

Deconstructing Events: The Neural Bases for Space, Time, and Causality

Alexander Kranjec^{1,2}, Eileen R. Cardillo², Gwenda L. Schmidt^{2,3},
Matthew Lehet², and Anjan Chatterjee²

Abstract

■ Space, time, and causality provide a natural structure for organizing our experience. These abstract categories allow us to think relationally in the most basic sense; understanding simple events requires one to represent the spatial relations among objects, the relative durations of actions or movements, and the links between causes and effects. The present fMRI study investigates the extent to which the brain distinguishes between these fundamental conceptual domains. Participants performed a 1-back task with three conditions of interest (space, time, and causality). Each condition required comparing relations between events in a simple verbal narrative. Depending on the condition, participants were instructed to either attend to the spatial, temporal, or causal characteristics of events, but between participants each particular event relation appeared in all three conditions. Contrasts compared neural activity during each condition against the remaining two and revealed how

thinking about events is deconstructed neurally. Space trials recruited neural areas traditionally associated with visuospatial processing, primarily bilateral frontal and occipitoparietal networks. Causality trials activated areas previously found to underlie causal thinking and thematic role assignment, such as left medial frontal and left middle temporal gyri, respectively. Causality trials also produced activations in SMA, caudate, and cerebellum; cortical and subcortical regions associated with the perception of time at different timescales. The time contrast, however, produced no significant effects. This pattern, indicating negative results for time trials but positive effects for causality trials in areas important for time perception, motivated additional overlap analyses to further probe relations between domains. The results of these analyses suggest a closer correspondence between time and causality than between time and space. ■

INTRODUCTION

Space, time, and causality hold a privileged place in the history of philosophy. Together, these abstract categories provide a basis for organizing our experiences. To make sense of a simple event, it is necessary to represent (1) the spatial relations among objects relative to the general topography of a scene, (2) the durations of actions or movements for particular objects relative to one another, and (3) the relations between causes and effects (or lack thereof). As such, space, time, and causality provide a general framework for organizing and structuring our knowledge about events as they unfold in the world. Attempts to describe and reconcile these domains represent a common thread linking ancient and early modern philosophies. For philosophers, questions regarding space, time, and causality have primarily related to ontology, or philosophical investigations concerning the underlying physical reality of these basic domains in and of themselves, and epistemology, or how knowledge is acquired and organized. Fundamental questions regarding the domains of space, time, and causality in philosophy often concern

the extent to which each has a reducible, independent structure and how they interact with one another, both in the world and the mind (Newton, 1687/2005; Hume, 1739/2003; Aristotle, 1999; Kant, 1781/1998; Locke, 1690/1995; Plato, 1977).

Complementing philosophical approaches, more recent cognitive and behavioral neuroscience methods have independently investigated the neural bases for *perceiving* the spatial, temporal, and causal components of simple arrays (Coslett, Wiener, & Chatterjee, 2010; Wencil, Coslett, Aguirre, & Chatterjee, 2010; Wiener, Turkeltaub, & Coslett, 2010; Eagleman et al., 2005; Fugelsang, Roser, Corballis, Gazzaniga, & Dunbar, 2005; Schubotz & von Cramon, 2001; Coull, Frith, Buchel, & Nobre, 2000; Colby, 1998; Milner & Goodale, 1995; Haxby et al., 1991). Importantly, experimental research has also explored the interaction of these three domains dynamically. Using what Michotte (1946/1963) first termed “launching events” (simple collisions between “billiard ball-like” stimuli), behavioral research, particularly in the area of infant cognition (Schlottmann, Allen, Linderroth, & Hesketh, 2002; Leslie, 1984), and more recent cognitive neuroscience approaches with adults (Straube & Chatterjee, 2010) have sought to define the cognitive and neural bases for separable spatial

¹Duquesne University, ²University of Pennsylvania, ³Hope College

and temporal components of perceived causality when viewing dynamic events. These studies tend to systematically vary the relative salience or strength of either spatial or temporal cues to determine how such manipulations influence the perception of causality.

However, the perception of spatial, temporal, and causal information can be distinguished from our *conception* of these basic categories of experience. For example, discriminating between a tone presented for either 100 or 6000 msec is different from knowing that a minute is shorter than an hour. Similarly, perceiving a causal relation between two balls during a simple collision event is different than understanding the causal mechanism for more complex events where the relation is not necessarily as concrete or transparent (e.g., understanding the germ theory of disease). The distinction between perceptual and conceptual levels of representation is particularly evident when information from each is at odds with one another. For example, if I sneeze and a nearby lamp in my office turns on, I may perceive a causal relation between these two events based on their spatial and temporal contiguity—although I know that sneezes do not turn on lights. That is, I experience the “perceptual illusion” of causality—or so it would seem. If, later in the day, my colleague informs me that our new office lamps are equipped with energy-efficient, motion-sensitive light switches, I can further revise my understanding of the relation between cause and effect with respect to these two events connected in space and time. Reasoning about events in such a flexible manner is made possible by our abilities to think abstractly.

Although studies of spatial, temporal, and causal concepts are in relatively short supply in cognitive neuroscience, work focusing on such abstract concepts is accumulating. At present, independent investigations of space and causality have received the most attention in this regard. For example, several investigations of spatial concepts (e.g., studies looking at the neural bases for spatial prepositions and categorical spatial relations) suggest that, in general, the left hemisphere, more than the right, processes abstract spatial relations across both verbal and nonverbal tasks and that, in particular, inferior parietal cortex and parts of dorsolateral pFC including inferior and posterior middle frontal gyri are important for representing spatial relational knowledge (Amorapanth, Widick, & Chatterjee, 2010; Noordzij, Neggens, Ramsey, & Postma, 2008; Emmorey et al., 2002; Damasio et al., 2001; Kosslyn, Thompson, Gitelman, & Alpert, 1998; Kosslyn et al., 1989; Kosslyn, 1987). In the domain of causality, nonverbal tasks have contrasted conceptual causality with perceptual causality, finding evidence for a dissociation between these distinct types of causal judgments, and a left hemisphere specialization for abstracting causal relations (Roser, Fugelsang, Dunbar, Corballis, & Gazzaniga, 2005). In particular, the medial frontal cortex appears to be involved (Fonlupt, 2003). Using a verbal task, Satpute and colleagues (2005) found that causal judgments, more than judgments about

semantic associations, increased activation in left dorsolateral cortex and right precuneus.

Neural studies examining temporal concepts are the least represented (Teuscher, McQuire, Collins, & Coulson, 2008; Kemmerer, 2005). Notably, these studies tend to investigate cross-domain mappings between time and space. For example, in a study investigating the relations between space, time, and the semantics of prepositions, Kemmerer (2005) found a double dissociation for impaired comprehension of spatial or temporal semantics in patients, wherein the left supramarginal gyrus was found to be important for representing the spatial meaning of prepositions. Regions important for representing temporal meanings were not as clearly defined. However, the left peri-sylvian cortex was implicated. The results suggest that the spatial and temporal meanings of locative prepositions are represented separately. These particular results, in conjunction with the relative scarcity of neural studies investigating temporal concepts in general, suggest that time may be more difficult to localize both functionally and anatomically as compared with other abstract concepts. It is not surprising that dissociating the unique neural components of spatial, temporal, and causal concepts should be difficult considering that several well-known current theories in the cognitive sciences point out that language frequently uses spatial and/or causal metaphors to represent other abstract concepts—especially time—and that these patterns of lexicalization reveal something important about the structure of our conceptual system (Casasanto, 2009; Wolpert, 2006; Boroditsky, 2000; Lakoff & Johnson, 1999).

Thus, it may be even more important to consider that research addressing the neural bases for abstract concepts like space, time, and causality, specifically as they relate to one another, is lacking (but see Speer, Reynolds, Swallow, & Zacks, 2009). Given the current general interest in interactions among these conceptual domains from cognitive linguistics, simulation theories, and other sensorimotor or embodied cognition accounts of semantic representation (Chatterjee, 2010; Casasanto, 2009; Zwaan, 2004; Glenberg & Kaschak, 2002; Talmy, 2000; Barsalou, 1999; Lakoff & Johnson, 1999; Jackendoff, 1996), the scarcity of neural data is less than ideal (Kranjec & Chatterjee, 2010). Therefore, one main goal of the current study concerns establishing a better understanding of how concepts of space, time, and causality share organization at the level of cognition and neural anatomy. Although philosophical investigations and experimental behavioral studies suggest that concepts of space, time, and causality may overlap (e.g., we think about durations [in time] in terms of length [in space]; an understanding of the asymmetrical nature of sequential relations is central to both causality and time) these concepts also clearly refer to distinct kinds of relations in the world.

Using fMRI and simple verbal materials, the present work attempts to both find overlap and delineate the boundaries between these fundamental conceptual domains in so far

as they might be independently represented in the brain. As discussed above, the neural basis for spatial and causal conceptualization has been looked at in isolation. Furthermore and despite much cognitive linguistic and experimental behavioral evidence, the idea that space and time share representational structure at the conceptual level has only begun to be investigated using neuroscience methods. With this in mind, two main objectives motivated the present study. First, and most generally, our study addresses the scarcity of neuroscience research on abstract conceptualization. Second, it specifically investigates space, time, and causality simultaneously—as a family so to speak—in recognition of the central and dynamic role these elemental abstract concepts play in representing events. As described in the Methods below, the present design aims at dissecting these three basic conceptual components of events. It does so by testing whether the spatial, temporal, and causal structure embedded within identical event narratives can be distinguished using functional imaging. That is, the present study concerns whether distinct forms of relational thinking are neurally dissociable *within* the same event.

METHODS

Stimuli

The question of how we structure events and relate one event to another with respect to their spatial, temporal, and causal characteristics is central to the design of the present study. Consider two simple events, first, with respect to their relative spatio-temporal properties: an ice cube melting at room temperature and a wine glass shattering to pieces. Compared with an ice cube melting, the shattering of a glass takes place over a greater extent of *space* but occurs over a lesser period of *time*. This example illustrates that particular spatial and temporal characteristics of simple events can be orthogonalized. However, with respect to “causality,” one would be hard pressed to find any clear causal relation between these two events:

the ice cube melted
the glass shattered

But if another simple event is inserted in between, a causal chain becomes apparent:

the ice cube melted
the waiter slipped
the glass shattered

This example illustrates that spatial, temporal, and causal judgments concerning any single event all involve a relation to another event. With this in mind, we varied the spatial, temporal, and causal relations between simple events to construct the stimuli used in the current study. The stimuli consisted of 252 sentences describing simple events like the ones in the example above. All sentences

were comprised of three words and took the form of “the [noun] [verbed]”. Forty-two blocks of six-sentence sequences were constructed. Each was organized around a general theme (see Appendix). In each block of thematically related sentences, the sequential order of the six sentences was fixed. During stimulus construction, in determining the placement of a sentence in the sequence of events, the spatial, temporal, and causal relations between contiguous pairs (that would serve as the basis for making comparisons) were varied. For the purposes of the present study, spatial comparisons between two adjacent events concerned their relative areas, and temporal comparisons their relative durations. Causal relations were restricted to nonintentional or mechanical ones; that is, causal relations that did not require participants to infer the mental states of goal-directed agents. In addition, 14 blocks of false font strings were constructed by converting a sample of the sentence stimuli into Wingdings. False font trials served as a baseline condition for the fMRI analyses (Kable, Lease-Spellmeyer, & Chatterjee, 2002). Balancing items on psycholinguistic variables (e.g., concreteness, frequency, etc.) was unnecessary as the sentences used across conditions of interest were identical (Table 1).

Participants

Eighteen paid volunteers (10 men, 8 women; mean age = 25 years; range = 21–32 years) from the University of Pennsylvania community participated. All were right-handed, native speakers of English with normal or corrected-to-normal vision. No participants had a history of neurological or psychiatric disorders. Each gave informed consent in accordance with the institutional review board of the University of Pennsylvania.

Experimental Design

Participants performed a 1-back task in a fixed-order blocked design. The instructions varied according to condition. There were four conditions: space, time, causality, and false font (baseline). A 5-sec instruction slide informing the participant of the current condition followed by a

Table 1. Example Stimuli and Probable Responses by Condition

	+Space?	+Time?	+Caused?
the burgers grilled	null	null	null
the wind gusted	yes	no	no
the leaf stirred	no	no	yes
the dog barked	yes	no	no
the squirrel froze	no	no	yes
the rain poured	yes	yes	no

1-sec blank screen began each block. There were six trials per block. Each trial was 5 sec separated by a 1-sec inter-trial interval (see Figure 1). The first trial of each block was treated as a null trial as no response was required. This left five critical trials for analysis in each block. Across all conditions, participants answered either yes or no by pressing one of two buttons. For critical trials in the space condition, participants judged whether the event described in the current trial would likely take place over a larger area compared with the preceding event. In the time condition, participants judged whether the current event would likely occur over a greater duration compared with the preceding one. In the causality condition, participants judged whether the current event was likely to be caused by the preceding one. The stimuli were designed such that judgments on a trial-by-trial basis required participants to determine the likelihood that two events were related in a particular way; depending on the block, attention was directed toward the spatial, temporal, or causal relations between them. There were no explicit correct answers for individual trials. The direction of responding (e.g., whether a participant judged a particular event as longer or shorter in duration as compared with another) was unrelated to the final analyses. Different conditions were intended only to constrain par-

ticipants to a domain-specific mode of thought (space, time, or causality), whereas the syntax (the [noun] [verbed]) and binary response (yes or no) was identical for all stimuli and conditions. For false font trials participants merely had to determine if a string of nonsense symbols was identical to a prior string. In this manner, false font trials did provide an absolute measure for accuracy and served a dual purpose as a baseline condition for fMRI analyses and as behavioral catch trials to ensure participants were sufficiently performing the task throughout the scanning procedure.

Participants saw each of the 252 sentences in either 1 of 14 space, 14 time, or 14 causality blocks, in addition to 14 false font blocks. The particular condition of interest (space, time, or causality) that a sentence/block was presented varied between participants. Pseudorandom lists were constructed such that across all subjects each sentence/block was presented approximately the same number of times in each of the space, time, or causality conditions. Both the order of condition blocks and the particular button press that served as a yes or no answer (left or right button), were counterbalanced between participants. The 56 total blocks were distributed across five runs. The first four runs had 12 blocks, whereas the last run consisted of eight blocks.

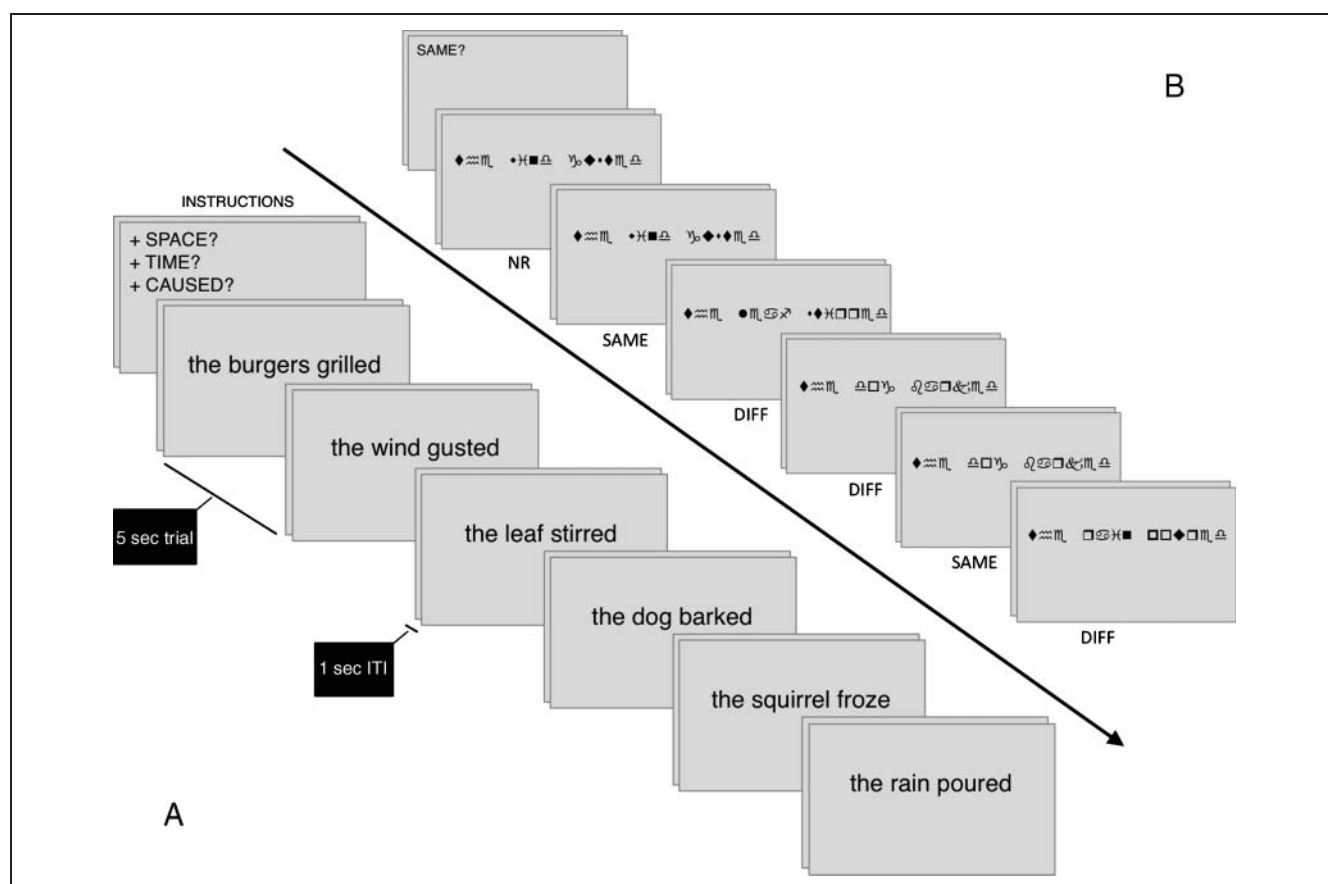


Figure 1. Experimental design for conditions of interest (A) and baseline task (B). Correct answers to an example block of the baseline task are shown.

Pilot Study

A pilot study ($n = 38$) using an identical design was conducted out of the scanner on a laptop computer. (Participants from the pilot study were not used in the fMRI study.) On the basis of demonstrating chance performance on the false font task, three participants' pilot data were excluded from all subsequent analyses. Of the remaining 35 participants, mean accuracy on false font trials was 91.2% and mean RT was 1454 msec. The pilot study revealed significant RT differences between the three conditions of interest ($M_{\text{time}} = 2745$ msec, $M_{\text{space}} = 2584$ msec, $M_{\text{causality}} = 2313$ msec; one-way ANOVA, $F(2, 68) = 51.52$, $p < .01$; post hoc tests, $t(1)$ all p s $< .01$). The RT differences were pervasive, such that removing particular blocks would not have equated conditions. Therefore, it was decided before carrying out the fMRI study that RT would be modeled as a covariate of noninterest during image analyses.

Imaging Procedures

Participants were familiarized with the task on a laptop computer before entering a 3-T Siemens Trio scanner using an eight-channel surface array head coil. For the experiment, the stimuli were back-projected onto a screen at the back of the scanner bore. In the scanner, participants lied down while viewing the stimuli through a mirror mounted on the head coil. BOLD-sensitive, T2* weighted functional images were acquired in 3-mm isotropic voxels using a gradient-echo, echoplanar pulse sequence (TR = 3000 msec, TE_{eff} = 30 msec). Forty 3-mm slices were acquired during each repetition, with each slice containing a 64° × 64° matrix within a 192 × 192 mm field of view. Head motion was minimized using foam padding, and the scanner performed both prospective (3-D Prospective Acquisition Correction, PACE) and retrospective motion correction on-line. High-resolution T1-weighted anatomical images were acquired for each participant using an MPRAGE pulse sequence (TR = 1620 msec, TE = 3 msec, TI = 950 msec). One hundred and sixty 1-mm slices were acquired, with each slice containing a 256° × 256° matrix within a 250 × 250 mm field of view. The participants performed no task during the first 6 sec of each functional scan (i.e., run) whereas the subsequent 6 sec represented the instruction screen for the first trial. This permitted 12 sec for steady-state magnetization to be attained. During the experiment, manual presses recording yes–no responses and RT were transmitted using a custom-designed fiber optic button box.

fMRI Data Analyses

Functional MRI data processing was performed using VoxBo software (www.voxbo.org) developed at the University of Pennsylvania's Center for Functional Neuroimaging.

After image reconstruction, the data were preprocessed in four main steps. First, the data were sinc-interpolated in time to correct for staggered slice acquisition (Aguirre, Zarahn, & D'Esposito, 1998). Second, using a procedure based on one used in the SPM analysis package (www.fil.ion.ucl.ac.uk/spm/), a six-parameter, least-square minimization motion correction algorithm was applied to realign all functional data to the first image acquired for each participant in a scanning session. Third, the data were automatically thresholded to exclude extraparenchymal voxels from subsequent analyses. The scan-wise global signals and power spectra were also derived and stored at this stage. Finally, parameters permitting normalization of the data to a standardized (Montreal Neurological Institute) space were automatically calculated.

A voxel-wise analysis was performed for each participant using a modified version of the general linear model for serially correlated error terms (Aguirre, Zarahn, & D'Esposito, 1997; Zarahn, Aguirre, & D'Esposito, 1997). This analysis identifies the voxels whose activity was significantly associated with the task covariate (i.e., the space, time, and causality conditions in contrast with the false font condition). The task covariate was a boxcar waveform convolved with an estimate of the BOLD hemodynamic transfer function empirically derived from the motor cortex in a large group of participants (Aguirre et al., 1998). For each participant, the voxels demonstrating a higher activity for space, time, and causality conditions relative to the false font baseline were identified. Random effects analyses were used for all group analyses and were measured by calculating the t value of the averaged time series of each voxel at the whole-brain level. Main effects for each condition identified significant voxels against baseline. For the subtractions, the distribution of t values across participants for a particular condition of interest was contrasted against the average of the remaining two to determine whether particular voxels showed preferential activation to the unique spatial, temporal, or causal attributes of an event relation (i.e., *space* > time and causality; *time* > space and causality; *causality* > space and time). False discovery rate (FDR) control was achieved with a procedure first described by Benjamini and Hochberg (1995) and adapted for brain image analysis by Genovese, Lazar, & Nichols (2002). The procedure allows one to identify a threshold that controls the expected proportion of false positives. In the current data set, selecting a false discovery rate (q value) of 0.05 yields a t threshold for each analysis of interest. This means that of the total number of voxels in an analysis with t values exceeding this threshold, the expected proportion of false positives is 5%. Cluster size was limited to those ≥ 10 , which in combination with FDR control makes for a conservative threshold. Overall, two analyses as described above were performed. In one, RT for each participant on a trial-by-trial basis was modeled as a covariate of noninterest. In the other, RT was not included in the model (see Imaging Results section).

RESULTS

Behavioral Results

Significant RT differences concordant with the pilot study were found between conditions of interest ($M_{\text{time}} = 2547$ msec, $M_{\text{space}} = 2334$ msec, $M_{\text{causality}} = 2097$ msec; one-way ANOVA, $F(2, 28) = 14.33$, $p = .02$; post hoc tests, $t(1)$, all p s $< .01$). Accuracy and RT on the false font task was also comparable to that for the pilot study (RT: $M_{\text{false font}} = 1231$ msec; accuracy: $M = 91.8\%$; range = 78.6–98.6%). Across all responses there were fewer yes trials (45.8%) as compared with no trials (54.2%) [$\chi^2(1, n = 3112) = 11.08$, $p < .001$]; however, the proportion of yes-and-no trials did not differ between conditions of interest [$\chi^2(2, n = 3112) = 0.77$, $p = .68$]. Additionally, the proportion of trials on which participants made a response compared with trials where no response was recorded was noted as an index of general alertness. These proportions did not vary significantly by condition, and overall, participants failed to make a response on fewer than 2% of all trials. Participant exclusion was determined on the basis of nine behavioral variables (four condition RT variables, four condition response–no response proportions, and false font accuracy) using a multivariate method for indentifying outliers (Mahalanobis distance, $p < .01$; Tabachnick & Fidell, 2001). Two participants were categorized as outliers, and their fMRI data were excluded from further image analyses. A third subject's fMRI data were not included in the group analyses as she self-reported doing a cognitively demanding secondary task during the baseline condition. Thus, a total of 15 participants' data (eight men and seven women; mean age = 25 years; range = 21–32 years) were included in the final fMRI analyses.

Imaging Results

Condition of Interest Subtractions

(See Figure 2 for statistical maps and Table 2 for anatomical descriptions, coordinate information, cluster sizes, and t values.)

Space. Relative to the other two conditions of interest, space trials elicited significant patterns of activation in bilateral occipitoparietal and frontal areas [$t \geq 3.197$, $q(\text{FDR}) = .05$]. Greater right-lateralized effects were observed in the superior frontal gyrus, middle frontal gyrus, and precuneus. Activations in inferior parietal lobules were relatively more prominent in the left hemisphere. Bilateral effects in the parietal lobule extended to the supramarginal gyrus in the left hemisphere and the angular gyrus in the right (or white matter directly undercutting these areas). The peak voxel for the *space* > time and causality subtraction was in the left supramarginal gyrus located within this largest cluster of active voxels in the inferior parietal lobule (see Figure 3).

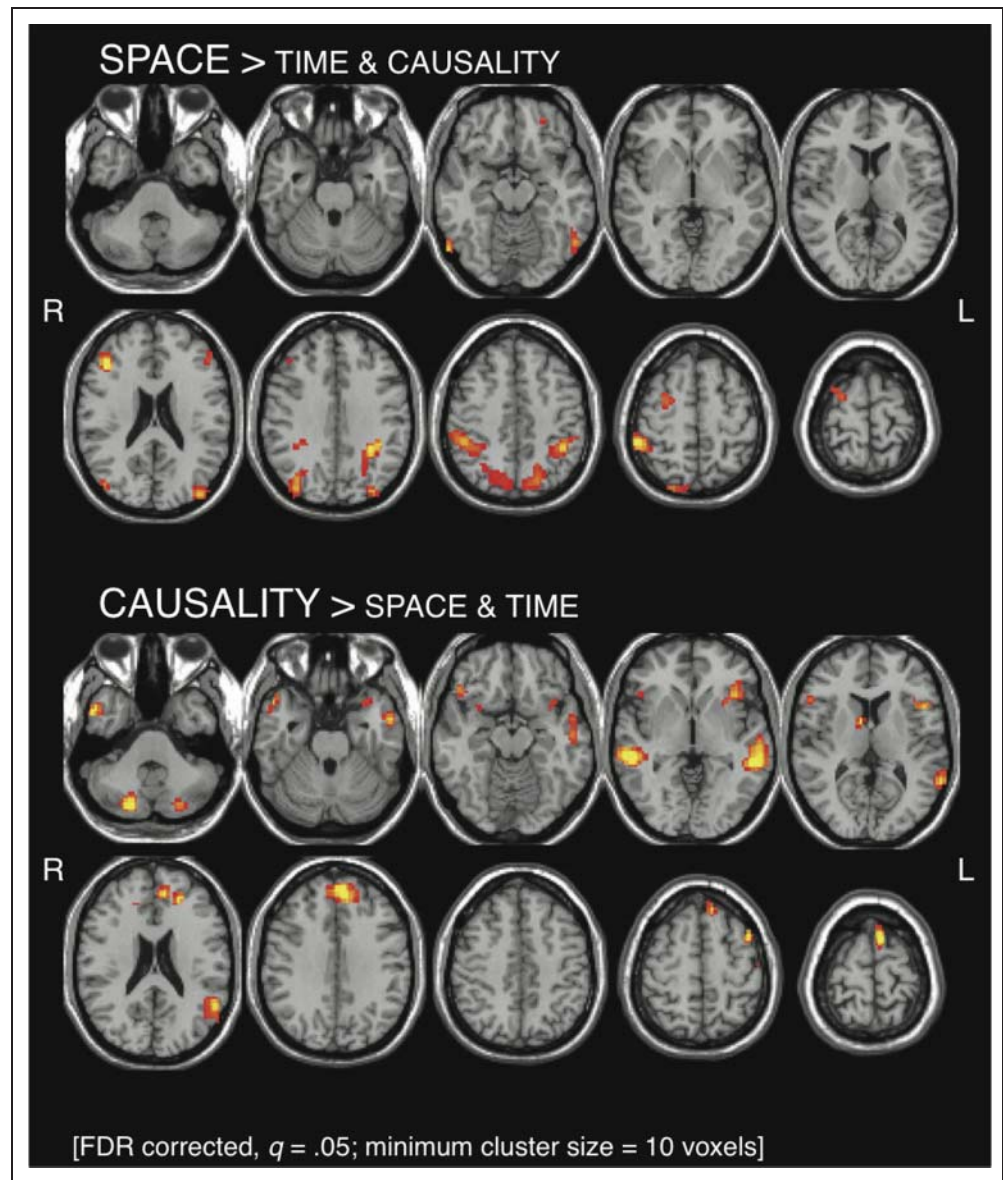
Causality. Most notably, causality more than space and time trials yielded significant cortical activations in bilateral frontal and middle temporal gyri, left insula, and a left-lateralized portion of the SMA [$t \geq 3.625$, $q(\text{FDR}) = .05$]. Although frontal activations included those in the right inferior frontal gyrus, left-lateralized activations were more robust. The peak voxel for the *causality* > space and time subtraction was located in a large cluster of active voxels in the left medial frontal gyrus. A cluster of significant voxels extending all the way from the anterior pole to posterior portions of the left middle temporal gyrus was the largest and remarkably cohesive (see Figure 3). A number of subcortical regions were also active for causality more than other trial types. These included effects in bilateral cerebellar areas and the left caudate of the BG.

Time. When contrasted with space and causality conditions, time trials produced no significant effects [$q(\text{FDR}) \leq .05$]. Because RT was modeled as a covariate of noninterest, initially we speculated that this feature of the analysis might have contributed to the lack of significant effects. RTs were longer for time trials relative to all other conditions. Therefore, when making direct contrasts, covarying for RT reduces the amount of activation for time trials while increasing the amount for space and causality (Kable et al., 2002). Also, we considered the possibility that mentally imaging the durations for time trial events might correlate with RT more so than for other conditions. For these reasons, we performed identical analyses between all conditions of interest, but without covarying for RT. However, although the pattern of results was remarkably similar for space and causality conditions, time still produced no significant effects. Further ROI analyses in SMA, BG, and the cerebellum (areas thought to be important for time perception; see references below) found no significant differences between time main effects and the other conditions.

Summary

Weighted contrasts compared each condition of interest (space, time, or causality) against the remaining two. There were several main findings. First, space trials recruited primarily bilateral frontal and occipitoparietal networks—areas traditionally associated with various forms of visuospatial processing (Berryhill, Fendrich, & Olson, 2009; Noordzij et al., 2008; Wu, Morganti, & Chatterjee, 2008; Wu, Waller, & Chatterjee, 2007; Emmorey et al., 2002; Damasio et al., 2001; Bisiach, 1999; Haxby et al., 1991; Vallar & Perani, 1986). Second, causality trials were associated with bilateral frontal activations and greater activity in the left middle temporal gyrus—areas previously found to underlie causal thinking and thematic role assignment, respectively (Wu et al., 2007; Satpute et al., 2005; Fonlupt, 2003). However, causality trials also produced activations in SMA, caudate, and cerebellum—cortical and subcortical sites associated with the perception of time at different timescales, suprasedond and subsecond durations,

Figure 2. Imaging results for condition of interest subtractions.



respectively (Coslett et al., 2010; Wencil et al., 2010; Wiener et al., 2010). Lastly, relative to space and causality conditions, time trials yielded no significant effects. The lack of effects for time trials was surprising, particularly because these trials were the most difficult as gauged by RT and self-reports in debriefings. Furthermore, causality trials, the easiest trials by the same measures, produced activations in regions important for time perception. This pattern of data, indicating negative results for time trials but positive effects for causality trials in areas important for time perception, motivated additional analyses intended to further probe relations between domains.

Overlap Relations between Time and Space or Causality

To further quantify relations between conditions, main effects for space, time, and causality (compared with base-

line; FDR corrected $q \leq .05$) were overlaid to create a combined overlap mask. For present purposes, the most interesting patterns regarding main effects concern those voxels that are either unique to a particular condition or overlap with only a single other condition. (For purposes of contrasting relations between representations of space, time, and causality, voxels implicated by all conditions are as uninformative as voxels activated by none of the conditions.) Therefore, we counted the total number of voxels unique to a particular main effect and the number involved in two-way overlaps between conditions. Of particular interest were the number of voxels activated by time trials that either overlapped with voxels activated by space *or* causality trials. Because no significant effects were found for the time versus space and causality subtraction, it was thought that an examination of main effects overlap in this manner could reveal something about conceptual

Table 2. Peak Activations by Condition Subtractions

<i>Anatomical Region</i>	<i>Brodmann's Area</i>	<i>Talairach</i>			<i>Size</i>	<i>t</i>
		<i>x</i>	<i>y</i>	<i>z</i>		
<i>Space > Time and Causality</i>						
Middle frontal gyrus (R)	4/6	27	-1	58	114	4.97
Middle frontal gyrus (R)	46	42	31	18	114	5.85
Inferior parietal lobule (R)	40	54	-39	54	297	5.45
Inferior temporal gyrus (R)	37	63	-65	-8	40	4.76
Precuneus (R)	19	36	-74	37	269	5.37
Middle frontal gyrus (L)	46	-48	33	18	21	3.72
Superior frontal gyrus (L)	11	-27	46	-16	10	3.99
Inferior parietal lobule (L)	40	-39	-39	38	394	6.15
Middle occipital gyrus (L)	37	-53	-62	-7	70	4.84
Superior occipital gyrus (L)	19	-39	-83	22	67	4.66
Posterior cingulate (L)	30	-21	-55	15	14	4.02
<i>Causality > Space and Time</i>						
Inferior frontal gyrus (R)	45	54	21	7	11	4.22
Inferior frontal gyrus (R)	13	30	12	-12	13	4.10
Middle temporal gyrus (R)	21	56	-29	-4	190	5.72
Anterior middle temporal gyrus (R)	21	53	4	-28	186	6.19
Cerebellum (pyramis) (R)	-	21	-77	-29	104	5.36
Superior frontal gyrus (L)	8	-9	38	53	21	5.37
Superior frontal gyrus (SMA) (L)	6/19	-6	15	63	62	5.47
Medial frontal gyrus (L)	9	-3	48	25	277	6.91
Middle frontal gyrus (L)	6	-42	15	52	39	5.00
Postcentral gyrus (L)	3	-54	-10	50	13	4.87
Insula (L)	13	-30	15	-6	198	5.05
Middle temporal gyrus (L)	22	-53	-38	-1	558	5.54
Caudate (L)	-	9	4	11	31	5.28
Cerebellum (pyramis) (L)	-	-24	-77	-31	37	5.08
<i>Time > Space and Causality</i>						
	Ø					

L = left; R = right.

relations between time and other domains. As displayed in Figure 4A, this analysis revealed that there were fewer overlapping voxels for space and time (S&T: 282 voxels) and space and causality (S&C: 267 voxels) compared with the numbers of voxels for causality and time (C&T: 1207 voxels). Figure 4B illustrates that for all voxels activated for time, the percentage of overlapping voxels activated for

both causality and time trials (C&T: 65.0%) is greater than the percentage of overlapping voxels activated on space and time trials (S&T: 15.2%; test of proportions, $p < .05$). Figure 5 illustrates the distribution of voxels activated by time trials overlapping with either space (pink) or causality (green) trials. The overall pattern suggests a left hemisphere bias. Most notable areas of overlap between causality

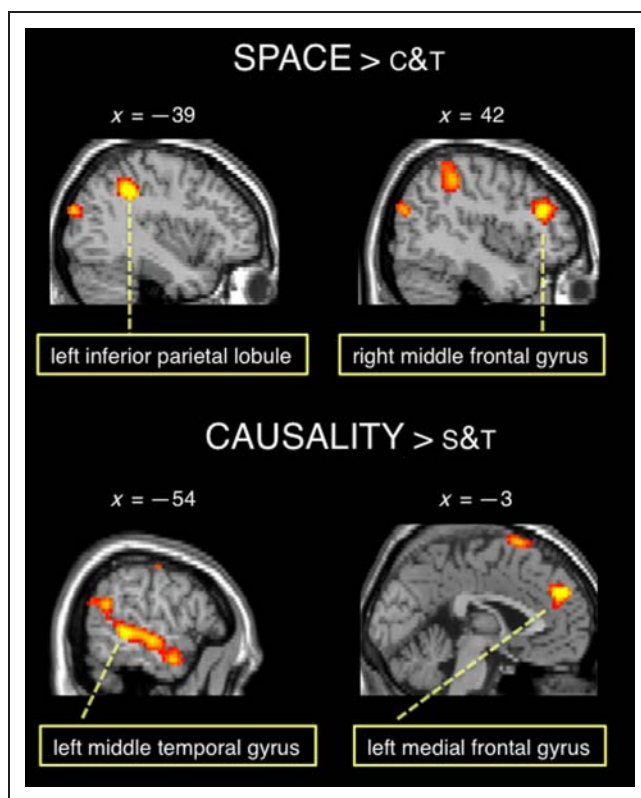


Figure 3. Sagittal views of main results (see Table 2) for condition of interest subtractions.

and time trials are in the left middle frontal gyrus and SMA. For space and time overlap, the most salient result is a negative one. There are few substantial areas of overlap, although some shared voxels for space and time in bilateral regions of the precuneus are noteworthy. Overall though, the overlap analyses suggest a much closer correspondence between time and causality than between time and space.

GENERAL DISCUSSION

Cognitive neuroscience has traditionally focused on the semantics of concrete objects (Forde & Humphreys, 1999; Caramazza & Shelton, 1998; Tranel, Damasio, & Damasio, 1997; Hillis & Caramazza, 1991; Warrington & Shallice, 1984), a trend that has carried over to more recent imaging studies (Thompson-Schill, Kan, & Oliver, 2006; Martin, Ungerleider, & Haxby, 2000; Chao, Haxby, & Martin, 1999; Cappa, Perani, Schnur, Tettamanti, & Fazio, 1998). Whereas much is currently known about how the brain categorizes things like faces, animals, and tools based on concrete features like *size*, *color*, and *shape*, much less is known about how abstract relations between events are neurally instantiated. Thus, we know relatively more about the “what” of semantic representation than we do about the “where,” “when,” and “why” of it. This more abstract level of repre-

sentation is particularly important for understanding event semantics and, more generally, the ability to think in terms of abstract relations may best explain “why we’re so smart” (Gentner, 2003). *Space*, *time*, and *causality* provide a natural framework for organizing the features of events. Only more recently has attention in cognitive neuroscience been paid to these abstract categories of experience (Amorapanth et al., 2010; Speer et al., 2009; Wu et al., 2007; Kemmerer, 1999, 2005, 2006; Satpute et al., 2005; Tranel & Kemmerer, 2004; Fonlupt, 2003; Chatterjee, 2001; Zacks & Tversky, 2001; Kemmerer & Tranel, 2000). Using fMRI, we tested whether the spatial, temporal, and causal structure embedded within identical narratives could be neurally distinguished *within* the same event relations. We found that spatial and causal relations were more easily segregated than temporal ones. However, time appeared to be more closely coupled with causality than space. To our knowledge, these three specific foundational components of events have not been looked at within an identical narrative structure before the current study.

We found that thinking about *space*, more than causality and time, activated bilateral frontal gyri, inferior parietal, and middle occipital regions, with a peak activation recorded in the left supramarginal gyrus. These activations are consistent with previous studies investigating spatial processing. Evidence in support of a general


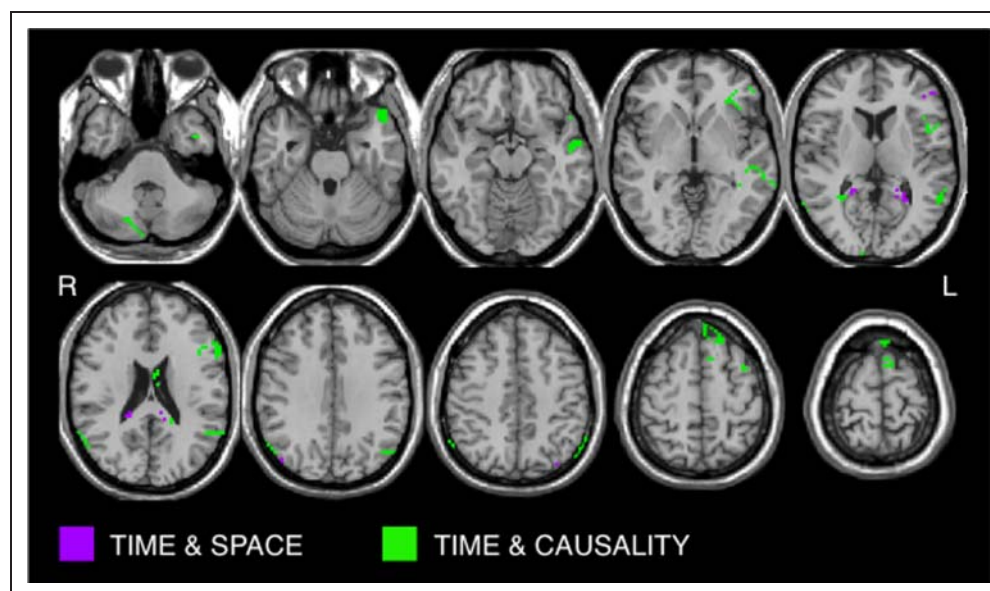


Figure 4. The absolute (A) and relative extent (B) of overlap between main effects for all conditions.

Figure 5. Distribution of voxels activated by time trials overlapping with either space (pink) or causality (green) trials.



right hemispheric bias for attending to spatial information was provided by frontal activations in primarily right-lateralized middle and superior gyri; anterorolandic activations—that is, in regions outside occipitoparietal areas more typically implicated in deficits of spatial attention (Vallar & Perani, 1986)—were right-lateralized. However, peak activation in the left supramarginal gyrus is consistent with work that finds a left hemisphere bias for representing categorical spatial relational information, particularly in inferior parietal areas (Amorapanth et al., 2010; Noordzij et al., 2008; Wu et al., 2007, 2008; Emmorey et al., 2002; Damasio et al., 2001; Kosslyn et al., 1989, 1998; Kosslyn, 1987). Remarkably, activations in middle occipital regions, along the “where” visual processing pathway (Haxby et al., 1991), suggest that, even during a verbal task such as the one used presently, spatial-specific information, as compared with other abstract concepts tested, was more closely associated with vision. This suggests that spatial relational thinking, although arguably supramodal in nature (Barsalou, 1999; Farah, Wong, Monheit, & Morrow, 1989; Farah, Hammond, Levine, & Calvanio, 1988), is more closely linked to vision than is thinking about causal and temporal relations.

Causality, more than space and time, was also associated with bilateral effects, but showed particularly large clusters of activation in the left middle temporal gyrus, and peaked in the left medial frontal gyrus. Frontal activations are common to the few fMRI studies investigating causality at the conceptual level. For example, Satpute et al. (2005) contrasted judgments of causal or associative relations between the same word pairs (e.g., “wind–erosion”) finding significant activation in the left dorsolateral pFC for causal, more than associative, judgments. Notably, this study also found activations in right precuneus for the same contrast, suggesting the possibility that abstracting the spatial

(event) structure evoked by certain word pairs could factor into thinking about causal relations during their task. Another fMRI study (Fonlupt, 2003) contrasted participants’ conceptual judgments of causality with their perception of causality using simple launching events. This study found increased activity in medial frontal cortex when participants were explicitly instructed to search for a causal relation, regardless of whether the launching event was actually construed as caused or uncaused. That is, simply thinking *about* causal relations activated frontal areas, even when the spatio-temporal details of the event did not result in the perception of causality. Moreover, passively perceiving caused events did not activate these areas more than passively perceiving noncaused events did.

Activations in middle temporal gyrus resonate with findings from a previous patient study from our laboratory. Wu et al. (2007) investigated thematic role assignment using tasks that required patients to map the thematic role (i.e., agent–recipient relations), or code *who does what to whom*, onto actors embedded within a very simple narrative structure. Patients had to either (1) match a sentence (e.g., “the circle kicked the square”) with a picture depicting the same simple event or (2) infer the consequences of similar actions involving agents and recipients in (A) verbal and (B) nonverbal tasks. Damage to middle temporal regions impaired performance on all three thematic relation tasks. These concordant results are sensible, given that assigning thematic role in such a fashion is essentially determining the cause and effect of an action.

Also noteworthy, the left SMA and bilateral cerebellar areas showed significant effects for causality more than other trial types. However, these activations are interesting primarily because both the motor and perceptual components of subsecond timing tasks have been localized to subcortical networks, including the cerebellum

and BG, and suprasedond timing tasks to SMA and pFC in a number of fMRI studies (Wencil et al., 2010; Wiener et al., 2010). That is, despite finding no significant activity for *time* relative to both space and causality, causality trials produced relatively larger effects in brain regions thought to be important for representing time. Although this result was somewhat unexpected, we sought to better understand it by means of examining the total number of voxels and patterns of overlap for each main effect relative to one another. The results of these analyses suggest that thinking about the temporal relations between simple events has more in common with thinking about their causal relations than either (time or causality) have in common with space. This is to say, the functional–anatomical overlap between time and causality suggest that these two domains are closely linked cognitively.

This finding is especially notable considering that some current theoretical accounts would predict a different pattern of results. For example, several cognitive models of cross-domain conceptual representation suggest that relatively abstract concepts (e.g., time) are structured in terms of more salient or relatively concrete concepts (like space) (Casasanto, 2009; Boroditsky, 2000; Lakoff & Johnson, 1999; Clark, 1973). More generally, spatial representations are thought to have a special involvement in structuring language and thought (Pinker, 2007; Chatterjee, 2001; Bloom, Peterson, Nadel, & Garrett, 1996). There is an especially fast growing literature demonstrating that we talk and think about temporal concepts using spatial abstractions (Kranjec & McDonough, 2011; Kranjec, Cardillo, Schmidt, & Chatterjee, 2010; Casasanto, 2009; Casasanto & Boroditsky, 2008; Nunez & Sweetser, 2006; Boroditsky, 2000). For example, *duration* is conceptualized in terms of *length*. The present task found little evidence in support of the idea that temporal concepts find grounding in neural areas associated with the representation of spatial relations. However, evidence that temporal thinking shares neural structure with causal thinking was found. Although investigations concerning the relations between causality and temporal concepts are less common than those between space and time, there are several reasons to think that causal concepts are particularly important for structuring language and thought in general, and that time in particular may play an important role in this.

The capacity to represent and invert the thematic roles of agents and recipients in relation to verbs is an important feature of language, allowing one to understand that *Man Bites Dog* is a newsworthy headline but *Dog Bites Man* is not. But thematic assignment also necessarily entails information about the beginning and end states of simple actions or events that depends on more basic representations of causal *and* temporal relations. Time may be thought to be even more basic than causality in this respect; to conceptualize causality—i.e., to comprehend that causes always precede effects—one must be able to per-

ceive duration and temporal order. A prior fMRI study from our laboratory, investigating the spatial and temporal components of perceptual causality, found that subjects more often used timing information when making causal judgments about Michotte-like launching events (Straube & Chatterjee, 2010). Sequence relations between time and causality can impact sentence comprehension; ERP evidence suggests that relations between two causally related events expressed in their actual temporal order of occurrence are easier to comprehend than when the same events are expressed in an inverted order (Munte, Schiltz, & Kutas, 1998). However, as *situation models* of discourse comprehension explain (Zwaan, 1999), temporal discontinuities in narratives are permitted (although with some cognitive cost) whereas a lack of causal cohesion can make a story very hard to follow or completely absurd.

Although our task used very simple verbal stimuli to focus on event processing at an intentionally abstract, conceptual level (the stimuli contained no explicit spatial, temporal, or causal language), a body of research concerned with complex text comprehension and higher-level inference making is relevant to the present discussion (for reviews, see Bornkessel-Schlesewsky & Friederici, 2007; Mason & Just, 2006; Mar, 2004). A general result found among these studies shows the involvement of fronto-medial and anterior-temporal regions during evaluative inferencing. Of particular interest to the current study, some research in this area has investigated neural regions associated with the spatial, temporal, and/or causal aspects of complex texts while participants read or listen to complex text (Cooper, Hasson, & Small, 2011; Speer et al., 2009; Ferstl & von Cramon, 2007; Virtue, Haberman, Clancy, Parrish, & Beeman, 2006; Ferstl, Rinck, & von Cramon, 2005; Mason & Just, 2004). These higher-level language comprehension tasks that attempt to contrast two or more domains tend to include materials with explicit spatial, temporal, or causal language to cue domain-specific processing (e.g., *the noises from “above” vs. the movie after “midnight”*). (In contrast, the present study uses identical language between conditions of interest *and* strips this language of domain-specific content using the standardized schematic construction the [noun] [verbed]). Despite differences in theoretical orientation and the materials used, some findings from this literature converge with those of the present study. For example, several studies demonstrate activations in superior and middle temporal gyri during “semantic integration” or when sentences are causally linked (Virtue et al., 2006; Mason & Just, 2004). Just and colleagues (Just et al., 2006; Just, Newman, Keller, McEleney, & Carpenter, 2004) found that high spatial imagery sentences activated the left intraparietal sulcus, hypothesizing a parietal “spatial network” for encoding spatial information relevant to situation models for text comprehension (Zwaan & Radvansky, 1998). Ferstl and colleagues (Ferstl & von Cramon, 2007; Ferstl et al., 2005) found bilateral fronto-parietal regions for spatial content during story comprehension and left prefrontal areas

for processing temporal information. Overall, the present results are consistent with much of the higher-level language comprehension literature showing increased language and inferencing demands while processing causal information (in middle temporal gyrus and frontal regions, respectively) and modality-specific effects in occipitoparietal regions while processing spatial information.

Conclusion

Thinking about causal relations, even when none are warranted, is arguably our most overarching conceptual bias. It may be peculiar to human beings and is argued by some to be responsible for driving the evolution of human reasoning and other species-specific behaviors like tool use (Tomasello, 1999). The tendency to construe two events as causally related is particularly powerful for humans. It is the same tendency that drives superstitious thinking and beliefs about other kinds of spurious correlations (Wolpert, 2006, 2009). In fact, inferring causal relations between events in a narrative may occur unintentionally and without awareness (Hassin, Bargh, & Uleman, 2002).

So where our laboratory finds a bottom-up bias toward using temporal (more than spatial) cues to represent the perception of causality (Straube & Chatterjee, 2010), the present work suggests a top-down bias where causal concepts may partly subsume temporal ones. However, both lines of research suggest that time, more than space, is linked to our judgments of causality. In relating time and causality, our findings suggest that time is more basic ontologically but causality is more salient phenomenologically. This characterization of the two domains as they relate to one another is mirrored in the way that time and causality are typically understood at perceptual and conceptual levels of representation. At the perceptual level, causality is portrayed as something (more abstract) that emerges from the (more concrete) spatio-temporal components of physical events (e.g., even a small temporal delay during a billiard ball “launching” eliminates the perception of a caused event). However, when thinking about event relations at a conceptual level, (more abstract) temporality is relatively vague and imprecise as compared with the (more concrete) causal relations that automatically link events in a narrative (e.g., causal inconsistencies in a narrative are very jarring, whereas temporal distortions are tolerated). Portraying time and causality as such predicts that neural areas important for time perception would be involved with understanding causality (consistent with an “embodied” view of abstract conceptualization), but that understanding causality would recruit relatively more resources when comparing the same events at a conceptual level. This is the conclusion supported by the pattern of observed neural results here. Sensorimotor areas for timing may be critical for representing causality, but, as evidenced by the lack of effects for the time condition sub-

traction and the relatively high degree of overlap between temporal and causal main effects activations, it is difficult not to experience causality when thinking about the temporal components of events.

Complementing previous patient work from our laboratory (Wu et al., 2007) and the neural literature on text comprehension (see Discussion above), the current study demonstrates that spatial and causal relations can be separated in the brain, with fronto-parietal and middle temporal cortical regions implicated as particularly important for representing these respective forms of relational thought. Temporal relations, however, were not as clearly segregated, but the present study also found perceptual motor areas activated for both spatial and causal relations. The pattern of results have interesting implications for simulation theories and other sensorimotor or embodied cognition accounts of semantic representation (Casasanto, 2009; Zwaan, 2004; Glenberg & Kaschak, 2002; Talmy, 2000; Barsalou, 1999; Lakoff & Johnson, 1999; Jackendoff, 1996). Looking at effects in neural areas subserving perception, we find that spatial relations were more strongly associated with activations in occipital areas critical for visual processing, and causal relations in the SMA, BG, and cerebellum—areas thought to be important for time perception at different scales (Wiener et al., 2010). Whereas prominent embodied views of abstract conceptualization suggest that we commonly talk and think about relatively abstract domains in terms of more concrete domains (Lakoff & Johnson, 1999), the present results suggest something slightly different. That is, at a neural level, it may be more useful to think about space and causality as our most central and overarching organizing concepts, with each being more or less embodied in a particular perceptual motor system, that is, vision and timing systems, respectively. So although space may be more concrete than time by some accounts (because of vision’s central place in structuring human experience), time may be more basic than a concept like causality according to others.

Much is known about how time *perception* is neurally instantiated. However, we know very little about how abstract *concepts*—especially time—are represented in the brain (Kranjec & Chatterjee, 2010). The present data indicate that our concepts of time and causality are not easily distinguished at a neural level. More generally, these data also suggest that embodied accounts should be careful when using terms like concrete and abstract when describing very basic categories of experience. In relating space, time, and causality to each other, a particular domain’s classification as relatively concrete or abstract can shift, depending on whether one is operating at a perceptual or conceptual level of representation. In this regard, the domain of time can be particularly indefinite. Saint Augustine (398/1961) famously wrote, “What then is time? If no one ask of me, I know; if I wish to explain to him who asks, I know not” (p. 301). Time may be very abstract when we need to think about it, but in some respects, nothing could be more basic at a neural level of representation.

APPENDIX

Example Stimuli

Block 1

the eels migrated
the oceans warmed
the iceberg melted
the seagull floated
the ship sank
the captain drowned

Block 2

the monkey jumped
the palm bent
the coconut bounced
the shell split
the milk oozed
the iguana napped

Block 3

the steeple chimed
the bats flapped
the termites chewed
the pulpit snapped
the priest tumbled
the bruise formed

Block 4

the village slept
the bird sang
the ape belched
the volcano erupted
the trees blazed
the sky darkened

Block 5

the umbrella tore
the clouds rained
the weed grew
the sidewalk cracked
the bikes raced
the puddle splashed

APPENDIX (*continued*)

Example Stimuli

Block 6

the shower poured
the tub filled
the sponge floated
the steam condensed
the mold grew
the bathroom smelled

Block 7

the taxi honked
the balloon burst
the ambulance wailed
the gun fired
the phone rang
the baby cried

Block 8

the farmer snored
the rooster crowed
the lightning hit
the silo toppled
the hay rotted
the cow chewed

Block 9

the thunder boomed
the dam broke
the river rose
the bank eroded
the metal rusted
the bridge collapsed

Block 10

the ice cube melted
the waiter slipped
the bottle shattered
the glass cut
the infection spread
the fever increased

Acknowledgments

The authors would like to thank Christine Watson and Olu Faseyitan for technical advice, Frank Pittenger for copyediting, and two anonymous reviewers for comments made on an earlier draft. This work was supported by the National Institutes of Health [RO1 DC004817, RO1 DC008779, T32 NS007413] and the National Science Foundation [subcontract under SBE0541957].

Reprint requests should be sent to Alexander Kranjec, Psychology Department, Duquesne University, 600 Forbes Avenue, 544 College Hall, Pittsburgh, PA 15282, or via e-mail: kranjeca@duq.edu.

REFERENCES

- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1997). Empirical analysis of BOLD fMRI statistics: II. Spatially smoothed data collected under the null-hypothesis and experimental conditions. *Neuroimage*, *5*, 199–212.
- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). The variability of human BOLD hemodynamic response. *Neuroimage*, *8*, 360–369.
- Amorapanth, P., Widick, P., & Chatterjee, A. (2010). The neural basis for spatial relations. *Journal of Cognitive Neuroscience*, *8*, 1739–1753.
- Aristotle. (1999). *Physics*. New York: Oxford University Press.
- Augustine, St. (398/1961). *The confessions*. New York: Penguin Books.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral & Brain Sciences*, *22*, 577–660.
- Benjamini, Y., & Hochberg Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society*, *57*, 289–300.
- Berryhill, M. E., Fendrich, R., & Olson, I. R. (2009). Impaired distance perception and size constancy following bilateral occipitoparietal damage. *Experimental Brain Research*, *194*, 381–393.
- Bisiach, E. (1999). Unilateral neglect and related disorders. In G. Denes & L. Pizzamiglio (Eds.), *Handbook of clinical and experimental neuropsychology* (pp. 479–495). East Sussex: Psychology Press.
- Bloom, P., Peterson, M., Nadel, L., & Garrett, M. (1996). *Language and space*. Cambridge, MA: MIT Press.
- Bornkessel-Schlesewsky, I., & Friederici, A. D. (2007). Neuroimaging studies of sentence and discourse comprehension. In G. Gaskell (Ed.), *Oxford handbook of psycholinguistics* (pp. 407–424). Oxford: Oxford University Press.
- Boroditsky, L. (2000). Metaphoric structuring: Understanding time through spatial metaphors. *Cognition*, *75*, 1–28.
- Cappa, S., Perani, D., Schnur, T., Tettamanti, M., & Fazio, F. (1998). The effects of semantic category and knowledge type on lexical-semantic access: A PET study. *Neuroimage*, *8*, 350–359.
- Caramazza, A., & Shelton, J. (1998). Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience*, *10*, 1–34.
- Casasanto, D. (2009). When is a linguistic metaphor a conceptual metaphor? In V. Evans & S. Pourcel (Eds.), *New directions in cognitive linguistics* (pp. 127–145). Amsterdam: John Benjamins.
- Casasanto, D., & Boroditsky, L. (2008). Time in the mind: Using space to think about time. *Cognition*, *106*, 579–593.
- Chao, L., Haxby, J., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, *2*, 913–919.
- Chatterjee, A. (2001). Language and space: Some interactions. *Trends in Cognitive Science*, *5*, 55–61.
- Chatterjee, A. (2010). Disembodying cognition. *Language and Cognition*, *2*, 79–116.
- Clark, H. (1973). Space, time, semantics and the child. In T. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 27–63). New York: Academic Press.
- Colby, C. L. (1998). Action oriented spatial reference frames in cortex. *Neuron*, *20*, 15–24.
- Cooper, E. A., Hasson, U., & Small, S. L. (2011). Interpretation-mediated changes in neural activity during language comprehension. *Neuroimage*, *55*, 1314–1323.
- Coslett, H. B., Wiener, M., & Chatterjee, A. (2010). Dissociable neural systems for timing: Evidence from subjects with basal ganglia lesions. *PLoS One*, *5*, e10324.
- Coull, J., Frith, C., Buchel, C., & Nobre, A. (2000). Orienting attention in time: Behavioral and neuroanatomical distinction between exogenous and endogenous shifts. *Neuropsychologia*, *38*, 808–819.
- Damasio, H., Grabowski, T. J., Tranel, D., Ponto, L. L., Hichwa, R. D., & Damasio, A. R. (2001). Neural correlates of naming actions and of naming spatial relations. *Neuroimage*, *13*, 1053–1064.
- Eagleman, D. M., Tse, P. U., Buonomano, D., Janssen, P., Nobre, A. C., & Holcombe, A. O. (2005). Time and the brain: How subjective time relates to neural time. *Journal of Neuroscience*, *25*, 10369–10371.
- Emmorey, K., Damasio, H., McCullough, S., Grabowski, T., Ponto, L. L., Hichwa, R. D., et al. (2002). Neural systems underlying spatial language in American Sign Language. *Neuroimage*, *17*, 812–824.
- Farah, M. J., Hammond, K. M., Levine, D. N., & Calvanio, R. (1988). Visual and spatial mental-imagery—Dissociable systems of representation. *Cognitive Psychology*, *20*, 439–462.
- Farah, M. J., Wong, A. B., Monheit, M. A., & Morrow, L. A. (1989). Parietal lobe mechanisms of spatial attention—Modality-specific or supramodal. *Neuropsychologia*, *27*, 461–470.
- Ferstl, E. C., Rinck, M., & von Cramon, D. Y. (2005). Emotional and temporal aspects of situation model processing during text comprehension: An event-related fMRI study. *Journal of Cognitive Neuroscience*, *17*, 724–739.
- Ferstl, E. C., & von Cramon, D. Y. (2007). Time, space and emotion: fMRI reveals content-specific activation during text comprehension. *Neuroscience Letters*, *427*, 159–164.
- Fonlupt, P. (2003). Perception and judgement of physical causality involve different brain structures. *Cognitive Brain Research*, *17*, 248–254.
- Forde, E., & Humphreys, G. (1999). Category-specific recognition impairments: A review of important case studies and influential theories. *Aphasiology*, *13*, 169–193.
- Fugelsang, J. A., Roser, M. E., Corballis, P. M., Gazzaniga, M. S., & Dunbar, K. N. (2005). Brain mechanisms underlying perceptual causality. *Cognitive Brain Research*, *24*, 41–47.
- Genovese, C. R., Lazar, N. A., & Nichols, T. E. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, *15*, 870–878.
- Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadows (Eds.), *Language in mind* (pp. 195–235). Cambridge, MA: The MIT Press.
- Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*, 558–565.
- Hassin, R., Bargh, J., & Uleman, J. (2002). Spontaneous causal inferences. *Journal of Experimental Social Psychology*, *38*, 515–522.

- Haxby, J. V., Grady, C. L., Horwitz, B., Ungerleider, L. G., Mishkin, M., Carson, R. E., et al. (1991). Dissociation of object and spatial visual processing pathways in human extrastriate cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *88*, 1621–1625.
- Hillis, A. E., & Caramazza, A. (1991). Category-specific naming and comprehension impairment: A double dissociation. *Brain*, *114*, 2081–2094.
- Hume, D. (1739/2003). *A treatise of human nature*. London: Dover.
- Jackendoff, R. (1996). The architecture of the linguistic-spatial interface. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and space* (pp. 1–30). Cambridge, MA: Cambridge University Press.
- Just, M. A., Newman, S. D., Keller, T. A., McEleney, A., & Carpenter, P. A. (2004). Imagery in sentence comprehension: An fMRI study. *Neuroimage*, *21*, 112–124.
- Kable, J. K., Lease-Spellmeyer, J., & Chatterjee, A. (2002). Neural substrates of action event knowledge. *Journal of Cognitive Neuroscience*, *14*, 795–804.
- Kant, I. (1781/1998). *The critique of pure reason*. Cambridge: Cambridge University Press.
- Kemmerer, D. (1999). “Near” and “far” in language and perception. *Cognition*, *73*, 35–63.
- Kemmerer, D. (2005). The spatial and temporal meanings of English prepositions can be independently impaired. *Neuropsychologia*, *44*, 1607–1621.
- Kemmerer, D. (2006). The semantics of space: Integrating linguistic typology and cognitive neuroscience. *Neuropsychologia*, *44*, 1607–1621.
- Kemmerer, D., & Tranel, D. (2000). A double dissociation between linguistic and perceptual representations of spatial relationships. *Cognitive Neuropsychology*, *17*, 393–414.
- Kosslyn, S. (1987). Seeing and imagining in the cerebral hemispheres: A computational analysis. *Psychological Review*, *94*, 148–175.
- Kosslyn, S., Koenig, O., Barrett, C., Cave, C., Tang, J., & Gabrieli, J. (1989). Evidence for two types of spatial representations: Hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 723–735.
- Kosslyn, S., Thompson, W., Gitelman, D., & Alpert, N. (1998). Neural systems that encode categorical versus coordinate spatial relations: PET investigations. *Psychobiology*, *26*, 333–347.
- Kranjec, A., Cardillo, E. R., Schmidt, G. L., & Chatterjee, A. (2010). Prescribed spatial prepositions influence how we think about time. *Cognition*, *114*, 111–116.
- Kranjec, A., & Chatterjee, A. (2010). Are temporal concepts embodied? A challenge for cognitive neuroscience. *Frontiers in Psychology*, *1*, doi: 10.3389/fpsyg.2010.00240.
- Kranjec, A., & McDonough, L. (2011). The implicit and explicit embodiment of time. *Journal of Pragmatics*, *43*, 735–748.
- Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh*. New York: Basic Books.
- Leslie, A. M. (1984). Spatiotemporal continuity and the perception of causality in infants. *Perception*, *13*, 287–305.
- Locke, J. (1690/1995). *An essay concerning human understanding*. Amherst, MA: Prometheus Books.
- Mar, R. A. (2004). The neuropsychology of narrative: Story comprehension, story production and their interrelation. *Neuropsychologia*, *42*, 1414–1434.
- Martin, A., Ungerleider, L., & Haxby, J. (2000). The sensory/motor model of semantic representation of objects. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 1023–1036). Cambridge, MA: MIT Press.
- Mason, R. A., & Just, M. A. (2004). How the brain processes causal inferences in text—A theoretical account of generation and integration component processes utilizing both cerebral hemispheres. *Psychological Science*, *15*, 1–7.
- Mason, R. A., & Just, M. A. (2006). Neuroimaging contributions to the understanding of discourse processes. In M. Traxler & M. A. Gernbacher (Eds.), *Handbook of psycholinguistics* (pp. 765–799). Amsterdam: Elsevier.
- Michotte, A. E. (1946/1963). *The Perception of Causality*. London: Methuen.
- Milner, A., & Goodale, M. (1995). *The visual brain in action*. New York: Oxford University Press.
- Munte, T. F., Schiltz, K., & Kutas, M. (1998). When temporal terms belie conceptual order. *Nature*, *395*, 71–73.
- Newton, I. (1687/2005). *The principia: Mathematical principles of natural philosophy*. London: Running Press.
- Noordzij, M. L., Neggers, S. F. W., Ramsey, N. F., & Postma, A. (2008). Neural correlates of locative prepositions. *Neuropsychologia*, *46*, 1576–1580.
- Nunez, R., & Sweetser, E. (2006). With the future behind them: Convergent evidence from Aymara language and gesture in the crosslinguistic comparison of spatial construals of time. *Cognitive Science*, *30*, 133–146.
- Pinker, S. (2007). *The stuff of thought: Language as a window into human nature*. New York: Viking.
- Plato. (1977). *Timaeus and Critias*. London: Penguin Books.
- Roser, M. E., Fugelsang, J. A., Dunbar, K. N., Corballis, P. M., & Gazzaniga, M. S. (2005). Dissociating processes supporting causal perception and causal inference in the brain. *Neuropsychology*, *19*, 591–602.
- Satpute, A. B., Fenker, D. B., Waldmann, M. R., Tabibnia, G., Holyoak, K. J., & Lieberman, M. D. (2005). An fMRI study of causal judgments. *European Journal of Neuroscience*, *22*, 1233–1238.
- Schlottmann, A., Allen, D., Linderoth, C., & Hesketh, S. (2002). Perceptual causality in children. *Child Development*, *73*, 1656–1677.
- Schubotz, R. I., & von Cramon, D. Y. (2001). Interval and ordinal properties of sequences are associated with distinct premotor areas. *Cerebral Cortex*, *11*, 210–222.
- Speer, N. K., Reynolds, J. R., Swallow, K. M., & Zacks, J. M. (2009). Reading stories activates neural representations of visual and motor experiences. *Psychological Science*, *20*, 989–999.
- Straube, B., & Chatterjee, A. (2010). Space and time in perceptual causality. *Frontiers in Human Neuroscience*, *4*, 10.
- Tabachnick, B., & Fidell, L. (2001). *Using multivariate statistics* (4th ed.). Needham Heights, MA: Allyn & Bacon.
- Talmy, L. (2000). *Towards a cognitive semantics: Concept structuring systems*. Cambridge, MA: MIT Press.
- Teuscher, U., McQuire, M., Collins, J., & Coulson, S. (2008). Congruity effects in time and space: Behavioral and ERP measures. *Cognitive Science*, *32*, 563–578.
- Thompson-Schill, S., Kan, I., & Oliver, R. (2006). Functional neuroimaging of semantic memory. In R. Cabeza & A. Kingstone (Eds.), *Handbook of functional neuroimaging of cognition* (pp. 149–190). Cambridge, MA: MIT Press.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge: Harvard University Press.
- Tranel, D., Damasio, H., & Damasio, A. R. (1997). A neural basis for the retrieval of conceptual knowledge. *Neuropsychologia*, *35*, 1319–1327.
- Tranel, D., & Kemmerer, D. (2004). Neuroanatomical correlates of locative prepositions. *Cognitive Neuropsychology*, *21*, 719–749.

- Vallar, G., & Perani, D. (1986). The anatomy of unilateral neglect after right-hemisphere stroke lesions—A clinical CT-scan correlation study in man. *Neuropsychologia*, *24*, 609–622.
- Virtue, S., Haberman, J., Clancy, Z., Parrish, T., & Beeman, M. J. (2006). Neural activity of inferences during story comprehension. *Brain Research*, *1084*, 104–114.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*, 829–854.
- Wencil, E. B., Coslett, H. B., Aguirre, G. K., & Chatterjee, A. (2010). Carving the clock at its component joints: Neural bases for interval timing. *Journal of Neurophysiology*, *104*, 160–168.
- Wiener, M., Turkeltaub, P., & Coslett, H. B. (2010). The image of time: A voxel-wise meta-analysis. *Neuroimage*, *49*, 1728–1740.
- Wolpert, S. (2006). *Six impossible things before breakfast: The evolutionary origins of belief*. London: Faber and Faber Limited.
- Wolpert, S. (2009). Cognition: Evolution does help to explain how minds work. *Nature*, *459*, 506
- Wu, D. H., Morganti, A., & Chatterjee, A. (2008). Neural substrates of processing path and manner information of a moving event. *Neuropsychologia*, *46*, 704–713.
- Wu, D. H., Waller, S., & Chatterjee, A. (2007). The functional neuroanatomy of thematic role and locative relational knowledge. *The Journal of Cognitive Neuroscience*, *19*, 1542–1555.
- Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin*, *127*, 3–21.
- Zarahn, E., Aguirre, G. K., & D'Esposito, M. (1997). Empirical analysis of BOLD fMRI statistics. *Neuroimage*, *5*, 179–197.
- Zwaan, R. A. (1999). Situation models: The mental leap into imagined worlds. *Current Directions in Psychological Science*, *8*, 15–18.
- Zwaan, R. A. (2004). The immersed experienter: Toward an embodied theory of language comprehension. *Psychology of Learning and Motivation: Advances in Research and Theory*, *44*, 35–62.
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, *123*, 162–185.