COGNITIVE SCIENCE 12, 485-527 (1988)

Combining Prototypes: A Selective Modification Model

EDWARD E. SMITH
University of Michigan

DANIEL N. OSHERSON
Massachusetts Institute of Technology

LANCE J. RIPS
University of Chicago

MARGARET KEANE

Massachusetts Institute of Technology

We propose a model that accounts for how people construct prototypes for composite concepts out of prototypes for simple concepts. The first component of the model is a prototype representation for simple, noun concepts, such as *fruit*, which specifies: (1) the relevant attributes of the concepts, (2) the possible values of each attribute, (3) the salience of each value, and (4) the diagnosticity of each attribute. The second component of the model specifies procedures for modifying simple prototypes so that they represent new, composite concepts. The procedure for adjectival modification, as when *red* modifies *fruit*, consists of selecting the relevant attribute(s) in the noun concept (color), boosting the diagnosticity of that attribute, and increasing the salience of the value named by the adjective (*red*). The procedure for adverbial modification, as in *very red fruit*, consists of multiplication-by-a-scalar of the salience of the relevant value (*red*). The outcome of these procedures is a new prototype representation. The third component of the model is Tversky's (1977) contrast rule for determining the similarity between a representation for a prototype and one for an instance. The model is shown to be consistent with previous findings about prototypes in general, as well as with specific findings about typicality judgments for adjective-noun conjunctions. Four new experiments provide further detailed support for the model.

Research on natural concepts, such as *apple* and *fish*, has led to the conclusion that part of the mental representation of a concept consists of a "prototype," roughly, a description of the best examples or central tendency of

We thank Susan Carey, Allan Collins, Dedre Gentner, David Israel, Michael Lipton, and Mary Potter for their many suggestions that helped shape this research. The research was supported by U.S. Public Health Service Grant MH37208 and by the National Institute of Education under Contract No. US-HEW-C-400-82-0030.

Correspondence and requests for reprints should be sent to Edward E. Smith, Department of Psychology, University of Michigan, 330 Packard Road, Ann Arbor, MI 48104.

a concept. Specifically, research has shown that the instances of any concept vary in how typical they are rated, and that such ratings predict how quickly and accurately an instance can be categorized, how readily it can be retrieved from memory, how early it can be learned, how efficiently it can be coded linguistically, and so on (see e.g., Mervis & Rosch, 1981). In light of such findings, it seems reasonable to posit that experience, direct or indirect, with exemplars of a concept gives rise to a prototype for that concept, that the rated typicality of an instance is a good predictor of its similarity to its prototype, and that similarity-to-prototype plays some role in categorization, memory and communication (see, e.g., Smith & Medin, 1981).

Such is the account commonly advanced within the prototype tradition for "simple" concepts, that is, concepts denoted by single words. But what about "composite" concepts, such as *striped apple* and *literary politician* (i.e., concepts formed from simpler constituents)?[1] For many such concepts, their instances seem to vary in typicality; yet we cannot have induced a prototype for each composite concept on the basis of experience with its exemplars, because many composite concepts are novel combinations and hence unfamiliar. There must therefore be some means of computing the typicality of an instance in a composite concept from knowledge about its constituents, some way of determining, say, that Lassie is not a very good example of *ferocious animal*, given what we know about Lassie and the concepts *ferocious* and *animal*.

More generally, if prototype theory is to be extended to composite concepts, principles of conceptual composition must be supplied. This is the concern of the present paper. In particular, we will focus on adjective–noun conjunctions such as *striped apple* and *not very red fruit*, and specify how prototypes for such conjunctions can be composed from prototypes for their constituents. While the specifics of our claims apply to only adjective-noun compounds, some of the broader principles we espouse may also characterize noun–noun compounds such as *dog house*.[2]

Our exposition is arranged as follows: First, we motivate and present a model of how adjectives modify noun prototypes to form prototypes for conjunctions. Second, we describe in detail an initial experimental test of this model. Third, we present two subsequent tests of assumptions of the model. Fourth, we extend the model to handle conjunctions involving adverbs. Fifth, we provide an experimental test of the extended model. Sixth and finally, we take up a number of outstanding issues.

---

[1] We use italics to indicate concepts, and reserve quotes for the words that denote these concepts.

[2] Fodor (1981) has argued that some composite concepts do not have prototypes, and hence there must be more to concepts than prototypes. This argument in no way conflicts with the present analyses. We are concerned only with those composite concepts that do have prototypes, and we have argued elsewhere that such prototypes do not exhaust the contents of a concept (see, e.g., Osherson & Smith, 1981).

A MODEL FOR ADJECTIVE-NOUN CONJUNCTIONS

Rationale for the Model

General Aspects of Prototypes. The term "prototype" has sometimes been used to mean a representation of the best example for a given concept (e.g., Medin & Schaffer, 1978; Mervis & Rosch, 1981). In this sense, a prototype for *apple* might be an image or a mental description of an especially good *apple* instance. Our own theory generalizes this idea to allow a prototype to be a more abstract description of the concept. In our view, a prototype is a prestored representation of the usual properties associated with the concept's instances (much as in schema or frame theory—see, e.g., Minsky, 1975; Rumelhart & Ortony, 1977). Thus, an *apple* prototype will include properties such as having seeds, properties that are part of our commonsense knowledge about apples. Earlier work on prototypes indicated that a concept's prototype includes properties that are not strictly necessary for concept membership (e.g., Rosch, 1973; Smith, Shoben, & Rips, 1974). The prototype of *apple,* for example, includes the nonnecessary properties of red, round, and smooth. Subsequent work has shown that the contents of a prototype must include far more than a list of properties.

For one thing, we need to decompose the notion of a property into two components: *attribute* and *value.* Thus the *apple* prototype includes attribute-value pairs such as color-red, shape-round, and texture-smooth. The reason for including attributes in prototypes is simply that there are numerous cases where people use attribute knowledge in categorization. Consider categorization with negative concepts, such as *nonred fruit.* Without the notion of an attribute, how can one ever know that a blueberry is an instance of *nonred fruit*? To know that blue counts as *nonred* while round does not, one must know that a certain set of values (the colors) constitutes an attribute.

A prototype also includes some indication of the salience of each relevant value. A couple of lines of evidence point to this conclusion. For one thing, when asked to verify that a property is true of a particular concept, people respond faster to properties that have previously been rated as more related or associated to the concept than to those rated less related (e.g., Glass & Holyoak, 1975). Thus, people are faster at deciding that *apples are red* than *apples are round,* suggesting that red is more salient than round in the prototype for *apple.* Another line of evidence for salience is that the nature of a value seems to be relative to a concept, the red in *apple,* for example, being redder than that in *brick* but less than that in a *fire engine* (Halff, Ortony, & Anderson, 1976). This suggests that the red in *apple* is more salient than that in *brick,* though less salient than the red in *fire engine.*

Finally, a prototype may also include some indication of the diagnosticity of each attribute, that is, a measure of how useful the attribute is in discriminating instances of the concept from instances of contrasting concepts. The importance of diagnosticity was demonstrated by Rosch and Mervis (1975); when subjects have to decide whether or not an item belongs to a target concept, they consider not only the item's attribute-by-attribute similarity to the target concept but also its attribute-by-attribute dissimilarity to concepts that contrast with the target.

In short, any model of prototype composition would do well to start with prototypes that include: (1) an attribute-value structure, (2) indications of value salience, and (3) indications of attribute diagnosticity.

Conjunction Effects. In addition to the above three general aspects, the development of our model was guided by three specific findings that involve typicality judgments for adjective-noun combinations. The most important of these findings we call the "conjunction effect." To illustrate the effect, consider the typicality of a particular red apple as an instance of the concepts *apple* and *red apple.* Several experiments have found that the rated typicality of the instance in the conjunction exceeds that in the simple concept. Our red apple is judged more typical of *red apple* than of *apple* (Hampton, 1982; Osherson & Smith, 1982; Smith & Osherson, 1984; Shafir, Smith, & Osherson, 1988; for a related effect, see Tversky & Kahneman, 1983).[3]

Smith and Osherson (1984) uncovered a second phenomenon of interest. They investigated "incompatible" conjunctions, such as *brown apple,* where the adjective denotes an unlikely value of the object denoted by the noun, and "compatible" conjunctions, such as *red apple,* where the adjective denotes a likely value of the object denoted by the noun. They found that the conjunction effect is greater for incompatible than compatible conjunctions. For example, the extent to which a brown apple is judged more typical of *brown apple* than of *apple* is greater than the extent to which a red apple is judged more typical of *red apple* than of *apple.* A third finding arises when the item to be categorized is not a true member of the conjunction (e.g., a brown apple paired with the conjunction *red apple*). Unsurprisingly, here there is a "reverse conjunction" effect, the item being judged less typical of the conjunction than of the noun constituent (Smith & Osherson, 1984).

The Selective Modification Model

The model that we propose is an extension of one discussed in Smith and Osherson (1984). The current model has three major components: (1) a prototype representation for simple noun concepts, (2) procedures for modifying such a prototype, and (3) a means for determining the typicality of an

[3] We follow Tversky and Kahneman (1983) in calling adjective-noun phrases "conjunctions." However, we do not mean to imply that phrases such as "red apple" are equivalent in meaning to explicit conjunctions such as "both red and apple." Indeed, in "red apple," the adjective seems to modify the noun concept rather than combine with it conjunctively (Oden, 1984), an intuition that lies at the heart of the model we develop.

instance vis-a-vis a prototype. We begin by describing the first and third components. The descriptions in this section are illustrative, more precision being supplied in the next section.

Typicality and Simple Concepts. A prototype for the concept *apple* is illustrated in the left-most panel of Figure 1, and it includes the three general aspects discussed earlier. The representation specifies: (1) A set of relevant attributes (color, shape, texture, etc.), and for each attribute a set of possible values that instances of the concept can assume (e.g., for color, the values include red, green, and brown); (2) The diagnosticity of each attribute for the concept, as indicated by the number to the attribute's left, and (3) The salience of each value of an attribute, as indicated by the number to the value's right (we refer to these numbers as "votes" for the value). With regard to a value's salience, we suspect that it reflects at least two contributing factors: the subjective frequency with which the value occurs in instances of the concepts, and the perceptibility of the value. Thus, red apples are encountered more frequently than green ones, and that may be why the red in *apple* is more salient (has more votes) than the green in *apple*. Also, the red of an apple is more perceptible than the red of a brick, and that may be why the red in *apple* is more salient (has more votes) than the red in *brick*.

The remaining panels of Figure 1 illustrate representations for two specific objects (or "instances"), I_1 and I_2, a typical red apple and a brown apple. We assume that an object representation is like a prototype except that it does not contain any indication of an attribute's diagnosticity. (We let the prototype alone determine attribute diagnosticity because it is being used as a standard against which the object is compared.) For simplicity, we have assumed further that for I_1 and I_2 all votes for an attribute are on one value (but this need not be generally true of object representations).

To determine an instance's typicality in a concept, we assume that typicality rests on similarity. To measure similarity, we use Tversky's (1977) "contrast" rule, which assesses similarity by a contrast between common and distinctive features. In our application of this model, each vote counts as a feature. In essence, n votes on a value (say, red) is equivalent to there being n copies of that value (n reds). The similarity between the features of a prototype ("P") and the features of an instance ("I") is given by:

$$\mathrm{Sim}(P,I) = af(P \cap I) - bf(P-I) - cf(I-P), \qquad (1)$$

where $P \cap I$ designates the set of votes or features common to the prototype and instance, P–I designates the set of features distinct to the prototype, and I–P designates the set of features distinct to the instance. In addition, f is a function that measures the importance of each of these three set of features, and a, b, and c are parameters that determine the relative contributions of the three sets. The basic idea is that similarity is an increasing function

Apple (A)		I_1		I_2	
1 color	red 25	color	red 30	color	red
	green 5		green		green
	brown		brown		brown 30
	—		—		—
	—		—		—
0.50 shape	round 15	shape	round 20	shape	round 20
	square		square		square
	cylindrical 5		cylindrical		cylindrical
	—		—		—
	—		—		—
0.25 texture	smooth 25	texture	smooth 30	texture	smooth 30
	rough 5		rough		rough
	bumpy		bumpy		bumpy
	—		—		—
	—		—		—

$$\begin{aligned} \text{Sim}(A, I_1) &= 1\,(25\text{-}5\text{-}5) \\ &+ 0.50\,(15\text{-}5\text{-}5) \\ &+ 0.25\,(25\text{-}5\text{-}5) \\ &= 15 + 2.5 + 3.75 \\ &\approx 21 \end{aligned}$$

$$\begin{aligned} \text{Sim}(A, I_2) &= 1\,(0\text{-}30\text{-}30) \\ &+ 0.50\,(15\text{-}5\text{-}5) \\ &+ 0.25\,(25\text{-}5\text{-}5) \\ &= -60 + 2.50 + 3.75 \\ &\approx -54 \end{aligned}$$

Figure 1. Illustration of attribute-value representations for a prototype (*apple*) and relevant instances (a red apple and a brown apple); beneath each instance-representation is the similarity between the instance and prototype.

of the features common to the prototype and instance, and a decreasing function of the features distinct to the prototype and of those distinct to the instance.

For purposes of making computations, it is convenient to use a version of Equation (1) that specifies the common and distinctive features on an attribute-by-attribute basis. This is given by:

$$\text{Sim}(P,I) = \Sigma_i[af_i(P \cap I) - bf_i(P\text{–}I) - cf_i(I\text{–}P)], \tag{2}$$

where i indexes the relevant attributes, and now $P \cap I$ designates the set of features or votes on attribute i common to the prototype and instance, P–I designates the set of features of attribute i distinct to the prototype, and I–P designates the set of features of attribute i distinct to the instance. Beneath each object representation in Figure 1, we have used Equation (2) to calculate the object's similarity to the *apple* prototype. We have assumed that a, b, and c are equal to one (just to keep things simple for now). We have further assumed that f_i multiplies the number of votes for attribute i in a set by the diagnosticity of i. To illustrate, to determine the similarity between the typical red apple (designated "I_1") and the prototype for *apple* (designated "A") on the color attribute, we note that *apple* and the red apple share 25 red votes, that *apple* has 5 distinct green votes, that the red apple has 5 distinct red votes, and that each component of the contrast is multiplied by the diagnosticity of 1.0 (see Figure 1). The computations are similar for the other attributes. For the examples provided in Figure 1, the contrast rule correctly predicts that the red apple, I_1, should be judged to be more typical of *apple* than is the brown apple, I_2.

Note that the only representational difference between I_1 and I_2 is on the color attribute. This difference eventuates in a large typicality difference between I_1 and I_2 because color has a substantial number of features (votes) and is a very diagnostic attribute. To appreciate the importance of diagnosticity, consider a third possible instance, I_3, which is identical to I_1 except that all its texture votes are on bumpy. The typicality of I_3 in *apple* is:

$$\begin{aligned} \text{Sim}(A,I_3) &= 1(25-5-5)+.5(15-5-5)+.25(0-30-30) \\ &= 15+2.5-15 \\ &= 2.5 \end{aligned}$$

Hence, I_3 is more typical of *apple* than is I_2, solely because I_3's mismatching attribute (texture) is less diagnostic than that of I_2 (color). This version of the contrast rule thus nicely captures differences in diagnosticity as well as differences in features or salience.

Adjective Modification. To extend this account to adjective-noun conjunctions, we need to specify how the adjective interacts with the noun. Two general ideas about the nature of this interaction have been proposed (see Cohen & Murphy, 1984, for discussion). One possibility is that the adjective and noun play symmetrical roles and that in forming a conjunction the features of the two concepts are somehow intersected. The other possibility is that the adjective and noun concepts play different and asymmetrical roles, the noun being the basic frame to be operated on and the adjective being the operator or modifier. We opt for the latter, "modification" approach. One reason for doing so is that there can be a striking change in meaning when the order of an adjective-noun combination is reversed—consider *red apple* versus *apple red*—and it is not obvious why this sort of change should occur if the conjunction is some kind of symmetrical intersection of the two prototypes. Other reasons for favoring a modification approach over an intersection one are discussed by Cohen and Murphy (1984).[4]

Our basic proposal about the modification process is as follows: Each attribute in the adjective concept selects the corresponding attribute in the noun concept; then, for each selected attribute in the noun, there is an increase in the salience (or votes) of the value given in the adjective, as well as an increase in the diagnosticity of the attribute. Consider *shriveled apple* as an example. Presumably *shriveled* contains attributes pertaining to shape and texture; accordingly, it would select these attributes in the *apple* prototype, boost their diagnosticities, and shift their votes away from round and smooth and toward irregular and bumpy.

In developing a precise account of the model, however, we will consider only those adjectives that presumably contain a single attribute, for example, *red* or *brown*. Figure 2 illustrates our specific assumptions for such cases. The adjective: (1) selects the relevant attribute in the noun (e.g., color), (2) shifts all votes on that attribute into the value named by the adjective, and (3) boosts the diagnosticity associated with the attribute. In the example at the top of Figure 2, most of the color votes already were on the value specified by the adjective, so few votes have to be shifted; this is the hallmark of compatible conjunctions. In the example at the bottom of Figure 2, all color votes have to be shifted, which is the hallmark of incompatible conjunctions.[5]

The above proposals hinge on two distinct intuitions about modification: roughly, that color is more important for determining typicality in *red apple*

---

[4] Another approach to modification is possible if one adopts the view that a prototype consists of the best examples of a concept. Modification might amount to a change in the examples; for example, in the prototype for *apple* two of the three best examples might be red, while in the prototype for *red apple* all three best examples might be red. A serious problem with this approach is that it offers no principled account of how the best examples are chosen.

[5] Instead of shifting all votes to the value specified by the adjective, the votes might be distributed so that the more similar a value to that specified by the adjective the more votes that value receives. For example, in composing *brown apple* most votes may shift to brown, but some may remain on red because that value is similar to brown. While this seems plausible, for purposes of simplicity we ignore this possibility in what follows.

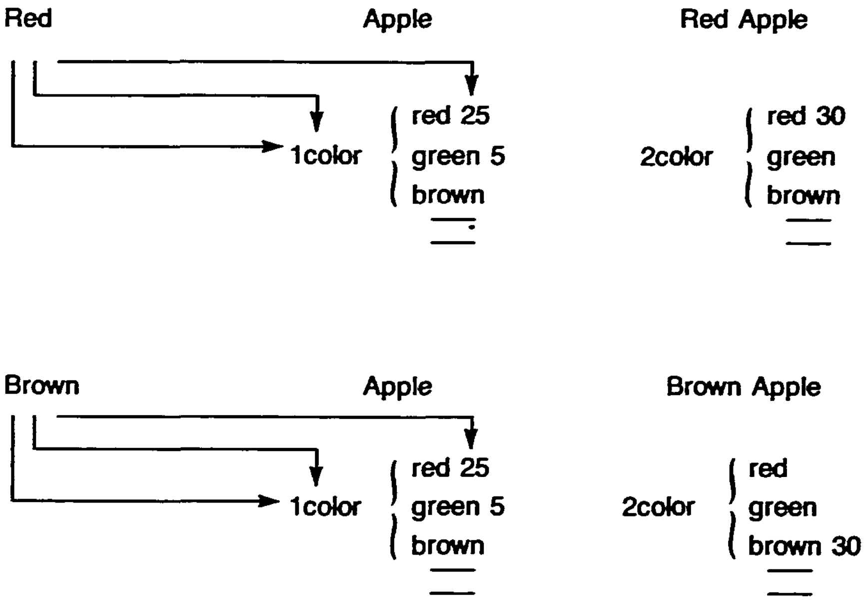

Figure 2. Illustration of three aspects of adjective modification.

than in *apple* (the diagnosticity boost), and that a typical *red apple* is redder than a typical *apple* (the salience change). The rationale for the salience change is obvious: The change from *apple* to *red apple* unequivocally signals a change in the color of typical instances. The rationale for the boost in diagnosticity is more subtle. The boost is likely mediated by a change in the perceived contrast class of the concept. As we change from *apple* to *red apple,* the contrast class may change from *oranges* to *green apple;* if so, then color is the only distinguishing attribute for the conjunction, and that is why its diagnosticity increases.

Figure 3 illustrates the implications of the above changes in salience and diagnosticity for typicality ratings with compatible conjunctions. The leftmost panel of the figure contains the prototype for *red apple.* The only differences between this representation and that of *apple* involve the color attribute. Now, all votes are on red, and the diagnosticity of color has increased by a factor of two (*two* being an arbitrary choice on our part). The effects of these differences for typicality are illustrated in the remaining panels of Figure 3. There we have repeated the representations for our red and brown apples, and computed the similarity for each of these objects in the conjunction. When these similarity scores are compared with those in Figure 1, the results are that (1) The red apple is more similar to *red apple* than it is to *apple,* while (2) the brown apple is less similar to *red apple* than it is to *apple.* Result (1) is the conjunction effect, result (2) is the reverse conjunction effect.

Red Apple (RA)		I_1		I_2	
2 color	red 30	color	red 30	color	red
	green		green		green
	brown		brown		brown 30
0.50 shape	round 15	shape	round 20	shape	round 20
	square		square		square
	cylindrical 5		cylindrical		cylindrical
0.25 texture	smooth 25	texture	smooth 30	texture	smooth 30
	rough 5		rough		rough
	bumpy		bumpy		bumpy

$$\begin{aligned} \mathrm{Sim}\ (RA, I_1) &= 2\ (30\text{-}0\text{-}0) \\ &+ 0.50\ (15\text{-}5\text{-}5) \\ &+ 2.25\ (25\text{-}5\text{-}5) \\ &= 60 + 2.50 + 3.75 \\ &\approx 66 \end{aligned}$$

$$\begin{aligned} \mathrm{Sim}\ (RA, I_2) &= 2\ (0\text{-}30\text{-}30) \\ &+ 0.50\ (15\text{-}5\text{-}5) \\ &+ 0.25\ (25\text{-}5\text{-}5) \\ &= -120 + 2.50 + 3.75 \\ &\approx -114 \end{aligned}$$

Figure 3. Illustration of attribute-value representations of *red apple* and relevant instances; beneath each instance representation is the similarity between the instance and prototype.

Figure 4 illustrates a comparable analysis for incompatible conjunctions. The prototype for *brown apple* differs from that of *apple* only in that now all color votes are on brown and color has doubled in diagnosticity. The effect of these changes for typicality are shown in the remaining panels. Now, in comparison with our original computations in Figure 1, we find that the brown apple is more similar to *brown apple* than it is to *apple,* while the red apple is less similar to *brown apple* than it is to *apple.* Again we have reconstructed the conjunction effect and its reverse. Note further than the conjunction effect predicted in this case exceeds that in the previous case, which reconstructs the third effect described earlier. That is, the extent to which the brown apple is judged more typical of *brown apple* than *apple* (66 – (– 54) = 120) is greater than the extent to which the red apple is judged more typical of *red apple* than *apple* (66 – 21 = 45).

Further Implications. The selective modification model incorporates the three general aspects of prototypes discussed earlier, and accounts in detail for the three conjunction effects. The model also has two additional implications that deserve to be spelled out.

One implication concerns a subtle prediction about conjunction effects. The typicality of, say, a brown apple in *brown apple* should be roughly equal to that of a red apple in *red apple.* That is, with regard to determining typicality in a conjunction, what matters is the attribute not the value. This prediction follows because the total number of brown votes in *brown apple* or red votes in *red apple* is simply the total number of color votes that *apple* has. The data reported in Smith and Osherson (1984) support the prediction.

The second implication of the model is our assumption that during modification a simple adjective, such as *red* or *long,* selectively influences a single attribute of the noun representation. All things considered, this "selective influence" assumption may be too strong. There are relations between attributes, and some of these relations may be part of a prototype (see, e.g., Malt & Smith, 1984). Among apples, for example, there are relations between color, shape, and sweetness—compared with a red apple, a brownish one is usually more shriveled and less sweet. Hence, *brown* applied to *apple* may change more than just the color attribute. Medin and Shoben (1988) have recently reported findings that appear to demonstrate such additional changes.[6]

[6] Medin and Shoben (1988) showed, for example, that a small spoon is judged more typical than a large spoon of *spoon,* while the reverse obtains for *wooden spoon.* This suggests that the adjective *wooden* has affected the size attribute. This argument seems plausible, but an alternative account should also be kept in mind. One's knowledge about the size of wooden spoons may have little to do with composition processes, but instead reflects prior experience with known instances of *wooden spoon.* That is, *wooden spoon* is not an unfamiliar concept, hence experience with its instances may figure in typicality judgments (Hampton, 1987). This consideration indicates that research on conceptual composition has not paid sufficient attention to the familiarity of the concepts involved, a criticism that applies to parts of the present paper as well.

Brown Apple (BA)	I_1	I_2
2 color { red, green, brown 30 }	color { red 30, green, brown }	color { red, green, brown 30 }
0.50 shape { round 15, square, cylindrical 5 }	shape { round 20, square, cylindrical }	shape { round 20, square, cylindrical }
0.25 texture { smooth 25, rough 5, bumpy }	texture { smooth 30, rough, bumpy }	texture { smooth 30, rough, bumpy }
	$\text{Sim}(BA, I_1) = 2\,(0-30-30) + 0.50\,(15-5-5) + 0.25\,(25-5-5) = -120 + 2.50 + 3.75 \approx -114$	$\text{Sim}(BA, I_2) = 2\,(30-0-0) + 0.50\,(15-5-5) + 0.25\,(25-5-5) = 60 + 2.50 + 3.75 \approx 66$

Figure 4. Illustration of attribute-value representations of *brown apple* and relevant instances; beneath each instance representation is the similarity between the instance and prototype.

These additional changes, though, may not be as profound as the primary change (e.g., in *brown apple,* only a few of the shape or taste votes may be shifted, and there may be only small increases in the diagnosticities of these attributes). Also, the additional changes may be made subsequent to the primary one so that the additional changes are in effect part of another process. That is, initial judgments about membership in conjunctive concepts may conform with our selective-influence assumption. This possibility is supported by the evidence that we present in subsequent sections where we demonstrate how well our model predicts typicality ratings for conjunctions. The issue of selective modification is sufficiently controversial, though, that we return to it at the end of the paper.

A FIRST TEST OF THE MODEL: STUDY 1

The preceding account was highly illustrative. We used but a single noun concept, had no basis for the attributes, values, or votes that were included in the noun's prototype, and no rationale for the diagnosticity weights nor for how much they were boosted by modification. Study 1 remedied these deficiencies. In part 1, subjects listed properties of various instances of the concepts *fruit* and *vegetable;* these listings were used to determine the attributes, values, votes, and diagnosticities for the instances and concepts. With such representations in hand, we could use our modification procedures to produce representations for conjunctions, and then employ the contrast rule to predict the typicalities of the instances in the simple concepts and conjunctions. These predicted typicalities were compared with actual ratings obtained in part 2 of the study.

Part 1

Method. The instances used were basic-level concepts, such as *apple, peach, carrot,* and *onion,* rather than specific objects like those employed by Smith and Osherson (1984) and used to illustrate the model in the previous section. This change in level allowed us to dispense with pictures or models of instances, in favor of one-word descriptions.

Thirty subjects listed properties for *fruit* instances and 30 listed properties for *vegetable* instances. All subjects were Harvard-Radcliffe undergraduates who were paid for their participation. Each subject was given a booklet that consisted of a page of instructions followed by 15 test pages, each of the latter containing the name of one *fruit* or one *vegetable* instance. The instructions were essentially the same as those used by Rosch and Mervis (1975). They informed subjects that, for each instance, they were to write down all its properties they could think of, and that they had 90 seconds to do this. The order of instances was randomly determined for each subject.

Results. First, any property that was mentioned by only one subject was eliminated. Then, for each instance, two raters (two of the authors) inspected the resulting set of listed properties for ones that intuitively seemed to be values of the same attribute. Wherever there was any disagreement between the raters, the properties were dropped from consideration. Only 12% of the listed properties was eliminated by these criteria.

For each attribute selected, the number of mentions of each property, or value, was taken as a measure of its number of votes. Part of this coding for the instance *carrot* is presented in Figure 5. While the figure shows only 4 attributes, there were in fact a total of 15 attributes for *carrot*. Most importantly, the attributes that emerged for this instance also appeared with other *vegetable* and *fruit* instances, as only 26 different attributes emerged across all 30 instances. This communality makes it reasonable to determine the attribute-value representation for the concepts *fruit* and *vegetable* by averaging over all relevant instances on each attribute, as illustrated on the right-hand side of Figure 5. Such averaging is in keeping with the idea of a prototype as a measure of central tendency. Our prototype for *fruit*, then, consisted of 25 attributes with an average of 7.16 values per attribute, while our prototype for *vegetable* contained 25 attributes with an average of 7.28 values per attribute. All of these attributes are listed in Table 1, along with the total number of votes cast for each, separately for *fruit* and *vegetable*. (Other ways of determining attribute-value representation for *fruit* and *Vegetable* are discussed in connection with Experiment 4.)

To estimate the diagnosticity weights for the concepts' attributes, we assumed that the diagnosticity of an attribute would be largely a matter of how useful it was for discriminating between fruits and vegetables. Accordingly, for each attribute we formed an n by 2 table; the two columns designated *fruit* and *vegetable*, the n rows designated every value of the attribute listed for any fruit or vegetable, and the cell entries were the numbers of votes for that value of *fruit* or *vegetable*. We then calculated the statistic v, a close cousin of chi-square (specifically, the square root of chi-square divided by the total number of votes in the table). This statistic varies between 0 and 1 and indicates the extent to which the values of the attribute are associated with *fruit* but not *vegetable*, or vice versa. We took the value of v as an estimate of the attribute's diagnosticity. These diagnosticity weights are given in the last column of Table 1.[7]

[7] The general formula for v is:

$$v = [X^2/N \min(I-1), (J-1)]^{1/2},$$

where X^2 is chi square, N is the total number of observations in the table, I is the number of rows, and J is the number of columns (Bishop, Feinberg, & Holland, 1975, Chap. 11). Because in our case there are always just two columns, the above formula reduces to:

$$v = [X^2/N]^{1/2}$$

Unlike X^2, v is not correlated with N.

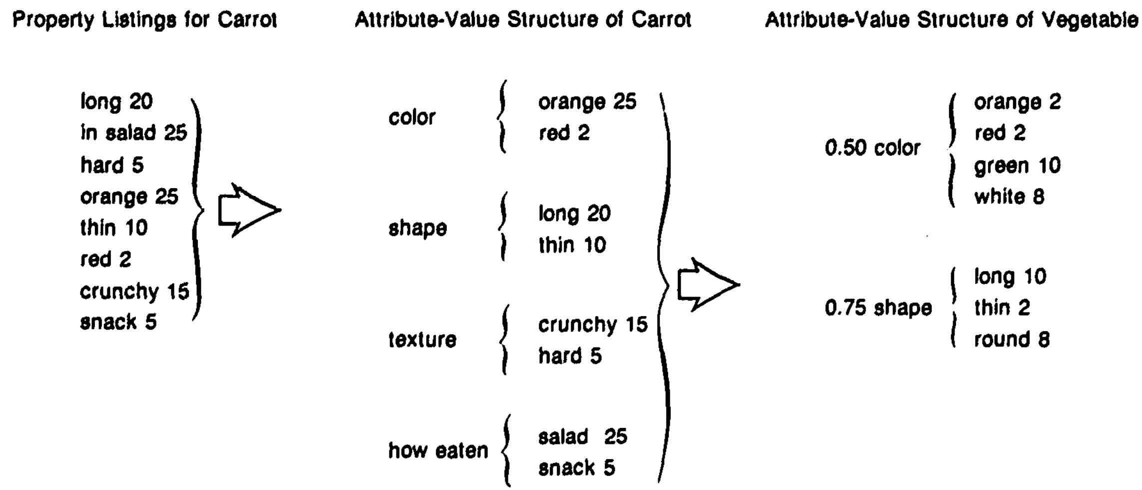

Figure 5. Illustration of coding of listed properties into attribute-value structures for instance *carrot* and concept *vegetable*.

TABLE 1
Attributes of *Fruit* and *Vegetable*, along with Total Votes and Diagnosticities (Study 1)

Attribute	Total Votes for Fruit	Total Votes for Vegetable	Diagnosticity v
Outside color	503	462	.44
Outside texture	261	206	.61
Taste	252	167	.69
How eaten	238	397	.84
How grown	203	90	.82
Seeds	191	48	.13
Shape	157	158	.70
Juiciness	146	34	.61
Inside color	119	30	.37
Size	109	34	.12
Pit	55	0	.43
Inside texture	51	44	.52
Original identity	44	18	.55
Where grown	41	19	.96
Skin	39	25	.50
Stem	38	15	.24
Varieties	37	83	.71
Side-effects	27	47	.95
When grown	22	2	.34
Container	21	12	.53
Nutritional value	18	138	.83
Leaf	14	68	.67
Favorite consumer	9	47	.90
Nonfood uses	7	10	.79
Cost	4	2	1.00
Smell	0	36	.41

We now have sufficient information to predict typicalities for each of the 15 instances in the simple concepts *fruit* and *vegetable*, as well as in certain adjective–noun conjunctions. With regard to the simple concepts, earlier we presented Equation (2), which, computes similarity on an attribute-by-attribute basis. However, the computing equation that we actually used iterates not only over attributes, but over values of an attribute as well:

$$\text{Sim}\ (\text{P,I}) = \sum_i v_i \sum_j [a \min(n_{ij}(\text{P}), n_{ij}(\text{I})) - b(n_{ij}(\text{P}) \dot{-} n_{ij}(\text{I})) - c(n_{ij}(\text{I}) \dot{-} n_{ij}(\text{P}))]. \tag{3}$$

Again i indexes the attributes, and now v_i is the diagnosticity of attribute i. Also, j indexes the values on an attribute, and $n_{ij}(\cdot)$ is the number of votes on value j of attribute i. The expression in brackets denotes a contrast between common and distinctive features for each value of each attribute, where the dot over the minus sign indicates the difference must be positive (i.e., $n_{ij}(\text{P}) \dot{-} n_{ij}(\text{I}) = n_{ij}(\text{P}) - n_{ij}(\text{I})$ if $n_{ij}(\text{P}) > n_{ij}(\text{I})$, and 0 otherwise). To illustrate, suppose the prototype had 20 red votes and the instance 15 red votes. Then the number of common features would be 15 (the minimum of the 2 vote counts), the number of features distinct to the prototype would be 5, the number distinct to the instance would be 0, these three numbers would be multiplied by *a, b,* and *c,* respectively, and the outcome of the contrast would be multiplied by the diagnosticity of color.

To predict typicalities in conjunctions, we have to augment our computing equation in the following way:

$$\mathrm{Sim}\ (\mathrm{P},\mathrm{I}) = \sum_i e_i v_i \sum_j [a \min(n_{ij}^*(\mathrm{P}), n_{ij}(\mathrm{I})) - b(n_{ij}^*(\mathrm{P}) \dot{-} n_{ij}(\mathrm{I})) - c(n_{ij}(\mathrm{I}) \dot{-} n_{ij}^*(\mathrm{P}))] \tag{4}$$

Equation (4) differs from its predecessor in two respects. First, for the prototype, the number of votes on value j of attribute i is now n_{ij}^* (P), which is defined as follows:

$$n_{ij}^*(\mathrm{P}) = \begin{cases} \sum_{ij} n_{ij}(\mathrm{P}), & \text{if the adjective in the conjunction matches value } j. \\ 0, & \text{if the adjective in the conjunction is a value of attribute } i \text{ other than } j. \\ n_{ij}(\mathrm{P}), & \text{otherwise.} \end{cases}$$

For example, if the adjective in the conjunction is "red," then the number of red votes in the prototype would be the sum of all color votes, the number of green votes in the prototype would be zero, and the number of votes on the value of any other attribute in the prototype would be as usual. The other novelty in Equation (4) is the addition of e_i, which multiples v_i, the diagnosticity of attribute i; e_i is defined as follows:

$$e_i = \begin{cases} d, & \text{if the adjective in the conjunction encodes } i. \\ 1, & \text{otherwise.} \end{cases}$$

Thus, the diagnosticity of an attribute is boosted by a factor of d ($d>1$) if that attribute is encoded by the adjective in the conjunction. Note that d, the "booster," is the only free parameter in (4) apart from the contrast weights *a, b,* and c. Equations (3) and (4) were used to predict the typicalities of the relevant instances in the simple concepts *fruit* and *vegetable* and in the eight conjunctions described below. (To estimate parameters, we used the program STEPIT [Chandler, 1969], and maximized average correlations between predicted and obtained ratings.)

Part 2

Method. Thirty subjects, drawn from the same population as in the previous part of the study, rated the typicality of instances in 10 different concepts. These concepts included the simple concepts *fruit* and *vegetable* and the conjunctions formed by combining them with *red, white, round,* and *long*. Every subject was given a booklet which included instructions and test pages; each of the latter contained the name of one of the 10 concepts followed by a list of 15 relevant instances. Subjects were instructed to rate each instance "...for how good an example it is of the category." The ratings were made on an 11-point scale, where *10* means the instance "is about as good an example as you can get of your idea or image of what the category is and *0* means you think the item does not fit at all with your idea or image of the category."

The same 15 instances were used with *fruit* and with the 4 conjunctions involving *fruit;* all of these were technically fruits though two—*tomato* and *pickle*—were sometimes classified as vegetables. There were also 15 instances that were used with *vegetable* and the 4 conjunctions involving *vegetable;* 13 of them were technically vegetables while the remaining 2 were *tomato* and *pickle.* (All *fruit* and *vegetable* instances are listed in Table 2). For a given subject, the instances were listed in the same order when they appeared with different concepts (to ease the subject's rating task), but the order of instances and concepts varied randomly across subjects.

Results: Evaluation of the Model. For each of the 10 concepts, we determined the average typicality ratings for the 15 instances. These ratings are presented in Table 2. Then, for each concept, we correlated the obtained ratings with those predicted by the model. These correlations are presented in the last row of Table 2. The model does a reasonable job with most of the concepts—the average *r* is about .70—and particularly with the *vegetable* concepts where the average *r* for the four conjunctions is .88.[8]

However, the model appears to fail with *white fruit* and *long fruit.* Further inspection of the obtained ratings for *white fruit* and *long fruit,* however, suggests that the problem is not in the model but in our selection of instances. First, and most important, our fruit instances showed little variation with respect to whiteness and length, which greatly limits the possible correlations. With regard to the obtained ratings for *white fruit* and *long fruit,* 11 of the 15 instances were rated less than 2.0 on our 11-point scale (see Table 2). Another problem with the instances is specific to *white fruit.* The three instances that subjects rated most typical of *white fruit* were *coconut, apple,* and *pear;* all three of these objects are white on the inside but not on the outside, yet it was outside color that subjects were instructed to

[8] There may have been some contribution of an instance's frequency to its predicted typicality (Nosofsky, personal communication, April, 1987). Because subjects tended to list fewer properties for very infrequent instances such as *pomegranate* and *avocado* than for other instances, more properties of frequent than infrequent instances appeared in the prototypes of *fruit* and *vegetable*. Consequently, infrequent instances may have ended up being deemed less similar to their prototype.

TABLE 2
Average Typicality Ratings for 10 Concepts
(Study 1)

	Vegetable	Red Vegetable	White Vegetable	Round Vegetable	Long Vegetable
String beans	8.70	.54	1.13	.54	8.03
Carrot	8.44	2.27	.50	.64	8.17
Spinach	8.44	.33	.30	.40	2.30
Lettuce	7.64	1.10	1.84	5.77	1.37
Squash	7.44	1.34	2.17	3.87	5.30
Cauliflower	7.37	.17	9.04	4.43	1.34
Potato	7.20	1.83	6.77	6.37	2.27
Lima bean	6.87	.84	2.80	2.80	1.97
Turnip	6.60	3.17	5.37	4.47	2.64
Onion	6.40	2.80	7.53	7.87	1.10
Tomato	6.40	8.50	.30	7.70	.60
Mushroom	6.07	.40	6.17	2.70	.67
Pickle	3.93	.60	1.37	1.60	4.93
Seaweed	3.10	1.04	.93	.17	3.50
Garlic	2.74	1.07	5.47	3.00	.87
Correlation between Observed and Predicted Typicalities	.40	.94	.84	.87	.85

	Fruit	Red Fruit	White Fruit	Round Fruit	Long Fruit
Apple	9.77	9.34	4.23	9.27	.67
Peach	9.23	3.33	.93	8.70	.67
Strawberry	9.14	8.57	.74	4.50	.97
Pear	9.04	1.37	3.84	3.67	2.94
Grape	8.40	3.60	3.23	8.14	1.04
Blueberry	8.34	.47	.53	8.60	.33
Watermelon	7.27	5.44	.80	4.57	5.44
Pomegranate	6.94	6.07	.70	7.00	1.37
Lemon	6.70	.30	1.20	4.93	1.64
Fig	4.73	1.07	.70	3.37	1.17
Raisin	4.64	.77	.87	3.30	1.07
Coconut	4.10	.24	5.37	6.90	1.14
Avocado	3.37	.44	.54	3.77	2.24
Tomato	3.33	6.37	.34	7.04	.50
Pickle	1.37	.27	.97	.73	2.17
Correlation between Observed and Predicted Typicalities	.75	.91	.35	.92	.11

rate and that our model considered in its predictions. The only other low correlation in Table 2 is for the simple concept *vegetable*. Again, the problem can be traced to a lack of variation in the instances: 12 of the 15 *vegetable* instances had typicality ratings between 6.07 and 8.70. The *fruit* instances showed substantially more variation.

The parameter values obtained in fitting the model were estimated separately for *fruit* and *vegetable* concepts. Three of the four parameters were intrinsic to the contrast model: a, the weight given to common features; b, the weight of features distinct to the concept; and c, the weight of features distinct to the instance. For *vegetable* concepts, $a = .88$, $b = .50$, and $c = .20$; for *fruit* concepts, $a = 1.84$, $b = .50$, and $c = .20$. The parameters are similar for the two kinds of concepts, and the ordering of the parameters is in agreement with prior results (Gati & Tversky, 1984; Tversky, 1977; Tversky & Gati, 1982). The fourth parameter, d, measures how much the diagnosticity of an attribute is boosted by modification. For *vegetable* concepts, $d = 8.36$, for *fruit* concepts, $d = 4.21$. Clearly, the attribute encoded by the adjective plays a major role.

Results: Conjunction Effects. Lastly, we want to examine the data for conjunction effects. Table 3 presents the relevant data and predictions for *vegetable* concepts; we ignore the data for *fruit* concepts in light of the problems with *fruit* instances noted above. The top half of the table contains the data for instances that were "good" members of their corresponding conjunctions; for *white vegetable*, for example, just those instances that had at least five color votes on white. The bottom half of Table 3 contains the data for instances that were "poor" members of the conjunction; for *white vegetable*, those instances that had zero white votes. The data are presented separately for each conjunction. The obtained data replicate all our previ-

TABLE 3
Obtained and Predicted (in parentheses) Typicality Ratings for Nouns and Conjunctions (Study 1)

Adjectives	Vegetable Rating		Adjective Vegetable Rating	
Good Members				
Red	6.40	(5.63)	8.50	(11.89)
White	5.84	(5.54)	6.72	(6.65)
Round	6.67	(5.79)	7.23	(7.99)
Long	7.02	(5.48)	7.04	(7.88)
Poor Members				
Red	6.31	(5.35)	0.86	(0.22)
White	6.71	(5.38)	1.68	(0.13)
Round	6.47	(5.35)	1.84	(2.56)
Long	6.65	(5.50)	1.86	(3.30)

ous results: For good members, an instance is judged more typical of the conjunction than of the noun constituent; the conjunctions effect is greater for the incompatible conjunction (*red vegetable*) than for the compatible ones (the other three conjunctions); and for poor members, reverse-conjunction effects occur as an instance is judged less typical of the conjunction than of the noun constituent. Moreover, the predicted data, presented in parentheses, show exactly the same effects.[9]

SUBSEQUENT TESTS OF THE MODEL: STUDIES 2 AND 3

In study 1, subjects rated the typicality of instances with respect to nouns and adjective-noun conjunctions. It is of interest, however, to also determine the typicality of the instances vis-à-vis adjective constituents (e.g., *red* and *round*), and we did this in studies 2 and 3. Obtaining adjective ratings (along with noun and conjunction ratings) allowed us to evaluate two important assumptions that were left implicit in the previous study. One is that adjectives, such as *red* and *round,* are represented by only a single attribute; if this is the case, we should be able to fit the model to the adjective ratings by assuming that *red,* say, contains only the attribute of color with all votes being on the value red. The second assumption of interest is that subjects base their typicality ratings for a conjunction (e.g., the typicality of *apple* in *red fruit*) on all attributes of the noun concept (e.g., color, shape, and texture), not just on the attribute singled out by the adjective (color). To check this, for each set of instances, we correlated the ratings in the conjunction separately with those in the noun and with those in the adjective, to determine whether each constituent was contributing to the conjunction's ratings.

The only differences between studies 2 and 3 is that in the latter study some noninstances of *vegetable* (*fruit*) were included among the items paired with the *vegetable* (*fruit*) concepts. For example, subjects had to rate the typicality of apple in *vegetable, red,* and *red vegetable.* The purpose of this change was to make the variability of the items paired with the noun concept more comparable with the variability of the items paired with the adjective concept; we need such comparability to compare the correlation between noun and conjunction ratings with that between adjective and conjunction ratings. Also, because the *fruit* instances in study 1 hardly varied with respect to *white* and *long,* in studies 2 and 3 we used the same instances, but only the conjunctions *red fruit, round fruit, red vegetable,* and *round vegetable.* Because studies 2 and 3 are similar and produced comparable results, we treat them together in what follows.

---

[9] The predicted ratings in Table 3 have been rescaled to have the same mean and standard deviation as the obtained ratings in Table 3. This rescaling preserves the correlations reported in Table 2 because in fitting our model we maximized the average correlations between predicted and obtained ratings rather than minimized predicted-observed deviations.

Method

Study 2. The subjects were 30 students drawn from the same population as in the previous study. All subjects rated the typicality of instances in three types of concepts: noun concepts, including *fruit* and *vegetable,* adjective concepts, including *red* and *round,* and conjunctions, including *red fruit, round fruit, red vegetable,* and *round vegetable.* The 15 instances used with the *fruit* concepts were the same as those used in study 1, and so were the 15 instances used with *vegetable* concepts. In addition, both kinds of instances were used with *red* and *round.* Subjects thus made 10 sets of ratings. The order of the instances within each of these sets was randomized anew for each subject.

Again subjects worked with booklets, where the first pages gave instructions and subsequent pages contained the names of concepts followed by lists of 15 relevant instances. Concepts were blocked, so that all concepts of one type—noun, adjective, or conjunction—appeared on consecutive pages. All possible orders of the concept types were used. Other procedural details —rating scale, general content of instructions—were the same as in the previous study.

Study 3. The subjects were 30 students drawn from the usual population. The only change from the preceding study was in the nature of the *fruit* and *vegetable* "instances." The set of *fruit* "instances" now included our usual 15 fruits, plus 8 vegetables. Similarly, the set of *vegetable* "instances" included our usual 15 vegetables, plus 8 fruits.

Results: Studies 2 and 3

Evaluation of the Model. For each of the 10 concepts, we determined the average typicality ratings for the 15 critical instances. Then, for each concept, we correlated the obtained ratings with those predicted by the model. Applying our model to the data for adjective-noun conjunctions and their noun constituents involved nothing new. To apply the model to the data for adjective concepts, we assumed that the representation for an adjective consists of a single attribute, with all its votes on the value named by the adjective. Consider *red:* The only attribute is color, and all 30 votes are on the value red (*30* because that is the maximum number of votes that any value had in our noun representations).

The results of these fits are in Table 4. In study 2, there is a relatively poor fit for the simple concept *vegetable* (just as we found in study 1), but now all other correlations between observed and predicted typicalities are very high. The average correlation for the adjective-noun conjunctions is a resounding .94. The results for study 3 lend further credence to the model, as the average correlation for the adjective-noun conjunctions is .81. Also, the correlations for the simple concepts, *vegetable* and *fruit,* have increased,

TABLE 4
Correlations between Obtained and Predicted Typicalities
(Studies 2 and 3)

Concepts	Study 2	Study 3
Vegetable	.37	.80
Red	.97	.98
Round	.89	.88
Red Vegetables	.96	.70
Round Vegetables	.88	.74
Fruit	.70	.87
Red	.98	.97
Round	.96	.92
Red Fruit	.97	.91
Round Fruit	.95	.88

presumably because of the increase in variability that resulted from including noninstances as well as instances of the concept.

In Study 2 the three contrast-rule parameters were as follows: for *fruit* concepts, $a = .80$, $b = .50$, $c = .20$; for *vegetable* concepts, $a = .92$, $b = .50$, $c = .20$. In Study 3, for *fruit* concepts, $a = 2.01$, $b = .50$, $c = .20$; for *vegetable concepts*, $a = 1.44$, $b = .56$, $c = .20$. The values are similar in the two studies and close to those obtained in the previous study. The remaining parameter is the booster, d. In study 2, d was 8.06 for *fruit* concepts and 8.46 for *vegetable* concepts; in study 3, d was 4.24 for *fruit* concepts and 2.70 for *vegetable* concepts. Again, these values are comparable with those obtained in study 1.

Correlations between Conjunctions and Constituents. The above findings support our claim that the problematical results of study 1 were due to a lack of variation in the instances, and that the adjectives *red* and *round* can be represented by a single attribute. What remains to be checked is our assumption that subjects base their typicality ratings for a conjunction on all attributes of a noun concept. We checked this for each conjunction by determining the correlation (across instances) between ratings in a conjunction and ratings in a particular constituent (with any contribution of the other constituent partialed out). These partial correlations are presented in Table 5. In study 2, ratings in a conjunction are more correlated with the adjective than the noun constituent, but the correlations with the noun are substantial and in two cases significant. In study 3, where the variability of instances in the noun concepts is comparable with that in the adjective concepts, all correlations are significant, and ratings in conjunctions are almost as correlated with the noun as the adjective constituents. Clearly, then, subjects ratings for conjunctions consider more than just the attribute singled out by the adjective.

TABLE 5
Partial Correlations between Typicality Ratings in Conjunctions and Typicality Ratings in the Constituents (Studies 2 and 3)

Concepts	Study 2	Study 3
Red Fruit—Red	.99	.93
Red Fruit—Fruit	.52	.65
Round Fruit—Round	.97	.78
Round Fruit—Fruit	.80	.87
Red Vegetable—Red	.99	.72
Red Vegetable—Vegetable	.35	.54
Round Vegetable—Round	.99	.72
Round Vegetable—Vegetable	.31	.77

Note. For $p<.05$, $r=.50$; for $p<.01$, $r=.62$

EXTENSION OF THE MODEL TO ADVERBS

In this section we extend the model to conjunctions that include the adverbs *very, slightly,* and *non.* Again the fundamental idea is that of modification. Now, however, the modifiers of interest are adverbs, and the frame that is altered is itself the outcome of a modification process (namely, the composite prototype that results when an adjective modifies a noun).

Single Adverbs

"Hedges" are a large class of adverbial modifiers whose major function seems to be that of qualifying predicates and which have previously figured in analyses of concept membership (e.g., Lakoff, 1973; Smith et al., 1974). As Lakoff (1973) notes, one subset of hedges includes terms like *very, slightly,* and *non,* where these terms seem to intensify aspects of the concepts or prototypes on which they operate (see also Clark & Clark, 1979; Cliff, 1959; Zadeh, 1971). While the principles by which such "intensifiers" operate are simpler than those characterizing most hedges (such as *technically speaking* or *loosely speaking*), intensifiers provide a useful starting point for extending a model of composite prototypes to include adverbs.

To illustrate how *very, slightly,* and *non* work, in *very red fruit, very* appears to augment the redness in *red fruit;* while in *slightly red fruit* or *nonred fruit,* the adverbs diminish the redness in *red fruit.* To capture these intuitions, we assume that:

1. *very* augments the modified value in a conjunction (e.g., the red in *red fruit*) by multiplying the votes on that value by some scalar greater than 1;
2. *slightly* diminishes the modified value in a conjunction by multiplying

the votes on that value by some scalar between 0 and 1, thereby ensuring that there is a decrease in votes on the value but still some votes left; and
3. *non* diminishes the modified value in a conjunction by multiplying the votes on that value by a scalar less than or equal to 0, thereby ensuring that there are no (positive) votes left on the value.

Possible scalars for the three adverbs are:.

very: k_v, where $k_v > 1$
slightly: $1 - k_s$, where $o < k_s < 1$
not (*non*): $1 - k_n$, where $k_n \geq 1$

Our reasons for using the format $1 - k_s$ and $1 - k_n$ will become apparent when we describe more complicated adverb combinations.

To illustrate the above scheme, suppose that the number of red votes in *red fruit* is 10. When *very* is applied to *red fruit,* the number of red votes is increased by $10\,k_v - 10$. When *slightly* is applied to *red fruit,* the number of red votes is decreased by $10\,k_s$ (since k_s must be less than 1, the resulting number of red votes will always be greater than 0). When *non* is applied to *red fruit* the resulting number of red votes is decreased by $10k_n$ (since k_n can never be less than 1, there can never be any red votes left). Thus *very red fruit* has more red than *red fruit; slightly red fruit* has less red than *red fruit* but more than zero red; and *nonred fruit* has no red at all. All of this is compatible with our intuitions.

A further comment is in order about our treatment of *non.* It might seem plausible that k_n should always be 1, thereby ensuring that $1 - k_n$ is always 0. As we will see, though, our results indicate that k_n exceeds 1, which means that $1 - k_n$ has a negative value. This in turn results in there being negative votes. Negative votes are to be treated as follows in computing similarity. Given a concept with negative votes on value j of attribute i, and an instance with some votes on value j' of attribute i, then the votes on j' can be converted to negative votes on j as long as $j' \neq j$. We can illustrate with the concept *nonred fruit; blueberry's* blue votes can be converted to negative red votes, thereby increasing its common color features with the concept; and the more salient the color of a particular nonred fruit, the more typical it will be of the concept *nonred fruit.*

Dual Adverbs

We can extend our model one step further by considering conjunctions that involve two adverbs such as *very nonred fruit* and *slightly nonround vegetable.* Combining adverbs comes down to combining scalars for single adverbs. One proposal for doing this is as follows:

slightly not: $1 - [(1 - k_s)k_n]$
very not: $1 - (k_v\, k_n)$
very slightly: $1 - (k_v\, k_s)$

In *slightly not, not* applies first to yield $1-k_n$, and then *slightly* operates directly on the value of k_n, resulting in $1-[(1-k_s)\ k_n]$. Hence the first-mentioned adverb has smaller scope than the second. Similarly, in *very not, not* applies first yielding $1-k_n$, and the *very* augments the value of k_n, resulting in $1-(k_v\ k_n)$. In *very slightly, slightly* applies first yielding $1-k_s$, and then *very* augments the value of k_s, resulting in $1-(k_v\ k_s)$. To illustrate, *slightly nonred fruit* has more red than *nonred fruit* because we have diminished the negation; in contrast, *very nonred fruit* has even less red than *nonred fruit* as we have augmented the negation, and *very slightly red fruit* has less red than *slightly red fruit* because we have augmented *slightly*. All of this seems in line with our intuitions.

There are other schemes for combining scalars that are consistent with out assumptions about *very, slightly,* and *non*. One obvious possibility is to simply multiply scalars; for example, the scalar for *slightly non* would be $(1-k_s)\ (1-k_n)$. It turns out, though, that this proposal does not do as well at predicting typicality ratings as the scheme we have proposed. Thus, part of the rationale for our proposal is *post hoc*. Still, our scheme captures basic intuitions about the interpretations of adverbs in conjunctions and offers some interesting claims about combining adverbs (e.g., the first-mentioned adverb has smaller scope than the second-mentioned one), in addition to doing as reasonable job of predicting typicality ratings in complex concepts as other schemes we have tried. (For some related proposals from a fuzzy-set theory perspective, see Hersh & Carmazza, 1976; Lakoff, 1973; and Zadeh, 1971; 1972).

To predict the typicality of an instance in a complex conjunction such as *nonred fruit*, we used a simple extension of Equation (4):

$$\text{Sim}\ (\text{P},\text{I}) = \sum_i e_i v_i \sum_j [a \min(A_i n_{ij}^*(\text{P}), n_{ij}(\text{I})) - b(A_i n_{ij}^*(\text{P}) \dot{-} n_{ij}(\text{I})) - c(n_{ij}(\text{I}) \dot{-} A_i n_{ij}^*(\text{P}))] \quad (5)$$

Now we multiply the number of votes on the modified value, $n_{ij}^*(\text{P})$, by the scalar associated with the adverb, A_i, where $A_i = k_v$ if the adverb is *very*, $A_i = 1-k_s$ if the adverb is *slightly*, and so on. (Strictly speaking, Equation (5) is correct only when $A_i\ n_{ij}^*(\text{P})$ is nonnegative). Equation (5) involves seven parameters: four are the same as in Equation (4), namely, *a, b, c,* and *e;* the three new parameters are embedded in the A_i term and are the scalars k_v, k_s, and k_n.[10]

[10] Lakoff (1973) has argued that characterizing *very* and *slightly* solely by numerical values runs into trouble when the adjective to be modified is *similar*. Thus, in contrasting "Richard Nixon and Warren Harding are similar" and "Richard Nixon and Warren Harding are very similar," *very* seems to do more than just raise the degree of similarity. In particular, *very* seems to increase the number of attributes that are taken into consideration. However, our numerical proposals for *very* and *slightly* seem to work well with most other adjectives.

A TEST OF THE EXTENDED MODEL: STUDY 4

In study 4, we obtained typicality ratings for conjunctions involving the adverbs of interest, and then compared these ratings with those predicted by Equation (5). Study 4 was performed along with study 1 and it used the same *fruit* and *vegetable* instances. As a consequence, the ratings for concepts involving *white fruit* and *long fruit* again showed hardly any variability, which once more resulted in artifactually low correlations between observed and predicted typicalities. In view of this, we will focus on the ratings for concepts that involved *red* or *round.*

Method

The subjects were drawn from the same population as in previous studies, and were divided into three groups of 30 each. Each group rated the typicality of instances in 16 different complex conjunctions. For group 1, the 16 conjunctions were generated by taking the basic 8 conjunctions from study 1 (*red fruit, white fruit. . ., long vegetable*), and then modifying each one by *very* or *non;* for group 2, the 16 conjunctions were formed by modifying the basic 8 conjunctions by *slightly* or *slightly non;* and for group 3, the basic 8 conjunctions were modified by *very slightly* or *very non.* The 15 instances used with *fruit* concepts were the same as those in study 1, and similarly for the 15 instances used with *vegetable* concepts. Again subjects worked with booklets, where the first pages gave instructions and subsequent pages contained the names of concepts followed by a list of the 15 relevant instances. Other procedural details—rating scale, general content of instructions—were the same as in previous studies.

Results

A Preliminary Look at the Data. For each concept, we determined the average typicality rating for the 15 relevant instances. Before assessing how well these ratings agree with those predicted by the model, it is instructive to display a sample of the obtained ratings. In Table 6, the first column lists the 15 *fruit* instances ordered by their number of red votes, while the second, third, and fourth columns give the obtained typicality ratings of each instance in three conjunctions: *red fruit* (these data are from study 1), *nonred fruit* (data from study 4), and *very red fruit* (study 4). The ratings for instances in *red fruit* increase roughly monotonically with the number of red votes in the instance (*watermelon* is a clear outlier, presumably because subjects rated inside rather than outside color). More importantly for present purposes, the ordering of the instances in *nonred fruit* is essentially the reverse of that in *red fruit.* This reversal is in line with our assumption that *not* multiplies the votes on red by a scalar less than zero, for then instances that were typical of *red fruit* (because they "had a lot of red") are likely to

TABLE 6
Obtained Typicality Ratings for 3 Sample Conjunctions
(Studies 1 and 4)

Instances	Red Fruit	NonRed Fruit	Very Red Fruit
Tomato	6.37	.07	9.24
Strawberry	8.57	.34	9.40
Apple	9.37	1.03	9.14
Pomegranate	6.07	3.34	6.60
Grape	3.60	5.24	4.20
Peach	3.33	5.60	3.37
[a]Pear	1.37	8.03	1.70
Fig	1.07	7.07	1.17
Raisin	.77	7.14	1.67
Coconut	.24	8.64	.07
Avocado	.44	8.70	.47
Watermelon	5.44	2.84	6.90
Pickle	.27	7.77	.30
Lemon	.30	9.44	.27
Blueberry	.47	9.00	.37

[a] All instances from Pear through Blueberry had zero red votes.

become atypical of *nonred fruit* (because they have too much red), and vice versa. Turning to *very red fruit*, the ratings are again roughly monotonic with the number of red votes in the instance. This is compatible with our assumption that *very* multiplies the votes on *red* by a positive scalar, for then instances with more votes on the adjective will be more typical of *very red fruit*.

Evaluation of the Model. To assess the adequacy of our model quantitatively, again we correlated the obtained ratings with those predicted by the model. (In fitting the model we included the data from study 1, as studies 1 and 4 essentially constitute a single experiment.) A summary of the correlations is presented in Table 7. The results are broken down by: (1) *vegetable* or *fruit;* (2) the form of the concept—adjective-noun conjunctions (these data are from study 1), or adverb-adjective-noun conjunctions, or adverb-adverb-adjective-noun conjunctions; and (3) the specific adverbs involved —*very, slightly,* and so on.

For *vegetable* concepts, the model captures a good portion of the data, as almost all the correlations are above .60 and the overall average correlation is .70. However, the goodness of fit depends on the complexity of the concept: The correlation drops from .90 to .73 when an adverb is added to an adjective-noun conjunction, and drops further to .59 when a second adverb is added. Another thing to note is how correlations vary with the specific adverbs. For conjunctions with single adverbs, the model does best with

TABLE 7
Correlations between Obtained and Predicted Typicality Ratings, Separately for Different Kinds of *Vegetable* and *Fruit* Concepts (Studies 1 and 4)

VEGETABLE CONCEPTS				
	Red / Round	Vegetable	.90	(.87)
Very	Red / Round	Vegetable	.86	(.87)
Slightly	Red / Round	Vegetable	.63	(.68)
Non	Red / Round	Vegetable	.70	(.60)
Average for Adverb	Red / Round	Vegetable	.73	(.72)
Very Slightly	Red / Round	Vegetable	.44	(.47)
Slightly Non	Red / Round	Vegetable	.64	(.48)
Very Non	Red / Round	Vegetable	.70	(.60)
Average for Adverb-Adverb	Red / Round	Vegetable	.59	(.52)
FRUIT CONCEPTS				
	Red / Round	Fruit	.92	(.54)
Very	Red / Round	Fruit	.90	(.57)
Slightly	Red / Round	Fruit	.74	(.52)
Non	Red / Round	Fruit	.70	(.39)

continued

TABLE 7 (continued)

FRUIT CONCEPTS			
Average for Adverb	Red Fruit Round	.78	(.49)
Very Slightly	Red Fruit Round	.28	(.20)
Slightly Non	Red Fruit Round	.70	(.41)
Very Non	Red Fruit Round	.64	(.34)
Average for Adverb Adverb	Red Fruit Round	.54	(.32)

[a] The entries in parentheses give the correlation when all concepts are included, i.e., those involving *white* and *long* as well as those involving *red* and *round*. While the correlations are much reduced, particularly for *fruit*, even these tainted data show the usual effects of complexity and specific adverbs.

concepts that involve *very* ($r = .86$) and worst with those that involve *slightly* ($r = .63$); for concepts with dual adverbs, the model does best with concepts that involve *very non* ($r = .70$) and worst with those that involve *very slightly* ($r = .44$).

A comment is in order about the decrease in correlation with the increase in concept complexity. There are two factors that contribute to this "complexity" effect. First, the effect is partly a consequence of the effects of specific adverbs. The model does poorest with concepts involving *slightly,* and the average correlation for single-adverb concepts includes one case with *slightly* while the average for dual-adverb concepts includes two cases with *slightly.* Inspection of Table 7 indicates that if we ignore concepts involving *slightly,* the complexity effect is substantially reduced. The second factor contributing to the complexity effect is reliability—subjects were less reliable in their ratings for dual-adverb concepts than for single-adverb or adjective–noun conjunctions. Split-half reliabilities were .93 and .95 for adjective–noun and single-adverb conjunctions, respectively, versus .88 for dual-adverb concepts. Apparently, subjects had some difficulty composing concepts that involved more than one adverb—possibly because of ambiguities in the scope of the adverbs—and when their performance became less systematic the model of course faltered.

Turning now to the data for *fruit* concepts in study 4, the results mirror the preceding ones in most important respects (see Table 7). The model captures a reasonable amount of the data, with the overall average correlation again being .70. Once more correlations decrease as the concepts become increasingly complex, dropping from .92 to .78 when a single adverb is added to an adjective–noun conjunction, and dropping further to .54 when a second adverb is added. Again this complexity effect can be attributed to the effects of *slightly* (see Table 7) and to a breakdown in reliability; with regard to the latter, the split-half reliabilities were .96 and .94 for adjective–noun and single-adverb conjunctions, respectively, versus .80 for dual-adverb conjunctions. With regard to the effects of specific adverbs, once more the model does best with *very* for single-adverb concepts, and worst with *very slightly* for dual-adverb concepts.

In sum, the selective modification model does an excellent job of accounting for ratings in adjective–noun conjunctions (the average correlation for the data in Table 7 is .91), a good job of accounting for ratings in adverb–adjective–noun conjunctions (the average correlation is .76), and at best a moderate job of accounting for ratings in dual-adverb concepts (the average correlation is about .56). There is, however, a blatant trouble spot for the model, namely concepts involving *slightly*, as correlations between predicted and observed ratings are routinely lower for such concepts than for other conjunctions. We return to the *slightly* problem soon, after we have taken a look at the parameters obtained in fitting the model.

Parameter Estimates. The top half of Table 8 gives the parameters obtained when the model is fit to *vegetable* concepts, and the bottom half gives the parameters for *fruit* concepts. The different rows give the parameters for the different studies (to facilitate comparisons, we have repeated the parameters obtained in studies 1–3). The three parameters of the contrast model are very stable. In virtually every case, the weight given to common

TABLE 8
Parameter Values, Separately for *Vegetable* and *Fruit* Concepts

	Parameters						
Vegetable Concepts	a	b	c	d	k_n	k_s	k_v
Study 1	.88	.50	.20	8.36	—	—	—
Study 2	.92	.50	.20	8.46	—	—	—
Study 3	1.44	.56	.20	2.70	—	—	—
Study 4 (& 1)	1.71	.50	.20	39.35	1.66	.74	1.04
Fruit Concepts							
Study 1	1.84	.50	.20	4.21	—	—	—
Study 2	.80	.50	.20	8.06	—	—	—
Study 3	2.01	.50	.20	4.24	—	—	—
Study 4 (& 1)	.41	.50	.20	12.92	1.59	.56	3.63

features, *a,* exceeds that given to features distinct to the concept, *b,* which in turn exceeds that given to features distinct to the instance, *c.* The booster parameter is somewhat more variable, *d* being quite low for *vegetable* concepts in study 3 and extremely large for *vegetable* concepts study 4. The remaining three parameters pertain to the adverbs, and generally these values are stable and reasonable. Thus k_n averages about 1.60. This results in a scalar for *not* of roughly $-.60$, which eliminates all positive votes on the modified value of a conjunction. Similarly, k_s averages about .65. This results in a scalar of about .35, which substantially reduces the votes on the modified value of a conjunction but does leave some votes there. The remaining parameter, k_v, is a bit more variable being 1.04 for *vegetables* and 3.63 for *fruits.* This reflects the fact that there was little difference between, say, concepts involving *very red* and those involving *red* or *vegetable,* but a noticeable difference between the two kinds of concepts for *fruit.*

Note that the parameters common to all four studies—*a, b, c,* and *d*—provide support for two fundamental assumptions of the model. The fact that the similarity parameters *b* and *c* were always greater than zero attests to the contrast rule's assumption that similarity depends on distinctive features, not just common ones. And the fact that the booster parameter *d* was always substantial supports the claim that an adjective modifier boosts the diagnosticity of the associated attribute in the noun.[11]

In addition to the above, there are many "hidden" parameters in our model—namely, the attributes, values, and votes of the instances and concepts that were estimated by empirical means. For these hidden parameters, the critical question is whether their values depend on the methods used to obtain them. In particular, are our attribute-value representations for *fruit* and *vegetable*—which enter into every single prediction of the model—biased by the fact that they were determined by averaging over instances? To answer this, we had a new group of 30 subjects list properties for *fruit* and *vegetable* and used these listings to determine directly the attribute-value structures for the two simple concepts. The new *fruit* and *vegetable* prototypes contained fewer attributes than the old ones (an average of 13 attributes rather than 25), but the vast majority of attributes in the new prototypes were among those attributes of the old prototypes that had the largest numbers of votes. Most importantly, when the new representations for *fruit* and *vegetable* were used to recompute all predictions for studies 1–4, there was hardly any decrease in how well the model fit the data (in study 4, for example, the average overall correlation changed from .70 to .68).

[11] Note that in Table 8 the parameters *b* and *c* virtually always equaled .50 and .20, respectively. These constancies reflect the fact that, in fitting the model, we relied on the outcomes of preliminary fits and set the initial values of *a, b, c,* and *d* to 1.00, .50, .20, and 6.0, respectively.

Further Analysis of Slightly. Because the model did relatively poorly with concepts involving *slightly,* it is useful to take a closer look at the data obtained with such concepts. An analysis of the ratings for *slightly* concepts is presented in Figure 6. We have plotted the observed typicality of an instance in a conjunction involving *slightly* as a function of the instance's votes on the modified value. Thus the data entering into the curve labeled "slightly" include the typicality of the *fruit* instances in *slightly red fruit* as a function of how many red votes they have, the typicality of the *fruit* instances in *slightly round fruit* as a function of how many round votes they have, and similarly for the typicality of the *vegetable* instances in *slightly red vegetable* and *slightly round vegetable.* Each curve therefore averages over results for conjunctions involving *red fruit, round fruit, red vegetable,* and *round vegetable* (with different instances contributing the different points). There are separate curves for the three kinds of conjunctions involving *slightly,* and a fourth curve for adjective-noun conjunctions that serves as a frame of reference.

Compared with the curve for adjective-noun, all functions involving *slightly* are relatively flat. This uniformity sheds some light on the problems we have had with *slightly*—though the model is compatible with more uniform ratings for *slightly* than for the other adverbs, the fact that the obtained ratings vary so little makes it difficult for correlations with predicted scores to emerge. This statistical problem notwithstanding, the functions for *slightly* and *slightly non* manifest several trends that are congruent with the modification model.

The function for *slightly* initially rises and then levels off. The model is compatible with the rising trend (e.g., because *slightly red fruit* has a few red votes, it is more similar to *peach* than to *blueberry*). The model further suggests that the *slightly* function should decline at instances that contain many votes on the adjective (e.g., because *slightly red fruit* has only a few red votes, it is quite dissimilar to *strawberry*), thereby resulting in an overall nonmonotonic curve. (Lakoff's, 1973, analysis of *sort of* provides a similar argument.) There is some evidence for this nonmonotonicity in the *slightly* function in Figure 6. There might have been more evidence for this nonmonotonicity were it not for the case that the word "slightly" has a pragmatic ambiguity common to all quantitative words (Gazdar, 1979; Horn, 1972). Quantitative terms can be used to mean either *N* or *at least N.* Thus, *slightly red* can mean either $(1-k_s)$*red* or *at least* $(1-k_s)$*red,* and under the latter interpretation even *strawberry* is *slightly red.*

The function for *slightly non* is remarkably similar to that for *slightly.* This similarity may seem surprising, but it is compatible with the model (because a particular relation obtained among the parameters for *slightly* and *non,* namely that $k_s = k_n/1 + k_n$). The remaining function in Figure 6 is

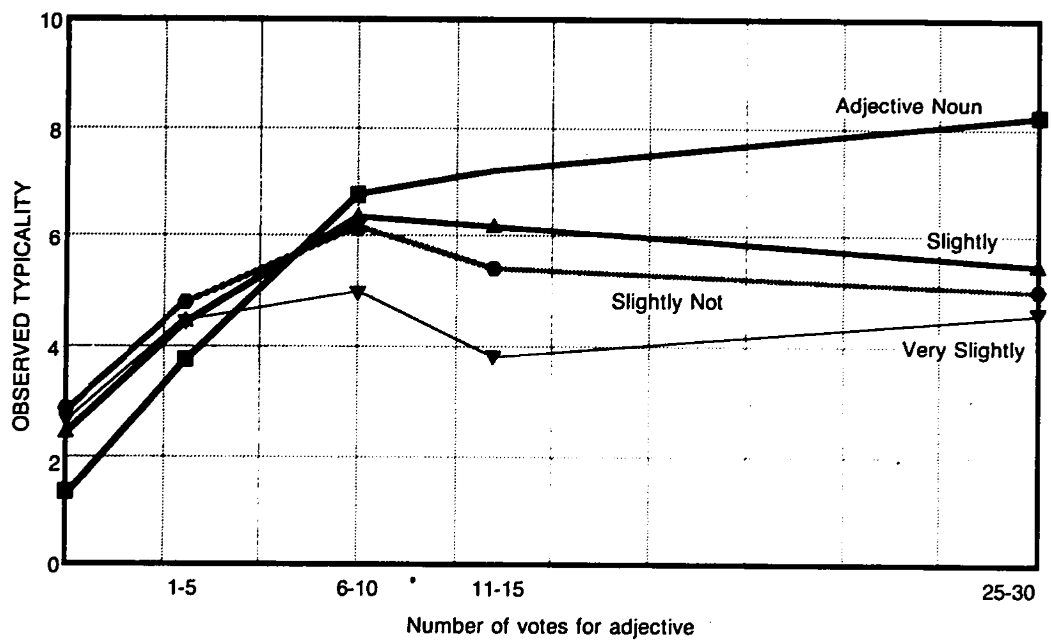

Figure 6. Obtained typicality of instances in conjunctions involving *slightly*, as a function of the instances' votes on the relevant value in the conjunction (Study 4).

for *very slightly* and here there is a serious discrepancy from the model. While *very slightly* resembles *slightly*, the model would have it behave more like *non*, being relatively high for instances that lack votes on the adjective and relatively low for instances that have many votes on the adjective.

The preceding qualitative analyses clarify the trouble spots that emerged in our quantatitive modeling. For all concepts involving *slightly*, typicality ratings were uniformly low; though compatible with our model, this uniformity restricts the correlations with predicted scores. Furthermore, because of the ambiguity of "slightly," the predicted nonmonotonicity of the function for *slightly* might have been obscured. And for concepts involving *very slightly*, the obtained data are at odds with some aspects of the model, which explains why these concepts yielded the lowest correlations between observed and predicted scores.

SUMMARY AND OTHER ISSUES

Summary

We began with three general aspects of a prototype—attribute-value structure, salience, and diagnosticity—and three specific findings about typicality effects in conjunctions—the conjunction effect, a greater conjunction effect for incompatible than compatible conjunctions, and the reverse conjunction effect. These aspects and findings guided our development of the selective modification model. One component of the model is a prototype representation for simple noun concepts, which specifies the relevant attributes and values for a concept along with numerical indicators of the diagnosticities of the attributes and the salience or votes for each value. The second component of our model specifies procedures for modifying prototypes. The procedure for adjectival modification consists of: (1) selecting the relevant attribute(s) in the noun prototype, (2) shifting all votes on that attribute(s) into the value(s) denoted by the adjective, and (3) boosting the diagnosticity of the attribute(s). The third component of the model is Tversky's (1977) contrast rule for determining typicality. The resulting package not only readily accounts for the three specific effects that started us off, but also allows us to elaborate each simple effect into a continuum; for example, the magnitude of the conjunction effect for an instance of *red fruit* depends on the number of red votes in the instance.

Study 1 provided an initial test of the model. First we used property listings to construct attribute-value representations for the noun concepts *fruit* and *vegetable* and for various instances of these concepts. Then we applied the model to produce representations for conjunctions, and to predict typicalities for the instances in the simple concepts and conjunctions. These predicted ratings were highly correlated with obtained ratings for most of the adjective–noun conjunctions that we tested.

There were, however, three concepts where the model's predictions failed to correlate highly with the obtained ratings. These concepts included *long fruit, white fruit,* and *vegetable.* We attributed the apparent failures of the model to a lack of variability in length, whiteness, and vegetableness among the instances paired with the relevant concepts (e.g., the instances paired with *long fruit* were all about the same size). To circumvent these variability problems, in studies 2 and 3 we used the same instances as in study 1 but employed only those conjunctions that involved *red* or *round.* Also, in study 3 we increased the variability of the items paired with *fruit* and *vegetable* concepts by including some noninstances of these concepts. In study 2, predicted scores were highly correlated with obtained scores for all concepts save *vegetable,* while in study 3 the model's predictions worked for all concepts. Moreover, the results of studies 2 and 3 provided support for two additional assumptions of the model; namely, that adjectives such as *red* and *round* contain only a single attribute, and that subjects base their typicality ratings for a conjunction on all attributes of the noun concept.

Next we extended the model to conjunctions that involve the adverbs *very, slightly,* or *non,* which seem to function as intensifiers. We added to the modification component of our model procedures for adverbial modification. The basic procedure consists of multiplication-by-a-scalar of the votes on a modified value. There are different scalars for the different adverbs, the scalar for *very* increases the votes on the values, that for *slightly* decreases the votes on the value but does not eliminate all votes on the value, and that for *non* eliminates all votes on the value. We extended the modification component further by proposing a means of combining scalars, thereby enabling the model to deal with concepts that involved dual adverbs.

Study 4 provided a test of the extended model. We generally found reasonable correlations between obtained typicality ratings for concepts involving single or dual adverbs and ratings predicted by the extended model. The exceptions were concepts involving *slightly.* Further analyses indicated that the obtained typicality ratings for *slightly* concepts were uniformly low, which restricted correlations with predicted scores, and that some aspects of the ratings for *very slightly* were not captured by the model.

Taken together these four studies provide evidence for many of the assumptions that make up the model. Moreover, the model is among the only theories of prototype combination to offer quantitative predictions about typicality, these predictions being based on data obtained from a paradigm (property listing) that is entirely different from rating typicalities. Of course, our case for the model is limited by our having dealt only with typicality ratings. But, as we noted at the outset, research with simple concepts has demonstrated that such ratings can be used to predict typicality effects on many kinds of performance, and there is good reason to expect the situation to turn out similarly with composite concepts.

Some Qualitative Implications

In addition to providing quantitative predictions about typicality, the selective modification model also offers qualitative insights about some important phenomena, including the generation of goal-derived concepts and apparent changes of concept structure with changing context.

Following Barsalou (1985), some concepts are created "on the fly' during one's efforts to achieve a goal. Examples of such "goal-derived" concepts include *foods not to eat on a diet* and *possessions to save in the event of fire.* Thus far researchers have treated goal-derived concepts as disconnected from natural concepts, such as *fruit* and *vegetable.* From the perspective of the selective modification model, however, many goal-derived concepts may be modified concepts like the ones discussed in this paper. *Foods not to be eaten on a diet,* for instance, may be roughly synonymous with *high-calorie foods,* a modification of *food* where all votes on the calories attribute have been shifted to the high end. Similarly, *possessions to save in the event of a fire* may be roughly synonymous with *valuable possessions,* a modification of *possession* where all votes on the worth attribute have been shifted to the high end.

Along with suggesting a mechanism for *how* goal-derived concepts are constructed, our approach provides a new account of some of Barsalou's (1985) major results. In particular, consider the finding that the typicality of an instance in a goal-derived concept is a function not of its similarity to other instances (or to a prototype summarizing those instances), but rather of its value on the dimension most relevant to the concept. For *foods not to eat on a diet,* for example, chocolate is more typical than bread and it also has a higher value on the relevant dimension of calories. All of this is intelligible in terms of the selective modification model. In composing *high-calorie foods,* if the diagnosticity of the calories attribute is boosted sufficiently high, this attribute will dominate typicality decisions and the effect of the instance's overall similarity to its prototype will be minimized. What looks like a qualitative switch in the process underlying typicality judgments may thus turn out to be just a quantitative change in one parameter of the process.

A similar story can be told about Roth and Shoben's (1983) proposal that context can change the basic structure of a concept. In one relevant experiment, subjects were timed as they read pairs of sentences in succession, such as (1a) and (2), or (1b) and (2):

(1a) Stacy milked the animal on the farm.
(1b) Fran wanted to ride the animal.
(2) She was very fond of the cow.

To understand (2), the reader must determine that "cow" refers to the same entity as "animal" does in the preceding sentence. Roth and Shoben found that this determination of coreference was easier when (2) was preceded by (1a) than (1b), and concluded that the context in (1) changed the meaning of the concept *animal.* The selective modification model, however, offers an alternative explanation. The different contexts in (1a) and (1b) lead the reader to construct different modified concepts, roughly, *milkable farm animal* and *ridable animal,* and it is these composites that the reader must relate to *cow* when reading the second sentence. (Alternatively, the modified concepts may strongly suggest particular instances—*milkable farm animal* suggests *cow*—which are then related to *cow* when reading the second sentence.) More generally, the basic claim of the selective modification model is that the meaning of a simple concept is relatively fixed and apparent meaning changes are due to modification.

Relations to Other Proposals

It is worth commenting briefly on the relation of the qualitative aspects of the selective modification model to some related proposals about prototype combination due to Thagard (1984) and Cohen and Murphy (1984). These researchers represent prototypes as "frames" (in the sense of Minsky, 1975). Our prototype representations can also be interpreted as frames. Our *attributes* are *slots,* our *values* are *slot-fillers,* and the distribution of votes over values is a *distribution of defaults.* Once this interpretation is made, it is easy to see that our model is compatible with the proposals of Thagard (1984) and Cohen and Murphy (1984).

Thagard (1984) treats simple concepts as frames, and notes that a slot-filler can be either a default value or an actual value. He then goes on to argue that in an adjective-noun conjunction, often there is at least one slot that is common to the two constituents (say, color), and the slot-filler for the adjective will dominate because it is interpreted as an actual value while the corresponding filler for the noun is only a default. Thagard's critical assumption, then, is that "actual values drive out defaults." All of this is completely in line with the modification component of our model. In the latter, it is the value named by the adjective (which corresponds to a slot-filler) that determines the final locus of the votes on the relevant attribute (slot) of the noun representation. In essence, we capture the same intuition as Thagard by letting the adjective direct the noun's votes.

Cohen and Murphy (1984) assume that: (1) Simple noun concepts are represented by frames like those discussed above; and (2) Slots are restricted to particular values—for example, at a minimum, the values that fill the shape slot for *fruit* cannot also fill its color slot. Adjective-noun conjunction is treated as a further restriction on values. In *red fruit,* for example, the color slot of *fruit* is restricted to red, and in *expert repair* the agent slot of *repair* is restricted to experts (these notions are adapted from the KL-ONE system of knowledge representation—see, e.g., Brachman & Schmolze, 1985). Cohen and Murphy's ideas have their obvious counterparts in the selective modification model: Attributes (slots) are restricted to the values (fillers) listed for them, and adjectival modification results in all votes for an attribute (slot) being restricted to one value (filler).

The point of the above is not that there are no distinct aspects of the work of Thagard or of Cohen and Murphy—Thagard considers context effects in adjectival modification, while Cohen and Murphy consider far more complex cases of modification than we do. Rather, the point is that there are no basic incompatibilities between our work and that of others using frame representations to model conceptual combination (see Murphy, 1988, for an extended comparison of models).

Limitations of the Selective Modification Model

Having made the case for our model, it seems judicious to close with some discussion of its limitations.

For one thing, there is a problem of "neutral" adjectives (as opposed to compatible and incompatible ones). Our modeling considered only those cases where the adjective encoded an attribute that presumably was part of the noun representation (*red* and *round,* say, encode the attributes color and shape, which are prestored in the prototype for *fruit*). What about cases like *upside-down fruit,* where the adjective encodes an attribute (orientation) that is unlikely to be part of the prototype for *fruit*? In Smith and Osherson (1984), we showed that such neutral conjunctions gave rise to the usual conjunction effects and their reverse. Our problem is how to model such effects. A possible solution is outlined in Cohen and Murphy (1984). In cases such as *upside-down fruit,* the relevant attribute must temporarily be added to the noun representation; then perhaps the value, votes, and boosted diagnosticity can be filled in (the value is that named by the adjective, the votes may be the maximum possible, and the boosted diagnosticity may be a constant). Of course this suggestion raises further questions—for example, how one determines the relevant attribute given only the adjective —but the suggestion seems worth pursuing.

A related difficulty is that the adjectives we have treated are all simple ones like *red,* which plausibly affect only a single attribute of a noun's representation. But many adjectives have more complex consequences for the noun phrases in which they appear. For starters, there are adjectives like *shriveled,* which we used in an earlier example and which require the model to operate simultaneously on two or more attributes (e.g., texture and shape). Although we believe the model is adequate for most multiattribute adjectives of this sort, we lack empirical evidence to support this claim. A more serious challenge comes from adjectives that produce more sweeping changes to the nouns they modify. Consider, for example, a combination such as *fake apple.* Although *fake* seems to leave some of the attributes of *apple* intact, it negates many others; thus, a fake apple might be of roughly the same color and shape as a real apple, but has a different texture, origin, and taste. R. Clark (1970) provides a useful catalog of these nonstandard adjectives (see also Kamp, 1975) that comprises the following types: *negators* (e.g., *fake*), *enlargers* (e.g., *possible*), *fictionalizers,* (e.g., *mythical*), *defictionalizers* (e.g., *simulated*), and *neutralizers* (e.g., *alleged*). Adjectives of each type can be distinguished by the inferences that they license; a *fake apple,* for instance, is necessarily a nonapple, whereas an *alleged apple* may be an apple. Presumably, the prototypes of these noun phrases differ in corresponding ways, but the task of characterizing these prototypes goes beyond what we can accomplish by means of our selective modification model.

Another limitation of our model is that it might be restricted to conjunctions of a certain syntactic form. Consider the contrast between (1) *red fruit* and (2) *red and fruit.* We have produced evidence that the adjective and noun play different roles (modifier and frame) in the first construction, but it is by no means obvious that this is true for the second construction (Oden, 1984, makes a similar point). In *red and fruit,* perhaps *red* is treated as a concept like *fruit,* rather than as a procedure for operating on *fruit,* and conceptual combination involves determining the intersection of the two concepts. This is an idea that we rejected for adjective-noun constructions, but it might work for the explicit *and* construction. These reservations about the generality of our proposals are amplified when we consider constructions even further removed, say explicit disjunction as in *red or fruit.* In short, complex concepts can be composed in different syntactic forms and it remains an open question whether different forms use different modification procedures.

Another limitation of our model is that it does not offer a convincing account of a phenomenon involving adjective-noun conjunctions that has figured centrally in previous discussions of conceptual combination (Osherson & Smith, 1982). Consider an object whose shape is midway between that of a block and that of a ball, and which is considered equally typical of the simple concepts *block* and *ball.* Intuitively, it seems that the object will be considered more typical of the conjunction *round block* than of the conjunction *round ball* because it is round for a block but not for a ball. This phenomenon rules out a large class of models of the fuzzy-set theory type. The most plausible way to account for this phenomenon in terms of the selective modification model is to assume that *ball* and *block* differ appreciably in their number of shape votes (to the extent these numbers are the same, the object in question should be equally typical of *round block* and *round ball*). But this assumption seems very dubious.

However, a slight revision of the model makes it compatible with the above phenomenon. When an adjective is applied to a noun, instead of all votes on the relevant attribute being shifted to the value denoted by the adjective, perhaps some votes are left in their original position. Thus, in forming *round block,* while most of the shape votes are shifted to round, some are left on square. Consequently, the object in question, which also has shape votes on both round and square, will find more matching shape votes in *round block* that in *round ball* (in the latter, all shape votes are round). While this revision of the model introduces another free parameter (the proportion of votes shifted), it leads to an interesting prediction. In an incompatible conjunction, typicality should be less for objects that maximally exemplify the adjective than for those that moderately exemplify it. The most typical *round block* is not perfectly round, and the most typical *square cantelope* is probably a bit round. This prediction remains to be tested.

Our last reservation about the model is that it deals with only the knowledge contained in prototypes, yet sometimes people bring to bear their general knowledge in making a decision about concept membership (see, e.g., Lakoff, 1987; Murphy & Medin, 1985; Rips, 1988). To illustrate with an example from Murphy and Medin, if at a party you see a person jump into a pool fully clothed, probably you would categorize him or her as *drunk.* This categorization is almost certainly not based on similarity to a prototype, because your prototype for *drunk* is unlikely to include any mention of jumping into pools clothed. Rather, your categorization is probably based on your general knowledge about parties, liquor, and erratic behavior.

In addition to influencing categorization processes, general knowledge may also affect the processes involved in constructing a composite prototype. Murphy (1988) points out that one may need general knowledge to determine which attributes of the noun are to be modified as well as to fill in the complex concept so that it is coherent (see also Hampton, 1987). Given the conjunction, *apartment dog,* for example, one needs general knowledge to: (1) know that *apartment* modifies the habitat attribute of *dog,* and (2) refine the composite by adding that an apartment dog is likely to be smaller, quieter, and better-behaved than dogs in general.

The above arguments indicate that the selective modification model offers an incomplete picture of the composition and use of modified concepts. Several considerations, however, suggest that the model will form an important component of any more general theory of conceptual combination. For one thing, the model may be adequate to handle many adjective–noun combinations, for the examples used to show the role of general knowledge typically involve noun–noun composites (such as *apartment dog*). Another matter is that even in cases where general knowledge is used, it may not come into play until after the procedures specified in the selective modification model. That is, our model may describe a rapid composition process, which is sometimes followed by a slower composition process that uses general knowledge. A third consideration is that, even if it turns out that the processing of general knowledge must be interleaved with the procedures of the selective modification model, the model still tells us the basic subtasks that modification must accomplish, such as finding the relevant attributes in the noun, boosting their diagnosticities and altering the saliences of their constituent values, and so on.

■ Original Submission Date: July 23, 1987

REFERENCES

Barsalou, L.W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 629–654.

Bishop, Y.M.M., Feinberg, S.E., & Holland, P.W. (1975). *Discrete multivariate analyses.* Cambridge, MA: MIT Press.

Brachman, R.J., & Schmolze, J.G. (1985). An overview of the KL-ONE Knowledge Representation System. *Cognitive Science, 9,* 171–216.

Chandler, J.P. (1969). STEPIT—Finds local minimum of a smooth function of several parameters. *Behavioral Science, 14,* 81–82.

Clark, H.H., & Clark, E.V. (1979). *Psychology and language: An introduction to psycholinguistics.* New York: Harcourt Brace Jovanovich.

Clark, R. (1970). Concerning the logic of predicate modifiers. *Noûs, 4,* 311–336.

Cliff, N. (1959). Adverbs as multipliers. *Psychological Review, 66,* 27–44.

Cohen, B., & Murphy, G.L. (1984). Models of concepts. *Cognitive Science, 8,* 27–60.

Fodor, J.A. (1981). The present status of the innateness controversy. In J.A. Fodor (Ed.), *Representations: Philosophical essays on the foundations of cognitive science.* Cambridge, MA: MIT Press.

Gati, I., & Tversky, A. (1984). Weighting common and distinctive features in perceptual and conceptual judgements. *Cognitive Psychology, 16,* 341–370.

Gazdar, G. (1979). *Pragmatics: Implicature, presupposition and logical form.* New York: Academic.

Glass, A.L., & Holyoak, K.J. (1975). Alternative conceptions of semantic memory. *Cognition, 3,* 313–339.

Halff, H.M., Ortony, A., & Anderson, R.C. (1976). A context-sensitive representation of word meaning. *Memory & Cognition, 4,* 378–383.

Hampton, J.A. (1982). A demonstration of intransitivity in natural categories. *Cognition, 12,* 151–164.

Hampton, J.A. (1987). Inheritance of attributes in natural-concept conjunctions. *Memory & Cognition, 15,* 55–71.

Hersh, H.M., & Carmazza, A. (1976). A fuzzy set approach to modifiers and vagueness in natural language. *Journal of Experimental Psychology: General, 105,* 254–276.

Horn, L.R. (1972). *On the semantic properties of logical operators in English.* Unpublished doctoral dissertation, University of California, Los Angeles.

Kamp, J.A.W. (1975). Two theories about adjectives. In E.L. Keenan (Ed.), *Formal semantics of natural language.* Cambridge, England: Cambridge University Press.

Lakoff, G. (1973). Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic, 2,* 458–508.

Lakoff, G. (1987). *Women, fire, and dangerous things.* Chicago: University of Chicago Press.

Malt, B.C., & Smith, E.E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior, 23,* 250-269.

Medin, D.L., & Schaffer, M.M. (1978). A context theory of classification learning. *Psychological Review, 85,* 207-238.

Medin, D.L., & Shoben, E.J. (1988). Context and structure in conceptual combination. *Cognitive Psychology, 20,* 158-190.

Mervis, C.B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology, 32,* 89-115.

Minsky, M. (1975). A framework for representing knowledge. In P.H. Winston (Ed.), *The psychology of computer vision.* New York: McGraw-Hill.

Murphy, G.L. (1988). Constructing complex concepts: A cognitive science case study. *Cognitive Science, 12,* 529-562.

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

Oden, G.C. (1984). *Everything is a good example of something, and other endorsements of the adequacy of a fuzzy theory of concepts* (WHIPP 21, Wisconsin Human Information Processing Progress Report). Madison, WI.

Osherson, D.N., & Smith, E.E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition, 9,* 35-58.

Osherson, D.N., & Smith, E.E. (1982). Gradeness and conceptual combination. *Cognition, 12,* 299-318.

Rips, L.J. (1988). Similarity, typicality, and categorization. In S. Voisniadou & A. Ortony (Eds.), *Similarity, analogy, and thought.* New York: Cambridge University Press.

Rosch, E.H. (1973). On the internal structure of perceptual and semantic categories. In T.E. Moore (Ed.), *Cognitive development and acquisition of language* (pp. 111-144). New York: Academic.

Rosch, E., & Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7,* 573-605.

Roth, E.M., & Shoben, E.J. (1983). The effect of context on the structure of categories. *Cognitive Psychology, 15,* 346-378.

Rumelhart, D.E., & Ortony, A. (1977). The representation of knowledge in memory. In R.C. Anderson, R.J. Spiro, & W.E. Montaque (Eds.), *Schooling and the acquisition of knowledge.* Hillsdale, NJ: Erlbaum.

Shafir, E.B., Smith, E.E., & Osherson, D.N. (1988). *Typicality and reasoning fallacies.* Unpublished manuscript, University of Michigan, Ann Arbor.

Smith, E.E., & Medin, D.L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Smith, E.E., & Osherson, D.N. (1984). Conceptual combination with prototype concepts. *Cognitive Science, 8,* 357-361.

Smith, E.E., Shoben, E.J., & Rips, L.J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81,* 214-241.

Thagard, P. (1984). Conceptual combination and scientific discovery. In P. Asquith & P. Kitcher (Eds.), *PSA,* 1. East Lansing, MI: Philosophy of Science Association.

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

Tversky, A., & Gati, I. (1982). Similarity, separability, and the triangle inequality. *Psychological Review, 89,* 123-154.

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

Zadeh, L.A. (1971). Quantitative fuzzy semantics. *Information Sciences, 3,* 159-176.

Zadeh, L.A. (1972). A fuzzy-set theoretic interpretation of hedges. *Journal of Cybernetics, 2,* 4-34.