Accessing Causal Relations in Semantic Memory

Daniela B. Sellner and Michael R. Waldmann

Department of Psychology, University of Göttingen, Germany

Keith J. Holyoak

Department of Psychology, University of California, Los Angeles

Address for proofs:

Dr. Daniela Sellner

Otto-von-Guericke University

Department of Neurology II

Leipziger Str. 44

39120 Magdeburg

Germany

Phone: ++49 391 6713431

Fax: ++49 391 6715233

e-mail: sellner@neuro2.med.uni-magdeburg.de

Abstract

Most studies investigating semantic memory have focused on taxonomic or associative relations. Very little is known about how other types of relations, such as causal relations, are represented and accessed. In a series of experiments we presented participants with pairs of words one after another either describing events that referred to a cause (e.g., spark) or an effect (e.g., fire). We manipulated the temporal order in which the words were presented in the task, and the question to which participants had to respond. The results revealed that questions referring to the existence of a causal relation are answered faster when the first word refers to a cause and the second word to its effect than vice versa. In contrast, no such asymmetry was observed with questions referring to the associative relation. People appear to distinguish the roles of cause and effect when queried specifically about a causal relation, but not when the same information is evaluated for the presence of an associative relation.

Keywords: semantic memory, causal relations, associative relations, causal-model theory

Introduction

Semantic memory is regarded as the long-term repertoire of our world knowledge (Tulving, 1972). Without world knowledge we would be incapable of understanding the world around us, and hence would be unable to communicate or to act in the service of goals (Hodges & Patterson, 1997). Semantic memory contains knowledge about categories and features that we use to represent the world, as well as knowledge about relations between categories and features (see Murphy & Medin, 1985). Semantic memory research has exploited paradigms such as sentence verification, typicality judgments, and priming of lexical decisions in efforts to understand the organization of semantic memory. The great majority of this work has focused on either taxonomic relations (e.g., verifying category statements, such as “A robin is a bird”) or general associative relations (e.g., priming a lexical decision about DOCTOR by first presenting an associate such as HOSPITAL). But although taxonomic and associative relations are certainly important for understanding cognition, other relations, such as temporal, functional or causal relations are also highly relevant for planning, predicting, acting, and reasoning. Research on categorization has recently increased attention to such relations (e.g., Lien & Cheng, 2000; Murphy & Medin, 1985; Waldmann, 1996; Waldmann, Holyoak, & Fratianne, 1995; Murphy, 2002); nonetheless, very few studies have addressed the question of how causal and similar functional relations are stored and accessed in semantic memory (but see Moss, Ostrin, Tyler, & Marslen-Wilson, 1995; Tyler & Moss, 1997; Krüger, Nuthmann, & van der Meer, 2001; van der Meer, Beyer, Heinze, & Badel, 2002).

The existence of multiple relations within semantic memory raises a particularly interesting question that has been neglected by theories of semantic knowledge: Assuming that different types of relational knowledge are relevant in different contexts, how are specific relations accessed within a network that contains many different kind of relations? The main goal of our study was to address this question by focusing on a particularly important class of relations, those that are causal in nature.

Causal Relations

Knowledge about causal relations is fundamental for our survival. Causal knowledge allows us to predict the future, achieve goals on the basis of actions, or explain why something happened. The human ability to acquire and retrieve causal knowledge is thus a central cognitive capacity (Cheng, 1993).

The nature of causality has for a long time been a hotly debated topic in philosophy. Many philosophers have struggled with the question of whether our causal impressions are solely based on associations between constantly paired events (Hume, 1739/1978) or whether it is necessary to assume metaphysical notions such as causal powers (Kant, 1781/1950). However, one aspect of causal relations seems undisputed: causal relations are asymmetric (Hausman, 1998; Waldmann, 1996). Hume noted in his famous definition of causality that in addition to spatio-temporal contiguity and constant conjunction, a necessary feature of a causal relation is that a cause must temporally precede its effect. Other philosophers have pointed out the related asymmetry that interventions can set causes to achieve effects, but generating these effects by other means does not produce their causes. For example, exposure to rhino viruses may cause a cold, but contracting a cold because of some other factor would not spawn rhino viruses in the body. In addition, proponents of the view that causal relations are based on a physical transmission of energy or information have pointed out that these processes are typical irreversible and generate specific statistical patterns that presuppose causal directionality.

Although causal asymmetry is an undisputed feature of the physical world, there has been sharp disagreement among psychologists as to whether this asymmetry is mirrored in human cognitive representations. Some researchers in the area of learning have claimed that causal asymmetry is not a feature that is represented when people learn about causal relations

(Shanks & López, 1996; Cobos, López, Cano, Almaraz & Shanks, 2002.). According to this *associative* view, learning leads to knowledge about the associative strength between cues and outcomes irrespective of what the cues and outcomes represent in the real world. By contrast, an alternative account, *causal-model theory*, has postulated that learners explicitly represent asymmetric causal relations and use this knowledge in learning (Waldmann & Holyoak, 1992; Waldmann et al., 1995; Waldmann, 1996, 2000, 2001). The main goal of the present research is to derive predictions from these two competing paradigms for the area of semantic memory for causal relations, and to test these predictions.

Causal-model Theory

According to causal-model theory, people use prior knowledge about causal directionality to interpret learning data when they are acquiring knowledge about causal relations. The goal of causal learning is the acquisition of knowledge about causal models that represent events as causes and effects. Thus, it is crucial to assign the learning events (i.e., cues and outcomes) to their causal roles (i.e., causes and effects) within directed causal relations.

Directed graphs provide a convenient way to model the structure of causal models, using nodes to represent causal events (causes, effects) and arrows to express the direction of the causal influence between events that are directly linked. Figure 1 displays an example that might approximate a fragment of people's knowledge of the common cold. As can be seen in Figure 1, the same event (e.g., infection) can play the role of both cause and effect with respect to different events. For example, infection is the effect of viruses and the cause of cough (among other effects). Thus, the causal role of an event is not a stable feature of the event in isolation, but rather depends on the context of other events in which it is presented.

Insert Figure 1 about here

One recent example of how causal-model theory has been tested is provided by studies reported by Waldmann (2000). In one of the experiments, participants had to learn to classify artificial diseases on the basis of information about the presence of substances in the patients' blood. Differential initial instructions were used to manipulate whether learners viewed the substances as potential causes of the disease (substances coming from food items) or effects of the disease (substances produced by the disease). Although cues and outcomes were identical in the learning task, learners appeared to induce different causal models and made different inferences. Specifically, a cue competition effect (blocking) was only shown when the cues represented causes of the outcome but not when they represented effects of the outcome. This finding is consistent with the view that learners attempt to acquire representations of causal models with directed links, rather than simple associations between cues and outcomes.

Asymmetry in Recognition of Causal Relations

The present research uses a different task that focuses on the retrieval of existing causal knowledge rather than learning. We used a *relation recognition* paradigm in which we presented participants with pairs of words one after another, each describing an event that referred to a cause (e.g., "spark") or an effect (e.g., "fire"). We manipulated the temporal order in which the words were presented, and the question to which participants had to respond. In the causal conditions we asked participants to judge as quickly as possible whether the two mentioned events were causally related. We compared two types of causal conditions. In the *predictive* condition, the cause (e.g., "spark") was presented temporally prior to the effect (e.g., "fire"), whereas in the *diagnostic* condition the order of presentation was reversed. Regardless of the ordering, the correct answer would be "yes" in both conditions.

Events such as spark and fire can be connected by a number of different relations (e.g., spatial, temporal, associative); accordingly, assessing whether a causal relation exists between them requires access to the specific set of causal relations stored in semantic memory. Since causal relations are asymmetric, it does not suffice to check whether the two events are merely linked; rather, it is necessary to map them to the roles of cause and effect and check whether they are connected by a causal relation. Causal-model theory predicts that this mapping process will yield an asymmetry between the recognition of causal relations in the predictive versus diagnostic order. A key assumption driving this prediction is that people (and other animals) typically experience causes prior to their effects. In fact, the possibility of experiencing effects prior to their causes appears to be a late achievement in evolution, depending on the availability of symbol systems. Symbol systems allow us to describe events and to learn on the basis of described events, rather than solely the real events that are being described. For example, in the real world one would first experience a specific food item and later the nausea that may be caused by it. However, a medical student may first read about nausea (i.e., a description of the ailment) and later be informed about the causes of nausea.

Because the natural order of events is predictive, we assume that we have an automatic tendency to infer that when one causal event is presented temporally prior to another, then the former event is the cause and the latter its effect (a kind of "congruity" effect). A variety of specific mechanisms might yield a congruity effect. One possible mechanism involves analogical mapping. In order to assign events to roles in a causal relation, each item in a pair must be mapped to a causal role. According to at least one major model of analogical mapping, Learning and Inference with Schemas and Analogies (LISA; see Hummel & Holyoak, 1997, 2003; Kubose, Holyoak & Hummel, 2002), mapping is based on distributed semantic representations. Mapping will be facilitated to the degree that the semantic features of a presented item overlap with the features of the role to which the item maps. Given the assumption that the semantic code for "temporally prior" overlaps with the semantic code for "cause", whereas the code for "temporally after" overlaps with that for "effect", a congruity effect of the form predicted by causal-model theory will be a natural consequence of the operation of such a mapping mechanism.

Evidence consistent with this prediction of a congruity effect was reported in a recent study by van der Meer et al. (2002). These researchers were interested in the role of temporal order relations in text comprehension. Their experiments tested the assumption that a text is easier to understand when the temporal order in which events are mentioned in the text corresponds to the temporal order in which these events occur in the real world. In their Experiment 1, van der Meer et al. (2002) used an item pool consisting of sequences of four events each (e.g., bite off-chew-digest-swallow). A norming study ensured that the event pairs were equated in terms of temporal relatedness (i.e., the association frequency based on the instruction to generate events on the basis of temporal succession), as well as typicality, duration, and temporal distance. Participants saw the prepositions “before” or “after” along with an event label (e.g., “before digest”). After a stimulus onset asynchrony (SOA) of either 200 ms or 1000 ms another word was presented that either described the event following or preceding the event initially mentioned within the normed sequences of four events. The results showed that at both SOAs, items that correctly expressed an “after” relation (e.g., “after bite off -chew”) were verified faster than items that correctly expressed a “before” relation (e.g., “before digest - swallow”). Relations based on “after”, unlike those based on “before”, preserved the mapping between temporal order of mention and temporal order of the actual events.

In a related study, participants read two words describing events one after another with an SOA of 1000 ms (Krüger et al., 2001). The word pair described either events that were ordered in the chronological direction or the reverse direction, or they were not semantically related at all. As in the study by van der Meer et al. (2002) described above, items were chosen on the basis of a norming study so as to be equated in terms of temporal relatedness and other factors. The task was to judge whether the two words within an item were *semantically related*. Similar to the previous study, this experiment also showed that information following the temporal arrow was accessed faster than items that were ordered in the reverse direction. The greater difficulty of accessing knowledge in the reverse temporal order could also be seen in pupillometric indices.

The basic finding that temporal order judgments are made more quickly when order of presentation matches order of the underlying events (Van der Meer et al., 2002) is consistent with our prediction for causal relations. However, Van der Meer and her colleagues appear to have shown that temporal relations affect judgments independently of whether the instructions requested access to temporal information (van der Meer at al., 2002) or to semantic relations in general (Krüger et al., 2001). By contrast, we predict that different relations can be selectively accessed depending on the task. We will discuss methodological issues that may have contributed to the finding of Krüger et al., (2001) in the Discussion of Experiment 3 below.

Dissociating the Causal-Model from the Associative View

In contrast to causal-model theory, the associative view claims that humans and animals learn about associative relations between cues and outcomes regardless of the causal roles of the events to which the cues and outcomes are referring. Indeed, one of the key features of associative links is that they do not explicitly represent causal directionality (see Cobos et al., 2002; Waldmann, 1996). Associative links in feedforward associative networks are directed from input cues to outcome representations, but this directionality follows the flow of information in the task context (cues are presented prior to effects), and not causal priority. In a diagnostic task in which effect information precedes cause information, causal directionality (from causes to effects) is pitted against associative directionality (from cues to outcomes).

Nonetheless, associative theories could potentially predict faster access to predictive as compared to diagnostic items in a relation detection task by assuming that associations in the predictive direction may tend to be stronger than associations in the reverse direction. Assuming that associative relations are based on the experience of spatio-temporal contiguity of event pairs in the real world or in texts, it is plausible to assume that people encounter event pairs more often in the predictive than the diagnostic direction. Friedman (1990) reports studies showing that if participants are asked to recall a sequence of events or talk about past events in their lives they typically order them along the timeline. Moreover, in a developmental study, Friedman (2002) found that four and eight month-old children showed a significant preference for the forward as compared to the reverse presentation of a video in which water was poured into a glass.

In order to dissociate the predictions of the causal-model versus associative views, in the present study we used materials for which the association strength between cause and effect was equated for the predictive and the diagnostic direction. Although it is likely that predictive relations in general tend to be associated more strongly than diagnostic relations, it is possible to find items in which both associations are roughly equal. According to the associative view there should be no difference in access time for these items in the two directions, whereas the causal-model view predicts that the mismatch between order of presentation and causal directionality should lead to faster access for predictive relations as compared to equally associated diagnostic relations.

Whereas Experiments 1 and 2 focus on the issue of whether there is an asymmetry in access time for predictive versus diagnostic orderings of the item pairs, Experiments 3 and 4 used an additional technique to dissociate the two views. We compared instructions that either required participants to check whether the evens mentioned in the item pairs were *causally linked* or whether they were *associated*. Since associative theories do not distinguish between these two types of relations, no difference between these two tasks is predicted by this theoretical account. In contrast, the causal-model view predicts shorter retrieval times for predictive than for corresponding diagnostic items (when association strength is equated in both directions) only when participants check for the existence of a causal relation, but not when they assess the associative relation. An associative relation, unlike a causal relation, need not distinguish the nature of the roles played by the two events that are linked. In addition to its theoretical importance, finding that the advantage of predictive over diagnostic pairs is eliminated when the question involves associative rather than causal relations would confirm that the initial norming procedures indeed succeeded in identifying items with symmetric association strength.

Experiment 1

The main goal of Experiment 1 was to test our prediction that predictive relations will be accessed more quickly than diagnostic relations even for items that are equally associated in both directions. The stimuli pairs were selected on the basis of the USF “Word Association Norm” list (Nelson, McEvoy & Schreiber, 1998), coupled with an additional norming study focusing on causal relations. We used the norming study to select causally-related item pairs that were equated in both directions in terms of the strength of statistical relations. Our procedure differed in a number of aspects from the previous related studies of van der Meer et al. (2002) and Krüger et al. (2001). Our focus was on causal rather than temporal relations. We selected item pairs that presented identical events in either the predictive or the diagnostic direction, whereas in previous studies of temporal relations the item pairs representing the forward direction differed in one event from item pairs representing the backward direction. In addition, whereas in previous studies items were chosen that were equated in terms of temporal relatedness, we used both associative and statistical relatedness as criteria for our item selection. This method of item selection allowed us to test the causal-model view against the view that causal relations can be reduced to associative relations.

Method

Participants and Design

Twenty-six undergraduate students from the University of California, Los Angeles (UCLA), participated for course credit. Their vision was normal or corrected to normal. Three participants had to be replaced (one due to computer failure, two with an overall error rate above 20%). Stimuli were presented in a within-subject design. Every participant saw every word pair once; the order of the trials and the type of stimuli within a trial (predictive vs. diagnostic order) were randomized and counterbalanced, respectively.

Stimuli

The stimulus material consisted of 68 causally related (e.g., “moon-tide”) and 68 weakly associated filler word pairs (e.g., “ring-emerald”; see Appendix for the item list). The causal and the associated filler word pairs shared the same low strength of association in each direction (Table A1). The word pairs were selected from the USF “Word Association Norm” (Nelson et. al., 1998). In that study, participants were asked to write down the first word that came to mind that was meaningfully related or strongly associated to a presented word. Each participant only produced a single word. Nelson et al. (1998) calculated the forward and backward strength between the cue and the target word for each of 5019 words. The forward strength (FSA in Table A1) was calculated as the number of participants who produced a particular target word (for a given cue) divided by the number of participants in that group. The backward strength was calculated the same way, except that now the target word served as the cue word (BSA in Table A1). The strength of association between the two words is represented by a number between 1.0 and 0 for each direction. We chose word pairs from this database that were connected by a causal relation. These word pairs had a low strength of association (0.00-0.2) in both directions (predictive/forward and diagnostic/backward). Only weakly associated words were chosen so as to reduce the possible influence of strength of association as much as possible. A total of 250 word pairs were selected.

As a next step, we conducted an additional norming study in which we presented these 250 word pairs to 80 participants (UCLA undergraduates) in a questionnaire task. Both causal and filler word pairs were presented, and for each pair the students were asked to judge whether there was a causal relation between the presented words. If the students detected a causal relation, the next step for them was to imagine that the event described by the first of the two words occurred 100 times. Their task was to estimate the conditional frequency for the event described by the second word (e.g., "*fire* occurs a 100 times, how often does *heat* occur?"). They indicated their rating on a scale from 0 to 100, with increments of 10. The order of items was randomized within the questionnaire and the version (diagnostic or predictive) of each item was randomly chosen, and counterbalanced over two versions of the questionnaire.

On the basis of the results of the norming study, we selected 68 pairs for our experiments that did not differ by more than 30 mean rating points between the two directions (predictive and diagnostic). Table A1 shows the difference of the ratings for the individual items (FR). The rating difference between the two directions summed up to 0 over all selected word pairs, which means that the overall conditional frequency rating for predictive causal relations was the same as the overall frequency rating for the diagnostic causal relations. The associated filler word pairs did not share a causal relation, but were roughly equal to the causal items in terms of bidirectional strength of association (see Table A1).

In the reaction time (RT) experiment, the test pairs were presented in font "Arial Black" and size 24 on a white background. All experiments were programmed in Superlab® and implemented on Macintosh iMac computers connected to a 15" screen with 1026 x 768 pixel resolution and 256 colors. The words were created as pct- files in Canvas 6.0 graphics software.

Procedure

Participants were shown written instructions. They were asked to decide if there was a *causal* relation between the events described by the two words presented on the computer screen. To make it clear that we wanted participants to assess the existence of a causal relation independently of the sequence of the item pairs, we additionally specified that the task was to assess “whether the event described by the first word causes or is caused by the event described by the other word.” After reading the instructions, participants in this and the other experiments were asked to repeat it in their own words to avoid any misunderstanding about their task.

Insert Figure 2 about here

At the beginning of each trial participants saw a fixation cross in the center of the screen (see Figure 2). After 1000 ms the cross disappeared and a blank screen was presented for 500 ms. The first word of the item pair was then presented for 1000 ms, followed by the second word (which replaced the first word). Thus, the interstimulus interval (ISI) was 0 and the stimulus onset asynchrony (SOA) was 1000 ms. The second word remained on the screen until the participant pressed one of the two response keys. If participants viewed the item as causal they were requested to press the letter "C" on the keyboard; otherwise they were to press the letter "N". Each participant completed 136 trials. Of the 68 causal trials, half were presented in the predictive direction (cause-effect) and half were presented in the diagnostic direction (effect-cause); the remaining 68 trials were non-causal filler word pairs. The response time for each trial was measured and recorded by the Superlab® software. The program also recorded errors.

Participants were required to respond to 20 practice trials (10 causal, 10 filler) at the beginning of the session. They received feedback in the practice trials; during the experimental trials no feedback was given.

Results and Discussion

The most important findings concern RTs for the predictive and diagnostic causal items. In all experiments, the analysis of reaction times included only trials that were answered correctly, excluding outliers that were more than two standard deviations above or below the mean. On average, this criterion led to the exclusion of 2 to 3 trials per participant.

Overall, RTs were 68 ms faster for predictive (mean RT of 1016 ms) than for diagnostic trials (mean RT of 1084 ms). A paired sample *t*-test revealed that the advantage for the predictive trials was significant, $t(25) = 2.64$, $p = .01$, with a small to medium effect size of $d = .34$. Error rates for the predictive and diagnostic pairs were 15% and 18%, respectively, and did not differ significantly, $t(25) = 1.74$, $p = .93$.

The results confirm our hypothesis that causal relations are accessed faster when the order of presented events follows the predictive rather than the diagnostic order. These results support the hypothesis that causal direction influences the retrieval of causal relations. The analysis of errors revealed no speed-accuracy tradeoff between the different conditions. The advantage of predictive over diagnostic orderings favors the causal-model view over the associative account because the effect was found with items that had symmetric association strengths in both presentation orders. Thus, the difference in RTs cannot be explained by asymmetries of associative strength.

Experiment 2

In Experiment 1 the predictive and diagnostic items were presented within a single block in a random order. Although we explicitly instructed participants that both predictive and diagnostic relations are part of the general class of causal relations, there may be a tendency to interpret a request to check for causal relations as primarily directed toward the more natural predictive direction. Participants may therefore first check for the predictive direction before considering the possibility of a diagnostic relation, a strategy that would yield longer RTs for diagnostic items. To rule out the possibility that participants primarily interpret causal questions as requests to assess predictive relations, we presented predictive and diagnostic items in separate blocks and explicitly specified the access direction in the task instructions. In two different blocks participants had to judge either whether the first word described the cause of the second event (predictive block), or whether the first word described an effect of the second event (diagnostic block). As in Experiment 1, the filler items were weakly associated words. The item pairs in the predictive block were either causally related (predictive order) or associated, whereas the item pairs in the diagnostic block were either causally related (diagnostic order) or associated.

Method

Participants and Design

Twenty-eight UCLA undergraduates with normal or corrected to normal vision received course credit for their participation. We replaced four participants because their overall error rate exceeded 20%. For the four replaced subjects the false alarm rate for those items that were associated but not causally related was particularly high, indicating that these participants were biased towards a causal response and therefore did not clearly fulfill the task demands to differentiate between word pairs that were causally related rather than only associated.

The stimuli were presented in a within-subject design: every participant saw every word pair once. The order of the two blocks and the version of the stimuli assigned to either block (predictive vs. diagnostic) were counterbalanced, creating four different counterbalanced versions.

Stimuli

The stimulus material consisted of 64 causal and 64 associated filler word pairs that were used in Experiment 1. We excluded four causal word pairs that were consistently missed in Experiment 1 to reduce the error rate and decrease the noise in the data. Their removal did not alter the overall balanced strength of association ratings. The overall frequency rating was virtually symmetric, with a miniscule bias towards the diagnostic direction (FR = -1). For each block a list of stimuli was created containing 32 causal word pairs in one of the two directions and 32 associated filler word pairs.

Procedure

Half of the participants received the predictive block first and the diagnostic block second; the other half were given the blocks in the opposite order. In the predictive block participants were asked to judge whether or not the first word described the cause of the event represented by the second word, whereas in the diagnostic block they had to decide whether the first word described the effect of the event represented by the second word. Prior to each block participants were given ten practice trials (five predictive and five filler before the predictive, and five diagnostic and five filler before the diagnostic block), for which feedback was given.

Results and Discussion

The results replicated the pattern observed in Experiment 1. Mean RT was significantly shorter (by 69 ms) in the predictive condition (817 ms) than in the diagnostic condition (886 ms), $t(27) = 2.35$, $p = .02$, corresponding to a medium effect size of $d = .44$. Participants made slightly more errors in the predictive block (13%) than the diagnostic block (12%), but this difference did not approach significance, $t(27) = .80$, $p = .43$.

Experiment 2 thus replicated the results of Experiment 1 using a design in which predictive and diagnostic items were blocked with specific instructions to either check for a predictive or a diagnostic relation. Moreover, each block was preceded by practice trials.

Accordingly, participants were clearly aware of the fact that in the diagnostic block they would only be confronted with diagnostic items and that their task was to say “yes” when they saw a diagnostic relation. Nonetheless, RTs were slower for the diagnostic than for the predictive block. Thus the detrimental effects of a mismatch between the order of presentation and causal order is not a simple bias that can be remedied by instructions and blocking of trials; rather, it seems to be an automatic consequence of the mechanism by which the cognitive system accesses causal knowledge.

Experiment 3

Experiments 1 and 2 tested the causal-model view against the associative view by asking participants to access causal relations using pairs of items that had symmetric association and similar statistical strengths in both causal directions. Experiments 3 and 4 extended the relation-verification paradigm by asking participants to check whether the events described in an item pair were *causally related* or whether they were *associated*. In terms of causal-model theory, representing and evaluating a causal relation requires a representation in which each event is mapped to a specific role, the “cause” and the “effect”. This mapping will be made more quickly if the temporally-prior item is the cause, yielding shorter RTs in the predictive than the diagnostic presentation order. In contrast, the general relation “associated” does not differentiate the roles of the two items, so no such mapping process is required. Accordingly, causal-model theory predicts an interaction, such that the advantage of the predictive over the diagnostic ordering should be observed only for queries about causality, and not for queries about association. In contrast, the associative view does not predict an order effect for either query (given that pairs are selected to have symmetric association strengths).

Method

Participants and Design

Forty-four UCLA undergraduate students participated in this experiment and received course credit. Their vision was normal or corrected to normal. Half of this group was assigned to the causal and half to the associative condition. Two participants in the associative and three in the causal condition had to be replaced because of high overall error rates (> 20%).

Stimuli

All 44 participants saw causal word pairs and unrelated filler word pairs. In both the causal and the associative condition participants were presented with 64 causal and 64 unrelated filler word pairs. The causal items were the same stimuli used in the previous experiments. The unrelated filler word pairs were made up of single words from causal relations that we did not select for the previous studies due to excessively divergent ratings in the norming study. The words were combined randomly to avoid any semantic or associative relationships. As in the earlier experiments, the causal items were presented in either the predictive or the diagnostic direction, with the version of the items being counterbalanced across participants.

For the associative condition, in order to disguise the fact that the study focused on causally-related items, we also included some associated, but non-causal items (e.g., “emerald-ring”, “vehicle-bicycle”). Additional unrelated items were also included in this condition to equate the proportion of correct “yes” and “no” responses. The associative condition included 64 causal, 32 associated but non-causal items, and 96 unrelated filler word pairs (see Table A2 in the Appendix for the list of unrelated words). The associative condition thus included more items than the causal condition. The causal and the associated word pairs were the same stimuli used in the previous experiments.

Procedure

In the causal condition participants were requested to judge whether there was a causal relation (as in Experiment 1). In the associative condition participants were instructed that their task was to judge whether “there is an association between the two words.” They were told to press the “A” key “if the words are related in some meaningful way”, and to press the “N” key if there was no relation. In other respects the procedure was the same as in the previous experiments.

Results

Figure 3 displays the mean RTs for the causal items for both the causal and the associative instructions. As in the previous experiments, with causal instructions diagnostic items led to slower RTs than did predictive items. In contrast, with associative instructions the very same items were responded to equally quickly in both orders. The analysis of the design for the within-subject factor of causal direction (predictive vs. diagnostic) and the between-subjects factor type of instruction (causal vs. associative) yielded a significant effect for causal direction, $F(1, 42) = 7.81$, $p < .01$, $MSE = 1231.4$, with an effect size of a partial eta square (η^2) = .16. There was no significant main effect of instruction, $F(1, 42) = .94$, $p = .34$, $MSE = 39283.3$. The interaction for the two factors was marginally significant, $F(1, 42) = 4.03$, $p = .05$, $MSE = 1231.4$ with a small effect size of $\eta^2 = .09$.

Insert Figure 3 about here

A post-hoc analysis using the Scheffé test showed a significant difference between the predictive and diagnostic causal word pairs for the causal instruction, $p < .05$. The difference for the causal word pairs for the associative instruction was not significant, $p = .96$. Moreover, the diagnostic items in the causal condition led to longer RTs than both the diagnostic items ($p < .01$) and the predictive items ($p < .01$) in the associative condition. In contrast, RTs for the predictive items in the causal condition were not significantly different from either type of item in the associative conditions.

We also conducted an analysis of the errors. The error rates for the predictive and diagnostic items were 11% and 15%, respectively, in the causal condition, and 13% and 14%, respectively, in the associative condition. An analysis of variance revealed an effect of causal direction, $F(1, 42) = 4.57$, $p = .04$, $MSE = 3.78$, with an effect size of $\eta^2 = .09$. No other factor was significant.

Discussion

The main finding in Experiment 3 was that participants proved capable of task-specific access to different types of relations when queried in different ways about an identical set of causally-related items. When participants were asked to judge whether a causal relation was present, the usual RT advantage of predictive over diagnostic relations was found. However, for the identical items, access was equally fast regardless of causal direction when the instructions asked participants to judge whether the words within the item pairs were associated. This interaction supports the assumption of causal-model theory that mapping into causal roles is required for assessing a causal relation, but not for assessing an associative relation.

On the surface, the observed difference between the role of temporal order in causal versus associative queries runs counter to findings from the studies by van der Meer and colleagues (Krüger et al., 2001; van der Meer et al., 2002), who appeared to have shown that temporal direction affects access to semantic memory both for judgments of temporal relations and judgments of more general semantic relatedness. However, in addition to possible differences between retrieval of causal versus temporal relations, differences in the materials may have led to this apparent divergence of findings. Participants in the studies by van der Meer and colleagues only saw temporally-related or unrelated materials, and hence may have adopted a strategy of assessing temporal relations under both sets of instructions. In contrast, the present experiment included non-causal associates in the associative-instruction condition, which would have blocked such a strategy. In addition, van der Meer and colleagues did not directly compare access for two instructions within the same experiment, making it impossible to determine whether an interaction was present.

Experiment 4

In Experiment 3 we presented identical causal items and manipulated instructions that directed participants in different conditions to access causal or associative relations. To prevent participants in the associative condition from noticing that all the related items were causal, we added items with other types of associations in this condition. In Experiment 4 we switched to a within-subjects design in which participants were presented with items in blocks, with instructions to either evaluate causal or associative relations in the subsequent block of items, in order to determine whether our findings can be replicated in a within-subjects design. The design of Experiment 4 was intended to serve as a feasibility study for a later neuroimaging study based on the same design. In the present experiment, all participants were presented with the same lists of items.

Method

Participants and Design

Twenty-four UCLA undergraduate students participated for course credit. Their vision was normal or corrected to normal. Four participants were replaced because of computer failure, and three because of their high overall error rates (> 20%).

The stimuli were presented in a within-subject design; every participant saw every word pair once. There were ten blocks of trials for which the instruction was varied (see Procedure for details). The order of the blocks and the version of stimuli within a given block (predictive vs. diagnostic) were randomized and counterbalanced, respectively.

Stimuli

The stimulus material consisted of 64 causal, 16 strongly associated but non-causal (see Table A3), and 30 unrelated filler word pairs. The causal word pairs were the same as those used in the previous experiments, and the unrelated pairs were selected from the item set used in Experiment 3. The strongly associated word pairs (strength > 0.3) were selected from the USF database (Table A3). Blocks with strongly associated word pairs were included to provide a further test that participants indeed used association-retrieval strategies in the associative blocks, as strongly associated pairs should yield faster responses than the weakly associated pairs selected as causal items.

Procedure

Participants were instructed that they would be presented with ten blocks, each containing eleven different trials. Prior to each block they were prompted by either the word “CAUSAL?” or the word ”ASSOCIATIVE?” If they were prompted with the word ”CAUSAL?” they had to decide whether or not there was a causal relation between the two words; if the prompt word was “ASSOCIATIVE?” they had to decide whether or not these words were related in a meaningful way (as in Experiment 3). For both instructions participants were to press the letter “Y” if the designated relation held for the pair; otherwise they were to press “N”.

The session began with 26 practice trials (6 predictively-related causal items, 6 diagnostically-related causal items, 3 associated items and 11 filler items), with feedback. After the practice trials, the experimental blocks were presented. Each block started with a prompt (causal or associative). Participants were then presented with eleven trials per block. If the task was to assess causal relations, eight of these eleven trials were causal and three of these items were unrelated. We presented filler items in the block to ensure that participants processed the prompted relation. The order of the trials within a block was randomized. There were two different kinds of causal blocks, predictive or diagnostic. Within each causal block only one type of causal relation (predictive or diagnostic) was presented. After each block there was a 16 second break; then a new prompt was presented.

If the prompt was to look for associations, then eight word pairs were related and three filler word pairs were unrelated. The relationships within an associative block were either only causal or only associative. The causal trials within an associative block were either predictive or diagnostic. Overall, ten blocks were presented, including four blocks in which participants were prompted to look for a causal relation. Half of these four blocks only presented predictive pairs (along with unrelated filler items), and half presented only diagnostic pairs. Four other blocks contained either predictive (two blocks) or diagnostic (two blocks) pairs (along with fillers), and participants' task was to judge whether the items were associated. There were two blocks in which highly associated pairs were presented along with filler items, with the prompt to look for associative relations. The ten blocks were organized into two sequences of five blocks, which each contained one block of each type in a random order.

Results and Discussion

We first analyzed mean correct RTs for the five types of blocks: causal-predictive (i.e., causal instruction with predictively ordered items), causal-diagnostic, associative-predictive (i.e., associative instruction with predictively ordered items), associative-diagnostic and associative-associated. The results are depicted in Figure 4. An overall repeated-measures design with the factor block yielded a statistically significant difference, $F(4, 92) = 10.15$, $p < .01$, $MSE = 9792.6$ with an effect size of $\eta^2 = .31$. A planned contrast analysis revealed that there was a significant difference between the two types of causal word pairs in the causal blocks (mean RT of 872 ms for predictive order vs. 940 ms for diagnostic order), $F(1, 23) = 6.42$, $p = .02$, $MSE = 8696.80$, whereas no difference was obtained between these item types in the associative conditions ($F<1$). The diagnostic items in the causal conditions yielded slower RTs than the diagnostic items, $F(1, 23) = 3.16$, $p = .09$, *MSE* = 14557.5, and also than the predictive items, $F(1, 23) = 5.05$, $p = .04$, *MSE* = 11437.3, in the associative conditions. Finally, as expected, the strongly associated items yielded dramatically faster reaction times than the weakly associated predictive and diagnostic items within the associative task, $F(1, 23) = 39.4$, $p < .01$, *MSE* = 4910.6.

Insert Figure 4 about here

An analysis of errors revealed a significant difference for the error rates in the five different type of blocks, $F(4, 92) = 4.85$, $p < .01$, *MSE* = 1.63. This effect is solely due to the fact that the error rate was the lowest for the highly associated word pairs (4%) relative to the other four conditions (mean of 11%). In particular, the difference between the weakly associated causal items and the strongly associated items in the associative task was significant, $F(1, 23) = 32.6$, $p < .01$, *MSE* = .77.

Experiment 4 thus replicated the pattern of results found in Experiment 3, this time using a within-subjects design. We again observed an RT advantage for predictive over diagnostic pairs when the instructions asked for the assessment of causal relations, whereas no significant difference was observed when participants judged whether the item pairs were merely associated. This pattern was found with material that was structurally identical for all participants. Experiment 4 included strongly associated items, which produced much shorter RTs under associative instructions than were observed for the more weakly associated causal pairs. This finding supports the interpretation that participants were indeed sensitive to associative relatedness when the task requested them to check for associative relations. Overall, the results of Experiment 4 support the hypothesis that people are capable of accessing relations in a task-specific manner.

General Discussion

Summary

The four experiments reported here showed that causal relations are retrieved from semantic memory and evaluated more easily in the predictive cause-effect order than in the reverse diagnostic effect-cause order. We demonstrated this effect with materials that had equal association strength and statistical strength in both directions, assuring that the effect is not driven by asymmetries of association or statistical strength. In addition, we were able to show that the RT advantage for the predictive order persisted even when participants knew in advance what type of causal relation (predictive or diagnostic) they were going to see. The latter result indicates that this type of congruity effect is not due to interpreting causal relations primarily as referring to predictive relations, but is rather grounded in a deeper characteristic of causal semantic memory.

In addition, in Experiments 3 and 4 we were able to demonstrate a dissociation between retrieval and evaluation of causal relations versus general associative relations. In particular, the RT advantage of pairs in the predictive compared to the diagnostic order disappeared when participants were asked to assess whether the items were associated, rather than causally-related. This dissociation is consistent with the hypothesis that assessing causal relations requires mapping of the two events into specific causal roles (“cause” and “effect”), a process that is facilitated in the predictive order because the semantic features of the “temporally prior” event overlap with those of the “cause” role, whereas the features of the “temporally after” event overlap with those of the “effect” role (Hummel & Holyoak, 1997, 2003). In contrast, the general “association” relation does not distinguish the roles of the two items in a pair; hence temporal order of presentation no longer has an effect on decision time. The present findings refute theories that treat causal relations as simply instances of general associative relations (Shanks & Dickinson, 1987).

Implications for Cognitive Neuroscience

The present experiments (especially Experiment 4) may provide a paradigm that can be extended to allow collection of neuroimaging data. Robin and Holyoak (1995) argued that the prefrontal cortex (PFC) is critical to processing relations, a hypothesis supported by neuropsychological findings with patients suffering from degenerative disorders of the frontal cortex (Morrison et al., in press; Waltz et al., 1999). Studies using functional magnetic resonance imaging (fMRI; Christoff et al., 2001; Kroger et al., 2002; Prabhakaran et al., 1997) have also implicated the PFC as the neural locus for the working memory required to manipulate structured relations in reasoning. In addition, a study by Münte, Schiltz and Kutas (1998) used event-related potentials to identify left prefrontal areas that respond to linguistic inputs that necessitate internal manipulation of temporally-related events. Accordingly, we predict that left dorsolateral PFC will be selectively activated during the verification of causal relations presented in the non-canonical effect-to-cause order. In addition, we hypothesize that causal judgments in both directions will activate neural areas distinct from those involved in processing simple semantic associations between words.

Other Future Directions

In the present study we found that predictive causal relations can be accessed faster than diagnostic ones. One interesting research question is whether this effect can be reduced, or perhaps even reversed, with expertise. Physicians, for example, are confronted with diagnostic questions all day long, so it seems likely that they will learn to access knowledge in the diagnostic direction very efficiently. Future research should address the question of whether the process of accessing and evaluating causal relations changes with diagnostic expertise.

Experiments 3 and 4 demonstrated a dissociation between causal and associative relations. Causal relations are a type of basic semantic relation in which the entities being related play specific roles (in contrast to associative relations, which do not distinguish roles).

Other semantic relations may also show similar dissociations. For example, Spellman, Holyoak and Morrison (2001) have shown that several basic semantic relations (e.g., category membership, antonymy) play a role in priming lexical access. Some of these, such as category membership, clearly distinguish specific roles (e.g., “all canaries are birds” is true, but “all birds are canaries” is false, reflecting the asymmetry of the roles of “instance” and “category”). The paradigm used in the present study could be readily adapted to investigate the cognitive representations of other relations in semantic memory. It is important to identify both the commonalities and differences among the varied semantic relations that constitute our knowledge of the world.

References

Cheng, P. W. (1993). Separating causal laws from casual facts: Pressing the limits of statistical relevance. In D. L. Medin (Ed.), *The psychology of learning and motivation*, Vol. 30 (pp. 215-264). New York: Academic Press.

Christoff, K., Prabhakaran, V., Dorfman, J., Zhao, Z., Kroger, J. K., Holyoak, K. J., & Gabrieli, J. D. E. (2001). Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *NeuroImage*, *14*, 1136-1149.

Cobos, P. L., López, F. J., Cano, A., Almaraz, J., & Shanks, D. R. (2002) Mechanisms of predictive and diagnostic causal induction. *Journal of Experimental Psychology: Animal Behavior Processes, 28,* 331-346.

Friedman, W. J. (1990). *About time.* Cambridge/MA: MIT Press.

Friedman, M. J. (2002). Arrows of time in infancy: The representation of temporal-causal invariances. *Cognitive Psychology, 44,* 252-296.

Hausman, D. (1998). *Causal asymmetries*. Cambridge: Cambridge University Press.

Hodges, J. R., & Patterson, K. E. (1997). Semantic memory disorders. *Trends in Cognitive Science*, *1*, 67-72.

Hume, D. (1978). *A treatise of human nature.* Oxford: Clarendon Press. (Original work published 1739).

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427-466.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220-264.

Kant, I. (1950). *Critique of pure reason* (N.K. Smith, Trans.). London: Macmillan. (Original work published 1781).

Kroger, J. K., Saab, F. W., Fales, C. L., Bookheimer, S. Y., Cohen, M. S., & Holyoak, K. J. (2002). Recruitment of anterior dorsolateral prefrontal cortex in human reasoning: A parametric study of relational complexity. *Cerebral Cortex*, *12*, 477-485.

Krüger, F., Nuthmann, A., & van der Meer, E. (2001). Pupillometric indices of temporal order representation in semantic memory. *Zeitschrift für Psychologie, 209*, 402-415.

Kubose, T. T., Holyoak, K. J., & Hummel, J. E. (2002). The role of textual coherence in incremental analogical mapping. *Journal of Memory and Language*, *47*, 407-435.

Lien, Y., & Cheng, P. W. (2000). Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology*, *40*, 87-137.

Morrison, R. G., Krawczyk, D. C., Holyoak, K. J., Hummel, J. E., Chow, T. W., Miller, B. L., & Knowlton, B. J. (in press). A neurocomputational model of analogical reasoning and its breakdown in Frontotemporal Lobar Degeneration. *Journal of Cognitive Neuroscience*.

Moss, H. E., Ostrin, R. K., Tyler, L. K., & Marslen-Wilson, W. D. (1995). Accessing different types of lexical semantic information: Evidence from priming. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 21,* 863-883.

Münte, T. F., Schlitz, K., & Kutas, M. (1998). When temporal terms belie conceptual order. *Nature, 395,* 71-73.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92,* 289-316.

Murphy, G. L. (2002). *The big book of concepts.* Cambridge, MA: MIT Press.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms. http://www.usf.edu/FreeAssociation/.

Prabhakaran, V., Smith, J. A. L., Desmond, J. E., & Glover, G. H. (1997). Neural substrates of fluid reasoning: An fMRI study of neocortical activation during performance of the Raven's Progressive Matrices Test. *Cognitive Psychology, 33,* 43-63.

Robin, N., & Holyoak, K. J. (1995). Relational complexity and the functions of prefrontal cortex. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 987-997). Cambridge, MA: MIT Press.

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H. Bowwer (Ed.). *The psychology of learning and motivation, 21* (pp. 229-261). New York: Academic Press.

Shanks, D. R., & López, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition, 24,* 511-22.

Spellman, B. A., Holyoak, K. J., & Morrison, R. G. (2001). Analogical priming via semantic relations. *Memory & Cognition*, *29*, 383-393.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization and memory* (pp. 381-403). New York: Academic Press.

Tyler, L. K., & Moss, H. E. (1997). Functional properties of concepts: Studies of normal and brain-damaged patients. *Cognitive Neuropsychology, 14,* 511-545.

Van der Meer, E., Beyer, R., Heinze, B., & Badel, I. (2002). Temporal order relations in language comprehension. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 28,* 770-779.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121,* 222-236.

Waldmann, M. R., Holyoak, K. J., & Fratianne, A. (1995). Causal models and the acquisition of category structure. *Journal of Experimental Psychology: General*, *124*, 181-206.

Waldmann, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak & D. L. Medin (Eds.). *The psychology of learning and motivation, 34: Causal learning* (pp. 47-88). San Diego: Academic Press.

Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *26,* 53-76.

Waldmann, M. R. (2001). Predictive versus diagnostic causal learning: evidence from an overshadowing paradigm. *Psychonomic Bulletin and Review, 8,* 600-608.

Waltz, J. A., Knowlton, B. J., Holyoak, K. J., Boone, K. B., Mishkin, F. S., de Menezes Santos, M., Thomas, C. R., & Miller, B. L. (1999). A system for relational reasoning in human prefrontal cortex. *Psychological Science*, *10*, 119-125.

Author Note

The experiments were conducted in partial fulfillment of the requirements for a Ph.D., which was obtained by the first author under the supervision of the second author at the University of Göttingen. The experiments were conducted during a visit of the first author at the Psychology Department of the University of California, Los Angeles, which was funded by the HSP II fellowship of the German Academic Exchange Service (DAAD). Portions of this research were presented at the 2001 Annual Meeting of the Psychonomic Society, Orlando, Florida, and the 2002 conference of the Deutsche Gesellschaft für Psychologie, Berlin.

Correspondence concerning this article should be addressed to D. Sellner, who is now at the Otto-von-Guericke University, Department of Neurology II, Leipziger Str. 44, 39120 Magdeburg, Germany, e-mail: sellner@neuro2.med.uni-magdeburg.de.

Figure Captions

Figure 1: Example of a causal model represented by a directed graph.

Figure 2: Example of a predictive trial.

Figure 3: Mean reaction times and standard error bars for causal stimuli in the associative and causal conditions for pairs presented in the predictive and diagnostic orders (Experiment 3).

Figure 4: Mean reaction times and standard error bars for the five different types of blocks (Experiment 4).

Fig. 1

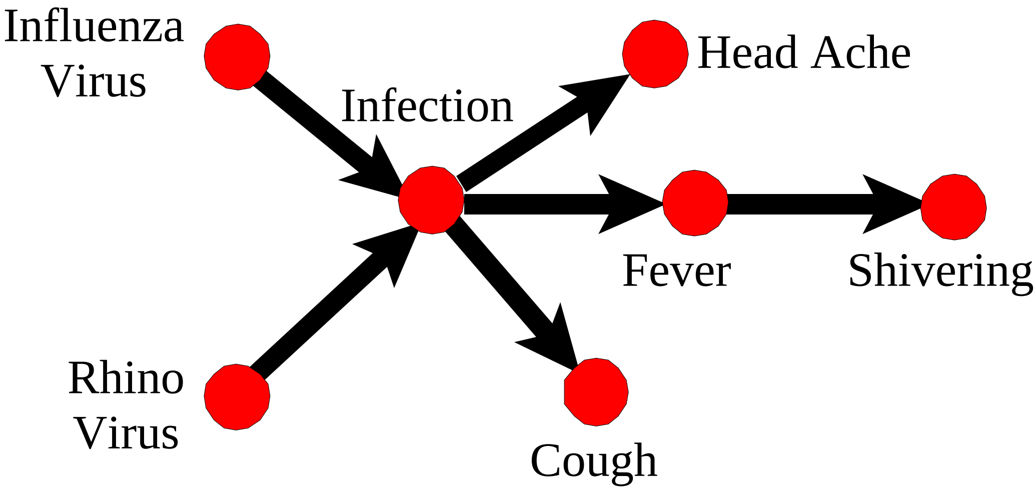
Influenza Virus
Infection
Head Ache
Fever
Shivering
Rhino Virus
Cough

Fig. 2

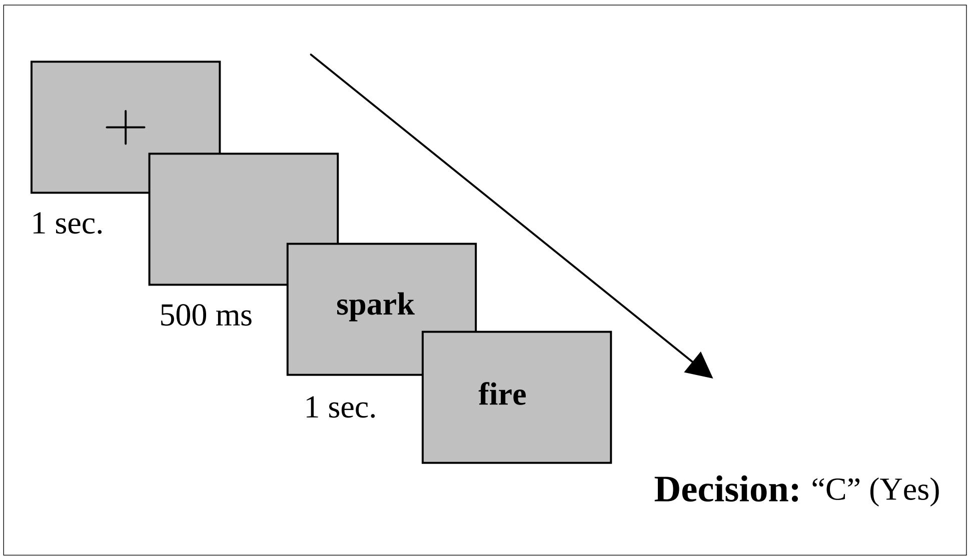

Fig. 3

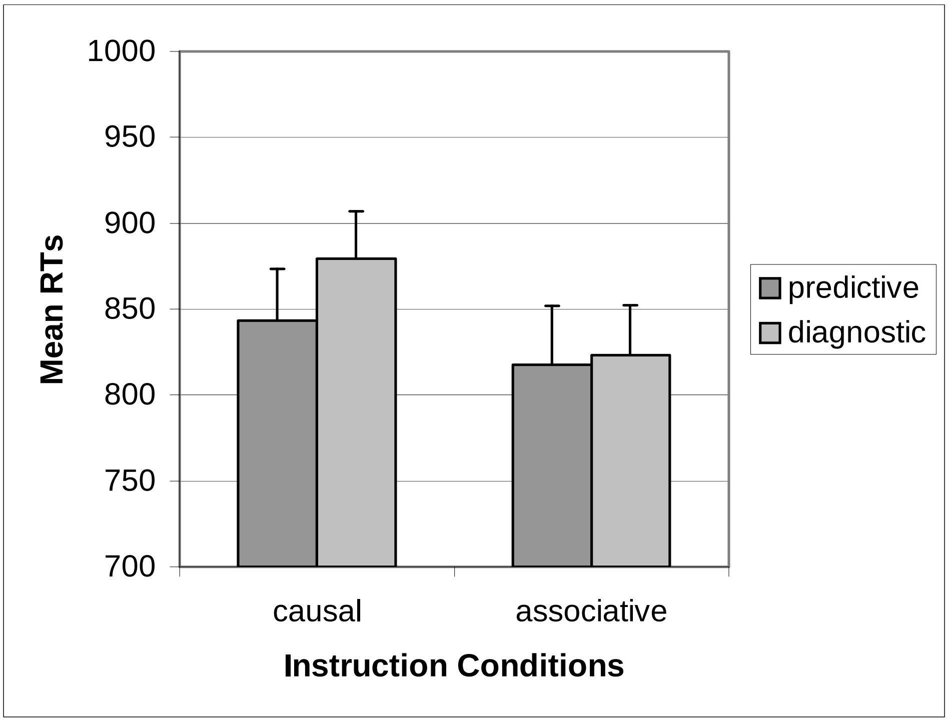
1000
950
900
850
800
750
700
Mean RTs
predictive
diagnostic
causal
associative
Instruction Conditions

Fig. 4

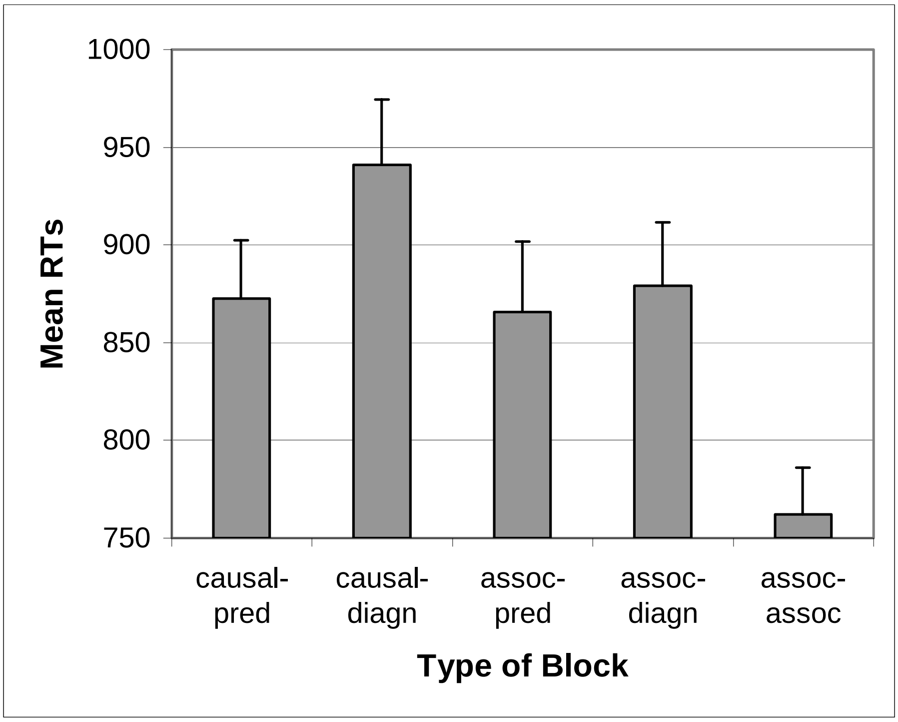
1000
950
900
850
800
750
Mean RTs
causal-pred
causal-diagn
assoc-pred
assoc-diagn
assoc-assoc
Type of Block

Appendix

Table A1:

Normed causal and associative filler word pairs that were used in the experiments. FSA and BSA indicate the forward and the backward strength of association in the USF data base (Nelson et al., 1998), respectively. FR describes the difference of the frequency ratings in our norming study. Positive values indicate a higher perceived frequency for the predictive direction, whereas negative values indicate a higher perceived frequency for the diagnostic direction.

Causal word pairs		FR	FSA	BSA	Associative word pairs		FSA	BSA
absence	withdrawal	-12	0.01	nn	acrobat	athletes	0.02	0
acid	corrosion [x]	5	0	nn	agency	firm	0.02	0
alcohol	accident	-4	0	0	ambulance	rush	0.01	0
attack	defense	1	0.05	0	antelope	gazelle	0.03	0.03
bacteria	infection	0	0.01	0	atlas	dictionary	0.01	0
bang	deafness	-11	0	nn	basketball	teams	0	0.02
beat	bruise	8	0	0.05	bedroom	furniture	0.01	0
betrayal	distrust	13	0	0	caffeine	mountain	0.01	0
birthrate	population	1	nn	0	car	plane	0.01	0.03
carcinogen	tumor	2	nn	0	chipmunks	acorn	0	0.01
chromosome	gender	2	0.03	0	claw	dogs	0.01	0
compliment	blush	-4	0	0	cocktails	fruits	0	0.01
crime	arrest	-20	0	0.02	computer	apple	0.01	0.02
crush	damage	3	0	0.01	control	volume	0	0.01
dairy	diarrhea	-7	0	nn	dagger	fight	0.01	0
diet	hunger	16	0.01	0	decency	respect	0.03	0

Causal word pairs		FR	FSA	BSA	Associative word pairs		FSA	BSA
disease	injection	13	0	0.01	elephant	zebra	0	0.01
drought	famine	-2	0	0.02	elevator	floor	0.01	0
drug	relief	23	0	0	email	attachment	nn	0
education	career	-7	0	0.03	engine	roar	0.02	0
espionage	treason [x]	-4	0	0	envy	admire	0.03	0.03
eyedrops	dilation	-20	nn	0	family	sibling	0	0.03
fertilizer	growth	15	nn	0	forecast	weather	nn	0
fracture	cast	12	0	0	girl	maid	0	0.01
frequency	pitch	-3	0	0	glands	pituitary	0.03	nn
frowning	wrinkles	24	0.02	0	glass	window	0.01	0.02
gang	riot	-6	0	0	grab	pull	0.03	0.11
gases	explosion	-24	0	0	graduation	gown	0.02	0
genes	baldness [x]	-10	0	0	graph	numbers	0.02	0
gold	wealth	7	0.02	0	ground	potatoes	0	0.02
hormones	mood	4	0	0	harbour	seaman	0	nn
humidity	sweat	27	0.02	0	insurance	estimate	0	0.01
illness	treatment	12	0	0.04	kill	theft	0	0.01
invitation	visit	10	0	0	kindness	sympathy	0	0.02
joke	amusement	-1	0	nn	latin	medicine	0.02	0
lamp	heat	-3	0	0	lettuce	vegetables	0.08	0
lesion	scar	7	nn	0	lime	corona	0.02	nn
lightning	fire	4	0	0	lover	girlfriend	0.03	0.02
magnet	attraction	6	0.1	0.04	money	groceries	0	0.02
moon	tide	-3	0	0.02	mother	wife	0	0.03

Causal word pairs		FR	FSA	BSA
movie	nightmare	0	0	0.02
mutation	cancer	-19	nn	0
nuts	allergy	-9	0	0
order	delivery	-9	0	0
pain	aggression [x]	10	0	nn
panic	escape	-3	0	0
period	cramps	4	0	0.07
pollution	asthma	-7	nn	0
pressure	bursting	-22	0	0
sadness	crying	-20	0.05	0.13
salt	thirst	21	0	0
scratch	blood	-7	0.01	0
shampoo	tears	-6	0	0
shock	scream	3	0	0
spice	flavor	-24	0.05	0.02
spill	stain	-7	0	0
sprain	swell	18	0	nn
stress	fatigue	6	0	0.02
study	pass	6	0	0
sunlight	freckles	5	0	0.02
sweets	cavity	-27	0	0
training	fitness	12	0	0
trash	stink	20	0	0

Associative word pairs		FSA	BSA
newspaper	gossip	0	0.01
office	employment	0	0.02
ounce	gallon	0.02	0
painting	wall	0.05	0.02
paper	envelope	0	0.03
patty	hamburger	nn	0.02
planter	farmer	nn	0.02
power	voltage	0.01	0.02
propeller	helicopter	nn	0.03
protestants	baptist	0	0.01
ring	emerald	0	0.03
round	screw	0	0.01
sandwich	tomatoes	0	0
security	force	0.01	0
session	course	0.01	0
shape	curve	0.01	0.02
shrimp	ocean	0.03	0
soup	cracker	0.03	0.04
spray	roach	0.03	0.03
story	passage	0	0
terms	meaning	0.01	0
test	hypothesis	0	0.03
towers	skyscraper	0.02	0.05

[x] These four items were presented in Experiment 1, and were excluded in Experiments 2 to 4.

Causal word pairs		FR	FSA	BSA	Associative word pairs		FSA	BSA
trauma	coma	-9	0.01	0	tuba	saxophone	0.03	0.01
UVlight	tanning	-9	nn	nn	umbrella	tote	0	0.01
vacuum	suction	14	0.2	nn	uniforms	officers	0	0.01
virus	epidemic	-24	0	nn	vehicle	bicycle	0.02	0.01
wind	erosion	10	0	nn	vessel	vein	0.01	0.02
Sum of rating differences		0						
Mean strength of association			0.01	0.01	Mean strength of association		0.01	0.01

Table A2:

Unrelated filler word pairs used in Experiment 3.

Unrelated word pairs		Unrelated word pairs		Unrelated word pairs	
ambulance	window	Girl	agriculture	patty	kitchen
ankle	farming	glands	sailor	phone	switch
archer	phonebook	Glass	rush	planter	power
basin	academy	Grab	screw	point	queen
basket	kite	grass	fist	posters	hamburger
beauty	compass	gymnastic	vegetables	potatoes	insult
bedtime	tomatoes	harbor	garage	printer	angel
brush	dices	insurance	icecream	radiation	jockey
bubble	velvet	investor	priming	report	bike
caffeine	sky	kill	clock	revolting	roach
chef	fear	king	ceiling	rise	bank
clown	map	lamb	bearing	roles	cats

Unrelated word pairs		Unrelated word pairs		Unrelated word pairs	
conductor	groceries	landscape	maths	roof	seaman
consulate	door	latin	disgust	round	roar
cookie	nose	lead	curve	salad	respect
couch	dough	leather	pull	savage	airport
dancer	liquid	lemon	soccer	security	floor
deer	pencil	leopard	river	shape	aluminum
diabetes	penny	lettuce	bars	ship	e-mail
diamond	gear	medicine	passage	smock	plug
disk	ground	mile	apron	spray	theft
doorbell	architect	miracle	ginger	square	indian
drill	guest	money	piano	store	session
eagle	child	mouse	light	story	dinner
eggs	liar	mousepad	justice	survivor	cup
elbow	pistol	mouth	actor	tea	graph
elevator	force	needle	currency	therapy	barrel
engine	glove	office	mirror	traffic	armrest
fabric	soup	onions	sphere	tree	maid
fairy	officers	page	tuxedo	truck	zebra
gate	cord	pants	bandage	voltage	mountain
gentlemen	chapter	parents	weather	water	boxer

Table A3:

Normed strongly associated stimuli used in Experiment 4.

Associated word pairs		FSA	BSA
checkers	chess	0.26	0.22
cigarette	smoke	0.45	0.32
corals	reef	0.38	0.40
cuss	swear	0.35	0.19
fist	knuckles	0.05	0.26
garden	landscape	0.00	0.01
halloween	pumpkin	0.14	0.03
leadership	team	0.03	0.00
nouns	adjectives	0.04	0.33
patient	doctor	0.37	0.03
positive	negative	0.63	0.60
pyramid	Egypt	0.35	0.34
razor	blade	0.29	0.24
reunion	gathering	0.00	0.02
robber	thief	0.36	0.22
town	city	0.53	0.31
Mean strength of association		0.26	0.22