Alac, M., Emmorey, K., 2010. Modulation of BOLD response in motion-sensitive lateral temporal cortex by real and fictive motion sentences. J. Cogn. Neurosci. 22, 2480–2490.

Speer, N.K., Reynolds, J.R., Swallow, K.M., Zacks, J.M., 2009. Reading stories activates neural representations of visual and motor experiences. Psychol. Sci. 20, 989–999.

Sridharan, D., Levitin, D.J., Chafe, C.H., Berger, J., Menon, V., 2007. Neural dynamics of event segmentation in music: converging evidence for dissociable ventral and dorsal networks. Neuron 55, 521–532.

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.

Sweetser, E., 1992. English metaphors for language: motivations, conventions, and creativity. Poetics Today 13, 705–724.

Tettamanti, M., Buccino, G., Saccuman, M.C., Gallese, V., Danna, M., Scifo, P., Fazio, F., Rizzolatti, G., Cappa, S.F., Perani, D., 2005. Listening to action-related sentences activates fronto-parietal motor circuits. J. Cogn. Neurosci. 17, 273–281.

Therriault, D.J., Rink, M., Zwaan, R.A., 2006. Assessing the influence of dimensional focus during situation model construction. Mem. Cogn. 34, 78–89.

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.

Walsh, V., 2003. A theory of magnitude: common cortical metrics of time, space and quantity. Trends Cogn. Sci. 7, 483–488.

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

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.

Zwaan, R.A., Magliano, J.P., Graesser, A.C., 1995. Dimensions of situation model construction in narrative comprehension. J. Exp. Psychol. Learn. Mem. Cogn. 21, 386–397.