

Relations Relating Relations

Robert L. Goldstone
Dedre Gentner
Douglas L. Medin
University of Illinois at Urbana-Champaign

Abstract

The aim of the current work is to incorporate structural information in judgments of similarity. According to the assumption of feature independence, how one feature affects similarity is independent of the values of the other features present. We present three violations of this assumption, all arising from the influence of relations between features and of relations between relations. A shared relation is more important for similarity judgments if it cooccurs with (A) relations that augment the first relation by "pointing in the same direction" as the first relation, (B) relations which are themselves salient, and (C) salient relations that involve the same objects as the first relation. We interpret these results as suggesting that relations do not have separately determined weights or saliences; the weight of a relation depends the relational structure in which it exists. Relations influence each other by creating higher-order relational structures, and also by affecting processing.

Introduction

Our current interest in structural similarity stems from two sources. First, similarity has been a pivotal, if under-explored, agent in many theories. Problems are said to be easier to solve if similar problems have been previously solved. The transfer of one skill to another is said to be proportional to their degree of similarity. One event reminds us another if they are similar. An object belongs to a category if it is similar to the category's examples or prototype. An account of what makes things similar is essential to our understanding of cognition. Second, it has long been argued, particularly in artificial intelligence, that simple, independent feature representations, inadequately capture most of our knowledge. Propositional representations, explanation based learning, and pattern recognition all point to the need for systems that represent how the units of analysis are related and interconnected in structural descriptions. Thus, separate fields have argued for the importance of structure and for the importance of similarity; at the most general level, our research goal is to consolidate structure and similarity into a single framework.

One of the most important and common assumptions associated with a wide variety of psychological models is that there exists a set of features that are independent of one another. Features are independent if the effect of the value on any one feature does not depend on the value of any other feature. A selection of prominent examples illustrates the prevalence of the feature independence assumption. Tversky's influential contrast model (1977) makes the explicit assumption that the joint effect of two components in determining similarity is independent of the fixed level of the third component. Posner and Keele (1968) assume that the psychological distance of a dot from its prototypical location is independent of the locations of other dots. Independent-cue prototype models propose that categorization of an example into a group is a function of the feature matches between the example and the group, additively combined in an independent manner. The assumption of feature independence is a powerful simplifying assumption, permitting analyses which would otherwise be unwieldy or impossible. Without the independence assumption, any single feature of a system with ten binary-valued features could produce 1024 different effects; with the independence assumption, the same feature can produce only two effects.

Unfortunately, it is not always true that what makes for a simple and powerful model also makes for an accurate model of human behavior. There is good reason to suppose that cognitive processes are often based on highly interactive features. The feature "gray" does not mean the same thing in "gray hair" as it does in "gray cloud" (Medin & Shoben, 1988). Pomerantz (1986) has shown that in displays such as ")(" the two parentheses are not independently perceived - although they are physically detached, they are psychologically fused, creating emergent properties. Gati and Tversky (1984) present examples in which monotonicity and feature independence are violated because adding certain features causes interactions with the features already present. Indeed, all context effects can be construed as cases of non-independence. In general, there is reason to think that feature independence is the exception rather than the rule - object features often mutually constrain and modulate each other.

It might be thought that admitting feature dependencies into a model would yield an overwhelming number of degrees of freedom and a loss of important constraints. Contrary to this position, we will argue that feature dependence need not be a "counsel of despair." The similarity judgments that this paper investigates are not constrained by feature independence, but they are constrained by specific principles. The central claim will be that the importance of a relation in a scene depends on the quality, quantity, and location of the other relations present.

The general representational system outlined in Dedre Gentner's Structure-mapping Theory (SMT) (Gentner, 1983, 1989) will be used. In Gentner's terminology, a first-order relation is any relation that takes two or more objects as arguments. The relation "DARKER-THAN" is first-order because it takes two objects, two wings for example, and establishes a relation between them. The relation can be propositionally represented as DARKER-




Figure 1

THAN (Wing1, Wing2) (see also Palmer, 1975). A relation is "higher-order" if it takes two or more relations as objects. According to SMT, a common relation is more important for an analogy if it is involved in a common higher-order relation. For the analogy between the solar system and an atom, the relation GREATER (MASS(sun), MASS(planet)) will be more important than the relation GREATER (BRIGHTNESS(sun), BRIGHTNESS(planet)) because it is involved in the higher-order relation CAUSE(GREATER (MASS(sun), MASS(planet)), REVOLVE-AROUND (planet, sun)) . Consequently, according to SMT, the goodness of an analogy depends on the relational correspondences between two domains, and on the relational structure in which these correspondences exist.

Higher-order relations illustrate one case of relations between relations. While the importance of higher-order relations has been shown for causal analogies (Clement & Gentner, 1988; Gentner, 1983; Winston, 1980), we will present demonstrations that suggest the importance of higher-order relations in determining perceptual similarity. All higher-order relations are not treated equally; some higher-order relations serve to highlight the relations that compose them, whereas other potential higher-order relations do not. As such, not only are object features involved in dependencies - the relations themselves display dependencies. A second case of relations between relations is also suggested - even when a true higher-order relation is not formed by two relations, the relations can still influence each other's importance by affecting how each is processed. For example, the perception of one relation may prime individuals to see other similar relations. In general, both cases argue that a person's sensitivity to a relation depends on the other relations with which it is composed.

Determining the Polarity of the Darker-than and Larger-than Relations

While Gentner's work with analogies typically uses causal relations as the higher-order relations that connect first-order relations, the higher-order relations used here will be comparisons of magnitude differences. In doing so, we preserve the basic vocabulary of propositional representation (Palmer, 1978). In Figure 1, the two blocks have the following first-order relations: GREATER-THAN (SIZE(triangle), SIZE(circle))¹ and GREATER-THAN (NUMBER-OF-DOTS (triangle), NUMBER-OF-DOTS (circle)). There is also a relation between these two relations - namely, in both relations, the triangle has the larger quantity. This higher-order relation will be called "AUGMENTATION" because each of the first-order relations augments the effect of the other by "pulling in the same direction" as the other; more technically, the two relations are magnitude relations of the same type, that take the same object arguments in the same order. The opposite of an AUGMENTATION relation is an "OPPOSITION" relation, whereby one relation cancels out the other relation by "pulling in the opposite direction"; more technically, the two relations are magnitude relations of opposite types, that take the same arguments in the same order. If the triangle were taller than the circle, but had fewer dots than the circle, then the relation between the dots and the relation between the sizes would be in an OPPOSITION relation.

To determine whether a higher-order relation is AUGMENTATION (both first-order relations in same objects pointing in the same direction) or OPPOSITION (first-order relations pointing in opposite directions) we must first determine the direction of greater magnitude for each relation. Following the general logic of Smith, Sera, and Goodrich's (in press) procedure for determining the "polarity" of a relation, we asked 18 subjects to label one side of a butterfly "positive" and the other side "negative." Subjects were told to use the words "positive" and

¹ GREATER-THAN (SIZE (triangle), SIZE (circle)) is equivalent to the representation LARGER-THAN (triangle, circle). The former can be viewed as an expansion of the latter.




Figure 2

"negative" in a loose or metaphorical sense. The butterflies in Figure 2 tested the polarity of the TALLER-THAN and DARKER-THAN relations.

To control for possible left-right biases, half the subjects received butterflies identical to A and B except that the left and right wings were switched. All 18 subjects thought the larger side of A was positive, supporting the intuition that large is (metaphorically) positive and small is negative. Further, The results of 2B show that for most adults dark is positive and light is negative.² When the relational structure of butterflies are later represented, we will treat large wings as being GREATER-THAN small wings, and dark wings as being GREATER-THAN light wings. For example, if the left wing of a butterfly is larger than and darker than the right wing, then these two relations augment each other; if the left wing is larger than and lighter than the right wing, then the relations oppose each other.

General Method


Subjects were shown computer displays containing three butterflies. The butterfly at the top was the comparison butterfly. Subjects were told to either choose the lower left or the lower right butterfly, whichever was more similar to the comparison butterfly. Subjects pressed keys to choose either A or B as more similar to T. The relations used were: SAME-SIZE, LARGER-THAN, SMALLER-THAN, SAME-SHADE, DARKER-THAN, LIGHTER-THAN, and DIFFERENT-SHADE. There is no DIFFERENT-SIZE relation because any pair of wings which are differently sized can be described by the LARGER-THAN or SMALLER-THAN relation. In all experiments, the left/right order of butterflies A and B was randomized.

Higher-order Relations that Increase Relational Responding AUGMENTATION vs. OPPOSITION

Thirty-six University of Illinois undergraduates were presented butterfly triads with two choices, one of which was superficially similar to T and one of which was relationally similar to T. An example of two such triads is shown in Figure 3. For both the left and right cases, if a subject responds that A is more similar to T than is B, then we take this as evidence that they are responding superficially - they are basing their similarity on the gray that both A and T have on their right wings. A response of "B" is labelled a relational response, based on the common relation that B and T share - their right wings are darker than their left wings. A relation which is shared by T and only one of the two choices is called a unique relation. DARKER-THAN (right wing, left wing) is a unique relation, possessed only by T and B. A relation which is possessed by all three butterflies is called an accompanying relation. LARGER-THAN (right wing, left wing) is an accompanying relation for the three butterflies on the left because each of them possesses this relation.

The unique relations used in the experiment were: SMALLER-THAN, LARGER-THAN, DARKER-THAN, and LIGHTER-THAN. Each unique relation was paired with an augmenting and an opposing accompanying relation, yielding a total of eight different butterfly triads. For each triad, the left/right order of the relational and superficial responses were randomized, as was the presentation order of the pictures. Forty-two

² These results may seem to be opposed to Smith et al.'s (in press) evidence that adults have some tendency to view lighter objects as positive. One possible resolution is that our stimuli heighten the "dark=positive" effect because of they appear on a computer screen with a white background.



Augmenting Relations

Figure 3

Opposing Relations

University of Illinois undergraduates were shown each of the eight pictures four times, and were required to make each response within 3.5 seconds.

Our interest is not in the absolute number of superficial and relational responses that subjects give; we are interested in variable that *shift* subjects' judgments. Here, the manipulated variable of importance is whether the accompanying relation augments or opposes the unique relation. For the three butterflies on the left, the accompanying relation LARGER-THAN (right wing, left wing) augments the unique relation DARKER-THAN (right wing, left wing) because they both are GREATER-THAN relations which take the same arguments in the same order. For the three butterflies on the right, the accompanying relation is SMALLER-THAN (right wing, left wing); this relation opposes the unique relation because SMALLER-THAN is a LESS-THAN relation. If the unique relation pertains to size, then four different accompanying relations are used - one for each shade relation. If the unique relation pertains to shade, then all three size relations are used as accompanying relations.

The question of primary interest is: What proportion of subjects give the relational response (as opposed to the superficial response) when the accompanying relation augments/opposes the unique relation. When the accompanying relation is augmenting, subjects give the relational response (choice B in the above figure) on 50.1% of trials; the proportion of relational responding drops down to 39.4% when the accompanying relation is opposing. This significant difference ($p < .01$) between the amount of relational responding suggests that the salience of a common relation in a similarity judgment is affected by the other relations present in a scene. Two simple explanations of our results will not suffice.

First, our results cannot be explained by simply assuming that some relations are inherently more salient than others. When subjects make choice B over A as most similar to T, they must be doing so because of the unique relation DARKER-THAN (right wing, left wing), since this is the only property that distinguishes A and B that also belongs to T. This relation is the same in the two triads. The only difference between the augmenting and opposing conditions is the wing size relation, which is shared by all three butterflies. Thus, how important the DARKER-THAN relation is for similarity depends on a relation which, by itself, does not distinguish between the two choices.

Second, our results also cannot be explained by saying: "Perhaps there is a general advantage to picking up other relations when the right wing is larger than the left wing." If the unique relation had been DARKER-THAN (left wing, right wing), then the opposite accompanying relation [LARGER-THAN (left wing, right wing)] resulted in a boost to relational responding. It is not the specific quality of a first-order relation considered by itself which

results in the increased relational responding - it is the relation between first-order relations which affects the degree of relational responding.

Nonspecific Increases in Responding to Unique Relations Due to Accompanying Relation

Putting aside the specific relation-relation interactions, other results also suggest that some accompanying relations can serve as "universal facilitators," increasing the likelihood of relational responses in general. When the accompanying relation shared by all three butterflies is SAME-SIZE or SAME-SHADE (dimensional identity relations), then responding on the basis of the unique relations is higher than when the accompanying relation is one of the other relations. Figure 4 depicts an example of this effect. For both the left and right triads, the relational choice (B) has the unique relation SMALLER-THAN(LEFT WING, RIGHT WING) that the superficial choice (A) does not. The tendency (non significant if Figure 4 alone is considered, but significant overall) is for subjects to choose the relational choice more when wings have a SAME-SHADE relation. The unique relations include equal numbers of each of the seven relation (four shade relations, three size relations) types. The same subjects used in the preceding study produced the following results:

Accompanying Relation	% of Responding based on Unique Relation
DIFFERENT-SHADE (left wing, right wing)	70%
SAME-SHADE (Left wing, right wing)	76% (significantly greater than 71%)
LIGHTER-THAN or DARKER-THAN (left wing, right wing)	71%
LARGER-THAN or SMALLER-THAN (left wing, right wing)	44%
SAME-SIZE (left wing, right wing)	50% (significantly greater than 44%)

The presence of SAME-SIZE and SAME-SHADE relations within the butterflies causes further relational correspondences between the butterflies to be noticed. It has been suggested (Goldstone, Medin, & Gentner, 1987) that increasing the importance of relational correspondences between two objects will bias the subject to look for further relational matches as opposed to superficial matches. The percentage of responses on the basis of a relation




Figure 4


can be increased by either (1) increasing the number and salience of other relations present, or (2) decreasing the number or salience of superficial similarities present. The current results support this suggestion if SAME-SHADE and SAME-SIZE are assumed to be stronger, more noticeable relations than DIFFERENT-SHADE, LIGHTER-THAN, and SMALLER-THAN, so that, as accompanying relations, they call attention to the unique relation. Independent support for this assumption comes from the fact that when the unique relation is SAME-SHADE or SAME-SIZE, the relationally similar butterfly is chosen 72% of the time. Relational responding drops to 61% when the unique relation is one of the other relations. In the current experiments, we find that that a relation which might not be thought to have any influence on judgment (because both of the choices possess it), still influences judgment by highlighting other relations in a non-specific manner. The highlighting is non-specific in that SAME-SIZE and SAME-SHADE relations increase the likelihood of responses based on any other relation. No good higher-order relation results from the juxtaposition of SMALLER-THAN and SAME-SHADE (as in Figure 4). SAME-SHADE's tendency to increase the importance of SMALLER-THAN is due to its general processing facilitation of all relations.

Increased Relational Responding with Relations Associated with the Same Objects

A final demonstration shows a specific influence of relations on relations. Whereas the first demonstration showed an influence specific to the nature of the relations involved, another influence is specific to the structural configuration of the relations. A common pattern is that relations that share objects facilitate each other more than relations that belong to different objects. Thus, a relation will count for more in a similarity judgment if it involves objects that share another strong relation. Subjects rated the similarity of two scenes on a scale from 1 to 9, with 1 referring to very low similarity and 9 referring to very high similarity. Six sets of pictures were given of the same abstract design as Figure 5, intermingled with 30 filler sets. In Figure 5, the left pair (A and B) and the right pair (C and D) have exactly the same shapes and shades in common. Beyond this, they also share the same first-order relations: one SAME-SHAPE, one SAME-SHADE, and several DIFFERENT-SHAPE and DIFFERENT-SHADE relations. The only difference between the left and right pair is that the two SAME relations cooccur in the same objects on the right, whereas they belong to different objects on the left. So, the two triangles of C and D instantiate both the SAME-SHAPE and the SAME-SHADE relation. For A and B, the SAME-SHAPE relation is located with the bottom shapes while the SAME-SHADE relation is stationed in the middle row of shapes. Consistent with the previous demonstration, it will be assumed that relations involving SAME are more salient than relations involving DIFFERENT for these stimuli.

As figure 5 shows, when both of the salient relations belong to the same objects, similarity is rated higher. Five out of the six picture sets show an advantage for the location-coupled relations, and the sixth showed no difference between the coupled and uncoupled relations. The mean difference between the similarity ratings for the coupled and uncoupled sets was .255 ($p < .05$).

Again, alternative explanations of the results do not seem to adequately account for the results. Criticisms such as "perhaps the SAME-SHADE relation is always more salient when it is on the bottom" are refuted by counterbalancing conditions in which the second and third rows of all the scenes are switched. Similar counterbalancings were performed for each of the scenes. The absolute location of the salient relation does not matter; what matters is whether the salient relation is in the same location as the other salient relation. In addition, it does not appear that the two relations have to be similar. In Figure 5, the two salient relations both involve identity - either of shape or shade. However, in two other sets that yield the same effect, one of the relations is DARKER-THAN while the other is SAME-SHAPE.



Conclusions

In explaining the three observed relation-relation interactions, we can either 1) posit particular higher-order relations, or 2) posit processing principles. The first observation, that augmenting relations are more likely to be the basis of similarity judgments than opposing relations, is naturally handled by the first strategy - by postulating that people are selectively sensitive to the particular higher-order relation AUGMENTS (LIGHTER-THAN (left wing, right wing), SMALLER-THAN (left wing, right wing)). Consequently, two butterflies with an AUGMENTATION higher-order relation will be perceived as highly similar, whereas two butterflies with an OPPOSITION higher-order relation will not be perceived to be as similar. On the other hand, a processing account might argue that perceiving a relation with one magnitude relation facilitates noticing/evaluating other relations with a relation of the same polarity, without requiring the psychological reality of the actual higher-order relation AUGMENTS.

At the other end of the continuum, the observation that a SAME-SHADE relation (as opposed to DIFFERENT-SHADE) increases the saliency of LARGER-THAN, seems best explained by the processing principle that a salient relation primes people to look for any other relation, even if these relations do not form a compelling higher-order relation. The only higher-order relation shared by SAME-SHADE and LARGER-THAN is DIFFERENT-RELATION, and it is unlikely that this higher-order relation would have psychological salience in our stimulus set. Relations take part in specific interactions with other relations, expressible as higher-order relations; they also interact non-specifically, with highly salient relations acting as "universal facilitators" for the perception of other relations.

To review, the results argue that how important a relation is for similarity depends on the other relations present -relations do not have absolute, intrinsic saliences. First, a relation becomes more important if it cooccurs with a relation which points in the same direction that it does. That is, in order to make a relation clear or noticeable, other relations should be introduced so as to augment, not oppose, the relation at issue. Second, there is also a non-specific interaction between relations, such that a prominent relation (SAME-SIZE or SAME-SHADE) increases the salience of all of the other relations in a scene. To make a relation more likely to be noticed, a good strategy is to put the person into a "relational frame of mind" by adding other salient relations to the object. Third, a relation is more important if it connects objects that are connected by another salient relation; unlike the second effect, the facilitative effect of a relation is focused on other relations with which it coincides; unlike the first effect, this specific facilitation is primarily based on the configuration and not the nature of the relations.

These three points argue for an account of similarity that is not based upon the independent analysis of features, or even the independent analysis of relations between features. The importance of relations between relations precludes any model which simply assigns each matching and mismatching feature/relation a static weight, and counts the weighted features/relations to determine the similarity of two objects. The weight given to a relation cannot be assigned without knowing what other relations are present. Rather than viewing similarity as based on lists of features, similarity is best viewed as being sensitive to the structural relations between features. Even in the domain of visual perception, where the postulation of simple feature detectors is most appealing, we find that the higher-order structural representation of a scene influences the use of lower-order relations, which in turn influence the use of object features.

References

- Clement, C., & Gentner, D. (1988). Systematicity as a selection constraint in analogical mapping. In *The Tenth Annual Conference of the Cognitive Science Society*, (pp. 412-418), Montreal.
- Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgments. *Cognitive Psychology*, *16*, 341-370.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity, analogy, and thought*. New York: Cambridge University Press.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155-170.
- Goldstone, R., Medin, D., & Gentner, D. (1987). Relational similarity and the non-independence of features in similarity judgments. Address to the Mid-West Psychological Association: Chicago.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411-429.
- Palmer, S. E. (1975). Visual perception and world knowledge. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in cognition*. San Francisco: W. H. Freeman.
- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Pomerantz, J. R. (1986). Visual form perception: An overview. In *Pattern recognition by humans and machines: Visual perception, Volume 2*. New York: Academic Press.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353-363.

- Smith, L. B., Sera, M., & Goodrich, T. (1984). A developmental analysis of the polar structure of dimensions. in press.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327-352.
- Winston, P. H. (1980). Learning and reasoning by analogy. In: Communications of Association for Computing Machinery, *23*, 689-703.