Commonalities, Differences and the Alignment of Conceptual Frames During Similarity Judgments

Arthur B. Markman
University of Illinois
Department of Psychology
Beckman Institute
603 E. Daniel St.
Champaign, IL 61820

Dedre Gentner
Northwestern University
Department of Psychology
Institute for the Learning Sciences
2029 Sheridan Rd.
Evanston, IL 60208

Abstract

Tversky demonstrated that the similarity of two objects increases with their commonalities and decreases with their differences. We believe that determining commonalities and differences is a complex task. Using analogical mapping as a guide, we propose the process of *frame-alignment* which can be employed to find the commonalities and differences of structured representations. We then test the predictions of this approach by asking subjects to list the commonalities and differences of word pairs that vary in their degrees of similarity. The results of this study support the predictions of the frame-alignment view.

Introduction[1]

Similarity is a central aspect of virtually all cognitive processes, whether of categorization (Smith and Medin, 1981), problem solving and transfer (Ross, 1989, Holyoak and Koh, 1987), or even skill acquisition (Logan, 1989). Thus, it is incumbent on cognitive psychologists to detail the algorithm by which similarity is determined.

Tversky (1977) provides evidence that the similarity of two objects increases with their commonalities and decreases with their differences. Given this assumption, the process of calculating the similarity of two objects reduces to the subtasks of finding their commonalities and differences and combining these elements. At first blush, this approach seems straightforward. Consider the pair of circles in Figure 1a. We could describe both circles using lists of features. The features 'circle' and 'medium-sized' are common to both objects, and will increase the similarity of the pair. However, one object has the feature 'shaded', while the other has the feature 'striped'. These distinctive features will decrease similarity.

If we consider the pair in Figure 1b, we see the need for a more complex account. Notice that each of these scenes can be described by exactly the same set of features, though clearly they are not identical. In this case, we must make a decision before counting commonalities and differences: either the two circles are placed in correspondence and 'shading' and 'location' are distinctive features, or the two leftmost objects are placed in correspondence and 'shading' is a common feature, while 'shape' becomes a distinctive feature. The point is, before we can determine what is common and what is distinctive, we must align the representations based on their structure.

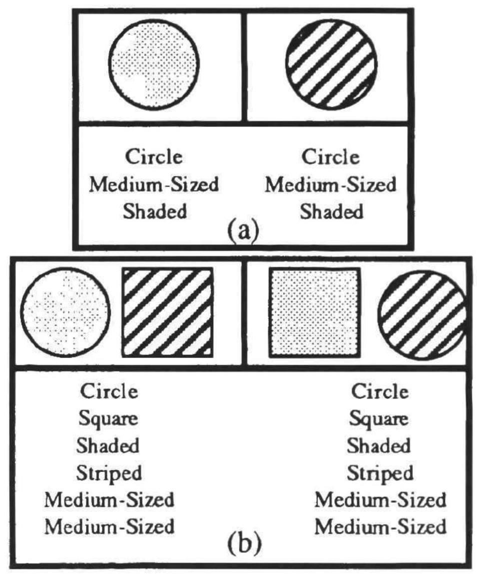

Figure 1: Examples of feature alignment.

There is empirical support for the claim that similarity is a process of alignment of structured representations. For example, Goldstone, Gentner and Medin (1989) demonstrated that pairs of objects with consistent relations between their parts are more similar than objects that do not have consistent relations between their parts. Furthermore, Gentner and Rattermann (1987) found that stories with similar plots and different characters (and hence similar structural elements) are often thought to be more similar than stories with different plots and similar characters (thus sharing few structural elements). Taken together, this evidence indicates that any theory of similarity must take relational structure into account.

[1]This work was supported by NSF grant BNS 87-20301 and ONR grant N00014-89-J1272 awarded to the second author as well as a University Fellowship awarded to the first author by the University of Illinois.

One format for structured representation is *conceptual frames* (Minsky, 1981, Schank and Abelson, 1977, Lenat and Guha, 1990, Barsalou, 1990, Norman and Rumelhart, 1972). In a conceptual frame, objects are represented as *slot-value pairs* with relations between them. Slots represent dimensions along which a particular object varies. For example, a car may have a color or a type of engine (see Figure 2). The value associated with a slot represents the way that role is filled for that object. This value may be a single entry (e.g. **powerDevice:** engine) or it may be a list of entries (e.g. **instances**: DatsunB210, ToyotaCorolla, LincolnTownCar...). Furthermore, relations between slots, such as mutual constraints between the values of two slots and information about the function of particular slots, are included. All slots, values and relations are themselves represented by frames.

Frame: Car
specializationOf: Vehicle
instances: DatsunB210, ToyotaCorolla, LincolnTownCar. . .
locomotionDevice: Wheels
numberOfWheels: 4
powerDevice: Engine
engineTypes: V-8, V-6, Rotary . . .
structureOfBody: EnclosedBody
colorOfBody: Red, Green, Blue. . .
fuel: Gasoline, Gasahol

Figure 2: Sample Frame

Lenat and Guha (1990) detail a long-term project in which all of common sense knowledge is to be encoded into an enormous knowledge-base composed entirely of frames. In this system, slots are organized in hierarchies. Through the creation of a consistent, global ontology, similar objects (which have similar slots) will have a common structure defined by the position of their slots within this organization. In addition, Barsalou (in press) reviews work demonstrating that frames may be used as the basis for diverse cognitive processes including mental-models, categorization, conceptual-combination and planning.

Given structured representations, we require a mechanism for comparison that can be used to determine similarity. We postulate that the commonalities and differences of representations can be found through a process of *frame-alignment*. During the alignment of two frames, we assume that slots are placed in correspondence if they are identical, or if they play the same role within the relational structures of both frames. One candidate for this process would be the local to global computation proposed to mediate analogical reasoning (Gentner, 1983, 1989, Holyoak and Thagard, 1989). Once the slots have been aligned, their values may also be compared.

Figure 3a depicts schematic object representations for two highly similar objects. Every slot in the first representation has a corresponding slot in the second representation (i.e. slot1 in Frame1 matches slot1 in Frame2). Further analysis of the values associated with the corresponding slots indicates that slots 1, 2 and 3 all have identical values in both frames. These matching slot-value pairs represent *commonalities* between the objects. In contrast, slot4 has value4 in Frame1 and value10 in Frame2. These mismatching values represent differences between objects fill corresponding (aligned) slots. Thus, we will refer to this kind of difference as an *alignable difference* (AD).

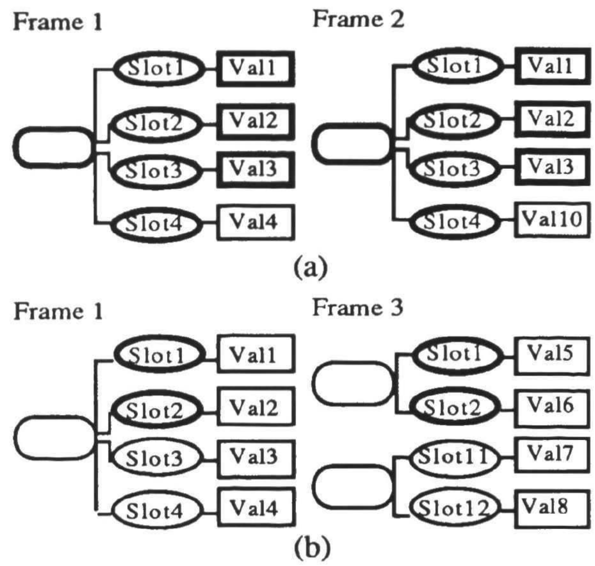

Figure 3: Schematic Frames

There is a third possibility, as shown in Figure 3b which depicts the frame representations of two dissimilar objects. Besides the two alignable differences [slot 1:(val1/val5)] and [slot2:(val2/val6)], there are two *non-alignable differences*. These differences between frames arise because the slots do not correspond. For example, slot3 in frame1 has no equivalent slot (and hence no corresponding slot-value pair) in frame2. A similar analysis may be made for slot4 in frame1 as well as slot11 and slot12 in frame2.

The frame-alignment view of similarity gives rise to the ternary distinction of commonalities, alignable differences and non-alignable differences.[2] Each of

[2]Other models of similarity make different partitions of commonalities and differences. Mental-distance models (Kruskal and Wish, 1978) are based on a unary scheme which does not differentiate commonalities from differences. Feature models (Tversky, 1977) are based on a binary distinction of commonalities and differences.

these factors is expected to play a discernably different role in the determination of similarity. Like the model proposed by Tversky (1977), the frame-alignment view assumes that the similarity of two objects increases with the number of their commonalities and decreases with the number of their differences (both ADs and NADs). However, the numbers of commonalities ADs and NADs is expected to differ for objects of high and low similarity.

For example, the comparison of similar objects, like those in Figure 3a, should yield many matching slot-value pairs (commonalities) as well as matching slots that have mismatching values (ADs). Similar pairs should have few slots that have no corresponding slot in the other frame (NADs). For dissimilar pairs the pattern should be quite different. As depicted in Figure 3b, we suggest that dissimilar objects have few common slots. Hence there should be few commonalities and few ADs. Rather, we expect that many slots will have no corresponding slot in the other representation, resulting in many NADs.

This analysis leads to four major predictions of the frame-alignment model of comparison. First, we expect that, as the number of commonalities between two objects increases, the number of ADs should increase as well. Second, because ADs arise from a matching slot, there should be a semantic link between ADs and commonalities. Third, similarity should increase with the number of commonalities and decrease with the number of differences (both ADs and NADs). Finally, if we make the simplifying assumption that the frames being compared are roughly the same size (have the same number of slots), there should be few NADs when there are many ADs and commonalities (and vice versa). Thus, the total number of differences listed need not increase with distance.

To test these predictions, we asked subjects to list commonalities and differences for various pairs of concepts in an effort to tap the output of the comparison process. In order to vary similarity systematically, we used an ontology tree similar to the ones used by Keil (1979) and Lenat (Lenat and Guha, 1990). This tree is shown in Figure 4. Eight basic level object names were listed at each leaf of the tree. These objects were taken from two different categories that fall in that area of the ontology. For example, we used four birds and four reptiles under the animal branch of the ontology. Sixty-four terms were spread evenly among the eight branches of the tree.

Subjects were presented with pairs that differed in their distance within the tree. Distance was calculated by the number of nodes that need to be traversed to get from one term to the other. Thus, two animals (i.e. salamander/robin) are at distance 0, an animal and a plant (i.e. salamander/oak) are at a distance of 1, an animal and an artifact (i.e. salamander/bus) are at a distance of 3 and an animal and any abstract term (i.e. salamander/meeting) are at a distance of 5.[3]

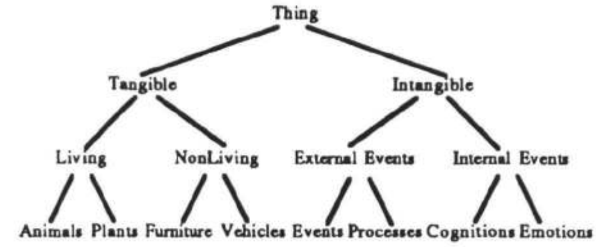

Figure 4: Ontology tree used here.

In the classic account, in which objects are represented as lists of features (and we retain the simplifying assumption that the representations are roughly the same size[4]), the prediction is that similar objects should have many commonalities and few differences, while dissimilar objects should have few commonalities and many differences. The frame-alignment model also predicts that similar objects will have many commonalities and few non-alignable differences, while dissimilar objects will have few commonalities and many non-alignable differences. However, an opposite pattern is predicted for alignable differences: similar objects should have many alignable differences, while dissimilar objects will have few. Thus, in this account all differences are not alike.

Experiment

Method

Subjects: Subjects were 44 students at the University of Illinois who received course credit for their participation. A second group of 40 subjects produced similarity ratings.

Stimuli: As described above, we use the ontology depicted in Figure 4. Eight basic level terms (4 from each of 2 categories subsumed under that node of the

[3]These a priori distances were highly related to the subjects' similarity ratings, with similarity decreasing as distance increased.For pairs of distance 0, mean similarity was 5.06. For pairs of distance 1, mean similarity was 3.02. For pairs of distance 3, mean similarity was 2.01. Finally, for pairs of distance 5, mean similarity was 1.59.

[4]In the context of this experiment, this assumption requires that size of frame and ontological distance be decorrelated. Although abstract and concrete objects probably differ in the number of slots in their frames, equal numbers of pairs of abstract and concrete items are used at distances 0, 1 and 3. Pairs of distance 5 are made of one abstract item and one concrete item, so this assumption may be less tenable for these items. If abstract or concrete items are analyzed alone the same pattern of results is obtained.

ontology) were listed at each node for a total of 64 terms. Two stimulus sets were created. Each set contained 16 pairs of distance 0, 8 pairs of distance 1, 4 pairs of distance 3 and 4 pairs of distance 5. Each word used in a pair of distance 0 in one set was used in a pair of distance greater than 0 in the other set.

Procedure: Subjects were run individually. Subjects were seated at a table in front of a pile of index cards (face down) and a tape recorder. Each card contained a word pair and instructions to list either commonalities or differences for that pair. Subjects were given one minute to make their listings aloud. If the subject did not make a response for 30 seconds, the trial was terminated and the next word pair was presented. All sessions were taped.

A separate group of subjects was given the pairs with the instructions to rate the similarity of the pairs on a scale from 1 (low) to 9 (high). Twenty subjects rated each stimulus set.

Design: All distance conditions were within subjects. Half of all trials were commonality listings, and half were difference listings. Thus, commonalities and differences were obtained between subjects. Two different stimulus sets were used in this experiment. Stimulus set was a between-subjects factor.

Scoring: All subject tapes were transcribed. The transcripts were scored as follows. For listings of commonalities, one commonality was counted for each distinct item listed (i.e. both a car and motorcycle have wheels).[5] However, grammatical commonalities, commonalities of the actual words and associations between the words were not counted as commonalities, and will not be considered further here.

For difference listings, the items were broken down into alignable differences (AD) and non-alignable differences (NAD). A listing was considered to be an AD if subjects mentioned a contrasting value for both objects in their protocol (i.e. a car has 4 wheels and a motorcycle has 2). Subjects need not have mentioned the exact values for both objects (i.e. cars and motorcycles have different kinds of engines). All other listings (except for differences in grammar, differences in the actual words or associations) were considered to be NADs (i.e. cars have trunks, motorcycles don't). An independent rater reviewed a random 20% subsample of the data. Interrater reliability was 0.89.

Results

Eleven subjects listed commonalities and differences for each of the 64 pairs in this study. Table 1 shows sample protocols. As predicted, this subject listed more commonalities for the highly similar pair <*Chair/ Dresser*> than for the dissimilar pair <*Pine/Motorcycle* > (6 vs. 3). Interestingly, the same number of differences was listed for the similar and dissimilar pair: however, the pattern of ADs and NADs differed. The differences for the similar pair (*Car/Motorcycle*) were mostly ADs (i.e. Cars carry more people than motorcycles) while the differences for the pair <*Sofa/Frog*> were mostly NADs (Frogs have bulging eyes, Sofas don't). [6]

This pattern of results was obtained across items as well. We averaged the responses for each item and performed analyses on the average scores for each of the 64 pairs. The first prediction that commonalities would decrease as similarity decreased was borne out. As shown in Figure 5a, the number of commonalities for each pair was greatest for pairs of distance 0 and decreased as distance increased. These differences were significant ($F(3,60)=11.65$ $p<.001$).

Consistent with the predictions of the frame-alignment view, more ADs were listed for similar items than for dissimilar items, while more NADs were listed for dissimilar items than for similar ones (see Figure 5b). A MANOVA on the number of AD and NAD listed for items indicates that these differences are significant ($F(6,118)=5.272$, $p<001$). Strikingly the total number of differences did not change with distance in the ontology ($F(3,60)=0.011$, $p>.5$).

Next, we examined the relationship between the number of commonalities listed and the numbers of ADs and NADs listed. As predicted, there was a positive correlation between the number of commonalities listed for a particular pair and the number of ADs listed for that pair ($r(62)=0.74$, $p<.001$). However, as expected, there is no positive correlation between the number of listed commonalities and the number of listed NADs for each stimulus ($r(62)=-0.12$, $p>.05$).

Finally, we assumed (in common with Tversky's featural account) that commonalities would increase similarity, while ADs and NADs would decrease similarity. We performed a multiple regression of similarity (SIM) on commonalities (COMM), AD and NAD, and found

$$SIM = 0.676\ (COMM) - 0.293\ (AD) - 1.528\ (NAD).$$

As expected, similarity increases with the number of listed commonalities and decreases with the number of listed ADs and NADs (although the regression coefficient for ADs is not significantly greater than 0). Overall, this equation accounts for 67.6% of the variance in the similarity judgments. Finally, a

[5] Listed commonalities did not have to be true of the items. Half of the subjects considered frogs as amphibians, but half considered them to be reptiles.

[6] The two commonality listings were taken from one subject and the two difference listings were taken from another.

Commonalities

Chair/Dresser (Similarity=5.40)

1	Both are pieces of furniture
2	Both found in the home
3	Both can be made of wood
4	Both might be in your room
5	Both can be bought at the furniture store
6	Both are common, but chairs are more common

Pine/Motorcycle (Similarity=1.35)

1	Both are kept outdoors
2	Both can be dangerous
3	Both touch the ground

Differences

Car/Motorcycle (Similarity=6.20)

AD	Car has 4 wheels/Motorcycle has 2
AD	Car has bigger engine than Motorcycle
AD	Car can carry more people
AD	Motorcycle can go faster than a car
NAD	Car can carry luggage/A motorcycle can't
NAD	Car protects from rain/Motorcycle doesn't
AD	Motorcycle's engine is exposed/Car's engine is under the hood

Sofa/Frog (Similarity=1.30)

AD	Frogs are greenish/Sofas come in many colors
AD	Frogs found outdoors/Sofas found indoors
AD	Sofas are furniture/Frogs are animals
NAD	Sofas can be sat in/Frogs can't
NAD	Frogs jump/Sofas can't
NAD	Frogs have bulging eyes/Sofas don't
NAD	Sofas can be jumped on/Frogs can't
NAD	Frogs can jump/Sofas can't

Table 1: Sample subject protocols

similar regression performed with only commonalities and the total number of differences (AD + NAD) accounted for only 59.8% of the variance.[7]

Discussion and Conclusions

From our analysis of the frame-alignment model, we derived four predictions about the nature of commonalities and differences. First, similarity should increase with commonalities and decrease with differences. Second, the number of commonalities and ADs should be positively correlated. Third, commonalities and ADs should be linked conceptually as well as numerically. Finally, the number of NADs should decrease with similarity. The results obtained in this experiment support all four of these predictions.

We found that the number of commonalities and ADs listed for each item is indeed positively correlated. Despite this relationship, commonalities serve to increase similarity, while differences (both ADs and NADs) decrease similarity. However ADs and NADs are distributed differently. Subjects list more ADs for similar items than for dissimilar items, while they list more NADs for dissimilar items than for similar ones. In fact ADs and NADs appear to play a different role in the calculation of similarity. A regression equation that predicts similarity from commonalities and differences accounts for more of the variance when the differences are broken down as ADs and NADs than when they are aggregated together.

In addition to the numerical relationship between commonalities and ADs, there are semantic relationships. Examination of subjects' protocols clearly reveals this link as well. For example, 91% (10/11) of subjects asked to list commonalities of the

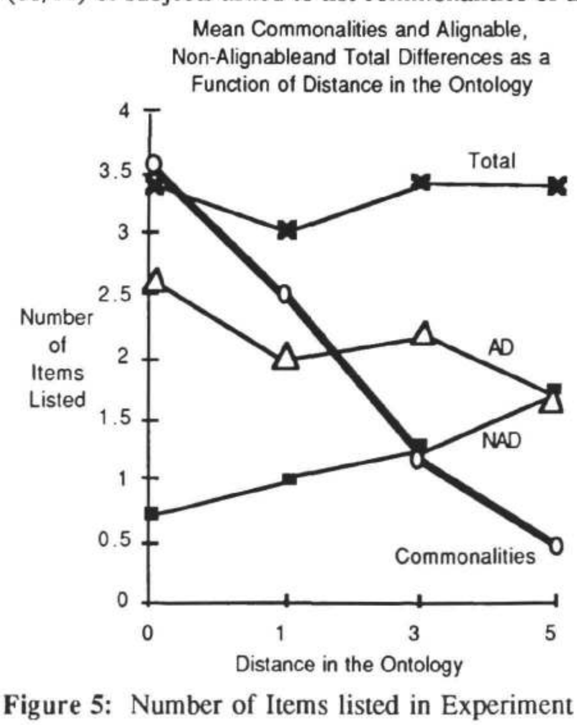

Figure 5: Number of Items listed in Experiment

pair <bus/bicycle> stated that both have wheels, while 82% (9/11) of the subjects asked to list differences for this pair noted that a bus has more wheels than a bicycle. In another example, 100% (11/11) of the subjects listing commonalities of the

[7]Since Total Differences=AD + NAD, we can conclude that merely redescribing the difference data as ADs and NADs accounts for more similarity variance than total differences alone. No statistical test is required.

pair <chair/car> noted that you can sit in both, while 73% (8/11) of the subjects listing differences stated that the primary function of a chair is sitting, while the primary function of a car is travelling. This pattern of responses arose often, and indicates that many of the commonalities and differences are strongly linked. This conceptual connection between commonalities and ADs is a direct prediction of the frame-alignment model. On the other hand, feature-based models, which are unambiguous in their prediction that the sets of common and distinctive features are independent, forbid such conceptual associations.

The frame-alignment model draws on earlier work on analogical mapping (Gentner, 1983, 1989, Holyoak and Thagard, 1989) in order to provide an account of how similarity is determined. Our results indicate that the seemingly simple task of partitioning representations into commonalities and differences requires that the alignment process be sensitive to the relations between features. The use of analogical mapping allows features to be placed in correspondence based on their common position within a relational structure.

The debt of the frame-alignment model to featural models of similarity is clear. Tversky's (1977) contrast model (a featural account of similarity) proposed that similarity increases with the commonalities of the pair and decreases with the differences, a result which was obtained in this experiment. However, the frame-alignment model represents an advance over the feature-matching view in three areas. First, the frame-alignment model predicts the relationship (both numerical and conceptual) between commonalities and alignable differences. Second, the ternary distinction of commonalities, ADs and NADs accounts for subjects' similarity ratings better than the binary distinction of commonalities and differences. Finally, and perhaps most importantly, the frame-alignment view provides a mechanism for comparing two conceptual representations that captures the role of common relational structures in determining what counts as a commonality and what counts as a difference.

Acknowledgments

The authors would like to thank Gary Dell and Brian Ross for helpful comments on earlier versions of this paper. We would also like to thank Evan Heit, Ken Forbus, Rob Goldstone, Doug Medin, Mary Jo Rattermann, Ed Wisniewski and the rest of the Similarity and Analogy Group for many discussions and suggestions during the evolution of this project. A special thanks to Linda May for transcribing the data. Finally, thanks to Bill Parcells for two Super Bowls.

References

Barsalou, L. (in press) Frames and the Construction of Semantic Fields. to appear in A. Lehrer and E. Kittay (eds.) The Structure of the Lexicon: Semantic Fields and Other Alternatives.

Gentner, D. (1983) Structure-Mapping: A theoretical framework for analogy. Cognitive Science, 7, 155-170.

Gentner, D. (1989) The Mechanisms of Analogical Reasoning. in S. Vosniadou and A. Ortony (eds.) Similarity and Analogical Reasoning. London: Cambridge University Press.

Goldstone, R.L., Gentner, D. and Medin, D.L. (1989) Relations Relating Relations in The Proceedings of the 11th Annual Conference of the Cognitive Science Society. Ann Arbor, MI.

Holyoak, K.J. and Koh, K. (1987) Surface and Structural Similarity in Analogical Transfer. Memory and Cognition. 15(4), 332-240.

Holyoak, K.J. and Thagard, P. (1989) Analogical Mapping by Constraint Satisfaction. Cognitive Science. 13, 295-355.

Keil. F.C. (1979) Semantic and Conceptual Development. Cambridge, MA, Harvard University Press.

Kruskal and Wish, M. (1978) Multidimensional Scaling. Newbury Park: SAGE Publications

Lenat, D. and Guha, R.V. (1990). Building Large Knowledge-Based Systems. San Francisco: Addison, Wesley Publishers, Inc.

Logan, G. (1990). Toward an Instance Theory of Automization. Psychological Review.

Minsky, M. (1981) A Framework for Representing Knowledge. in J. Haugeland (ed.) Mind Design. Cambridge, MA: The MIT Press.

Norman, D. and Rumelhart, D.E. (1972) Explorations in Cognition. San Francisco, Freeman Publishers, Inc.

Rattermann, M.J. and Gentner, D. (1987) Analogy and Similarity: Determinants of Accessibility and Inferential Soundness. in The Proceedings of the Ninth Annual Conference of the Cognitive Science Society. Seattle, WA.

Ross, B. H. (1989) Distinguishing Types of Superficial Similarities: Different Effects on the Access and Use of Earlier Examples. Journal of Experimental Psychology: Learning, Memory and Cognition. 15(3), 456-468.

Schank, R., and Abelson, R. (1977) Scripts, Plans, Goals and Understanding. Hillsdale, NJ: Lawrence Erlbaum Associates.

Smith, E.E. and Medin, D.L. (1981). Concepts and Categories. Cambridge, MA, Harvard University Press.

Tversky, A. (1977) Features of Similarity. Psychological Review. 84(4), 327-352.