

Defining Structural Similarity*

Dedre Gentner, Arthur B. Markman

Department of Psychology, Northwestern University

Department of Psychology, University of Texas

gentner@northwestern.edu

markman@mail.utexas.edu

There is general agreement that structural similarity — a match in relational structure — is crucial in analogical processing. However, theories differ in their definitions of structural similarity: in particular, in whether there must be conceptual similarity between the relations in the two domains or whether parallel graph structure is sufficient. In two studies, we demonstrate, first, that people draw analogical correspondences based on matches in conceptual relations, rather than on purely structural graph matches; and, second, that people draw analogical inferences between passages that have matching conceptual relations, but not between passages with purely structural graph matches.

Introduction

The discovery of common structure is a central aspect of analogical processing (Gentner, 1983; Gentner & Markman, 1997; Hofstadter, 2001; Holyoak & Thagard, 1995; Ramscar & Yarlett, 2003). But what exactly enters into this process? Most theories of analogy agree that analogical processing involves finding a correspondence between the conceptual structures of the

Journal of Cognitive Science 6: 1 - 20, 2005.

© 2005 *Institute for Cognitive Science, Seoul National University.*

*This work was supported by grant N00014-92-J-1098 from the ONR Cognitive Science Program to D. Gentner, by NSF-ROLE grant 21002/REC-0087516 to D. Gentner, and by NIDA Grant R21-DA015211 and a Fellowship in the IC² Institute to Art Markman. We thank Kathleen Braun for help in conducting and analyzing the studies and Mark Keane and Miriam Bassok for helpful comments on an earlier version of this paper.

two domains compared. As Kokinov and French (2003) put it, “Experimental work has demonstrated that finding this type of structural isomorphism between base and target domains is crucial for mapping.” For example, when comparing

- (a) Mary hugged John because she loves him.
- (b) Rover nuzzled Sarah because he loves her.

people seek a *structurally consistent match* between the two representations: that is, a match that preserves the constraints of *one-to-one mapping* and *parallel connectivity* (Forbus, Gentner & Law, 1995; Gentner & Markman, 1997). One to one mapping requires that each element in one domain match to at most one element in the other domain.¹ For example, if Mary in (a) is placed in correspondence with Rover, she cannot also be placed in correspondence with Sarah. Parallel connectivity states that if two predicates are placed in correspondence, then their arguments must also be placed in correspondence; that is, there must be like bindings between the two analogs (See Figure 1). For example, if the two causal relations in (a) and (b) are matched, then their arguments must also be matched: Loves → Loves and Hug → Embrace. Likewise, if Loves → Loves, then Mary → Rover and John → Sarah. Almost all current theories of analogy agree on the importance of structural consistency, although models vary in whether it is implemented as a hard constraint (Falkenhainer, Forbus, & Gentner, 1989; Gentner & Markman, 1997) or as a soft constraint (Halford, 1992; Holyoak & Thagard, 1989; Hummel & Holyoak, 1997; Mitchell, 1993).

Although there is universal agreement that structural similarity is crucial in analogical processing, there is a dismaying lack of agreement on exactly what is meant by “structural similarity.” In particular, theories disagree as to whether conceptual similarity in the relations is required, or whether a pure graph match is sufficient. To see this issue, contrast the initial sentence pair (a) and (b) with

¹ One principled exception occurs when two or more arguments are relationally equivalent, and can be collapsed into one: e.g., in “The hydrogen atom is like the solar system”, the nine planets can be treated as one and placed in correspondence with the electron.

the additional pair

- (a) Mary hugged John because she loves him.
- (c) Fred heated the sandwich before he ate it.

Both pairs — (a)/(b) and (a)/(c) — have identical graph structure, as shown in Figure 1. However, in pair (a)/(b), the corresponding relations in the graph structure are conceptually similar, whereas in pair (a)/(c), they are not. The crux of the disagreement is whether pure graph-structure matches like pair (a) and (c) are processed as analogies. According to Gentner (1983), structural similarity involves conceptual similarity between corresponding relations. Thus pair (a)/(b) is analogous, but pair (a)/(c) is not. A contrasting view is that structural similarity need only entail graph matches (Holyoak & Thagard, 1989; Hummel & Holyoak, 1997). Theories in this camp concede that pure graph-matching analogies such as (a)/(c) are more difficult to process than relationally similar pairs, but hold that the processes are nonetheless fundamentally the same.

This may sound like a simple matter of terminology, but there is more at stake. The problem of placing two structural representations in correspondence is one of matching two directed acyclic subgraphs. This kind of graph matching is known to be in the class of NP-hard problems; the size of the computation required increases exponentially (or worse) with the size of the representations. Even if one settles for approximate rather than exact solutions, the computational burden still grows rapidly with the number of elements. Thus, any psychologically plausible process for finding analogical correspondences must control the computational burden of the match. This issue is particularly important because analogical comparison is often assumed to be a subcomponent of many other cognitive processes, such as problem solving and categorization.

Computational models of analogy can reduce the complexity of the graph match in one of two ways: by using conceptual similarity to constrain the match, or by using selective projection to reduce the size of the search set. In the conceptual similarity method, an initial parallel matching process is carried out that finds all similar pairs (both predicates and elements) in the base and target, without regard to the structure of the match. Only these conceptually

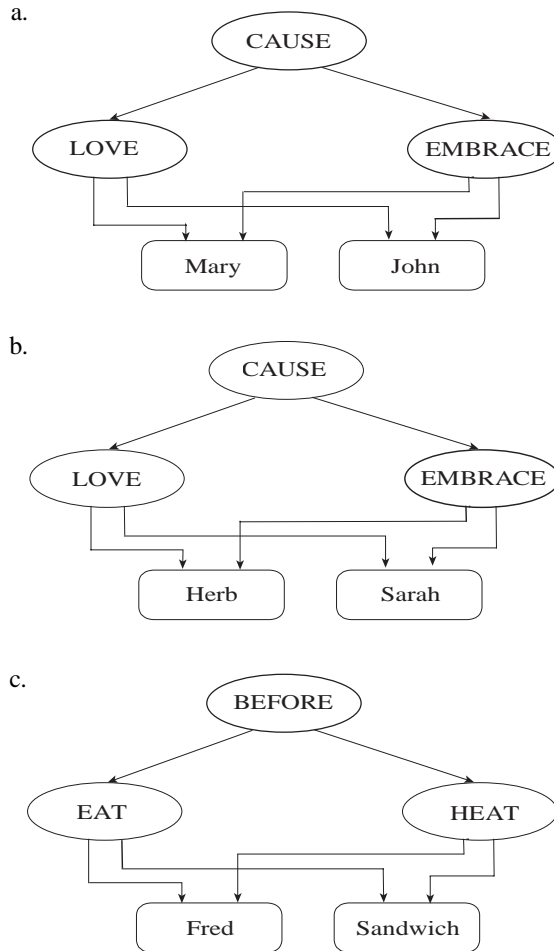


Figure 1. Graph representations for sentences (a), (b) and (c).

matching pairs enter into further processing, thus reducing the effective size of the problem. This method is used in the Structure-mapping Engine (SME) (Falkenhainer, Forbus, & Gentner, 1989; Forbus, Ferguson, & Gentner, 1994; Forbus, Gentner, & Law, 1995).

This initial set of local conceptual matches in SME includes both object matches and relational matches. It typically includes n-to-1 matches and other

matches that will later be rejected. In the next stage, structural consistency is imposed, with the effect of sorting the matches into structurally consistent kernels. Then these kernels are combined into one or more global mappings. The systematicity bias is implemented by a trickle-down computation, in which each matching predicate passes down a fraction of its evidence to its arguments. Finally, inferences are drawn by a kind of pattern completion from base to target (For details, see Forbus, Gentner, & Law, 1995; Gentner & Markman, 1997). Structural alignment and inference are computationally intensive. However, because the computation is restricted to the set of conceptually matching components (rather than processing the entire graphs for the two domains), the process is rendered tractable. Thus, the initial parallel conceptual match stage is crucial to SME's operation.

The second method for simplifying the graph-matching problem is to first select a particular set of assertions or structures in one domain and project only that set to the other domain. The order of projection may be chosen on grounds of systematicity or centrality, or it may simply mirror input order. In either case, the set of potential matches is reduced, because only the selected assertions in the first domain need to be considered. Once a predicate is selected, the best match in the other domain can be determined on the basis of a purely structural match. In models that use this directional projection method — e.g., Greiner's (1988) NLAG, Keane's (1997) IAM and Hummel and Holyoak's (1997) LISA — conceptual similarity is not necessary, though it may be used to facilitate processing or to decide ambiguous cases. Rather, what most matters is the selection and ordering of predicates to be mapped.

To summarize, while all extant theories of analogy consider structural similarity central, this key concept is defined differently in different theories. On the relational similarity view, as in structure-mapping theory, structural similarity requires conceptual similarity between corresponding relations. On the pure graph-isomorphism view, as in IAM and LISA, structural similarity requires only graph isomorphism; although conceptual similarity between relations can be helpful, it is not central to the algorithm, nor is it considered fundamental to the nature of analogy. Returning to our example, both views would agree that the relationally similar pair [(a) and (b)] is processed as an analogy. However, the views diverge on the graph-isomorphic pair [(a) and (c)]. On the relational similarity view, such nonconceptual matches are not

processed as analogies, but are matched (if at all) as some kind of logical puzzle. Indeed, SME cannot process such matches in its normal mode of processing, because it requires conceptual similarity to find potential correspondences in its initial parallel matching stage. In contrast, on the graph-isomorphism view, pair [(a) and (c)] is simply a very difficult analogy. For example, LISA can process such non-conceptually-similar matches as analogies, although its processing is more efficient when given conceptually similar analogs (Hummel & Holyoak, 1997, p. 450; Kubose, Hummel & Holyoak, 2002). Thus what's at stake here is the kinds of algorithms that are viable in a computational simulation of analogy.

Comparing graph isomorphism with relational similarity

The question we ask is what constitutes structural similarity in analogy: that is, what kind of match is used in human analogical processing. To answer this we need an operationalization of how to detect analogical processing. We focus on two key phenomena: (1) how people align the representations, as discussed above; and (2) whether and how people draw inferences from the match. The generation of candidate inferences is a core aspect of analogical reasoning (Gentner & Markman, 1997; Kokinov & French, 2001; Holyoak & Thagard, 1995; Spellman & Holyoak, 1996). The idea is that candidate inferences are generated during analogical processing by mapping information from the base to the target that is structurally consistent with the match (Bowdle & Gentner, 1997; Clement & Gentner, 1991; Markman, 1997; Ramscar & Yarlett, 2002; Spellman & Holyoak, 1996). Algorithmically, this can be done either by completing the pattern match achieved in the alignment process by mapping across further information connected to the matching relational structure (Falkenhainer et al., 1989) or by copying information from the base into a queried slot and substituting in corresponding elements of the target (Holyoak, Novick, & Melz, 1994).

We present two studies aimed at clarifying the nature of structural similarity in analogy. The logic of our studies is to give people pairs of passages that permit different kinds of matches — e.g., a pure graph isomorphism vs. a relational similarity match — and see (1) which correspondences people select; (2) whether they consider the match a good analogy; and (3) whether the

match leads to any inferences.

We operationalized graph structure in two different ways across the two studies. In Experiment 1, we took graph structure to be given by subject-verb-object sentence grammar. In Experiment 2 we used conceptual representations to determine the graphs. These two methods can be exemplified by a simplified example from Keane’s (1997) investigation. Keane asked people to find correspondences between sets of sentences with non-conceptually similar predicates: e.g.,

Joel sees Marge.	Ruby motivates Doris.
Joel hugs Marge.	Ruby knows Doris.
Bert sees Marge.	Lana motivates Doris.

Although participants made some errors, most could decide that Joel corresponds to Ruby and Doris to Marge.² One way they could have done this is by using grammatical matches — e.g., by noting that Joel and Ruby are each the subjects of two different verbs, and that Doris with Marge are the sentential objects in all three sentences. A second way people might arrive at a graph match (which we take up in Experiment 2) is to set up conceptual graphs and match the parallel arguments of corresponding relations. For example, one might decide sees corresponds to motivates (because there are two instances of each in their respective sets); this correspondence then dictates that their arguments (Joel → Ruby; Marge → Doris; Bert → Lana) must match in the corresponding order. In Experiment 1, we operationalized graph structure as subject-verb-object sentence grammar. This has the advantage of requiring few assumptions beyond ordinary syntactic rules.³ In Experiment 2, we used

² The fact that people can derive correspondences between such sets, if told to find a match (Keane, 1997), has been used as evidence that these sets qualify as (very difficult) analogies (Hummel & Holyoak, 1997; Keane, 1997). But our evidence suggests that these nonconceptual matches are processed very differently from analogies.

³ The terms “structure” and “syntax” are often used interchangeably in discussions of representation. That is, the term “syntax of a match” is used to refer to the structure of the representation. However, this can lead to confusion with the sense of “syntax” as the grammatical syntax of the sentences that describe the domains. While graph

conceptual graphs. In both cases, we pitted graph isomorphism against relational similarity.

Experiment 1

In order to explore the distinction between pure graph isomorphism — here operationalized as grammar-matching — and relational similarity, we gave people the sentence sets shown in Table 1. These sentences are constructed so that the correspondences dictated by the semantics of the relations are different from those dictated by the grammatical roles of the elements in the sentences. For example, if the sentence pair “Freddy chases after Fido” and “Rover runs from Bobby” is matched on the basis of grammar, the correspondences are Freddy → Rover (the grammatical subjects) and Fido → Bobby (the grammatical objects). However, if people seek to find relational similarity between conceptual representations, they should represent the two sentences roughly as *chases (Freddy, Fido)* and *chases (Bobby, Rover)*. In this case, the correspondences will be Freddy → Bobby and Fido → Rover.

The questions are, first, whether people base their correspondences on conceptual-relation matching or on grammatical graph matching, and, second, whether analogical inferences would follow the conceptual match or the purely grammatical graph match. To answer these, participants were asked, first, to state the best correspondences between the sentence sets; and, second, to make an inference about the second set based on using the first set.

Method

Participants

Participants in this study were 76 undergraduates at Northwestern University who were paid or received course credit for their time. Of these participants, 12 were dropped because they failed to follow instructions, leaving 64 participants whose data were analyzed.

structure and grammatical syntax may sometimes be in registration (as in our Experiment 1), they need not be [e.g., Jackendoff, 2002, p. 2237]. In this paper we use the term “structure” as the general term and reserve “syntax” for linguistic grammatical structure.

Table 1. Materials used in Experiment 1

Analogy	
A	B
Freddy is searching.	Rover is hiding.
Freddy catches a glimpse of Fido.	Rover attracts the notice of Bobby
Freddy chases after Fido.	Rover runs from Bobby.
Freddy fails to catch Fido	
Matches	
Please state which items in B best match these items from A:	
Freddy	_____
Fido	_____
Inference	
Please make a new inference about B, based on A.	

Materials and Procedure

The materials for this study were the sentence sets shown in Table 1. The two sets were constructed to have sentences that were conceptually similar but used complementary verbs, such that the grammatical subject of one was the grammatical object of the other. The sets were presented on a single page with corresponding sentences on the same line as in Table 1. Sentence set A had four sentences, and sentence set B had three sentences.

After reading the sentence sets, participants were asked to state which item in set B matched each of the two protagonists from set A: “Freddy” and “Fido.” Then participants were asked to make a new inference about set B based on its match to set A.

Results and Discussion

Of the 64 participants in this study, 60 (94%) placed Freddy in

correspondence with Bobby and Fido with Rover, $\chi^2(1) = 49.00, p < .05$. Thus, people overwhelmingly placed the objects in correspondence based on the conceptual similarities between the sentence sets, even though making this match required going against the correspondences suggested by the syntax of the sentences. The inference results were also striking. Of the 60 participants who made conceptual matches, 86% (52/60) made inferences consistent with their match. In contrast, not one of the four participants who made a graph match chose to draw any inference. This result suggests a deep difference between the process of deriving a purely graph-structural match and the process of deriving a relational similarity match.

These results provide evidence that in analogical matching people align conceptual relational structures and generate inferences from them, even when there are readily apparent grammatical matches that oppose these relational similarity matches. However, striking as they are, the results leave some issues open. First, the finding that graph isomorphism failed to yield inferences, whereas relational similarity matches almost always did, is potentially extremely important, because inference projection is a signature phenomenon in analogical mapping. However, because only four participants chose the graph-isomorphic match, this finding is based on a very small pool. Second, although the results argue against a graph-matching process based on sentence grammar, they leave open the possibility of pure graph-matching processes based on conceptual event representations. For example, suppose people in Experiment 1 had encoded the sentences as chasing events, and had encoded the arguments as “chaser” and “chasee” rather than in terms of SOV grammar. In this case the correspondences arrived at by pure graph matching would mirror those arrived at via relational similarity matching. In Experiment 2, we consider the other possible realization of pure graph isomorphism, namely, isomorphism in conceptual graphs.

There are other reasons to replicate the results of Experiment 1. The use of complementary verb pairs like chase/flee created a direct and obvious conflict between the conceptual matches and the grammatical isomorphism, and this may have led to some special processing strategy or in some other way have depressed the rate of isomorphic matching. Finally, the relational match may have been aided by the match between human names and dog names.

Although the use of object matches is consistent with a conceptual matching account, it is important to separate it from the use of relational similarity.

Experiment 2

In Experiment 2 we again compared relational similarity matches with graph-isomorphic matches, with two major changes. First, we used a between-subjects design, so that the graph isomorphism could be processed on its own, without competing conceptual matches. In addition to clarifying the status of graph matches in generating correspondences, this will also permit a better assessment of whether graph matches lead to inferences. Second, we assumed conceptual representations (rather than sentence-grammatical structure) as the basis for the graph-matches. We also made other improvements such as avoiding complementary verbs and using neutral object names.

We focused on the generation of inferences in this study, both because it is a central issue and because whether and which inferences people draw is diagnostic of their correspondences. To the extent that pure graph isomorphism can be shown to yield humanlike analogical inferences, it becomes a more viable prospect for a psychological mechanism of analogy. In order to assess whether participants' inferences fit human patterns, we included in the base passage an extra piece of information not present in the target, that was causally connected to the rest of the base passage. This allowed us to investigate not only whether participants would draw inferences, but whether the inferences would be drawn from shared systems of matches. There is considerable evidence that people prefer interpretations based on connected systems of matches, and that their inferences are drawn from shared systems rather than from isolated matches (Clement & Gentner, 1991; Keane, 1996; Markman, 1997). Thus, to the extent that participants' inferences show this preference for systematicity (Gentner, 1983), this will be evidence for the psychological plausibility of their processing method.

In Experiment 2 we used the materials in Table 2. Passage 1 describes a battle between the Fox corporation and the Time-Warner corporation. Participants compared this passage with either Passage 2a or Passage 2b. Comparing Passage 1 and Passage 2a affords a match based on relational similarity: e.g., just as Time-Warner owns a cable system and a news network, Jason owns an insurance company and an office building. Just as the Fox Corporation wants space on Time-Warner's cable system, Arthur wants an

Table 2. Materials for Experiment 2**Passage 1**

Currently, the Fox corporation is feuding with the Time-Warner corporation. Time-Warner owns a cable system and a news network. Fox owns a news network, and wants space on the Time-Warner system to air its programs. Time-Warner does not want to allow the Fox news network on its cable system. To make matters worse, Time-Warner is trying to hire away Fox's best employees. Time-Warner offers stock options to its employees making it an attractive place to work.

Passage 2a

Arthur and Jason are fighting. Jason has an insurance company and an office building. Arthur has an insurance company. Arthur wants to lease an office in Jason's building so he can conduct business.

Passage 2b

Arthur has always admired Jason. Jason often goes to casinos and baseball games. Arthur goes to casinos and he likes to play baseball at games that Jason is watching to show off his skills.

office in Jason's building. In contrast, the comparison between Passage 1 and Passage 2b affords a pure graph-isomorphism, with no relational similarity. In this case, Jason corresponds to Time-Warner and Arthur to Fox on the basis of the parallel syntactic structure between the relations in the two domains. For example, 'Time-Warner *owns* a cable system and a news network' in Passage 1 corresponds to 'Jason *goes to* casinos and baseball games' in Passage 2b.

The key question was the degree to which these two kinds of matches would support inferences. To test this, we asked participants to make an analogical inference and assessed whether they did so: that is, did they propose any new information about second passage on the basis of the first. To increase the chances of such inferences, and also to permit a more specific question, we placed an extra, nonmatching piece of information in Passage 1: namely, that Time-Warner does not want to allow Fox on its cable system. This is a *shared system fact*, because it is connected to the system of matches between the two passages. Thus this manipulation also allows us to test whether the two kinds of match are sensitive to systematicity. By design, the shared system fact can be unambiguously mapped to both of the second passages. When carried over

from Passage 1 to 2a, it leads to the inference that Jason does not want to lease space in his building to Arthur's insurance company; carried from Passage 1 to 2b (by making substitutions for the corresponding (but dissimilar) relations in the domains), it leads to the inference that Jason does not want Arthur to watch his games. Each of these is a reasonable continuation of the target passage. Thus, the question is how the relational similarity group and the graph isomorphic group will fare with respect to (1) whether they generate inferences; and (2) whether they generate shared-system inferences.

Method

Participants

The participants were 54 undergraduates at the University of Texas at Austin who received course credit. Of these, 25 were randomly assigned to the relational similarity condition and 29 to the graph match condition.

Procedure

Participants were given a sheet containing either Passages 1 and 2a or Passages 1 and 2b (Table 2). They were told "Please read the following two passages. In the space below, please write out anything that you would predict about what might happen in Passage 2 based on what happened in Passage 1. If you cannot think of any predictions, just write 'None.'" After writing their predictions, they rated their confidence in these predictions on a scale from 1 (low) to 7 (high).

Results and Discussion

The responses were first categorized as analogical inferences (inferences drawn from the base passage as instructed) and extraneous inferences not drawn from the base. (Typically, these last were plausible completions of the target based on general background knowledge.) We next scored each analogical inference as to whether it was a shared system inference, a nonshared system inference or an inference unrelated to the base domain. (Participants' responses could fall into more than one category.) As shown in Table 3, relational similarity participants ($M = .92$) made five times as many

Table 3. Results of Experiment 2: Number (and mean number) of subjects making no inferences or only extraneous inferences, and number (and mean per subject) of analogical inferences by type.

	Percent with no Inferences	Percent with only Extraneous Inferences	Mean Number of Analogical Inferences Per Subject		
			Shared System	Other Inferences from Base	Total
Relational Similarity Match (N=25)	4%	28%	0.64	0.28	0.92
Graph Match (N=29)	41%	45%	0.14	0.03	0.17

analogical inferences as did graph match participants ($M = .17$). More specifically, relational similarity participants ($M = .64$) made four times as many shared-system inferences as did graph match participants ($M = .14$), $\chi^2 = 14.51$, $p < .05$. The rate of shared system inferences in the relational similarity condition is comparable to what has been observed in other studies of analogical inference (e.g., Clement & Gentner, 1991; Markman, 1997). Other analogical inferences (nonshared system inferences) were also far more prevalent in the relational-similarity condition ($M = .28$) than in the graph match condition ($M = .03$), $\chi^2 = 6.41$, $p < .05$.

If participants failed to see any connection between the two passages, we would expect them either to draw inferences based only on the target passage (and unrelated to the base passage), or else no inferences at all. Among graph match participants, 45% drew only extraneous (nonanalogical) inferences and 41% drew no inferences. The corresponding figures are 28% and 4% for relational similarity participants. Thus, participants in the relational similarity condition made more shared-system inferences than either nonshared-system inferences or extraneous inferences, showing the typical sensitivity to systematicity observed in analogical mapping. In contrast, subjects in the graph

match condition made more extraneous inferences than either shared or nonshared-system inferences.

This pattern is further illuminated by participants' confidence ratings. If participants are sensitive to systematicity (as both the SME and IAM models predict), they should be most confident about shared-system inferences. This prediction was borne out for the relational similarity participants. For the shared system inferences made by relational similarity participants, the mean confidence rating was 5.13 (out of 7), as compared with a mean rating of 3.09 for their extraneous inferences. Graph match participants showed no such preference: the mean confidence rating for their shared-system inferences was 2.50, as compared with 3.96 for their extraneous inferences.

These results show a strikingly different pattern of inferencing between relational similarity matches and graph matches. Whereas 68% of the relational similarity participants drew an analogical inference, 86% of the graph match participants either drew no inferences or drew only extraneous inferences. Thus, even when explicitly encouraged to draw inferences from the analogy, most of the people in the graph match condition either declined entirely or simply invented facts about the target story. The confidence ratings were consistent with this pattern. Whereas relational similarity participants had high confidence in shared system analogical inferences, even those graph match participants who did make inferences showed no preference for analogical inferences over nonanalogical inferences (extraneous inferences drawn purely from the target).

General Discussion

Our goal here is to clarify the nature of structural similarity in analogical comparison. In Experiment 1, pure graph matches based on grammatical structure were pitted against relational similarity matches within-subjects. Participants overwhelmingly made correspondences based on relational similarity, not graph isomorphism. Further, nearly all the 60 participants who made relational similarity correspondences went on to draw an inference; in contrast, none of the 5 participants who drew graph-isomorphic correspondences did so. In Experiment 2, we compared conceptual graph matches with relational similarity matches in a between-subject design to allow

for the possibility that the poor showing for graph matches in Experiment 1 resulted from competition with the conceptual match. In this study, every effort was made to ‘level the playing field’ between relational similarity and graph matching. The graph-isomorphic passages contained no competing conceptual similarity alignment; the only good alignment was the graph match. Further, the passages were designed such that the graph-isomorphic correspondences yielded a plausible inference, just as did the correspondences in the relational similarity match. Yet when given instructions to seek an analogical inference, the two groups showed a clear divergence. Among participants given relationally similar passages, 92% drew analogical inferences — inferences from the base to the target. In contrast, among participants given the graph isomorphic passages, only 17% made analogical matches. In fact, 41% failed to draw any inference, despite instructions encouraging them to do so. Further, even when participants actually made an inference, most of the time this inference was simply a plausible completion of the target story rather than one based on an analogy between the base and target.

Finally, another key difference between the two groups was that the similarity participants showed the systematicity bias that is typical of analogical mapping. They not only tended to draw the inference that belonged to the connected system of predicates shared by the two passages, but they were also highly confident of this shared-system inference when they did so. In contrast, only a small minority of the graph matching participants drew the shared system inference, and their confidence in that inference was low. In sum, people are remarkably unwilling to draw inferences from nonconceptual mappings, even when instructed to do so; and when they do draw inferences, they do not show the normal systematicity bias. All of this suggests that when dealing with these nonmeaningful pairs, people are not engaging in normal analogical processing, but rather are solving a rather artificial puzzle.

These findings support the theoretical claim that conceptual similarity is crucial in human analogical processing. One implication of these results is that it is unnecessary, and possibly counterproductive, to test computational simulations of analogy on graph-matching pairs that have no conceptual similarity. We suggest that a model’s ability to handle pure graph matches unsupported by conceptual similarities is irrelevant to its ability to capture human analogical mapping processes. Attempts to improve a model’s

performance on graph matches may result in distortions of its central process account.

These findings invite further questions. For example, while relational similarity may be the chief contributor to structural alignment, other kinds of conceptual similarity also enter into analogical processing. However, the outcome of an analogical comparison can be influenced by object similarities as well as by relational similarity (e.g., Bassok, Wu & Olseth, 1995; Bassok & Medin, 1997; Keane, Hackett, & Davenport, 2001; Markman & Gentner, 1993; Rattermann & Gentner, 1998). A natural question is whether conceptual similarity matters more for relations than for objects, as structure mapping-claims or whether all matches count equally. Recent research by Gentner and Kurtz (in press) bears out the claim that relational similarity matters more than object similarity. They asked participants to judge whether two events, each expressed in a single sentence, were or were not analogical. The degree to which participants considered two events analogical varied strongly with the similarity of their relations, as expressed in the verbs, but hardly at all with the similarity of the objects, as expressed by the object nouns.

The Gentner and Kurtz research also addressed another question that follows from this research is what degree of synonymity is necessary to achieve a conceptual match? They found that pairs with closely synonymous relations (e.g., Greg built the deck/Chad constructed the deck) were almost always accepted as analogies, and pairs with very dissimilar relations (e.g., Greg built the deck/Chad swept the deck) were almost never accepted. Intriguingly, pairs whose relations were of intermediate similarity — e.g., Greg built the deck/Chad repaired the deck — were accepted about half the time; and further, participants required much longer to decide for these pairs than for the very close or very far pairs — suggesting that participants may have re-represented the meanings to seek conceptual overlap.

In sum, these findings support the view that analogical processing is driven by the presence of conceptual similarities — particularly relational similarities — between the base and target domain. As Polya (1954, p.13) put it, "... two systems are analogous, if they agree in clearly definable relations of their respective parts."

References

- Bassok, M., & Medin, D. L. (1997). Birds of a feather flock together: Similarity judgments with semantically rich stimuli. *Journal of Memory and Language*, 36, 311-336.
- Bassok, M., Wu, L., & Olseth, K. L. (1995). Judging a book by its cover: Interpretative effects of content on problem solving transfer. *Memory & Cognition*, 23(3), 354-367.
- Bowdle, B., & Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology*, 34(3), 244-286.
- Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, 15, 89-132.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, 41(1), 1-63.
- Forbus, K. D., Ferguson, R. W., & Gentner, D. (1994). Incremental structure mapping. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, 313-318.
- Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19(2), 141-205.
- Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 61-68.
- Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. & Kurtz, K. J. (in press). Relations, objects and the composition of analogies. *Cognitive Science*.
- Gentner, D., & Markman, A. B. (1997). Structural alignment in analogy and similarity. *American Psychologist*, 52(1), 45-56.
- Greiner, R. (1988). Learning by understanding analogies. *Artificial Intelligence*, 35, 81-125.
- Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. *Human Development*, 35(4), 193-217.
- Hofstadter, D. R. (2001). Epilogue: Analogy as the core of cognition. In Gentner, D., Holyoak, K. J., & Kokinov, B. N. (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 499-538). Cambridge, MA: MIT Press.
- Holyoak, K. J., Novick, L. R., & Melz, E. R. (1994). Component processes in analogical transfer: Mapping, pattern completion, and adaptation. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural*

- computation theory: Vol. 2. Analogical connections* (pp. 113-180). Norwood, NJ: Ablex.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3) 295-355.
- Holyoak, K. J., & Thagard, P. R. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3), 427-466.
- Jackendoff, R. (2002). *The foundations of language*. New York: Oxford University Press.
- Keane, M. T. (1996). On adaptation in analogy: Tests of pragmatic importance and adaptability in analogical problem solving. *The Quarterly Journal of Experimental Psychology*, 49/A(4), 1062-1085.
- Keane, M. T. (1997). What makes an analogy difficult? The effects of order and causal structure on analogical mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(4), 946-967.
- Keane, M. T., Hackett, D., & Davenport, J. (2001). Similarity processing depends on the similarities present: Effects of relational prominence in similarity and analogical processing. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 477-482). Mahwah, NJ: Lawrence Erlbaum Associates.
- Kokinov, B. & French, R. M. (2003). Computational models of analogy-making. *Encyclopedia of Cognitive Science*.
- Kubose T. T., Holyoak K. J., Hummel J. E. (2002). The role of textual coherence in incremental analogical mapping. *Journal of Memory and Language*, October 2002, 47(3), 407-435.
- Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science*, 21(4), 373-418.
- Markman, A.B. & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, 25, 431-467.
- Mitchell, M. (1993). *Analogy making as perception*. Cambridge, MA: MIT Press.
- Polya, G. (1954). *Induction and analogy in mathematics: Vol. 1. Of mathematics and plausible reasoning*. Princeton, NJ: Princeton University Press.
- Ramscar, M. & Yarlett, D. (2002) Semantic grounding in models of analogy: An environmental approach. *Cognitive Science*, 27, 41-71.
- Rattermann, M.J. & Gentner, D. (1998). More evidence for a relational shift in the development of analogy: Children's performance on a causal mapping task. *Cognitive Development*, 13, 453-478.

- Spellman B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush?: Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, *62*(6), 913-933.
- Spellman, B. A., & Holyoak, K. J. (1996). Pragmatics in analogical mapping. *Cognitive Psychology*, *31*, 307-346.
- Yarlett, D., & Ramscar, M. (2000). *Structure mapping theory and lexico semantic information*. Paper presented at the Proceedings of the 22nd Annual Conference of the Cognitive Science Society, Philadelphia, PA.