Title
More than Feature Comparison : Processes Underlying Similarity and Probability Judgment

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 17(0)

Authors
Chase, Valerie M .
Bassok, Miriam

Publication Date
1995

Peer reviewed

More than Feature Comparison: Processes Underlying Similarity and Probability Judgment

Valerie M. Chase
Department of Psychology
University of Chicago
5848 S. University Ave.
Chicago, IL 60637
vmc5@ccp.uchicago.edu

Miriam Bassok
Department of Psychology
University of Chicago
5848 S. University Ave.
Chicago, IL 60637
mbmb@ccp.spc.uchicago.edu

Abstract

Explanations of many cognitive processes, including probability judgment, rely on the construct of similarity. The present paper is concerned with the similarity-based explanation of reasoning in the conjunction task. Although high positive correlations have been found between similarity and probability judgments in this task, these alone cannot validate the assumption that similarity is judged by a process of feature comparison or that similarity judgment is an explanation of probability judgment. Preliminary results from a study in which we collected written justifications from subjects who made both types of judgment suggest that these assumptions are not tenable. Subjects cited considerations of causality and statistics -- not just feature overlap -- when judging both similarity and probability, indicating that (1) feature comparison is only one way in which people judge similarity and (2) similarity judgment can involve processes usually associated with probability judgment. These findings suggest that the role of similarity in explaining other cognitive processes needs to be revised. It is proposed that the power of similarity and probability to predict one another can be exploited for the purpose of making either type of judgment.

Introduction

Explanations of many cognitive processes, including categorization, probabilistic reasoning, and analogical transfer, rest on the construct of similarity at some level. Similarity-based explanations of cognition generally make two crucial assumptions about similarity judgment mechanisms and the role of similarity judgment in explaining other cognitive processes. The first assumption is that similarity is the outcome of a process of *comparison*, usually mediated by a mechanism that evaluates matches and mismatches between features representing the compared entities. The second assumption is that similarity judgment is a fundamental cognitive process that underlies many other cognitive processes and serves as an *explanation* of them.

In the present paper, we present evidence from a protocol study in which we investigated both of these assumptions. We collected similarity and probability estimates in a classic probabilistic reasoning task, the conjunction task. As in previous studies, we found that the two types of estimate are highly positively correlated. However, subjects' written justifications of their estimates revealed that, in addition to feature comparison, judgments in both tasks were also mediated by causal and statistical inferences. We chose to study the conjunction task because the most prevalent explanation of reasoning in this and other probabilistic reasoning tasks -- the "representativeness heuristic" -- rests squarely on the similarity-as-explanation and similarity-as-comparison assumptions.

Similarity and probability judgments

The representativeness heuristic has been defined as a rule of thumb by which subjects make a similarity judgment in lieu of the requested probability judgment (Kahneman & Tversky, 1972). For example, they might judge the similarity between a person and a particular category of people in order to derive the probability that the person is a member of the category. Because similarity is often positively correlated with probability, it is argued that judgment by representativeness is a plausible and useful way to make probability judgments. Thus researchers of probabilistic reasoning have made the similarity-as-explanation assumption by proposing that similarity judgment underlies the judgments made by subjects performing a variety of probabilistic tasks (e.g., Shafir, Smith, & Osherson, 1990). The same researchers have adopted the similarity-as-comparison assumption by operationalizing the representativeness hypothesis in terms of traditional feature-matching models of similarity judgment (e.g., Smith & Osherson, 1989).

Feature-matching models of similarity judgment make simplifying assumptions at the levels of both representation and processing. First, they assume that categories such as "spoon" can be represented as lists of feature slots (e.g., color, size, purpose) that take on different feature values (e.g., brown, large, used for mixing food). Next, they posit that the similarity between two entities is a weighted additive function of their common and distinctive features. Tversky's (1977) classic contrast model of similarity judgment assumes such a feature-matching process and has served as the basis for many subsequent cognitive models involving similarity judgment. Feature-matching models have fared reasonably well in predicting human judgments of similarity between schematic figures (Tversky, 1977; Tversky & Gati, 1978) and between simple adjective-noun conceptual combinations such as "brown apple" and "green apple" (Smith & Osherson, 1984; Smith, Osherson, Rips, & Keane, 1988).

Despite their ubiquity, feature-matching models of similarity judgment have drawn considerable criticism for their lack of constraints and untenable assumptions (e.g., Goodman, 1972; Medin, Goldstone, and Gentner, 1993). These assumptions become increasingly compromising as one moves from relatively simple, artificial categories to naturally occurring object categories such as "spoon," all the way up to the kinds of complex concepts most relevant to probabilistic reasoning, such as "friend" and "conservative." There is considerable evidence that modeling categories in terms of features and feature slots does not capture the structure of people's conceptual representations (Murphy & Medin, 1985), particularly when it is assumed that features are independent (Goldstone, Medin, & Gentner, 1991). Moreover, it has been demonstrated that in similarity judgment, relations as well as features are aligned for the purposes of comparison (e.g., Markman & Gentner, 1993). However, even models that take account of the role of relations cast similarity judgment as a process of comparison.

In this paper, we explore the processing that underlies similarity judgment in the conjunction task as well as in a typicality judgment task based on it. In a typical conjunction task, subjects receive a description of a person (e.g., "Linda is 31, single, majored in philosophy and was concerned with issues of discrimination and social justice in college") and are asked to judge the probability that the person described is a member of three categories: two constituent categories (e.g., bank teller, active in the feminist movement) and a conjoint category comprised of the two constituents (e.g., bank teller who is active in the feminist movement). When the description is informative with respect to the categories and especially when the description differentially "points to" the constituent categories as in the Linda task, most subjects (typically about 80%) rank the probability of the conjoint category to be higher than the probability of one of the constituent categories (i.e., they tend to think it is more probable that Linda is a bank teller who is active in the feminist movement than that she is a bank teller). This is considered to be a violation of the conjunction rule.

The similarity-based explanation of conjunction rule violations states that subjects assess the given person's similarity to each category representation in order to judge the probability that the person is a member of each category. The claim that similarity judgment underlies probability judgment in the conjunction task has received support from studies in which typicality[1] ratings were found to be highly positively correlated with probability ratings on the same items (Shafir *et al.*, 1990; Tversky & Kahneman, 1983). Shafir *et al.* (1990) explained the high positive correlation in qualitative terms as the result of a feature-comparison process responsible for both typicality and probability judgment. Elsewhere, Smith and Osherson (1989) proposed a formal feature-matching model of reasoning in the conjunction task based on Tversky's (1977) contrast model.

The positive correlation between typicality and probability judgments used to support the similarity-based explanation of probabilistic reasoning leaves two questions unanswered: (1) Does typicality judgment underlie probability judgment, does probability judgment underlie typicality judgment, or are both supported by another process? (2) Assuming that typicality judgment underlies probability judgment, is feature comparison the only process by which typicality is computed? If not, what other processes are involved?

The present study was designed to investigate these questions using items and measures adapted from Shafir *et al.* (1990). In order to assess the reasoning behind the estimates, we asked subjects to provide written justifications describing their thinking as they solved the task. This is the first study we know of in which justifications were collected in a correlational study of probability and typicality judgment. Although such justifications will not necessarily capture the underlying cognitive processing, they should reveal something about subjects' understanding of what information is relevant to the task and how this information should be used.

Method

Subjects. Subjects were 160 University of Chicago undergraduates who completed the experiment in small groups (2-8) immediately after various classroom lectures.

Materials. The experiment was based on Experiment 1 in Shafir *et al.* (1990). Each subject received a questionnaire consisting of an instruction page and a response page. The instruction page informed the subject that the task was to make either probability or typicality estimates. (Estimate type was varied between subjects.) A sample item was given so that subjects knew what to expect in the task. The response page showed one item consisting of a personality description and three categories to be judged.

Procedure. Each subject received one of eight possible personality descriptions. Depending on the condition to which they were assigned, subjects were required to judge either the probability that the person described was a member of each of three categories or to judge the typicality of the person in each of the same categories. They were instructed to express their estimates on a 0-to-1 scale, where 0 meant minimal probability (typicality) and 1 meant maximal probability (typicality). Each category triplet consisted of two single categories and a third category representing their conjunction. For example, a subject in the probability condition who read the description of Linda summarized earlier might have to judge the probability that Linda is (1) a bank teller, (2) active in the feminist movement, and (3) a bank teller who is active in the feminist movement, while a corresponding subject in the typicality condition judged Linda's typicality in the categories "People who are bank tellers," "People who are active in the feminist movement," and "People who are bank

[1]Ordinarily, the term "similarity" is used when the comparison is between instances, and "typicality" when the comparison is between an instance and a category. The terms will be used interchangeably here.

tellers and who are active in the feminist movement." Order of constituent categories was randomized, with the conjoint category always appearing last.

Half of the conjunctions, such as "bank teller who is active in the feminist movement," were "incompatible." The other half of the conjunctions, such as "teacher who is active in the feminist movement," were "compatible." Shafir *et al.* (1990) defined an incompatible conjunction in feature-matching terms as one whose constituents share few properties and a compatible conjunction as one whose constituents share many properties. Although we prefer not to define compatibility in this way, it should be intuitively clear that the two types of conjunction differ in how well they "fit together." (The compatibility factor did not affect the results reported in this paper and therefore will not be discussed further.) Three of the eight personality descriptions and their associated conjunctions (which will be used for the purpose of illustration later) appear below. Compatible conjunctions are labeled "C" and incompatible conjunctions "I."

Linda is 31, single, outspoken, and very bright. She majored in philosophy. In college, she was concerned with issues of discrimination and social justice and participated in antinuclear demonstrations.

C: teacher who is active in the feminist movement
I: bank teller who is active in the feminist movement

Jack began his job immediately after completing high school. He frequently talks on his CB radio and goes to sporting events when he can.

C: truck driver who plays softball for a hobby
I: truck driver who watches birds for a hobby

Richard is 50 years old. He loves his job, but is not very well liked by his colleagues. He is single, shy, and does not like to go out on social events.

C: engineer who collects stamps for a hobby
I: engineer who plays volleyball for a hobby

Data Analysis and Results

We first report on subjects' quantitative estimates and then present our qualitative analysis of their written justifications.

Quantitative estimates

Consistent with previous studies, we found that the mean probability and typicality estimates for a given category in a given item were positively correlated. The Spearman correlations between the two types of estimates for matching items were .71 for constituents and .78 for conjunctions. We computed the parallel correlations from the data reported by Shafir *et al.* (1990) and found them to be .91 and .73 for constituents and conjunctions, respectively. Thus, we replicated the general finding that probability and typicality estimates in the conjunction task are positively correlated. In the next section we report results that shed light on the nature of this correspondence at the level of processing.

In the probability condition, 24% of subjects (19 out of 80) estimated the probability of the conjunction to be greater than one or both of the constituents, i.e., violated the conjunction rule. In the typicality condition, 46% of subjects (37 out of 80) estimated the typicality of the instance in the conjoint category to exceed its typicality in one or both constituent categories. Although our finding of 24% violations may seem to contradict the results of other studies, in which up to 87% of subjects violated the conjunction rule in the Linda problem, Tversky and Kahneman (1983) themselves found that as few as 36% of subjects violated the conjunction rule in the Linda problem under some conditions.[2] Because Shafir *et al.* (1990) did not report the percentages of violations by item, we could not directly compare our results on this measure to theirs in either condition.

We believe that some probability and typicality subjects whose conjoint estimates did not exceed either constituent estimate allowed the conjunction rule to guide their judgments. From the binary measure of conjunction rule violations, it is impossible to tell whether the probability subjects who did not violate the rule followed it systematically or inadvertently. Similarly, one cannot tell if the typicality subjects intentionally applied the conjunction rule. However, we can begin to answer this question indirectly by looking at the estimates to see what percentage of subjects gave conjoint estimates exactly equal to the lower of their two constituent estimates. A high percentage would suggest that these subjects allowed the lower of the two constituent estimates to constrain the conjoint estimate.

In support of the notion that some subjects intentionally used the conjunction rule, 39% of probability subjects who did not violate the rule (24 out 61) gave a conjoint estimate equal to the lower of the two constituent estimates. Of the typicality subjects whose conjoint estimates did not exceed the lower of the two constituent estimates, 47% (20 out of 43) made the two estimates equal. It may not be surprising that some probability subjects seem to have applied the conjunction rule. It is surprising, however, that some typicality subjects seem to have done so. This quantitative result in the typicality condition is striking for two reasons: (1) One cannot reasonably argue that the conjunction rule applies to typicality judgment and it seems highly unlikely that people are ever instructed that it does; (2) Relative to the less typical constituent, the conjoint category was always more similar to the instance from a feature-matching perspective. The idea that typicality subjects applied statistical strategies to their judgments will receive further support from the analysis of justifications presented below.

[2] This low percentage of violations was found in a statistically sophisticated population of graduate student subjects who used rating scales rather than rankings to make their responses. We think that the University of Chicago undergraduate subjects in our study may be comparable to that population in statistical knowledge and that the 0-1 estimation response mode is more analogous to a rating scale than a ranking response mode, making this the most appropriate study to which to compare the present one.

Justifications of estimates

In general, the content of the justifications challenged feature-matching assumptions about both representation and processing. As for representation, we found that subjects in both conditions often went beyond the information given in the personality description and reasoned using their enriched representations of the instances. In particular, subjects inferred features of the instance that they then used to make their judgments. These inferred features naturally affected the selection of category features that were brought to bear on the judgment. The fact that subjects inferred features in the conjunction task highlights a weakness of feature-matching models first pointed out by Tversky (1977) himself: They fail to specify the features of instance and category that enter into the feature comparison. In addition, our finding that the "features" inferred were in some cases events rather than properties seems incompatible with the slot-value representations of categories assumed by feature-matching models.

To study processing, we had one judge categorize the justifications from both conditions using a scheme developed on justifications collected in a pilot study. The judge was trained on the pilot justifications and was blind both to condition (probability or typicality) and the purpose of the experiment. It was possible for a justification for a single estimate to fall in more than one category.

The four coding categories are explicated below. Examples are shown after each category definition (except Other) and are labeled "T" or "P" to indicate the condition (typicality or probability) from which they were drawn. The justifications for 17% of the constituent estimates and 40% of the conjoint estimates were omitted from analysis because subjects did not write anything or wrote statements that could not be categorized.

Feature comparison. Feature comparison involved mentioning a feature of the instance or category and either pointing out or implying the presence or absence of that feature in the other entity. Features of the instance could be directly quoted, paraphrased, or inferred from the given description. Examples of feature comparison justifications:

P: It appears very likely that Linda is an active feminist, because she is outspoken and interested in debates over social issues. This could mean that she is an ultraconservative antifeminist, but it doesn't seem logical because her opposition to nuclear power seems to connect her with left-wing rather than conservative ideas.

P: [Jack is] too aggressive to watch birds.

T: People who play volleyball play on a team, which implies that they are social. Richard is not like this.

T: Truck drivers use CB radios while on the job and often are known (or stereotyped) to watch sports.

Causal reasoning. When a justification specified how a feature of the person or a hypothesized set of circumstances could cause the person to be or not to be a member of the category, then it was considered to show evidence of causal reasoning. The most common type of causal reasoning involved construction of a hypothetical scenario about the instance that was then evaluated for its plausibility. Whether or not the causal scenario was accepted as plausible or rejected as implausible by the subject, the justification was put in this category if a causal mechanism was specified. Examples of causal reasoning justifications:

P: If she was a teacher, she would appear to be a rebel and more than likely be fired for teaching in what the school would believe to be an inappropriate way of teaching.

T: Volleyball is a team sport. A shy person is unlikely to play a sport that requires finding 9 other people in order to play.

T: This could be a temporary job for Linda.

Statistical and mathematical reasoning. We identified four types of statistical and mathematical justifications: base-rate, "general," conjunction rule, and averaging rule.

Base-rate and general statistical or mathematical justifications were used for both constituents and conjunctions. Base-rate justifications took into account the subject's estimate of the percentage of instances in a reference class (e.g., all people, men, 31-year-olds) that fall in a particular category (e.g., people who are teachers). General justifications mentioned statistics or mathematics as a consideration but failed to specify how statistics or mathematics was used to make the estimate. Examples of base-rate and general justifications:

P: I think because an engineer is independent of collecting stamps and I guess half of the people collect stamps.

T: While it is not unreasonable for him to be one or the other, I imagine the percentage of bird-watching truck drivers is low, so Jack has a small category in which to fit.

P: This would be the lowest probability because of the statistics involved.

T: Two qualifications are less likely to be had than either alone.

The remaining two types of statistical and mathematical justification cited rules for combining the constituent estimates and thus could only be used to justify conjoint estimates. Conjunction rule justifications were verbal expressions of the constraint that the lower of the two constituent estimates places an upper limit on the conjoint estimate or algebraic or arithmetic statements of the multiplication rule for independent events. Averaging rule justifications stated that the conjoint estimate was derived by quantitative or qualitative averaging of the constituent estimates. For combining probabilities, the averaging rule was incorrect from a normative standpoint but was still considered to be a mathematical strategy. Examples of conjunction rule and averaging justifications:

P: Mathematically, the odds of her being a bank teller and active in the feminist movement must be less than (or equal to) the odds of her just being a bank teller.

P: I multiplied the probabilities #1 and #2 [the constituent estimates] together to come up with this estimate. I guess it can be said that very low X very low = even lower.

T: I multiplied my first estimate by the second estimate because it seemed the most logical thing to do.

T: [The conjoint estimate is] the average of the first two.

Other. This catch-all category included indifference and conceptual combination justifications. If a justification said that it was impossible to make an estimate or that the person was as likely as not to be in the category *and* the quantitative estimate was .50, then it was categorized as indifference. Judgments of indifference were qualitatively different from all others in that they represented reactions to the task rather than responses to the question posed in the task (see also the Dick problem in Kahneman & Tversky, 1973). Conceptual combination justifications made assertions about the plausibility of the conjunction independent of the instance. The latter often had the flavor of base-rate justifications (e.g., "I don't think this category exists," "What graduate student goes to fashion shows?"), but were tallied separately for the sake of conservatism.

The percentages of justifications that fell in each category are shown in Table 1. Because the possible justification types differed for constituents and conjunctions, the results are broken down by category type (constituent or conjunction) as well as by experimental condition (probability or typicality). For constituents (top panel), the statistical and mathematical justifications are all of the base-rate subtype. Conjunction rule justifications accounted for 65% and 59% of statistical and mathematical justifications in the probability and typicality conditions, respectively. Conceptual combination justifications accounted for 77% and 78% of Other justifications in the two conditions, respectively.

Examination of the Table 2 shows that for constituent categories, feature comparison is by far the most common type of justification given for both probability and typicality judgments. The distributions of the other types of justification are comparable in the two conditions, with base-rate justifications and statements of indifference slightly more common in the probability than in the typicality condition. For conjunctions, however, feature comparison justifications were not in the majority in either condition. In both conditions, the distribution for conjunctions is concentrated in the statistical and mathematical and Other categories.

Table 1: Percentages of justification types.

Constituents

	Feature	Causal	StatMath	Other
prob N=147	74	8	9	9
typ N=139	84	8	3	6

Conjunctions

	Feature	Causal	StatMath	Other
prob N=50	18	10	46	26
typ N=50	38	10	34	18

The shift from feature comparison to statistical and mathematical justifications as one moves from constituents to conjunctions suggests that the underlying processing depends on the complexity of the category being considered. While comparing the features of an instance to a constituent category may be straightforward because the category features can be retrieved from memory, comparing them to a conjoint category first requires representing the conjoint category. As has been pointed out by many theorists (e.g., by Murphy and Medin, 1985), conjoint categories are not formed by merging feature lists, but by more complex cognitive processes such as causal reasoning. One way of approaching the conjunction task is to form a conjoint category representation by such processes and then to compare features of instance and category. Another way is to borrow combination rules from statistics and mathematics to derive a conjoint estimate. The justifications collected in our study suggest that some subjects took the latter approach to generating conjoint estimates.

Discussion

The results of our analysis of subjects' justifications speak to the similarity-as-comparison and similarity-as-explanation assumptions in a way that quantitative estimates alone cannot. To the extent that the justifications reveal underlying processing, probability and similarity judgment involve more than feature comparison. Particularly when judging the probability or typicality of instances in conjoint categories, some subjects justified their estimates by citing causal or statistical and mathematical considerations that cannot be incorporated into traditional feature-comparison models of similarity judgment.

One way to preserve the notion of similarity judgment strictly as a process of comparison might be to argue that processes other than feature comparison are distinct from similarity judgment but are sometimes used in addition to similarity to solve a particular task. But in that case, it is difficult to argue categorically that similarity judgment is made in lieu of probability judgment, since even in the typicality condition feature comparison was only one of several types of justification. Indeed, one could argue from the finding that some subjects borrowed statistical and mathematical rules to make their typicality estimates that probability judgment can also underlie similarity judgment. Our results validate similarity-based explanations of probabilistic reasoning to the extent that the similarity and probability judgment tasks elicited similar inferential processes, including comparison of features. However, feature comparison was only one of the processes of which we found evidence.

Not only do our results suggest the need to rethink the role of similarity as an explanation of other cognitive processes, but the need to change our conceptualization of similarity judgment itself. The justifications collected from typicality subjects in our study suggest that they reason about similarity in terms of plausibility and probability as well as in terms of feature comparison. In the case of causal scenarios, they evaluated the plausibility of an event sequence, beginning with the instance as described and ending with the instance as category member or nonmember (see also Bassok & Medin, this volume). In an even clearer departure from similarity judgment as feature comparison, some subjects borrowed statistical and mathematical strategies such as base-rate estimation and the multiplication rule for independent events to estimate typicality. It should be noted that causal and statistical strategies for solving both judgment tasks were not used to the exclusion of feature comparison. Further analysis of the written justifications is expected to shed light on how individual subjects deployed the different strategy types.

In principle, the ecologically valid positive correlation between similarity and probability can be exploited for the purpose of making either type of judgment: Similarity can be used as a heuristic for judging probability and probability can be used a heuristic for judging similarity. The ambiguity of both judgment tasks may lead subjects to look to other dimensions about which they are knowledgeable for guidance in defining the tasks and in making their judgments. Although the statistical and mathematical knowledge of subjects in the present study may be greater than that of subjects in most previous conjunction task studies, the fact that they relied on such knowledge to guide their typicality judgments nevertheless has broad implications because they also have knowledge of feature comparison. The point is that subjects drew on knowledge outside of feature comparison to make both similarity and probability judgments, using a variety of knowledge sources to guide the judgment they were required to make. We conclude that, because it involves inferential processes other than feature comparison, similarity judgment as it has been defined cannot necessarily serve as a unidirectional explanation of other cognitive processes.

Acknowledgments

The present study was supported by a grant from the University of Chicago School Mathematics Project to Miriam Bassok. We thank Boaz Keysar and Beth Hagen for their helpful comments on an earlier draft of this paper.

References

Bassok, M. & Medin, D. L. (1995). Structural and thematic alignments in similarity judgments. In *Proceedings of the Seventeenth Annual Cognitive Science Society* (this volume). Hillsdale, NJ: Lawrence Erlbaum.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology,* 23, 222-262.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and projects* (pp. 437-447). New York: Bobbs-Merrill.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology,* 3, 430-454.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review,* 80, 237-251.

Markman, A. B., & Gentner, D. (1990). Similarity involving attributes and relations: Judgments of similarity and difference are not inverse. *Psychological Science,* 1, 64-69.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review,* 100, 254-278.

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

Shafir, E. B., Smith, E. E., & Osherson, D. N. (199). Typicality and reasoning fallacies. *Memory and Cognition,* 18, 229-239.

Smith, E.E., & Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science,* 8, 337-361.

Smith, E.E., & Osherson, D. N. (1989). Similarity and decision making. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 60-75). Cambridge: Cambridge University Press.

Smith, E.E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science,* 12, 485-527.

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

Tversky, A., & Gati, I. (178). Studies of similarity. In E. Rosch & B. Lloyd (Eds.), *Cognition and Categorization* (pp. 79-98). Hillsdale, NJ: Erlbaum.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review,* 90, 293-315.