ELSEVIER

Cognition 75 (2000) 1–28

COGNITION

www.elsevier.com/locate/cognit

Metaphoric structuring: understanding time through spatial metaphors

Lera Boroditsky*

Department of Psychology, Stanford University, Jordan Hall, Bldg. 420, Stanford, CA 94305-2130, USA

Received 17 December 1998; received in revised form 22 September 1999; accepted 23 October 1999

Abstract

The present paper evaluates the claim that abstract conceptual domains are structured through metaphorical mappings from domains grounded directly in experience. In particular, the paper asks whether the abstract domain of time gets its relational structure from the more concrete domain of space. Relational similarities between space and time are outlined along with several explanations of how these similarities may have arisen. Three experiments designed to distinguish between these explanations are described. The results indicate that (1) the domains of space and time do share conceptual structure, (2) spatial relational information is just as useful for thinking about time as temporal information, and (3) with frequent use, mappings between space and time come to be stored in the domain of time and so thinking about time does not necessarily require access to spatial schemas. These findings provide some of the first empirical evidence for Metaphoric Structuring. It appears that abstract domains such as time are indeed shaped by metaphorical mappings from more concrete and experiential domains such as space. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Metaphoric structuring; Understanding time; Spatial metaphors

1. Introduction

How do we come to represent and reason about abstract domains like time, love, justice, or ideas? There are at least two interesting puzzles here. First, how do we

* Fax: +1-650-725-5699.
E-mail address: lera@psych.stanford.edu (L. Boroditsky)

PII: S0010-0277(99)00073-6

learn about abstract domains despite the dearth and vagueness of sensory information available about them? And second, how are we able to coordinate our mental representations of these domains enough to agree (at least some of the time) on the fairness of a decision, the strength of someone's love, or the worth of an idea? As a potential solution, Lakoff and Johnson (1980) proposed that the human conceptual system is structured around only a small set of experiential concepts – concepts that emerge directly out of experience and are defined in their own terms. These fundamental experiential concepts include a set of basic spatial relations (e.g. up/down, front/back), a set of physical ontological concepts (e.g. entity, container), and a set of basic experiences or actions (e.g. eating, moving). According to this view, all other concepts that do not emerge directly out of physical experience must be metaphoric in nature. Lakoff further proposes that these metaphoric, or abstract concepts are understood and structured through metaphorical mappings from a small set of fundamental experiential concepts (Lakoff & Johnson, 1980).

As evidence for this view, Lakoff and colleagues have pointed out that people often use metaphors to talk about abstract domains, and that in the majority of these conventional metaphors, language from a concrete domain is used to talk about the more abstract domain (Lakoff & Johnson, 1980; Lakoff & Kovecses, 1987). These conventional metaphors often reveal a particular source-to-target mapping, e.g. MIND IS A CONTAINER, and IDEAS ARE FOOD. To illustrate the IDEAS ARE FOOD schema, for example, readers might be reluctant to 'swallow Lakoff's claim' because they haven't yet gotten to 'the meaty part of the paper', or because they 'just can't wait to really sink their teeth into the theory'.

Such linguistic patterns suggest that there may be some systematic metaphoric relationships between abstract and concrete domains. However, the psychological reality of the proposed metaphoric relationships remains an open question. Lately, the metaphoric view of representation has been the subject of rigorous scrutiny and debate (Gibbs, 1996; Murphy, 1996, 1997). Two main criticisms have been put forward (Murphy, 1996, 1997). First, the majority of evidence in support of metaphoric representation has been of the purely linguistic form.[1] Strictly linguistic evidence can have only limited import for theories of mental representation as it would be scientifically imprudent to assume that patterns in language are necessarily a reflection of patterns in thought. Second, the theory has not been specified in enough detail to serve as a testable psychological model.

This paper aims to provide a more rigorous empirical treatment of metaphorical representation. To this end, it will be necessary to (1) propose one detailed account of how abstract concepts are learned, represented, and reasoned about, (2) provide psychological evidence in support of this proposal, and (3) show that the current evidence is not consistent with a plausible non-metaphoric account. The view proposed in this paper – the Metaphoric Structuring View – is derived from the

[1] For a discussion of the limitations of linguistic evidence see Murphy (1996).

Metaphoric Representation View set forth by Lakoff, but can be evaluated independently.[2]

1.1. The Metaphoric Structuring View

The Metaphoric Structuring View proposes that metaphors are used for organizing information within abstract domains. Those aspects of abstract domains that are evident from world experience may be represented in their own right. The job of the metaphor is to provide relational structure to an abstract domain by importing it (by analogy) from a more concrete domain. The mechanism for this type of metaphoric structuring may be the same as that used to understand analogies (Gentner & Wolff, 1997). Just like analogies, metaphors import the relational structure and not the surface features of the base domain to the target domain. When considering the IDEAS ARE FOOD metaphor, for example, we are not fooled into thinking that fried ideas are especially tasty or that thinking too much makes one fat. We can, however, infer that taking in a good idea can satisfy our intellectual appetite. In this case, the metaphor uses the relationship between food and hunger to describe the relationship between ideas and intellectual needs. The Metaphoric Structuring View proposes that metaphors provide relational structure to those domains where the structure may not be obvious from world experience.

This paper will focus on the abstract domain of time and consider whether time is structured through spatial metaphors. I will highlight a set of relational similarities between the conceptual domains of space and time, consider several explanations of how these similarities may have arisen, and describe three experiments that distinguish among these explanations. The described experiments will test the psychological validity of the claim that abstract conceptual domains such as time are structured by metaphorical mappings from more concrete experiential domains such as space.

1.2. Spatial metaphors for time

How is the domain of time learned, represented, and reasoned about? Certainly some elements of time are apparent in our experience with the world. From experience, we know that each moment in time only happens once, that we can only be in one place at one time, that we can never go back, and that many aspects of our experience are not permanent (i.e. faculty meetings are not everlasting, but rather begin and end at certain times). In other words, our experience dictates that time is a phenomenon in which we, the observer, experience continuous unidirectional change that may be marked by appearance and disappearance of objects and events.

[2] Because the Metaphorical Representation theory lacks the detail of an explicit psychological model, it is difficult if not impossible to extract a single, definitive, and testable statement of the theory. The Metaphorical Structuring View proposed in this paper is a view derived from the more general theory endorsed by Lakoff and colleagues, and aims to capture and make testable only some of the assumptions of this general theory. In constructing this more explicit psychological model of metaphorical structuring, it was necessary to make certain assumptions about the details of how metaphorical concepts might be acquired, represented, and used.

These aspects of conceptual time should be universal across cultures and languages. Indeed this appears to be the case. In order to capture the sequential order of events, time is generally conceived as a one-dimensional, directional entity. Across languages, the spatial terms imported to talk about time are also one-dimensional, directional terms such as *ahead/behind*, or *up/down*, rather than multi-dimensional or symmetric terms such as *shallow/deep,* or *left/right* (Clark, 1973; Traugott, 1978). Aspects of time that are extractable from world experience (temporally bounded events, unidirectional change, etc.) may be represented in their own right. However, there are many aspects of our concept of time that are not observable in the world. For example, does time move horizontally or vertically? Does it move forward or back, left or right, up or down? Does it move past us, or do we move through it? All of these aspects are left unspecified in our experience with the world. They are, however, specified in our language – most often through spatial metaphors. Whether we are looking *forward* to a brighter tomorrow, falling *behind* schedule, or proposing theories *ahead* of our time, we are relying on spatial terms to talk about time. The correspondences between space and time in language may afford us insight into how the domain of time is structured and reasoned about.

The Metaphorical Structuring View proposes that those aspects of time that are specified through spatial metaphors will be shaped by the metaphors used (see Boroditsky, 1999 for cross-linguistic evidence to this effect). The spatial schemas invoked by these metaphors will provide the relational information needed to organize events in time. This view can be formulated in several different strengths.[3] The weak version maintains that spatial metaphors play a role in shaping the domain of time. However, with frequent use, an independent representation is established in the domain of time, and so spatial schemas may no longer need to be accessed in thinking about time. This view is supported by recent findings showing that whereas novel metaphors are processed as on-line metaphorical mappings, conventional or frequently used metaphors tend to have stored meanings (Bowdle & Gentner, 1995, 1999). If a metaphorical mapping is frequently set up between two domains, the result of this mapping may eventually become stored in the target domain to avoid future costs of carrying out the same mapping.

Unlike the weak version, the strong version of Metaphorical Structuring maintains that spatial schemas are always necessary to think about time. On this view, relational information necessary to organize events is imported on-line from the domain of space and is not stored in the domain of time. Therefore, thinking about time requires accessing not only the temporal components of a scenario, but also the spatial schemas necessary to organize these temporal components. Experiments reported in this paper will test these two different formulations of the Metaphorical Structuring View and will attempt to establish (1) whether the domains of space and time are conceptually related, (2) whether spatial schemas can be used to understand time, and (3) whether spatial schemas are necessary to understand time.

[3] The weak/strong division proposed here is similar to that made by Murphy (1996) regarding the Metaphorical Representation view.

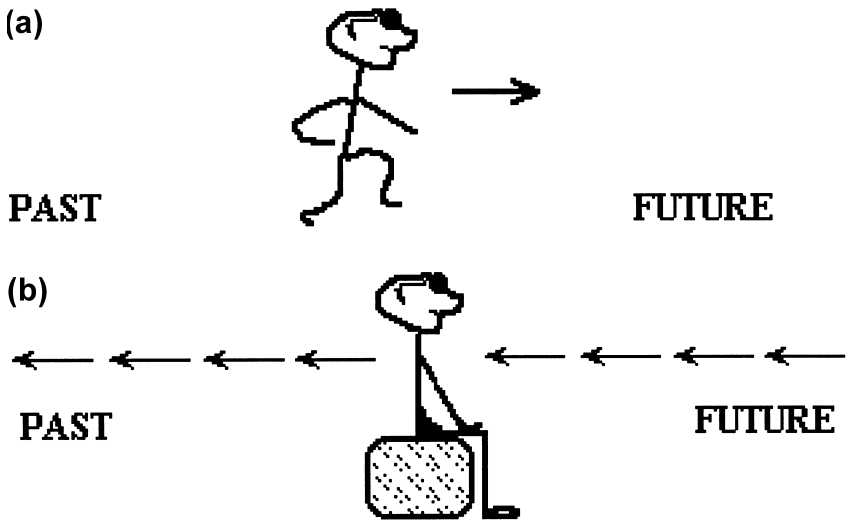

Fig. 1. (a) Schematic of the ego-moving schema used to organize events in time. (b) Schematic of the time-moving schema used to organize events in time.

1.3. The ego-moving and time-moving metaphors

Like most abstract domains, time can be described through more than one metaphor. This paper will focus on the event-sequencing aspect of conceptual time, that is, the way events are temporally ordered with respect to each other and to the speaker (e.g. 'The worst is *behind* us' or 'Thursday is *before* Saturday'). In English, two dominant spatial metaphors are used to sequence events in time (Clark, 1973; Lakoff & Johnson, 1980; McTaggart, 1908). The first is the ego-moving metaphor, in which the 'ego' or the observer's context progresses along the time-line toward the future as in 'We are coming up on Christmas' (see Fig. 1a). The second is the time-moving metaphor, in which a time-line is conceived as a river or a conveyor belt on which events are moving from the future to the past as in 'Christmas is coming up' (see Fig. 1b). These two metaphors lead to different assignments of front and back to a time-line (Clark, 1973; Fillmore, 1971; Lakoff & Johnson, 1980; McTaggart, 1908; Traugott, 1978).[4]

[4] This paper is focused mainly on the ego-moving and time-moving schemas, but several other ways of organizing events in time are possible. For example, what happens if both *time* and *observer* are stationary? One could imagine a system where events are organized according to their proximity to the observer. In this *ego-centric* system, front is assigned to parts of the time-line closest to the front of the observer. For items that are in front of the observer (in the future), those that are closer to the past (and also closer to the observer) are said to be in front. For items behind the observer, those that are closer to the future (and also closer to the observer) are said to be in front. Another system might simply rely on an absolute ordering of events from the past to the future. In this case, time may be like a number-line, ordered, but not moving. Also, many other metaphors are used to talk about time for purposes other than sequencing events (e.g. 'time is running out', 'time is a healer', 'time is money'). Further investigations of these and other metaphors will provide a more complete overall picture of the conceptual domain of time.

In the ego-moving metaphor, front is assigned to a future or later event (e.g. 'The revolution is *before* us'). In this example, the 'revolution' is a later or future event, and is said to be *before* because it is further along in the observer's direction of motion. An analogous schema exists for ordering objects in a line (see Fig. 2a). When an observer moves along a path, objects are ordered according to the direction of motion of the observer. In Fig. 2a, the dark can is said to be *in front* because it is further along in the observer's direction of motion.

In the time-moving metaphor, front is assigned to a past or earlier event (e.g. 'The revolution was over *before* breakfast'). Here, the 'revolution' is the earlier event, and is said to be *before* because it is further along in the direction of motion of time. Once again an analogous system exists for ordering objects in space (see Fig. 2b). When two objects (without intrinsic fronts) are moving, they are assigned fronts based on their direction of motion. In Fig. 2b, the light-colored widget is said to be *in front* because it is further along in the widgets' direction of motion.

A priori, there is no reason to believe that the linguistic distinction between the ego-moving and time-moving metaphors has any psychological implications regarding how these metaphors are processed. In the absence of further evidence, a more parsimonious view is that the distinction between these two different ways of talking about time is only language-deep. A skeptic might argue that linguistic evidence such as that provided by Lakoff and colleagues is at best an imaginative cataloguing of etymological relics with no psychological consequences. This skeptical position will be dubbed the Dubious View. The first challenge, then, is to empirically establish that the linguistic distinction between the ego-moving and time-moving metaphors has psychological consequences.

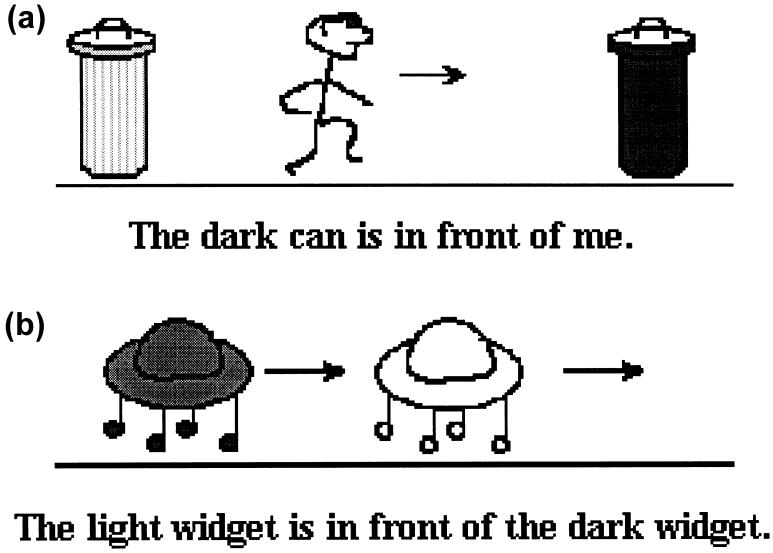

Fig. 2. (a) Sample ego-moving scenario used as a spatial prime in Experiment 1. (b) Sample object-moving scenario used as a spatial prime in Experiment 1.

1.4. Evidence for two distinct event-sequencing schemas

To investigate whether ego-moving and time-moving expressions are actually understood through different conceptual schemas, Gentner, Imai and Boroditsky (1999) measured processing time for temporal expressions presented either consistently or inconsistently with respect to either the ego-moving or the time-moving schema. They reasoned that if temporal expressions were processed as parts of globally consistent conceptual schemas, then processing should be fluent if the expressions are kept consistent to one schema (processing time should remain constant). If the schemas are switched, however, processing should be disrupted, and processing time should increase as it would take extra time to discard the old conceptual structure and set up a new one.

Participants were presented with a block of temporal statements that were either consistent with one schema, or switched between the ego-moving and time-moving schemas. For each statement (e.g. Christmas is six days before New Year's Day), participants were given a time-line of events (e.g. Past…New Year's Day…Future), and had to place an event (in this case Christmas) on the time-line. Response time data showed that switching schemas did indeed increase processing time.

In another study conducted at Chicago's O'Hare airport, participants were passengers not aware of being in a psychological study (Gentner et al., 1999). Participants were approached by the experimenter and asked a priming question in either the ego-moving form (Is Boston ahead or behind us in time?) or the time-moving form (Is it earlier or later in Boston than it is here?). After the participant answered, the experimenter asked the target question (So should I turn my watch forward or back?) which was consistent with the ego-moving form. The experimenter measured response times for the target question with a stopwatch disguised as a wristwatch. Once again, response times for consistently primed questions were shorter than for inconsistently primed questions. Switching schemas caused an increase in processing time. These results suggest that two distinct conceptual schemas are involved in sequencing events in time.

Converging evidence comes from studies that used a disambiguation paradigm (McGlone & Harding, 1998). Participants answered blocks of questions about days of the week phrased in either the ego-moving metaphor (e.g. 'We passed the deadline two days ago') or the time-moving metaphor (e.g. 'The deadline passed two days ago'). For each statement, participants indicated the day of the week on which the event in question had occurred or would occur. At the end of each block, participants read an ambiguous temporal statement such as 'The meeting originally scheduled for next Wednesday has been moved forward two days', and were asked to perform the same task. The 'moved forward' statement is ambiguous because it could be interpreted using one or the other schema to yield different answers. Participants in the ego-moving condition tended to disambiguate the 'moved forward' statement in an ego-moving-consistent manner (thought the meeting was on Friday), whereas participants in the time-moving condition tended to disambiguate in a time-moving-consistent manner (thought the meeting was on Monday).

These studies provide strong evidence for the psychological reality of two distinct, globally consistent schemas for sequencing events in time.

Since the linguistic distinction between the ego-moving and time-moving metaphors appears to be psychologically real, the Dubious View (that claims that differences between the two metaphors are only language-deep) can be rejected. However, the evidence described so far is not sufficient to conclude that time is understood as a metaphor from space. Just because space and time are talked about in a similar way, does not necessarily mean that they share deeper conceptual similarities. To claim that our understanding of time was shaped by our understanding of space, it is at least necessary to demonstrate that space and time have similar relational structure, and that spatial schemas could, in principle, be used to organize time.

So, can people use spatial schemas to think about time? If they can, then it should be possible to differentially prime particular spatial schemas to affect how people think about time. The following experiment examines whether making people think about spatial relations in a particular way might affect how they then think about time. First, participants answered several priming questions about spatial relations of objects in pictures. These pictures used either the ego-moving or the object-moving spatial schemas. Then, participants interpreted an ambiguous temporal statement such as 'Next Wednesday's meeting has been moved forward two days'. If the above statement is interpreted using the ego-moving schema, then *forward* is in the direction of motion of the observer, and the meeting should now fall on a Friday. In the time-moving interpretation, however, *forward* is in the direction of motion of time, and the meeting should now be on a Monday.[5]

If space and time do share some relational structure, then participants primed in the ego-moving spatial perspective should be able to reuse this perspective for time, and should thus think that the meeting will be on Friday. Participants primed in the object-moving perspective should prefer the time-moving interpretation and think that the meeting will be on Monday. However, if the domains of space and time do not share any relational structure, then spatial primes should have no effect on the way participants think about time.

2. Experiment 1

2.1. Method

2.1.1. Participants

Ninety-eight Stanford University undergraduates participated in this study as part of a course requirement.

[5] Because events are free to move in either direction in time (we can move meetings forward or back), the motion of events in itself cannot be used to organize the domain of time. This explains why statements like 'Wednesday's meeting has been moved forward two days' are ambiguous. These sorts of statements specify only that *events* are moving, but are indeterminate with regard to the motion of *observers* or *time*. This makes it is impossible to singly determine the direction that the event is being moved.

2.1.2. Materials and design

A two-page questionnaire was constructed. The first page contained four TRUE/FALSE priming questions. Priming questions were spatial scenarios consisting of a picture and a sentence description. The scenarios used either the ego-moving frame of reference (see Fig. 2a), or the object-moving frame of reference (see Fig. 2b), and were shown to different participants. These two frames of reference were predicted to map onto (and bias the use of) the ego-moving and time-moving perspectives in time, respectively. Half of the priming questions depicted motion to the left, and half to the right. Also, half of the questions were TRUE and half were FALSE. All of the objects depicted in the primes were frontless, and vertically symmetrical. Cans, trees, houses, and stools were used in the ego-moving primes, and widgets, wheels, carts, and blickets (an object similar to the widget) were used in object-moving primes. All of the objects depicted in the object-moving primes looked mobile (they were either round or had wheels).

On a separate page that immediately followed the primes, participants read an ambiguous temporal sentence (e.g. 'Next Wednesday's meeting has been moved forward two days') and indicated to which day the meeting had been rescheduled. A control group of participants responded to the above target sentence without having seen any primes. All participants also provided a confidence score for their answer to the target question on a scale of 1 to 5 (1, not at all confident; 5, very confident).

2.1.3. Procedure

Participants completed the two-page questionnaire individually with no time restrictions. The two pages of the questionnaire were imbedded in a large questionnaire packet which was distributed to an entire class in introductory psychology and contained many questions unrelated to this study. No special connection was made between the two pages of the questionnaire beyond that implied by their immediate adjacency. For the control group, only the page with the target question was included.

2.2. Results

As predicted by the Metaphoric Structuring View, people used primed spatial information to think about time. Overall, 71.3% of the participants responded in a prime-consistent manner. Of the participants primed in the ego-moving frame of reference, 73.3% thought that the meeting was on Friday, and 26.7% thought it was on Monday. Participants primed in the object-moving frame of reference showed the reverse bias. Only 30.8% of the participants primed in the object-moving frame of reference thought the meeting was on Friday, whereas 69.2% thought it was on Monday. Seven of the participants did not respond to the prime questions correctly and their responses were omitted from all analyses. A χ^2 statistic confirmed the effect of consistency, $\chi^2(1, N = 56) = 5.2, P < 0.05$. Control participants (who had not seen any primes) were about evenly split between Monday (45.7%) and Friday (54.3%).

Participants' confidence scores also confirmed this consistency bias. A confidence

score was computed for each participant by scoring a prime-consistent response as a +1, a prime-inconsistent response as a −1, and multiplying by the confidence rating that had been provided by the participant on the 1-to-5 scale. The mean observed confidence score for the primed conditions was 2.14, significantly higher than zero which would indicate no bias ($t = 2.81$, $P < 0.01$). This again confirmed the consistency effect. For the unprimed control condition, the mean confidence score (−0.23) did not differ from the null prediction. Participants in Experiment 1 spontaneously used the structural information made available by spatial primes to answer the target time questions.

2.3. Discussion

Making available different spatial schemas affected how people thought about time. Participants in Experiment 1 chose to disambiguate a sentence about time in a manner that was consistent with a recently used spatial schema. These findings confirm that the domains of space and time are similar in conceptual structure, and not just in language.

However, it is still too early to conclude that time is understood and structured online as a metaphor from space. Experiment 1 established that spatial schemas can be used to organize events in time, but this does not necessarily mean that they are necessary to do so. So far, only the effect of spatial thinking on thinking about time has been examined. But what if the experiment was reversed? Would making people think about time in a particular way affect how they think about space? If spatial schemas are necessarily accessed in thinking about time (as proposed by the strong Metaphoric Structuring View), then solving a problem about time should necessarily access and prime the appropriate way of thinking about space.

Another possibility is that spatial schemas are no longer necessary to organize time (as proposed by the weak version of Metaphoric Structuring). Since spatiotemporal metaphors are used so frequently, the spatial information used to organize time may become stored independently in the domain of time and would no longer need to be accessed through metaphorical mappings from space. Further, the structural schemas stored in the domain of time may differ in several ways from their spatial parents. Since the domain of time is less rich than the domain of space, spatial schemas imported into the domain of time may be simplified to include only those elements needed for time. For example, space has three dimensions, while time is generally thought of as one-dimensional. In space, objects have intrinsic fronts, and can face and move any which way in a layout. The domain of time, once again, is more restricted; events don't have intrinsic fronts per se and can only move in one dimension. Because the domain of time is restricted in several ways, it is reasonable to suppose that temporal schemas will be simpler, bare-bones versions of their spatial parents. If this is the case, then these restricted temporal schemas should only be useful for thinking about time. One could not use them to think about space as they would not include all the necessary details to construct a full spatial scenario.

Thus, the different strengths of the Metaphorical Structuring Hypothesis make different predictions. The strong view predicts that, since spatial schemas are neces-

sary to think about time, solving problems about time should necessarily access and prime the appropriate spatial schemas. According to the weak view, however, spatial schemas may no longer be necessary to think about time. Therefore, solving problems about time would not necessarily prime the appropriate spatial schemas. Further, people might not be able to use primed temporal schemas to think about space because temporal schemas might not include enough detail to construct a full spatial scenario.[6] Experiment 2 was designed to test these predictions.

3. Experiment 2

In Experiment 2 participants answered ambiguous questions about spatial and temporal scenarios. Each target question followed several prime questions that used either the ego-moving schema or the object/time-moving schema. For some of the participants, spatial primes preceded target questions about time. For others, temporal primes preceded target questions about space. This manipulation was designed to investigate whether spatial schemas are necessarily accessed in thinking about time. There were also two control groups for whom spatial primes preceded spatial targets, and temporal primes preceded temporal targets. These conditions were necessary as manipulation checks; the stimuli must produce an effect of consistency within a domain before consistency effects across domains can be interpreted. The primes (see Fig. 3) were designed to minimize the superficial differences between the ego-moving and object-moving scenarios used in Experiment 1.

3.1. Method

3.1.1. Participants

Three hundred two Stanford University undergraduates participated in this study as part of a course requirement.

3.1.2. Materials and design

A two-page questionnaire was constructed. The first page contained TRUE/FALSE schema priming questions, and the second page contained the ambiguous target question. Overall, the experiment was a 4 (transfer type) × 2 (prime schema type) fully crossed between participants design. The four levels of transfer type were: (1) 'space-to-space' – transfer from spatial primes to spatial targets; (2) 'space-to-time' – transfer from spatial primes to temporal targets; (3) 'time-to-time' – transfer from temporal primes to temporal targets; and (4) 'time-to-space' – transfer from temporal primes to spatial targets. The two levels of prime schema type were ego-moving, and object/time-moving.

[6] This is not to say that space could not in principle be thought of as a metaphor from time. In fact, there are several examples of temporal language being used to talk about space. One striking example is that spatial positions are commonly described by the blind as points on a clockface (12 o'clock meaning straight ahead and 6 o'clock meaning directly behind).

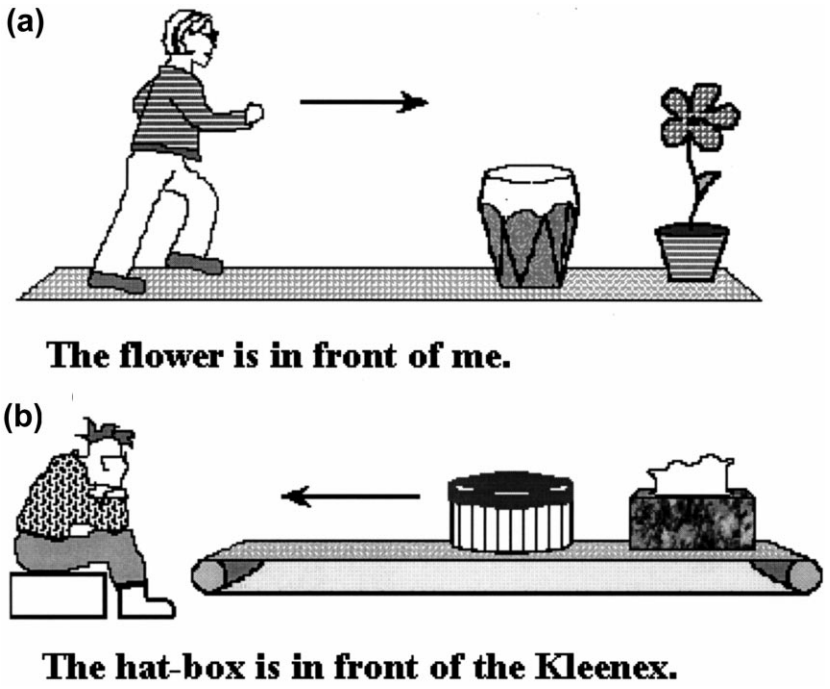

Fig. 3. (a) Sample ego-moving scenario used as a spatial prime in Experiment 2. (b) Sample object-moving scenario used as a spatial prime in Experiment 2.

3.1.2.1. Prime questions There were four types of priming questions. Half of the participants saw spatial priming questions, and half saw temporal priming questions. Half of the spatial priming questions employed the ego-moving schema (see Fig. 3a), and half employed the object-moving schema (see Fig. 3b). Likewise, half of the temporal priming questions employed the ego-moving schema (e.g. 'On Thursday, Saturday is before us'), and half employed the time-moving schema (e.g. 'Thursday comes before Saturday'). Each set of spatial primes contained three TRUE/FALSE questions, two of which were TRUE. The direction of motion depicted in the primes alternated between left and right, and all of the objects used (walnut, hatbox, drum, tissue-box, stool, can, and flower-pot) were frontless. Each set of temporal primes contained five TRUE/FALSE questions, three of which were TRUE. In each set, four of the questions asked about events that were 'before', and one of the questions asked about events that were 'behind' or 'after'. All of the temporal priming questions were about relationships between days of the week. Each set of priming questions was followed on the next page by the target question.

3.1.2.2. Target questions Two types of target questions were used: half were ambiguous time questions (e.g. 'Next Wednesday's meeting has been moved forward two days. Which day is the meeting now that its been moved?'), and half were ambiguous space questions (see Fig. 4). The widgets in Fig. 4 were arranged

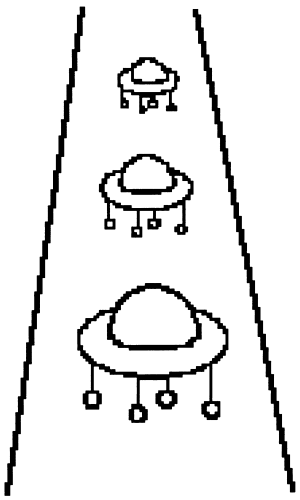

Which one of these widgets is ahead?
(please circle one)

Fig. 4. Ambiguous spatial target used in Experiment 2.

vertically from closest to farthest so as not to introduce any left/right bias. The widgets were designed to be frontless, vertically symmetrical, and mobile-looking. This was done so that a widget's 'aheadness' could not be inferred from any intrinsic properties of the widget, but rather required the observer to impose either an object-moving or an ego-moving perspective on the picture. If the participants imagined the widgets as moving out of the page (the object-moving perspective), then the bottom or closest widget should be 'ahead'. On the other hand, if the participants imagined themselves moving into the page toward the widgets (the ego-moving perspective), then the widget furthest along in the direction of motion of the observer (the top or farthest widget) should be 'ahead'.

To summarize, each participant answered either the spatial or the temporal ambiguous question after just having answered a set of either ego-moving or object/time-moving spatial or temporal priming questions. Of interest is the extent to which participants would disambiguate the target question in a prime-consistent manner across the four different transfer types.

3.1.3. Procedure

Just as in Experiment 1, participants completed the two-page questionnaire individually with no time restrictions. The two pages of the questionnaire were imbedded in a large questionnaire packet which was distributed to an entire class in introductory psychology and contained many questions unrelated to this study. No

special connection was made between the two pages of the questionnaire beyond that implied by their immediate adjacency.

3.2. Results

Results are summarized in Fig. 5. Participants were influenced by spatial primes when thinking about time (63.9% consistent), but were not influenced by temporal primes when thinking about space (47.2% consistent). Within-domain consistency effects were also observed for both the space-to-space transfer condition, and the time-to-time transfer condition (64.8 and 69.7%, respectively). Overall, the results are consistent with the weak Metaphorical Structuring prediction that, even though spatial schemas can be used to think about time, they are not necessary to do so.

Responses of 21 participants were excluded from all analyses because they did not respond correctly to all of the prime questions. There were no differences in error rates across conditions.

3.2.1. Within-domain schema consistency

The space-to-space and time-to-time conditions were necessary as manipulation checks. It was necessary to establish that the ambiguous targets used in this experiment were susceptible to consistency bias and could in principle be disambiguated by priming the ego-moving and object/time-moving schemas. This was especially important for the ambiguous spatial target (see Fig. 4) since it had not been used

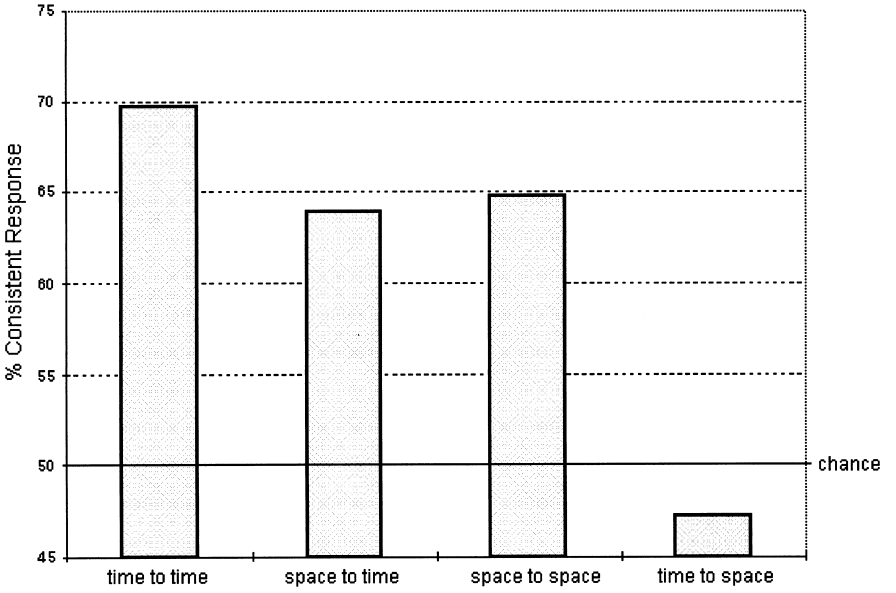

Fig. 5. Results of Experiment 2 are shown in terms of % Consistent Response (plotted on the ordinate) as a function of Transfer Type (plotted on the abscissa). There was a significant effect of consistency in all but the time-to-space condition. Chance is at 50%.

previously. Both the manipulation checks were successful. The spatial ego-moving and object-moving primes did indeed cause people to disambiguate the question in Fig. 4 in a schema consistent manner. The same was true for the time-to-time condition. The relevant statistical analyses are described below.

3.2.1.1. Space-to-space Overall, 64.8% of the participants in this condition responded in a prime-consistent manner. When given ego-moving primes (see Fig. 3a), 63.2% of participants said that the top widget in Fig. 4 was the one 'ahead' (this was the ego-moving consistent response). When given object-moving primes (see Fig. 3b), 66.7% of participants said that the bottom widget in Fig. 4 was the one 'ahead' (this was the object-moving consistent response). The prime consistency bias was significant in a 2×2 χ^2 analysis, $\chi^2(1, N = 71) = 6.28$, $P < 0.01$.

3.2.1.2. Time-to-time Overall, 69.7% of the participants in this condition responded in a prime-consistent manner. When given ego-moving primes (e.g. 'On Thursday, Saturday is before us'), 66.7% of participants said that Wednesday's meeting had been moved to Friday (the ego-moving consistent response). When given time-moving primes (e.g. 'Thursday comes before Saturday'), 71.4% of participants said that Wednesday's meeting had been moved to Monday (the time-moving consistent response). The prime consistency bias was significant in a 2×2 χ^2 analysis, $\chi^2(1, N = 66) = 9.07$, $P < 0.01$.

3.2.2. Cross-domain schema consistency

As predicted by the weak view, there was an effect of consistency in the space-to-time condition (63.9% consistent response), but not in the time-to-space condition (47.2% consistent response). A χ^2 comparison confirmed this difference between the two conditions, $\chi^2(1, N = 72) = 8.02$, $P < 0.01$.

3.2.2.1. Space-to-time Overall, 63.9% of the participants in this condition responded in a prime-consistent manner. When given ego-moving primes (see Fig. 3a), 61.2% of participants said that Wednesday's meeting had been moved to Friday (the ego-moving consistent response). When given object-moving primes (see Fig. 3b), 66.7% of participants said that Wednesday's meeting had been moved to Monday (the object/time-moving consistent response). The prime consistency bias was significant in a 2×2 χ^2 analysis, $\chi^2(1, N = 72) = 5.71$, $P < 0.05$. This schema-consistency effect suggests that there was relational transfer from the domain of space to the domain of time. This finding replicates the results of Experiment 1, and corroborates the hypothesis that people can use spatial schemas to think about time.

3.2.2.2. Time-to-space Only 47.2% of the participants gave prime-consistent responses in this condition which does not significantly differ from the chance prediction of 50%, $\chi^2(1, N = 72) = 0.09$. A main effect of response type appeared in this condition with 69.4% of the participants saying that the top widget in Fig. 4 was 'ahead', $\chi^2(1, N = 72) = 28.3$, $P < 0.01$. However, there

was no evidence of relational transfer from the domain of time to the domain of space.

It is important to note that the lack of a consistency effect in this condition cannot simply be dismissed as a failure of the spatial target to conform to the ego-moving or object-moving schemas. The same target showed a robust effect of consistency in the control space-to-space condition. Nor can the present lack of consistency be dismissed as a failure of the temporal primes to invoke the ego-moving or time-moving schemas. The same primes showed a robust effect of consistency in the control time-to-time condition. It appears that spatial schemas were not necessarily accessed in solving the temporal prime questions, and that people did not use the primed temporal schemas to think about space.

These findings support the weak Metaphoric Structuring prediction that people can use spatial schemas to think about time, but not the strong prediction that spatial schemas are necessary to think about time.

3.3. Discussion

In Experiments 1 and 2, participants were influenced by spatial primes when interpreting a question about time. This suggests that space and time do share conceptual similarities beyond similarities in language. Furthermore, in Experiment 2, participants were not influenced by temporal primes when interpreting a question about space. These findings support the weak Metaphorical Structuring View and contradict the strong view. It appears that spatial schemas are useful, but not necessary to think about time. Further, information sharing between these two domains appears to be asymmetric; people can use spatial information when thinking about time, but not temporal information when thinking about space.

Still, it may be premature to reject the strong view. There are two concerns. First, since all of the data so far have come from questionnaire studies, there are only measurements of the products of processing, not of the processing itself. The effects of temporal thinking on spatial thinking may become apparent if a more traditional measure of priming (e.g. reaction times) is used.

Second, it may still be possible to construct an alternative explanation for the results obtained in Experiments 1 and 2. It could be that space and time both use a set of generic, domain-independent (neither spatial nor temporal) schemas that can be used to mentally organize objects in space as well as events in time. Let's call this alternative the Generic Schema View. According to the Generic Schema View, time is not thought of in spatial terms. Rather, both spatial and temporal reasoning is accomplished by referencing the same generic schemas. If this is the case, then the effect of consistency observed in Experiment 1 is the result of spatial primes activating the appropriate generic schema which makes it more likely to be used by the domain of time (since time makes use of the same schemas as space). The Generic Schema View might also be able to explain the asymmetry in transfer between space and time observed in Experiment 2. The domain of space might be more strongly associated with the generic schema than the domain of time (perhaps because spatial thinking is more common than temporal thinking). This type of asymmetry could

explain why there was better transfer in the space-to-time than in the time-to-space condition; the domain of time may be too weakly associated with the generic schemas to produce enough priming to be noticeable in our paradigm.

One counterintuitive prediction that follows from this, however, is that space should prime time better than time should prime itself (because spatial primes should prime the generic schema more strongly than temporal primes). It is important to note here that the prediction is not that space should be a better prime for all aspects of time. Rather, the prediction is that space should be better at activating the relational schemas needed to structure the domain of time. Although this prediction was not borne out in the results of Experiment 2 – the effect of consistency was not bigger in the space-to-time condition (63.9%) than in the time-to-time condition (69.7%) – perhaps a more sensitive test, one that could separate out the effects of schema-consistency from simple semantic priming, would show the effect.

Experiment 3 was designed to examine participants' on-line processing, and to ascertain whether the Metaphoric Structuring or the Generic Schema View provide a better description of the data. Experiment 3 measured participants' response times to consistently and inconsistently primed questions about spatial and temporal relations. Each target question followed two prime questions that used either the same relational schema as the target (a consistent trial) or a different relational schema (an inconsistent trial). Just as in Experiment 2, the domains of the target and prime questions were varied in a 2×2 manner so that spatial primes preceded spatial or temporal targets on half of the trials, and temporal primes preceded spatial or temporal targets on half of the trials. The weak Metaphorical Structuring View predicts a pattern of results parallel to those found in Experiment 2. Participants should respond faster to consistently primed target questions for all conditions except when temporal primes precede spatial targets. The Generic Schema View predicts a greater effect of consistency from space to time than from time to time.

4. Experiment 3

4.1. Predictions

To be consistent with the results of Experiment 2, there should be an asymmetry in consistency effects between space and time. There should be greater effects of consistency when the transfer is from space to time, than from time to space.

These predictions can be accommodated by both the weak Metaphorical Structuring View and the Generic Schema View. However, according to the Generic Schema View, the asymmetry in relational priming between space and time is due to space being more closely associated to a generic schema underlying both domains (and therefore being a better prime in general). According to this view, in addition to any asymmetries in transfer between space and time, there should also be a greater effect of schema consistency when the transfer is from space to time than when the transfer is from time to time (because space should always prime the generic schema better than time).

The weak Metaphorical Structuring View makes a different prediction in this regard. According to weak Metaphorical Structuring, time can be structured in terms of its own schemas (which were at some point imported from the domain of space), or it can be structured using the actual spatial schemas. If spatial schemas are indeed functionally identical to the temporal schemas, then the effect of consistency should be the same whether the transfer is from space to time or from time to time.

4.2. Methods

4.2.1. Participants

Fifty-three Stanford University undergraduates participated in this study in order to fulfill a course requirement. Participants were tested individually in a computerized laboratory.

4.2.2. Materials

The experiment used 128 prime questions and 32 target questions. All questions had TRUE/FALSE answers. Each prime question appeared only once. Each target question appeared twice: once primed consistently, and once primed inconsistently. The order was randomized across participants.

4.2.2.1. Time questions Sixty-four statements about months of the year were constructed to use as primes. Half of these statements used the ego-moving schema (e.g. 'In March, May is ahead of us'), and the other half used the time-moving schema (e.g. 'March comes before May'). Also, half of the statements were TRUE and half were FALSE. Half of the statements referred to months that are 'ahead' or 'before', and half of the statements referred to months that are 'behind' or 'after'. All of these variations were fully crossed into eight types of primes, thus insuring that the task was too difficult for participants to develop a simple heuristic for answering the questions. In addition, 16 statements about months of the year were constructed to use as target questions. These statements were always TRUE, used either the ego-moving, or the time-moving schema, and always referred to months that are 'ahead' or 'before'.

4.2.2.2. Space questions Sixty-four spatial scenarios were constructed to use as primes. Sample items are shown in Fig. 6 and also in Appendix A. Each scenario consisted of a picture and a sentence. Half of these scenarios used the ego-moving schema, and half used the object-moving schema. Also, half of the sentences were TRUE descriptions of the spatial relations portrayed in the picture and half were FALSE. Half of the statements referred to objects that were 'in front', and half referred to objects that were 'behind'. All of these variations were fully crossed into eight types of primes. Also, left/right orientation of the pictures was counterbalanced across these variations. Only vertically symmetrical letters (M, V, W, T, H, O, X, and A) were used to identify objects in the spatial scenarios.

In addition, 16 spatial scenarios were constructed to use as target questions. Sentences in these scenarios were always TRUE descriptions of the picture, used

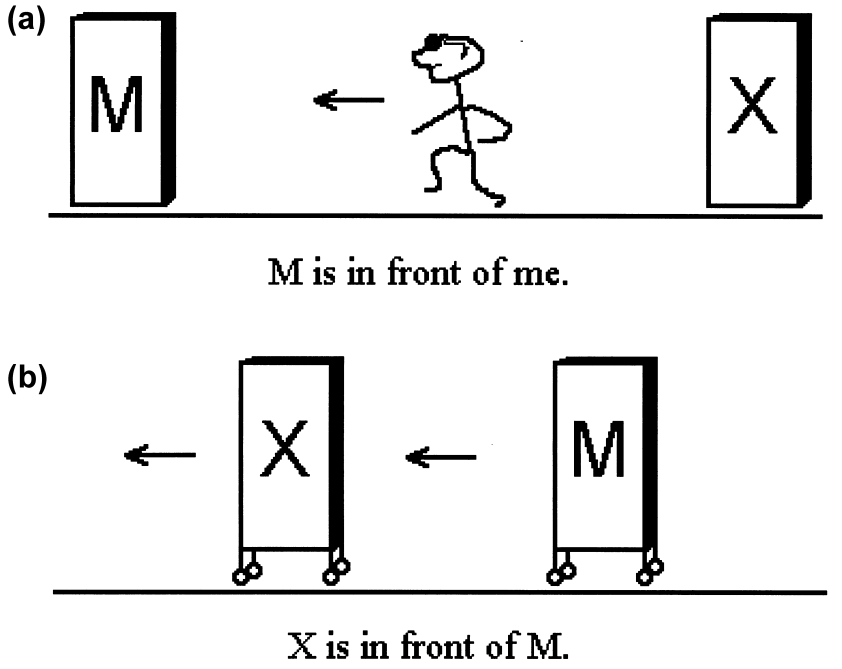

Fig. 6. (a) Sample ego-moving spatial scenario used in Experiment 3. (b) Sample object-moving spatial scenario used in Experiment 3.

either the ego-moving or the object-moving schema, and always referred to objects that were 'in front'.

4.2.3. Design

Each participant completed a short practice session followed by 64 experimental trials. Each trial was composed of two prime questions followed by one target question. Across the 64 trials, each target was presented twice, once in a consistent trial, and once in an inconsistent trial. In consistent trials, the prime questions and the target question belonged to the same schema (e.g. ego-moving prime, ego-moving target). In inconsistent trails, the prime questions and the target question belonged to different schemas (e.g. ego-moving prime, time-moving target). The critical measure was the effect of consistency on the response time to the same target question by the same participant. The order of trials was randomized. For each participant, consistent and inconsistent items appeared first and second equally often.

The design involved three factors fully crossed within participants, with factor levels of 4 (transfer type) × 2 (consistency) × 2 (target type). Just as in Experiment 2, the four levels of transfer type were: (1) space-to-space; (2) space-to-time; (3) time-to-time; and (4) time-to-space. The two levels of consistency were: (1) consistent – the primes and targets belonged to the same schema; or (2) inconsistent – the primes

and targets belonged to different schemas. The two levels of target type were: (1) ego-moving; and (2) object/time-moving.

4.2.4. Procedure

Participants were tested individually. Questions were presented on a computer screen one at a time and the participants' task was to answer TRUE or FALSE as quickly as possible by pressing one of two keys on a keyboard. In each trial, participants answered two prime questions followed by one target question. Participants were unaware that the experiment was divided into such three-part trials, nor did they figure it out just from participating in the experiment. For each question, participants needed to respond before the response deadline of 6 s. Participants received feedback during the practice session, but not during the 64 experimental trials.

4.3. Results

Results are summarized in Fig. 7. Just as in Experiment 2, people were influenced by spatial primes when thinking about time, but were not influenced by temporal primes when thinking about space. Within-domain consistency effects were also observed for both space-to-space and time-to-time trials. Contrary to the Generic

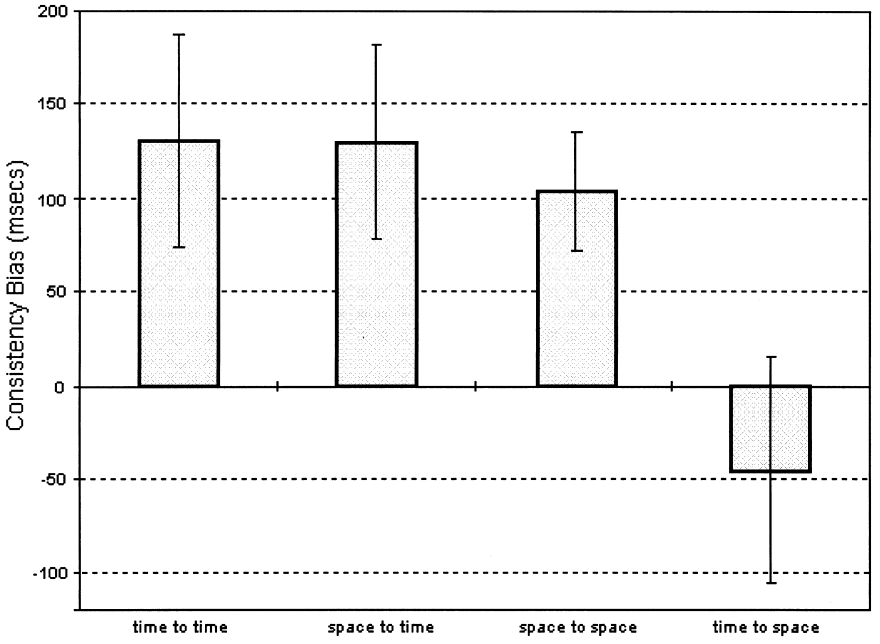

Fig. 7. Results of Experiment 3 are shown in terms of Consistency Bias in ms (plotted on the ordinate) as a function of Transfer Type (plotted on the abscissa). Consistency Bias is the difference between the mean RT for inconsistently primed targets and the mean RT for consistently primed targets. There was an effect of consistency for all but the time-to-space condition. The error bars represent two standard errors of the mean.

Schema prediction, the effect of consistency was not any different when the transfer was from space to time (129 ms), than when the transfer was from time to time (130 ms). This finding contradicts the Generic Schema explanation of the asymmetry between space and time, and gives reason to prefer the weak Metaphorical Structuring View.

The relevant statistical analyses are described below. Response times exceeding the deadline, incorrect responses, and those following an incorrect response to a priming question were omitted from all analyses.

4.3.1. Within-domain schema consistency

In both the time-to-time and space-to-space conditions, participants benefited from consistency (a benefit of 130 and 103 ms, respectively). When answering the target questions, people were able to reuse the relational information made available by the schema-consistent primes.

4.3.1.1. Time-to-time Participants responded faster to consistently primed targets (1846 ms) than to inconsistently primed targets (1976 ms), $F(1, 52) = 5.18$, $P < 0.05$. Ego-moving and time-moving targets benefited equally from consistency.

4.3.1.2. Space-to-space Participants responded faster to consistently primed targets (1619 ms) than to inconsistently primed targets (1722 ms), $F(1, 52) = 10.46$, $P < 0.01$. Ego-moving and object-moving targets benefited equally from consistency.

4.3.2. Cross-domain schema consistency

Whereas reasoning about time was facilitated by consistent spatial primes (a benefit of 129 ms), reasoning about space was not facilitated by consistent temporal primes (a non-significant decrement of 46 ms). This difference in effect of consistency between the two conditions was confirmed by a 2×2 ANOVA, $F(1, 52) = 4.35$, $P < 0.05$. Further statistical analyses are described below.

4.3.2.1. Space-to-time Participants responded faster to consistently primed targets (2086 ms) than to inconsistently primed targets (2215 ms), $F(1, 52) = 5.74$, $P < 0.05$. Response times did not differ by target type, and both ego-moving and time-moving targets benefited equally from consistency. When solving problems about time, people were able to reuse the relational information made available by consistent spatial primes.

There was also an interesting dissociation between the effect of consistency, and the effect of priming purely temporal information. Not surprisingly, people were overall faster to solve temporal targets after temporal primes (1911 ms) than after spatial primes (2150 ms), $F(1, 52) = 20.48$, $P < 0.001$. However, the effect of schema consistency was the same regardless of whether the primes were spatial or temporal. This suggests that while purely temporal information was better activated by temporal primes, the relational information needed to organize these temporal components was equally useful whether it came from spatial or temporal primes. This dissociation between the effect of consistency and the effect of simple

priming of temporal information is consistent with the Metaphorical Structuring assumption that the meat of abstract domains (the purely temporal information) is separable from their structural skeleton (the relational information used to structure the temporal components).

These findings corroborate the findings of Experiments 1 and 2, and once again support the hypothesis that people can use spatial schemas to think about time.

4.3.2.2. Time-to-space There was no transfer from the domain of time to the domain of space. Response times to consistently primed targets (1673 ms) did not differ from those to inconsistently primed targets (1627 ms), $F(1, 52) = 0.55$, $P = 0.46$. Target type did not interact with consistency indicating that ego-moving and object-moving targets were equally unaffected by consistency. Further, the effect of consistency in this condition was statistically different from that in the control space-to-space condition as confirmed by a 2×2 ANOVA, $F(1, 52) = 4.40$, $P < 0.05$. When thinking about space, people were only aided by consistent spatial and not by consistent temporal primes (presumably because the temporal primes accessed stored temporal schemas that were not detailed enough to fully represent the spatial scenarios).

These results support the weak Metaphoric Structuring claim that temporal scenarios can be understood and structured in terms of on-line mappings from the domain of space, and contradict the strong claim that spatial schemas are necessary to understand time. These findings are also consistent with the results of Experiments 1 and 2. Apparently, space and time can share structured relational information on-line, but this sharing is asymmetric; spatial schemas can be used to think about time, but temporal schemas cannot be used to think about space.

4.4. Discussion

Consistent with the Metaphoric Structuring View, people spontaneously used available spatial schemas to think about time. The structural information made available by spatial primes was just as useful for thinking about time as the structural information made available by temporal primes. However, when spatial information was not already available, people relied on separate schemas stored in the domain of time. Further, the relational information made available by these temporal schemas was not useful for thinking about space.

Just as in Experiment 2, these findings support the weak Metaphorical Structuring View, and contradict the strong view. It appears that while spatial schemas can easily be used to think about time, they are not necessary to do so.

5. General discussion

Results of three experiments show that the domains of space and time share relational structure. It appears that spatial schemas can be used as easily as temporal schemas to think about time; however, access to spatial schemas is not required for

thinking about time. These findings support the weak Metaphoric Structuring View, and contradict several plausible alternatives.

First, based on previous evidence from Gentner et al. (1999), and McGlone and Harding (1998), it was possible to reject the Dubious View that linguistic differences between ego-moving and time-moving statements have no implications for processing. In Experiment 1, priming particular spatial schemas changed the way participants thought about time. This demonstration showed that space and time have deeper conceptual similarities beyond just similarities in language. In Experiment 2, the effect of schema consistency was found to be asymmetric; although people were influenced by spatial primes when thinking about time, they were not influenced by temporal primes when thinking about space. These findings suggested that spatial schemas are not necessarily accessed to think about time (which contradicts the strong version of the Metaphorical Structuring View). Experiment 3 corroborated the results of the first two experiments, and further demonstrated that the asymmetrical sharing of information between space and time can be observed in real-time processing. Overall, these findings provide support for the weak Metaphoric Structuring View.

Although the results described above are all consistent with the weak Metaphoric Structuring View, it may still be possible to construct an alternative, non-metaphoric explanation. One such view is discussed below.

5.1. The Structural Similarity View

An alternative to metaphoric representation accounts of the origins of metaphorical language – the Structural Similarity View – was proposed by Murphy (1996). The Structural Similarity View maintains that all domains are represented directly, not metaphorically, and that metaphorical language arises when people notice preexisting structural similarities between domains. The domains of space and time, for example, though structurally similar, are represented separately.

5.1.1. Structural Similarity and cross-linguistic evidence

One claim made by Structural Similarity is that linguistic metaphors play no causal role in shaping abstract domains (Murphy, 1996). This implies that even if languages differ in the metaphors they use to describe abstract domains, speakers of these languages should not differ in their mental representations of these domains. Recent evidence suggests that this is not the case (Boroditsky, 1999). English and Mandarin speakers talk about time differently; English speakers use predominantly horizontal terms to talk about time, while Mandarin speakers use both horizontal and vertical terms. A Metaphorical Structuring account would predict that Mandarin speakers would be more likely to rely on vertical spatial schemas when thinking about time than English speakers. This is indeed what was observed. When answering TRUE/FALSE questions about time (e.g. ‘March comes earlier than April’), Mandarin speakers were faster after vertical spatial primes than after horizontal spatial primes. This result implies that Mandarin speakers were relying on vertical representations of time to answer the time questions. The reverse was true for the

English speakers. English speakers were faster after horizontal spatial primes than after vertical spatial primes. This difference is particularly striking since both groups performed the task in English, and all of the Mandarin speakers had had at least 10 years of 'contaminating' English experience. Further, English speakers who were briefly trained to talk about time using vertical metaphors produced results that were statistically indistinguishable from those of Mandarin speakers. This is strong evidence that metaphorical language plays an important role in shaping abstract thought. These new findings contradict the Structural Similarity position that metaphors play no role in shaping mental representations.

5.1.2. Structural Similarity and asymmetry in cross-domain transfer

Can Structural Similarity account for the findings presented in this paper? A proponent of Structural Similarity might argue that the schema consistency observed in Experiment 1 reflects not the sharing of relational schemas between space and time, but rather a simple case of relational priming. That is, while using a spatial schema to answer the spatial priming questions, activation is spread to a separate, but structurally similar schema in the domain of time. Thus activated, the time schema is in readiness to answer the critical time question. While this explanation can account for the results of Experiment 1, it might not be able to explain the results of Experiments 2 and 3.

Why would space be a better prime for time than vice versa? One seemingly plausible explanation is that space and time may be asymmetrically associated because one of the domains is richer, more elaborated, or more familiar than the other. Indeed, many examples of asymmetries have been documented between items that differ in familiarity, prototypicality, or salience (e.g. Rips, 1975; Tversky, 1977). There are two serious problems with this explanation.

First, it is generally the case that the less familiar member of an asymmetrically associated pair (e.g. 'leopard' of 'leopard–tiger') is more likely to elicit the more familiar member (e.g. 'tiger') than the reverse. Of space and time, space appears to be the richer, more elaborated, and more familiar domain; space has more dimensions than time, is more flexible with regard to direction of motion, and is more readily perceptible. From all this, one should predict that time (the smaller, less common domain) should remind us of space (the larger, more common domain) more than space should remind us of time. This prediction is, indeed, exactly the opposite of what was observed in Experiments 2 and 3 which showed that spatial thinking affected temporal thinking but not the reverse.

Second, although examples of asymmetries in similarity judgments and association frequencies are plentiful, it would be reckless to assume that asymmetries in such explicit measures lead to asymmetries in mutual priming. In fact, recent studies have found that even between asymmetrically associated items (e.g. 'bar–drink'), the effects of mutual semantic priming are symmetrical (Thompson-Schill, Kurtz & Gabrieli, 1998). Even though 'bar' calls to mind 'drink' much more often than 'drink' calls to mind 'bar', the semantic priming between the two words is symmetrical. If priming is symmetrical even between asymmetrically associated items, there appears to be no reason to expect asymmetries in simple priming between space and time.

5.1.3. Structural Similarity and the directionality of metaphors

One challenge for any non-metaphoric theory of representation is to explain the directionality of the linguistic metaphors pointed out by Lakoff. In a great majority of conventional metaphors, an abstract domain is described in terms of some concrete domain, and not vice versa. For example, we talk about ideas in terms of food, but not food in terms of ideas (while it is possible to say 'I contemplated that piece of meat' to mean 'I ate it slowly', such talk is not common, nor has it been conventionalized in the way that the IDEAS ARE FOOD metaphor has been). If all domains are represented directly, and concrete domains have no special influence over abstract domains, why should we expect this overwhelming directionality in linguistic metaphors?

One might argue that this directionality is simply a matter of pragmatics – perhaps people talk about love in terms of journeys (and not vice versa) simply because ''people wish to talk about love much more than they wish to talk about journeys'' (Murphy, 1996). Although this analysis seems to work fine for love and journeys, it becomes a bit more *hard to swallow* when we come to, say, ideas and food. Surely any survey of American households would reveal far more talk of meatloaf than of mental representation. And what if we find ourselves *seduced* by the pragmatic account, or suppose we even decide to *buy* into it? Would we also have to buy that people are more interested in discussing the nature of thought than they are in food, sex, or shopping? The simple pragmatic account appears to be insufficient. Although it seems natural to describe Princess Diana's marriage as a *stormy* affair, it would be an improbable group of Englishmen who wish to talk about romance much more than they wish to talk about the weather (M. J. A. Ramscar, pers. commun., 1 November 1998).

Clearly, factors other than the interestingness of a domain (such as how much sensory information is available about a domain, for example) need to be taken into account to be able to explain the striking directionality of conventional metaphors. The more general point that emerges from this discussion is that there is a dire need for more rigorous empirical testing in this area. Although there may be a wealth of anecdotal evidence in support of any particular claim, it appears that anecdotal counter evidence is just as easy to come by. Empirical approaches (such as the one presented in this paper) will help shape a more definitive view of metaphoric representation.

Recent empirical findings pose serious challenges to the Structural Similarity View. The Structural Similarity claim that metaphors play no causal role in shaping thought is contradicted by new cross-linguistic evidence. Further, Structural Similarity appears to have difficulty accounting for the asymmetry in transfer between space and time observed in Experiments 2 and 3. The findings presented and reviewed in this paper give reason to prefer the weak Metaphoric Structuring View over Structural Similarity.

6. Conclusions

Results presented in this paper suggest that similarities between space and time in

language have deeper conceptual underpinnings. Three experiments showed that space and time share enough relational structure to allow spatial schemas to be used as easily as temporal schemas to organize events in time. Further, recent cross-linguistic evidence shows that if spatiotemporal metaphors differ, so do people's conceptions of time. This suggests that using spatiotemporal metaphors causes spatial relational structure to be imported (as by analogy) to the domain of time. However, there was no evidence that spatial schemas are necessary to think about time. This may be because frequent mappings between space and time come to be stored in the domain of time. Taken together, these findings lend support to a metaphorical theory of concept learning. It appears that abstract domains such as time are indeed shaped by metaphorical mappings from more concrete and experiential domains such as space.

Finally, it is important to point out that the Metaphorical Representation View is in itself a metaphor. Namely, the 'MENTAL REPRESENTATION IS A METAPHOR' metaphor is used to explain how abstract conceptual domains might be represented. Although this *metaphor metaphor* may be productive, it is still largely underspecified and so has clear limitations as a cognitive model (see Murphy, 1996 for discussion of challenges for Metaphoric Representation). At present, little is known about how conceptual metaphors are processed, how a particular metaphor is chosen, or how (or even whether) conflicts between inconsistent metaphors are reconciled. Further, it may be worth keeping in mind that the MENTAL REPRESENTATION IS A METAPHOR metaphor is only one of the many possible metaphors that might be used to characterize conceptual representation. If domains like TIME and LOVE are characterized in terms of many different metaphors, then surely a detailed understanding of something as complex as MENTAL REPRESENTATION itself will require more than just this one.

Acknowledgements

The research presented in this paper was supported by a Graduate Research Fellowship from the National Science Foundation. Partial support was also provided by NIMH research grant MH-47575 to Gordon Bower. I would like to thank Gordon Bower, Herbert Clark, Dedre Gentner, Michael Ramscar, and Barbara Tversky for many valuable discussions and comments on earlier versions of this paper. Portions of this paper have been presented at the 19th annual meeting of the Cognitive Science Society at Stanford University, and at the Analogy '98 Workshop at the New Bulgarian University in Sofia, Bulgaria.

Appendix A

A.1. Examples of stimuli used in Experiment 3

Examples of spatial stimuli are shown in Fig. 8 and examples of temporal stimuli are shown in Table 1

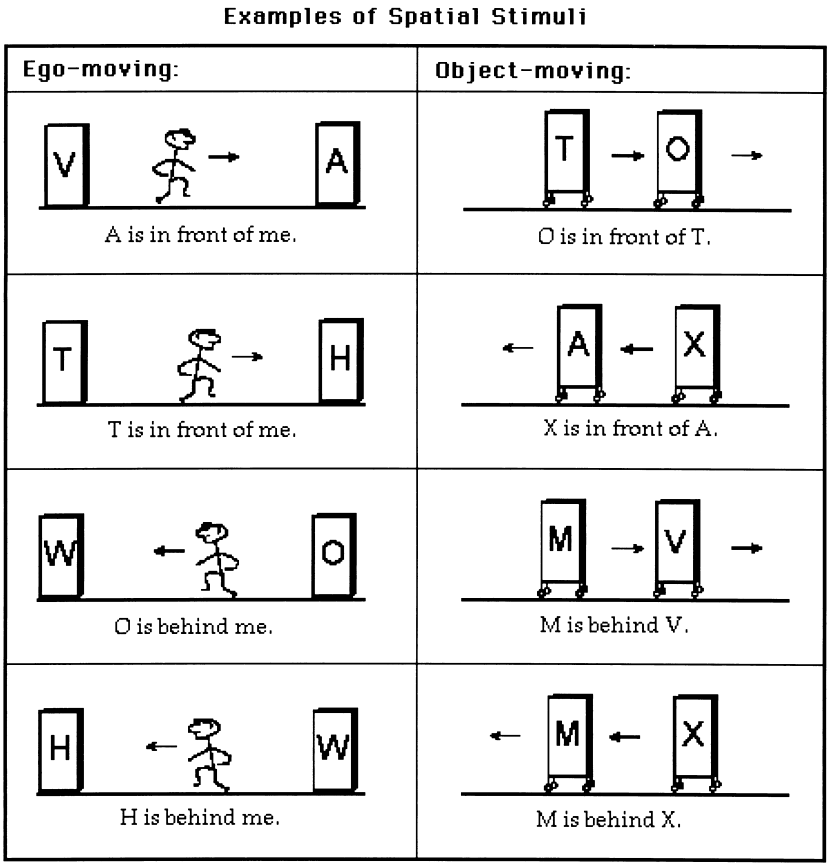

Fig. 8. Examples of spatial stimuli.

Table 1
Examples of temporal stimuli[a]

	Ego-moving	Time-moving
Ahead/before		
True	In March, May is ahead of us	March comes before May
False	In May, March is ahead of us	May comes before March
Behind/after		
True	In May, March is behind us	May comes after March
False	In March, May is behind us	March comes after May

[a] Only one month-pair is shown. Other month-pairs used were April–May, May–June, April–June, July–August, August–September, July–September, August–October, September–October, May–July, June–August, November–December, September–November, June–July, March–April, and October–December.

References

Boroditsky, L. (1999). First-language thinking for second-language understanding: Mandarin and English speakers' conceptions of time. *Proceeding of the Twenty-first Annual Meeting of the Cognitive Science Society*, Vancouver, BC.

Bowdle, B. & Gentner, D. (1995). The career of metaphor. *Paper presented at the meeting of the Psychonomics Society*, Los Angeles, CA.

Bowdle, B. & Gentner, D. (1999). Metaphor comprehension: from comparison to categorization. Paper to be presented at the Twenty-First Annual Meeting of the Cognitive Science Society, Vancouver, BC.

Clark, H. H. (1973). Space, time, semantics, and the child. In T. E. Moore, *Cognitive development and the acquisition of language*, New York: Academic Press.

Fillmore, C. J. (1971). *The Santa Cruz lectures on deixis*, Bloomington, IN: Indiana University Linguistic Club.

Gentner, D., & Wolff, P. (1997). Alignment in the processing of metaphor. *Journal of Memory and Language*, *37* (3), 331–355.

Gentner, D., Imai, M., & Boroditsky, L. (1999). As time goes by: understanding time as spatial metaphor. submitted.

Gibbs, R. W. (1996). Why many concepts are metaphorical. *Cognition*, *61*, 309–319.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*, Chicago, IL: University of Chicago Press.

Lakoff, G., & Kovecses, Z. (1987). The cognitive model of anger inherent in American English. In D. Holland, & N. Quinn, *Cultural models in language and thought*, (pp. 195–221). Cambridge, England: Cambridge University Press.

McGlone, M. S., & Harding, J. L. (1998). Back (or forward?) to the future: the role of perspective in temporal language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24* (5), 1211–1223.

McTaggart, J. (1908). The unreality of time. *Mind*, *17*, 457–474.

Murphy, G. L. (1996). On metaphoric representation. *Cognition*, *60* (2), 173–204.

Murphy, G. L. (1997). Reasons to doubt the present evidence for metaphoric representation. *Cognition*, *62* (1), 99–108.

Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, *14* (6), 665–681.

Thompson-Schill, S. L., Kurtz, K. J., & Gabrieli, J. D. (1998). Effects of semantic and associative relatedness on automatic priming. *Journal of Memory and Language*, *38* (4), 440–458.

Traugott, E. C. (1978). In J. H. Greenberg, *On the expression of spatiotemporal relations in language*, *Word structure, Universals of human language, vol. 3*. (pp. 369–400). Stanford, CA: Stanford University Press.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352.