The combination of prototype concepts

James Hampton

City University London

Chapter in "The Psychology of Word Meanings", ed. P. Schwanenflugel.

How does the meaning of a sentence relate to the meaning of the individual words within it? This question is one of the central issues in the semantics of natural language. In the absence of any strong contextual influence, the answer must be that the meaning of a sentence is composed or derived from the meanings of the individual words within it. On this view, semantics needs to provide a description of individual lexical meanings, together with combination rules for combining words into phrases, clauses, and sentences.[1]

This chapter is concerned with a particular problem in compositional semantics - namely the way in which the meaning of complex noun phrases is derived from the meaning of their component words. Shoben (this volume) describes some of the intricacies involved with different forms of adjectival modification of nouns. In forming adjective-noun phrases, a wide range of mechanisms can be seen. In fact relatively few cases are simply intersective in the way that CARNIVOROUS MAMMAL is the intersection of the classes of CARNIVOROUS and MAMMALS.

My aim in this chapter is to look at how people construct noun phrases involving two different nouns - either in a direct compound (like OFFICE FURNITURE) - or in various `logical' combinations such as relative clause modification (A SPORT WHICH IS A GAME), negation (A SPORT WHICH IS NOT A GAME) or disjunction (SPORTS OR GAMES). The interest in such examples is that at first sight they would appear to be relatively straight forward. The category of sports which are also games would seem to be simply those members of the sports category which are also in the category of games - that is the intersection of the two sets. I aim to show how people's interpretation of such apparently logical constructions is sometimes at odds with a classical interpretation in terms of Boolean set logic. For example SPORTS WHICH ARE GAMES is not coextensive with GAMES WHICH ARE SPORTS, and both categories may include items that are not judged to be GAMES simpliciter. The reasons for this divergence from set logic lie, I believe, in the nature of the concepts underlying the meanings of many common nouns - the so-called prototype structure of concepts. To understand the compositionality of meaning, we need to delve into the structure of the underlying lexical concepts, and look at how concepts combine.

The organization of the chapter is as follows. First, I provide a formulation of prototype theory (Rosch & Mervis, 1975) which argues that most of our everyday concepts like ART, CHAIR or SPORT are not defined in terms of clear cut common elements, but are based instead on clustering together similar objects around a `prototype' form. As well as attempting to provide a clear specification of my understanding of this model, I provide a brief review of some of the empirical evidence supporting the theory. I then look at the important question of how conceptual combination might work if concepts themselves are inherently vague, as prototype theory would predict. There follows a summary of research that I have conducted over the last few years into the question of how people interpret different ways of combining prototype concepts, and an outline of a proposed Composite Prototype Model for forming conjunctive concepts. There are finally some concluding remarks.

Prototype Theory - a formalization

Prototype Theory aimed to replace the so-called "classical view" of concepts according to which all instances of a concept shared some set of defining features in common (Bruner, Goodnow and Austin, 1956). The new theory suggested that concepts are represented in memory by some ideal prototype which represents the central tendency of the class. Taking up Wittgenstein's idea of family resemblances (Wittgenstein, 1953), Rosch and Mervis (1975) proposed that the meaning of a word such as FRUIT could be represented by a prototype of the ideal or typical fruit, and that other objects would be classified as fruit (that is correctly referred to using the word "fruit") depending on how similar they were to the prototype. Objects in the world as we perceive them fall naturally into clusters, based around the correlation of attributes across different objects. Creatures with feathers tend also to have wings, fly, have two legs and build nests. Hence we form a prototype of such creatures - a kind of generic example, and label it with a name - "bird". Other creatures are then considered to be birds to the extent that they have enough of the cluster of attributes that constitute the prototype of the category - effectively this means that we classify the world on the basis of similarity, where similarity is determined by some function of common and distinctive features (e.g. Tversky, 1977). A great deal has been written about prototype theory (see for example, Smith and Medin, 1981, and Lakoff, 1987). For the arguments which will follow, however, it is necessary to state clearly just what the theory involves.

Formally, prototype theory involves the following four axioms □:

A: a prototype concept is constituted by a set of attributes, each with a particular weight corresponding to its `definingness' or contribution to the concept's definition. The instantiation of the full set of attributes in an individual instance would therefore produce a prototype example of the concept. However I take the set of weighted attributes to be the basic representation of the concept, rather than any particular instance, since an instance would also be likely to instantiate many other irrelevant attributes, not defining for the concept. There may also be incompatibilities amongst the set of attributes defining the prototype, so that not all attributes can be instantiated at the same time. For example furniture can have a number of different functions, while apples can have a number of different typical colors. An attribute set is thus a more flexible representation than is an instantiated prototype.

B: a similarity scale is defined for the domain of possible instances based on the set of weighted attributes. For example, within the domain of creatures, objects can be placed on a scale of similarity to BIRDS. Usually a linear combination rule is assumed, such as the sum of attributes possessed by an item, weighted by the definingness of each attribute for the prototype, and also weighted by the degree to which the attribute is typically true of the item. Other non-linear functions would however be possible within the general prototype framework.

C: ratings made by subjects of item "representativeness" or "typicality" are monotonically related to the similarity scale. (Below a certain level, items will be seen as merely related to the concept, rather than being representative or unrepresentative of it.)

D: a criterion is placed on the similarity scale such that membership in the category is positive for items with similarity above criterion, and negative for items below criterion.

Two quick points can be made about this characterization. First, note that the use of a criterion in (D) means that items can vary in rated typicality without necessarily varying in the probability with which they are judged to belong in the class. Membership and typicality are both derived from the underlying similarity scale, but are not identical. So, for example thrushes and ostriches may differ widely in their typicality as birds, but yet a (reasonably well-educated) group of people may all consistently categorize both kinds of creature as birds. In this case the range of similarity within the category does not fall low enough to come within the range of variation of the criterion - unless extinct species like the pterodactyl or the archeopterix are considered.

Second, note that while the model can readily account for disagreement and uncertainty in category membership judgments, it does not require that all categorization is unclear (or `fuzzy' to use Zadeh's (1965) term). If both the criterion in (D) and the similarity measure in (B) are highly stable for a particular concept, or else if there are no items in the world corresponding to points in the similarity space that are near the criterion (as in the BIRD category), then there may be no items whose membership is fuzzy. The degree of fuzziness in categorization for a particular concept will therefore come from four sources:
- the variability across subjects and across situations of the weights attached to different attributes (leading to variability in perceived similarity),
- the variability across subjects and situations in the degree to which an item (or subclass of items) is assumed to possess each attribute,
- the variability across subjects in where the criterion is placed (a commonly observed individual difference), and
- the availability in the world of borderline instances.

Prototype theory - the evidence

The evidence for this model of concepts came from four main sources. To start with, the categories named by common nouns like FRUIT, SPORT, or CHAIR, show effects that can not readily be explained with a classical common-element definition. First, they often show fuzziness in category membership decisions. There are items at the borderline of the category which cannot be easily classified as in or out of the set. Tomatoes are a well-known case, where classification as a fruit seems to be a matter of debate, and as McCloskey and Glucksberg (1978) demonstrated experimentally, similar fuzzy boundaries can be seen in a wide variety of semantic categories, both as disagreement between different subjects, and, more significantly, as inconsistency within the same subject on different occasions. It is important to note that the classification is not fuzzy simply because of uncertainty or a lack of information about an instance - tomatoes are, after all, a very familiar part of our diet - but because of an inherent vagueness in the category definition itself, as captured by the prototype theory. The second aspect of categories that suggests a prototype structure is the variation in typicality or representativeness amongst category members. Apples are just much more typical as fruits than are olives. If all category members have the same attributes in common, then why should this be? There has subsequently been much research devoted to the question of what determines typicality of category items (e.g. Barsalou, 1985, Hampton & Gardiner, 1983), and one strong determinant for common natural categories does turn out to be the degree to which an item possesses the attributes shown by other category members (Rosch and Mervis, 1975, developed a "family resemblance score" based on this notion, which they found to correlate well with rated typicality).

The remaining two sources of evidence come from a consideration of the attributes themselves. First, when subjects are asked to describe the meaning of a term like BIRD, by listing different descriptive attributes, they appear to make no distinction between attributes which are common to the whole set (such as two legs and feathers) and those that are only true of typical birds (such as flying). The most commonly listed attributes may often be true of only the more typical members of a category. Second, when one looks through all the attributes generated, and ignores those which are not common to all category members, it is frequently not possible to construct a sufficient common element definition out of the remaining attributes (Hampton, 1979, McNamara & Sternberg, 1983 Rosch & Mervis, 1975).

Prototype theory has had considerable success, both in application to other domains (Hampton, 1981, Cantor et al., 1980) and in the demonstration of a wide range of psychological measures that are sensitive to typicality differences in category items. However, surrendering the neat logic of the classical common element definition left a serious gap in the theory of concepts, and it was to this flaw that Osherson and Smith (1981) drew attention in a highly influential article, in which they challenged the whole basis of prototype theory as a theory of concepts.

Combining concepts

With a classical theory of concept definitions, one can define the conjunction of two sets either extensionally (that is in the external world of things) - as the set intersection of the two categories of exemplars - or intensionally (in the internal world of ideas) - as the concept defined by the set union of the two sets of defining attributes. For example, within the general domain of people, the concept of British Sportsman can be defined extensionally as the intersection of the set of British people, and the set of Sportsmen, or intensionally as the set union of the defining features of being British and the defining features of being a Sportsman. Conjunction thus involves grouping together the two sets of defining attributes to make a new, more specific, set, and with common element definitions the two aspects - the extensional and the intensional - are in perfect accord. How could this process work however, if we give up common element definitions, and only have similarity-based prototypes to work with? Prototype theory itself had no answer.

Initially an answer was sought using the notion of Fuzzy Logic, developed by Zadeh (1965) for representing the vagueness of natural language sentences such as "John is tall" or "There is a pile of sand". The limits of size for tallness, or for constituting a pile, are naturally vague, and fuzzy logic provided a way of giving truth values to vague statements and explained how those truth values were changed by logical operators such as negation, conjunction and disjunction. The basic premise was to suppose that truth was a continuous variable that could be represented by some number between 1 (for True) and 0 (for False). One could therefore speak of a fuzzy set such as TALL MEN which had clear cases at its centre, and continuous gradation of degrees of membership, where degree of membership for any instance P in a fuzzy set S could be equated with the truth value of the proposition `P is a member of the set S'. It seemed a natural extension of the approach then to use fuzzy logic to model the logic of vague categorization statements. Just as `John is a tall man' may be true to a varying extent, as a function of John's height, so `X is a FRUIT' could be considered true to a varying extent as a function of X's similarity to the prototype definition of FRUIT. The proposition `A tomato is a fruit' might then be represented as having a truth value of 0.6, for example. (No mechanism was ever suggested for actually specifying numerical truth values as a function of any empirical measure, but one could clearly do so.)

As a brief example of this approach, consider the case of conjunction. Given that proposition p is true to extent cp and that q is true to extent cq, how true is the proposition p&q (both p and q)? Zadeh chose combination rules for conjunction that would yield a classic truth table definition in the case where the values of cp and cq are restricted to be simply TRUE (1) or FALSE (0). Two such rules that he proposed for conjunction were:

The Minimum rule - p&q is as true as the least true of p or q considered separately.

The Product rule - the truth of p&q is found by multiplying the individual truth values of p and q. This rule he termed `interactive'.

In both these rules, if either cp or cq has a value of zero, then so does the conjunction. Hence statements with clear-cut truth values will produce the standard truth table definition for conjunction - something that is false if either of the component statements is false, and true otherwise. Evidence for the rules was provided by Oden (1977) in a study in which he showed that people's judgments of the combined truth of two independent category statements (such as `A penguin is a bird and a tomato is a fruit') could be predicted from the individual judged truth values of each statement, and a fuzzy logic conjunction rule.

Initially, fuzzy logic was seen as a direct way of representing the degree of similarity of an instance to a prototype concept. Thus both differing typicality and fuzzy boundaries were to be captured by the value of cp for some p such as `Instance X is in category C'. Unfortunately for this general approach, Osherson and Smith (1981, 1982) were able to demonstrate quite neatly that fuzzy logic would not work as an account of the typicality of instances in conjunctions of semantic categories. This effect is illustrated by the so-called `guppy effect'. Osherson and Smith pointed out that although the guppy is a poor example of a pet, and a poor example of a fish, yet it is a good example of a pet fish. None of Zadeh's fuzzy logic rules for conjunction can allow this, since neither the minimum nor the product rules can assign truth values to a conjunction that are higher than those of a constituent proposition. One therefore has to abandon at least the notion that typicality judgments can be captured by a fuzzy logic formalism. In any case, fuzzy truth values may not be well suited to represent instance typicality, since as I argued earlier, the use of a criterion for determining set membership in the prototype approach means that it is quite consistent for items to vary in typicality even though they are equally clear members of the set, and thus have equal truth value for the proposition `X is in set C'.

The problem (as a number of researchers realized at about the same time) was that fuzzy logic operates on sets of items in categories, without any recourse to the semantic "intensional" information in the concept definitions. When words are combined, then there is frequently an interaction between their semantic contents, such that the complex concept no longer shares all the attributes of its constituent concepts. Pets are typically furry, home-based objects, while fish are typically scaly creatures living in the wild. Since one cannot be at the same time furry and scaly, or home-based and live in the wild, the conjunctive concept PET FISH can only have one or the other attribute - that is, some attributes are kept, and some are lost. Pet fish are scaly and home-based rather than furry wild creatures. Guppies have just these attributes and so they are good examples of pet fish, whereas they lack the furriness of pets, and the wildness of fish, and so are not good examples of either category alone.

The theoretical move away from a fuzzy logic applied to extensional set membership and towards intensional attribute-based concept definitions leads prototype theory to make a more radical prediction. It is normally taken as a given of rational behavior that people will only allow an item to belong in a conjunction of two categories if it also belongs in both categories individually. That is, the conjunction of concepts involves the set intersection of the two categories. However, if we abandon the extensional approach to combining prototype categories, as exemplified by fuzzy logic, then this logical constraint on conjunction will no longer hold. The intensional approach says that we take the two sets of attributes for, say, Pet and Fish, and then combine them interactively to create a new set of attributes - a composite prototype as I have termed it (Hampton, 1987, 1988b) - for the concept Pet Fish. If membership in the Pet Fish category is then determined as for other prototype concepts, by similarity to the prototype, (defined by some function of possession of the attributes in the composite prototype) then what is to prevent inconsistencies of classification? The guppy effect shows how some items, like guppies, may possess more of the attributes of the composite prototype than they do of either original constituent set. Unless the criterion for membership in the conjunction is set very high, this will mean that people may allow some items to be in the conjunction which were not in either one, or even both of the constituent sets.

Osherson and Smith (1982) found this a disturbing corollary of prototype theory - sufficiently disturbing to cast doubt on the adequacy of the theory as an account of concept use. Theories of concepts that embody apparently `irrational' combination mechanisms are a threat to the need for an account of our ability to reason logically, and the normative use of concepts (Hampton, 1989). The only way to use the intensional combination of attribute information to produce a set intersection of the categories, is the classical system of having defining attributes, which are common to all category members - a solution that Osherson and Smith (1982) therefore argued for. They suggested that prototypes may be involved simply in the determination of typicality within a category, while some `core' of defining features would be involved in any operation involving the logic of set membership (a position which has had considerable support: e.g. Armstrong, Gleitman & Gleitman, 1983, Miller & Johnson-Laird, 1976, Keil & Batterman, 1984, Smith, Shoben & Rips, 1974).

Experimental evidence

It was against this background, that I decided to collect some data on how people actually judge membership and typicality in conjunctive and other kinds of set combinations. In an earlier study (Hampton, 1982) following up some work in the anthropological literature (Randall, 1976, Kempton, 1978) I had shown that the set relation corresponding to natural language categorization, was not in fact class inclusion as one might suppose. Thus when people claim that the following statement is true:

(1) A chair is a type of furniture

they do not thereby imply the logical proposition:

(2) For all x, if x is a chair, then x is furniture.

Examples like carseats, and ski-lifts show that natural language categorization is based on intensional information, and not on extensional set inclusion. Typical chairs possess enough of the prototypical attributes of furniture to make (1) generally true. It is in this sense of "generic truth" that many natural language statements are understood. It may then be quite possible that other apparently extensional statements are treated intensionally with similar `inconsistent' results. To this end I set out to examine the logical consistency with which people normally interpret a number of other common language constructions - noun-noun compounds, relative clause qualification, disjunction using "or", and negation with relative clauses. The following sections briefly review the results of these studies.

Noun-noun compounds

If one interprets PET as a noun, then the phrase Pet Fish is an example of a noun-noun compound. The traditional linguistic treatment of such compounds (eg. Levi, 1978, Shoben, this volume) is that the first noun is a modifier, while the second is the head noun. It is the head noun that defines the class in which the concept falls, while the modifier selects a particular subset of that class. For instance an OCEAN DRIVE is a particular kind of drive, a CAR REPAIR is a kind of repair, and so forth.

Compound noun phrases have received a lot of attention of late (Murphy, 1988, Medin & Shoben, 1988), as they provide good examples of non-compositional combinations of word meanings - that is, examples where the result of the combination can not be predicted solely in terms of the constituent meanings. Consider the following pairs:

(3) A hand wound
A gun wound

The `mediating relation' specifying how the concept WOUND is modified in the case of hand is the PATIENT role - wound done to a hand, and in the case of gun it is instrumental - wound done with a gun. It may then be tempting to propose that the word HAND looks for a patient role to fill whereas GUN looks for an instrumental. However they can easily be reversed, as in the following pair of compounds:

(4) A hand repair
A gun repair

Here it is now the hand that is the instrument, and the gun that is the patient of the head noun.

Even though it is readily apparent from this kind of example that the role of the modifier noun is highly flexible and context dependent (see Cohen and Murphy, 1984), most treatments see the head noun as relatively fixed in meaning. However such need not be the case. If the modifier acts to change enough of the default attributes of the head noun class, then it should be possible that an item which is not normally similar enough to the head noun prototype to belong to the head noun class, could nevertheless be similar enough to the modified concept to belong in the category defined by the noun-noun compound. Hampton (1988b, experiment 1) showed evidence for this semantic interaction. Subjects were asked whether items were examples of furniture, and also if they were examples of school furniture, garden furniture, church furniture, and so forth. There was a significant tendency for some items to be included in the compound categories which had been excluded from the head noun category on its own. Blackboards are not furniture, but are considered to be school furniture. What is more, in a follow-up test, subjects assented that it was possible for both the following statements to be true:

(5) School furniture is a type of furniture

(6) Some school furniture is not furniture

It is clear therefore that for subjects in this task, categorization of a class as a subtype of another class did not bring with it the extensional implication of class inclusion.

Relative clause conjunctions

Given that most noun-noun compounds are not simple intersections of the two noun categories (school furniture is not the intersection of schools and furniture), they may not provide the best testing ground for studying the conjunction of prototype categories. I therefore looked for a more obviously conjunctive natural language phrase, and chose relative clauses for a further set of experiments. The phrase "Sports which are also Games" seems at first blush to be uncompromisingly the subset of sports which also happen to belong to the class of games. As such the phrase should be equivalent to the converse phrase "Games which are also Sports", since extensionally both phrases pick out the intersection of the sport and game categories.

How do subjects rate items in these categories? The basic method used to study this question was to present them with a list of objects. At the head of the list was a category name - either a simple category such as SPORTS, or a conjunctive category like "Games which are also SPORTS". Their task was to work down the list of items deciding if each was in the category or not. They were additionally asked to rate category members for their typicality in the category, and to judge the relatedness of non-members to the category. The experiment has been done both within subjects (where subjects first rate the list of items as Sports, then as Games, and then in the conjunction), and also between subjects (with different subject groups providing each of the judgments) with essentially the same result. The basic finding was that subjects overextended the conjunctive categories to include in them items which were judged not to belong to one of the two constituent categories. For example they would rate CHESS as a GAME but not as a SPORT, but would later judge it to be a SPORT which is a GAME (or a GAME which is a SPORT). Seven different pairs of categories were examined, with similar results in each case.

The overextension could be tested in three ways. First, individual subjects' sets of three category judgments for each item were analyzed and shown to have a much higher frequency of inconsistent overextended responses than would be expected by chance alone (even allowing for the unstable nature of category judgments close to the category boundary). Second, category membership scale values were calculated for each item, based on the rating scale and ranging from most typical member, through to least related non-member. This scale can be thought of as an operationalization of the similarity-to-prototype scale. Membership scale values for the conjunction were in the majority of cases intermediate between membership for one category and membership for the other. Indeed a very regular relationship obtained between the three sets of item means, suggesting that a geometric mean of the two constituent membership values could accurately predict the membership value for the conjunction. A final way of testing the overextension was to apply a probability model to the pattern of positive or negative categorization judgments, using data from different groups of subjects for each classification judgment (Hampton, 1990). Once again, the probability of an item being classed in the conjunction was generally predictable by a geometric average of the two constituent category probabilities.

The conclusion drawn from these studies was that when a noun occurs modified by a relative clause, or is itself a modifier in such a clause, the semantic interpretation of the concept is modified. More specifically the original concept is liable to lose some of its prototypical default features, through conjunction with another concept. The result is that some items will be more similar to the conjunction than to the simple concept, and hence may be included in the conjunction category although not included in the simple concept category.

We will return to this account in the final section which presents a model for composite prototype formation for conjunctive concepts.

Disjunction

I have so far tried to argue that categorization statements in language are not tied to a class inclusion logic, and that conjunctions of semantic categories do not follow the logic of set intersection. In a further set of studies (Hampton, 1988a) I considered the disjunction of semantic categories. If people are asked to consider whether an item falls into either of two categories, will they follow the logical rule of set union - that is give a positive answer to all and only those items that they would place in one or the other category (or both) considered individually?

The short answer is again no. The experimental procedure was once again to ask people to judge category membership and to rate typicality or relatedness, first for a list of items as members of one category, then as members of the other category, and finally as members of either (using expressions such as TOOLS OR INSTRUMENTS, or FRUITS OR VEGETABLES). A between-groups experiment was also carried out using the same tasks. In order to demonstrate the effects of intensional representation on classification judgments, it is necessary to choose categories to be disjoined that are close together in the same semantic domain. We may then expect on the basis of the previous results that subjects would be inclined to judge membership in a disjunction such as HERBS or SPICES, on the basis of some overall average membership in the two constituent categories. For a disjunction, a good example of either category is sure to be a member, regardless of membership in the other category, so one would expect any deviation from a simple linear averaging to be in the opposite direction from the case of concept conjunction. When the data were collected and analyzed, it was indeed found that both degree of rated membership in the disjunction, and probability of a positive response, could be accurately predicted by a regression equation involving the two component category membership values, and a negative interaction term. That is membership in a disjunction could be predicted with an equation of the general form of (7).

(7) disjunction = c1 + c2.A + c3.B - c4.A.B

where A and B represent degree of membership in the constituent categories, c1 is a constant, and c2, c3 and c4 are positive constants. Note that (7) has the same form as the rule for calculating the probability of a disjunction of two events (8).

(8) p(A or B) = p(A) + p(B) - c.p(A).p(B)

where c is a positive constant reflecting the degree of interdependence between the two events.

In terms of responses to individual items, the effect of the semantic interaction was to lead to a pattern where some items were included in the disjunction which had been rejected from both categories individually, while others were excluded, even though they had been included in one of the constituent categories. Examples of the former case of overextension occurred with cases like FRUIT or VEGETABLES, where the two categories being combined were close to being contrasting sets, which together covered most of the domain. In such a situation, a number of other edible plant stuffs - like nuts, herbs, rhubarb etc. - were not included in either category alone, but were included in the disjunction. As one subject put it - "I don't really think it is a Fruit, and I don't think it is a Vegetable, but I feel it has to be in one category or the other". That overextension was not just an effect of statistical dependence between the decisions was shown by the fact that for many decisions the probability of a positive disjunction response was greater than the sum of the two constituent response probabilities (breaking the constraint on disjunctions of probabilities - see (8) above).

Examples of underextension - the exclusion of items from the disjunction which belonged to one or other constituent category - could often be attributed to a contextual effect operating between the two categories. Considering HOUSEHOLD APPLIANCES on its own, subjects were prepared to give equal weight to all machine-like artifacts in the home as category members. When disjoined with KITCHEN UTENSILS, the feature "used in the kitchen" appeared to affect the overall decision, to the extent that items like electric toothbrushes were no longer considered to belong.

Negation

The final study in this series investigated the effect of negation on complex categorization (see Hampton 1990). In order to test the effect of the word "NOT" on the interpretation of noun phrases, it is necessary to place the phrase in some clearly defined domain. One cannot go around asking people to what extent a bicycle is an example of the category "Not a Fruit", since the category has no prototype. The problem could be easily solved by using the relative clause constructions studied earlier. By placing the NOT operator in the relative clause modifier, the presumed scope of the negator could be clearly specified. In an unpublished study (Hampton, 1990) subjects were thus asked to categorize items as SPORTS, GAMES, and then as SPORTS WHICH ARE NOT GAMES, or GAMES WHICH ARE NOT SPORTS. The analysis of mapping constituent membership onto membership in the complex noun phrase category was then carried out in the same way as for conjunctions and disjunctions.

For the mean membership ratings, the pattern was very much consistent with the earlier results. By placing a negative sign in front of membership in the negated category, the results for conjunctions (A which are B) could be translated into results for complement conjunctions (A which are not B). The form of the regression equations thus took the form (9):

(9) A which are not B = c1 + c2.A - c3.B + c4.A.B

where the sole change from conjunction is the negating of the constants affecting the B term.

As regards the `inconsistency' and overextension observed for conjunctions, the same pattern was observed for complements, with in some cases a more extreme degree of overextension. Items were included in the category "A which are not B" which had either been left out of category A, or which had been included in category B. For example, in a between-groups study, a majority of each group was in favor of including DESK LAMPS in each of the following categories:

(10) HOUSEHOLD APPLIANCE that is FURNITURE
FURNITURE that is a HOUSEHOLD APPLIANCE
HOUSEHOLD APPLIANCE that is NOT FURNITURE
FURNITURE that is NOT A HOUSEHOLD APPLIANCE.

It was also found for the complement sets, that there were some items which were much more likely to be included in the complex set, than in either of the two constituent sets (if one assumes membership in NOT B to be 1 minus membership in B). This double overextension was very rare in the case of un-negated relative clause conjunction, where items which were overextended for one category were almost always very typical examples of the other category. Double overextension is the phenomenon that would be predicted by the "guppy effect" discussed in the introduction. Since guppies are purportedly better examples of PET FISH than of either PET or FISH, then for items near the borderline one might expect degree of membership to follow a similar pattern. Failure to find it for simple conjunctions may of course reflect the materials used for those experiments.

An inheritance model for composite prototypes

The relationship between category membership in individual concepts and in different logical combinations of them appears to be a very regular one, at least for the range of categories used in the studies. The temptation might then be to try to revive a version of fuzzy set logic to account for the results. Given appropriate scaling factors, the data approximate reasonably well to Zadeh's "interactive" or multiplicative combination rules, by which degrees of category membership are combined in the same way as event probabilities. Indeed Zadeh himself (Zadeh, 1982) made a similar proposal.

There are at least two problems with using an extensional fuzzy logic approach. First, Osherson and Smith (1982) were able to demonstrate that no general fuzzy logic function can deal with concept conjunction. The degree to which pairs of categories allow the "guppy effect" seems to depend on the degree of overlap between them. Pets and fish are categories with very low overlap (most pets are not fish, and most fish are not pets). Other categories with much higher overlap, like Sports and Games, will show little or no double overextension. The combination function needs to reflect the overlap between the categories.

The second problem relates to an unexpected phenomenon that appeared in the studies described above on conjunction. When two categories are combined in conjunction with a relative clause, it turned out that one of the categories tended to have a much stronger effect overall than the other. For example the extent to which an activity is a SPORT had much more influence that the extent to which it is a GAME in determining its degree of membership in the conjunctive categories "SPORTS WHICH ARE GAMES" or "GAMES WHICH ARE SPORTS". This pattern of dominance was found in most of the category pairs studied, and was replicated in different studies with different lists of category items. An extensional account would have to explain the effect in terms of some extensional aspect of the category - such as its set size. However set size does not appear to predict the dominant category.

On the basis of these arguments against the extensional approach, the following intensional model was developed, which has the advantage of readily explaining why double overextension depends on semantic overlap, and also why some categories should dominate others.

The composite prototype model

The model makes the following initial assumptions:

1. Concepts are represented as sets of attributes, in much the same way as Rosch's prototype theory proposed. The attributes may however be interconnected by higher level "theory driven" relations. Malt & Smith (1982) and Murphy and Medin (1985) pointed out that as well as knowing that birds have wings and can fly, we also know that having wings is an enabling condition for flight. That is there are higher levels of understanding, such as a naive physics of flight, which provide additional semantic relations between concept attributes. Such relations will typically involve notions like CAUSE and ENABLEMENT, as for example in:

(Pets live in homes) CAUSES (Pets need litter trays)
or
(Pets are tame) ENABLES (Pets are cuddled)

2. The attributes are associated with a quantitative "degree of definingness" which I will term Importance. Those attributes which are most central or criterial for category membership will be the most important, and those that are least relevant to the categorization of any item will be the least important.
3. At the top end of the scale of attribute importance there may be some attributes which are so important as to be necessary for category membership. For example HAS GILLS may be treated as a necessary attribute of FISH.
4. Attributes are organized in such a way that at least some of them are represented as particular sets of values on particular dimensions. Obvious examples would be dimensions such as COLOR and SIZE. Smith et al. (1988) make a similar proposal in their concept modification model. For example the fact that Pets are domesticated may be represented as a dimension (or frame-slot) labelled HABITAT taking a particular value [DOMESTIC].
5. There is a "consistency" checking procedure which can tell whether different values for a dimension are incompatible with other values for that dimension, and also with other attributes in the prototype. This procedure will make use of the higher level attribute relations - for instance to determine that if a bird has a value of [20 kilos] for the dimension WEIGHT, it is unlikely to take the values [FLIES] for the dimension Means of Locomotion.

So far the model bears strong similarity to the prototype model, with some additional refinements in terms of higher order relations between attributes, necessity of the most highly weighted attributes, and consistency among attribute values.
A conjunctive concept is then represented semantically by a composite prototype which is formed from the two constituent concept prototypes according to the following rules (a genealogical metaphor of attribute inheritance will be used - with the constituent concepts as mother and father, and the conjunctive concept as offspring):

6. Initially the "genotype" composite prototype is formed as the union of the sets of attributes from both "parent" concepts. Thus initially the concept PET FISH will have all the attributes of both PET and FISH prototypes.

This new set of attributes is then modified, and assigned importance values according to the following rules:

7. Where an attribute is maximally important for at least one parent concept (i.e. a Necessary attribute), it will also be maximally important for the offspring, regardless of the importance value for the other parent. For example if HAS GILLS is necessary for FISH, then it will also be necessary for PET FISH. This places a Necessity Constraint on attribute inheritance.

8. For other attributes, the importance of each attribute for the offspring is determined as a monotonic positive function of importance for each parent - for example a simple average of the two constituent concept importance values could be taken. To take the PET FISH case again, supposing that LOVABLE is fairly important for PETS and irrelevant for FISH, it will be of intermediate importance for PET FISH.

9. Those attributes with low importance as determined in (8) above are dropped from the representation (or filtered out in the process of constructing it). Thus, if LOVABLE is now of relatively low importance, a subject may simply exclude it from the prototype for PET FISH.

10. A consistency checking procedure is applied to the new set of attributes. Where there are incompatible attributes, a choice has to be made to delete certain attributes. This provides a Consistency Constraint.

- Where the conflict of attributes is between a necessary attribute of one parent and a non-necessary attribute of the other, the necessary attribute will always win. This has the effect of predicting that any attribute which is considered to be "impossible" for one parent - that is any attribute which conflicts with a necessary attribute of the other parent - will also be impossible for the offspring. For example if PETS typically breathe air, but this is inconsistent with living underwater, which itself is necessary for the concept FISH, then breathing air will not be possible for PET FISH.
- Where the conflict is between two necessary attributes, then the process of constructing the prototype fails, and the conclusion is drawn that the conjunction is an empty set - a "logical impossibility". So, if asked to describe a FISH that is also a BIRD, subjects may say that such a creature is not possible since FISH must have GILLS, while BIRDS must have LUNGS. If the subject is pressed to continue into the realm of science fiction, or if the linguistic context is supportive of a non-literal interpretation, then one or other of the necessary attributes will be deleted. Hence one can arrive at interpretations for phrases such as

Robot Cats, or Stone Lions (see Franks, 1990), or speak more loosely of WHALES as FISH that are also MAMMALS, or perhaps (very loosely) of PENGUINS as BIRDS that are also FISH.

- Where the conflict involves two non-necessary attributes, then the choice of which to delete will depend on their relative importance, on the overall consistency that can be achieved with respect to the other inherited attributes, and on the context in which the phrase is being used.

It should be noted that, as stated, the model is far from a complete theory of conceptual modification. In particular, a great deal more specification is needed for the consistency checking process - how is information about the incompatibility of attributes represented and accessed during this stage? It is also worth noting however, that the proposed model could be applied to the conjunction of well-defined concepts with a core of common element defining features, with the desired results. The necessity constraint would ensure that all defining features of each concept remain criterial for the conjunctive concept, and the consistency constraint would ensure the correct identification of non-overlapping sets. Well-defined concepts would therefore require no different treatment in the model, and would behave intersectively.

Evidence for the model

To test the model, some recent experiments (Hampton, 1987) obtained attribute listings for the same pairs of categories studied in the conjunction experiments described earlier. Each attribute was then rated by subjects for its importance in defining each of the parent concepts, and for its importance in defining the conjunctive concept (on an intuitive scale ranging from "necessarily true of all possible examples of the concept" through "very important" to "not usually true" and "necessarily false of all possible examples of the concept"). The resulting data were then analyzed to test the predictions of the model.

In brief, the pattern of attribute inheritance shown by the subjects' judgments of importance supported the model well. Importance for the conjunction was predictable as a linear weighted average of importance for each constituent (see (8) in the model above), with Multiple R for the regression equations averaging around 0.8. As a result there were a number of attributes which although important for a parent concept were judged relatively unimportant for the conjunction (i.e. were not inherited). There was also clear evidence that with very few exceptions, attributes that were considered either necessary, or impossible, for either parent concept were also considered necessary or impossible respectively for the conjunction. To demonstrate the effects of consistency, an additional analysis was run on the conjunction PETS WHICH ARE BIRDS, in which each attribute that was distinctively true of one category or the other, was rated for consistency with the attributes of the other category. For example, distinctive attributes of BIRDS were singing and migrating. The extent to which each of these attributes (among others) was consistent with the attributes of being a PET was rated by subjects, and then entered into a regression predicting importance for the conjunction. Obviously singing is much more consistent with being a pet than is migrating. Indeed it was found that degree of consistency predicted additional variance in the conjunctive importance measure, over and above that predicted by importance for the two constituent concepts.

Overextension and Dominance

How does the model explain the pattern of typicality and membership ratings obtained in the categorization studies of conjunction? I shall consider three aspects - the "guppy effect" of double overextension, the more pervasive phenomenon of the overextension due to compensatory averaging of constituent typicalities, and finally the dominance relation between pairs of categories.

The start of the whole project was the claim by Osherson and Smith that GUPPIES are better PET FISH than they are either PETS or FISH alone. The Composite Prototype model accounts for this phenomenon by showing how attributes of the parent concepts fail to get inherited as a result of the necessity constraint and the consistency constraint. The typical pet fish just doesn't have many of the attributes normally associated with pets or fish alone. In more detail, consider a few of the attributes of pets and fish and how they are inherited by the conjunction:

PET	FISH	PET FISH	GUPPY
HABITAT:			
Domestic	Sea, lake, river	Domestic	Domestic
EDIBILITY:			
Not eaten	Often eaten	Not eaten	Not eaten
BREATHING:			

With lungs	With gills	With gills	With gills
LIMBS:			
Four legs	Fins	Fins	Fins

Attribute overlap:		
	Guppy - Pet	= 2
	Guppy - Fish	= 2
	Guppy - Pet Fish	= 4

Pet fish inherit habitat and not being eaten from the pet concept, because the attribute values for PET are more criterial for PETS than are the attribute values of FISH criterial for being a FISH. On the other hand, the possession of gills and fins is more criterial for fish-hood than is the possession of lungs and legs for pet-hood. Hence the composite prototype for PET FISH is a hybrid - having some attributes and lacking others of each parent. Instances like guppies and goldfish will obviously have more attributes in common with the pet fish prototype than they will with either pets or fish alone - hence they are considered more typical as pet fish.

As regards overextension through compensation, this phenomenon is a clear prediction of the intensional approach. So long as the necessary attributes are not together sufficient to determine category membership, but the definition also relies on a similarity metric defined over "characteristic" attributes (i.e. the prototype view), then intensional combination of attributes will tend to lead to inconsistency in the treatment of conjunctions. Strict set intersection requires (as in Zadeh's minimum rule, but not in his product rule) that if membership falls below criterion for either concept, then degree of membership in the other can have no positive influence at all on the degree of membership in the conjunction. If this requirement is not met, then inconsistency will result. For example Zadeh's original product rule would lead to serious underextension, since an item lying above criterion for two categories (e.g. c(A) = c(B) = 0.6) would be very likely to be excluded from the conjunction (c(AB) = 0.36). Rescaling the conjunction by lowering the membership criterion could remove the underextension, but only at the expense of increasing overextension. Thus if the criterion is set now at 0.35 for the conjunction, items with c(A) = 1.0 and c(B) = .4 would have to be included in the conjunction.

The dominance effect was the observation that some categories exert a much stronger influence on conjunctions than do others. SPORTS was dominant with respect to GAMES, and BIRDS with respect to PETS. The model outlined above can account for dominance very naturally, if it is assumed that some concepts have a greater number of salient and important attributes than others. In such a case, the concept with the higher number of important attributes would tend to contribute a greater proportion of the attributes that constitute the final composite prototype for the conjunction - because of the necessity constraint, and because of the average importance rule. Consider the case of two concepts X and Y, with features as below:

Concept X	Concept Y
f1	f1
f2	f4
f3	f5
	f6
	f7

Suppose the simplest case - that the composite concept is created without any consistency problems arising, so that the composite prototype for the concept `X that is a Y' has all the features of both concepts - that is f1 through f7.

Now the degree to which an instance is typical of the conjunction will depend on the number of features f1 through f7 that it possesses. Knowing its typicality as an X will tell us about the degree to which it possesses 3 of those 7. On the other hand knowing its typicality as a Y will tell us about its possession of 5 of the 7. It follows that typicality of items for concept Y will be a better predictor of typicality in the conjunction than will typicality in concept X, other things being equal. In other words, the more semantic specification a concept has, the more dominant a part it will play in the final composite.

There is a second possible mechanism within the model for accounting for dominance. If one concept has stronger higher order links between attributes than the other, then the conflict resolution process would favor such a concept simply because one could not remove one attribute from the composite without affecting the others to which it is linked. Dominance should therefore depend both on the relative number of important attributes, and the degree of internal coherence that results from inter-attribute links.

The attribute rating study (Hampton, 1987) found good evidence for the first account, and some indirect evidence for the second. The categories identified as dominant in the experiments on categorization (Hampton, 1988b) did indeed have a significantly higher average importance across all attributes, and a significantly higher number of important attributes than those that were non-dominant. In addition, the regression weight in the equation predicting importance for the conjunction from importance for each constituent also reflected the dominance of the categories. The importance of an attribute for the dominant category was a better predictor of conjunctive importance than was importance for the non-dominant category. Since this analysis applies to individual attributes, it could not be explained simply in terms of the relative numbers of important attributes for each concept. This second effect could reflect the coherence effect, although it can only be taken as indirect evidence at present.

Non-compositional effects

So far the story has been of a relatively simple mechanism, by which the only attributes found in representations of conjunctive concepts have been derived through inheritance from one or other parent. The picture is however not quite so simple. In the attribute inheritance study, there were a small number of emergent properties that subjects listed as important for the conjunction, but which were absent in the parent concepts. For example, PET BIRDS talk and live in cages - both attributes that are conspicuously absent in pets and birds when considered as separate categories. One way for such attributes to appear is through the retrieval of additional information from memory. The use of the noun phrase PET BIRD when first encountered may indeed set off the process of creating a new composite prototype. (Consider the novel concepts of PET SCORPION or GLASS BICYCLE.) But it is also possible that the phrase will identify a concept that is already available in memory - either as a result of an earlier encounter with the concept, or even as a result of direct experience with a number of instances of the conjunctive category. Someone when faced with the concept Black Politician may retrieve one or two familiar instances, and then base a prototype on their perceived common attributes.

A second possible source of emergent attributes would be from the coherence constraint as appled to the creation of the composite concept. If a particular inconsistency is identified, then one way to resolve it may be to make an additional modification to the concept. Suppose, while travelling in Georgia, that you hear of a PET SKUNK for the first time. The behavior of skunks is not in every respect compatible with a home environment. The problem can be resolved however with a surgical intervention, which you may then decide to add to your composite prototype as a likely attribute.

Final remarks - testing concepts to destruction

This chapter has focussed on a specific research program aimed at showing that prototype theory can provide a solid basis from which to build an account of how the meaning of compound noun phrases is composed from the meanings of constituent parts. I have argued that the rejection of the classical approach to concept definition predicts inconsistency in categorization, and that just such inconsistencies do occur. It should of course be understood that these results do not preclude the possibility of people responding in a logically consistent way if they chose to do so. Rather than reflecting a non-normativity in behavior, I would argue that the results reflect the flexibility inherent in the interface between language and the world. Natural language uses phrases such as "A is a kind of B" or "X is an A which is not a B", which appear in their syntactic form to entail class inclusion or set intersection and complementation. However because meanings are combined intensionally, and because concepts have prototype structure, the result is quite frequently non-Boolean. There are many other circumstances in which we operate in an intensional mode. The so-called "conjunction fallacy" demonstrated by Tversky and Kahneman (1983) is one example. The moral is that semantics has to be studied through analysis of the intensional meaning of terms. Attempts to build extensional model-theoretic accounts in which word meanings are tied to the sets of objects to which they refer in a model of a possible world will fail to capture the interesting interactions that occur in people's use of words in combination.

As a postscript, I will describe a pilot study in which subjects were asked to combine the uncombinable. The idea of the study was to exaggerate the conflict resolution part of the model by requiring subjects to imagine objects that combined two highly different categories, with no overlap in normal experience. There should then be plenty of opportunity for emergent attributes to be produced as subjects struggled to find a consistent composite prototype to represent the new type of object.

For example, subjects were asked to describe attributes for the following objects:

- A teapot that is a computer
- A bicycle that is a stove
- A fruit that is a kind of furniture
- A fish that is also a Vehicle

Subjects were encouraged to used drawings to help illustrate the typical appearance of such an object.

Apart from being too idiosyncratic for easy analysis, the study did provide some points of interest. For example, there was clear evidence of conflict resolution (something for which there was a pressing need in most instances!) achieved through the invention of new attributes, not seen in either parent alone. Thus the need for the FISH-VEHICLE to be directable at will led some subjects to postulate control wires implanted in the fish's brain, while others had the fish conditioned to respond to particular signals. One problem with FRUIT-FURNITURE was that fruit decays, while furniture should be relatively durable. The resolution was achieved by either replacing the furniture once a week (inheriting the attribute from FRUIT) or else by developing strains of fruit that were stable for very long periods (taking an attribute of FURNITURE). Others had the fruit still attached to a vine, so that it could be prevented from decay by still being alive. Again the process of maturation of the fruit would be extremely slow.

The impossible objects study brings out the similarity between the process of concept combination and the construction of scenarios, or mental models. People take one of the objects and try to adapt it to reflect as many of the attributes of the other concept as possible. Changes are introduced, and the consequences of those changes are then propagated through the object to see what else would be entailed. A fish that is a vehicle would have to be able to carry people or things (a highly important attribute of vehicles). Hence some kind of compartment or attachment must be added. Some subjects put a saddle on the back of the fish, others slung a basket underneath like a hot-air balloon would have, while still others surgically implanted a pressurized compartment within the fish. One can see that as each change is implemented, the subject has to think through the consequences, and make further adaptations. Different people will end up with quite different answers to the problem. Such processes are obviously an extreme form of concept modification, likely to be seen only in creative design and planning decisions (see for example Barsalou, Usher & Sewell, 1985). In less exaggerated form, however, they appear to be at work in everyday language understanding, whenever we hear a novel combination of words, and have to construct a new concept of what is being referred to.

Conclusion

In this chapter I have tried to argue for the interactive nature of conceptual combination. The result of considering two concepts in conjunction can only be properly explained with reference to the intensional semantic contents of those concepts, together with higher-order mechanisms for determining coherence and conflict amongst the attributes of the composite prototype. More generally, the results on intransitivity of categorization, and inconsistencies in the treatment of conjunction, disjunction and negation all point to the common conclusion that the way in which we use noun terms in common parlance only approximates to the corresponding interpretation of such constructions in set logic. To the extent that our concepts are defined around prototypes, with a corresponding fuzziness about just what a term should refer to, the same fuzziness carries over into our interpretation of logical connectives like relative clause constructions, or the phrase `A or B'. Most critically, the demonstration of the range of effects described here provides strong evidence for the lack of `core' common element definitions for a considerable number of concepts. The widespread occurence of non-normative effects of the kind that Osherson and Smith (1981) originally identified as false predictions of the prototype model, must be seen as the best kind of support for that model. I have tried to show how the basic prototype model has to be extended in order to take account of the more complex range of effects thrown up by the study of conceptual combinations.

References

Armstrong S.L., Gleitman, L.R., & Gleitman, H. (1983). What some concepts might not be. Cognition, 13, 263-308.
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.
Barsalou, L.W., & Billman, D. (1989) Systematicity in concepts. Unpublished manuscript. Georgia Institute of Technology.
Barsalou, L.W., Usher, J.A., & Sewell, D.R. (1985). Schema-based Planning of Events. Paper presented at the Annual Convention of the Psychonomic Society, Boston, November.
Bruner J.S., Goodnow, J.J., & Austin,G.A. (1956). A study of thinking. New York: Wiley.
Cantor, N., Smith, E.E., French, R., & Mezzich, J. (1980). Psychiatric Diagnosis as Prototype Categorization. Abnormal Psychology, 89, 181-193.
Chater, N., Lyon, K., & Myers, T. (1990). Why are Conjunctive Categories Overextended? Journal of Experimental Psychology: Learning, Memory & Cognition, 16, No.3. (in press).
Cohen, B. & Murphy, G.L. (1984). Models of concepts. Cognitive Science, 8, 27-58.
Franks, B. (1989). Concept Combination: Towards an Account of Privatives. In G.Dunbar, B.Franks, & T.Myers, (eds.), Papers from the 1989 Edinburgh Round Table on the Mental Lexicon. Edinburgh: Edinburgh University, Centre for Cognitive Science.
Hampton, J.A. (1979). Polymorphous Concepts in Semantic Memory. Journal of Verbal Learning and Verbal Behavior, 18, 441-461.
Hampton, J.A. (1981). An Investigation of the Nature of Abstract Concepts. Memory and Cognition, 9, 149-156.
Hampton, J.A. (1982). A demonstration of intransitivity in natural concepts. Cognition, 12, 151-164.
Hampton, J.A. (1987). Inheritance of attributes in natural concept conjunctions. Memory and Cognition, 15, 55-71.
Hampton, J.A. (1988a). Disjunction of Natural Concepts. Memory & Cognition, 16, 579-591.
Hampton, J.A. (1988b). Overextension of conjunctive concepts: Evidence for a Unitary Model of Concept Typicality and Class Inclusion. Journal of Experimental Psychology: Learning, Memory and Cognition, 14, 12-32.
Hampton, J.A. (1989). Concepts and Correct Thinking. Mind and Language, 4, 35-42.
Hampton, J.A. (1990). Conceptual Combination: Conjunction and Negation of Natural Concepts. Unpublished manuscript. City University, London.
Hampton, J.A., & Gardiner, M.M. (1983). Measures of Internal Category Structure: a correlational analysis of normative data. British Journal of Psychology, 74, 491-516.
Keil, F.C., & Batterman, N. (1984). A Characteristic-to-Defining Shift in the Development of Word Meaning. Journal of Verbal Learning and Verbal Behavior, 23, 221-236.
Kempton, W. (1978). Category grading and taxonomic relations: A mug is a sort of cup. American Ethnologist, 5, 44-65.
Lakoff, G. (1987). Women, Fire and Dangerous Things. Chicago: University of Chicago Press.
Levi, J. (1978). The syntax and semantics of complex nominals. New York: Academic Press.
McCloskey, M., & Glucksberg, S. (1978). Natural categories: Well-defined or fuzzy sets? Memory and Cognition, 6, 462-472.
McNamara, T.P., & Sternberg, R. (1983). Mental models of word meaning. Journal of Verbal Learning and Verbal Behavior, 22, 449-474.
Medin, D.L., & Shoben, E.J. (1988). Context and Structure in Conceptual Combination. Cognitive Psychology, 20, 158-190.
Miller, G.A., & Johnson-Laird, P.N. (1976). Language and Perception. Cambridge, MA: Harvard University Press.
Murphy, G.L. (1982). Cue Validity and Levels of Categorization. Psychological Bulletin, 91, 174-177.
Murphy, G.L. (1988). Comprehending Complex Concepts. 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. (1977). Integration of fuzzy logical information. Journal of Experimental Psychology: Human Perception and Performance, 3, 565-575.
Osherson, D.N., & Smith, E.E. (1981). On the adequacy of prototype theory as a theory of concepts. Cognition, 11, 35-58.
Osherson, D.N., & Smith, E.E. (1982). Gradedness and conceptual conjunction. Cognition, 12, 299-318.

Randall, R.A. (1976). How tall is a taxonomic tree? Some evidence for dwarfism. American Ethnologist, 3, 543-553.
Rosch, E., & Mervis, C.B. (1975). Family resemblances: studies in the internal structure of categories. Cognitive Psychology, 7, 573-605.
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, 337-361.
Smith, E.E., Osherson, D.N., Rips, L.J., & Keane, M. (1988). Combining Prototypes: A Selective Modification Model. Cognitive Science, 12, 485-527.
Smith, E.E., Shoben, E.J., & Rips, L.J. (1974) Structure and process in semantic Memory: A featural model fpr semantic decisions. Psychological Review, 81, 214-241.
Tversky, A. (1977) Features of similarity. Psychological Review, 84, 327-352.
Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. Psychological Review, 90, 293-315.
Wittgenstein, L. (1953). Philosophical Investigations. New York: MacMillan.
Zadeh, L. (1965) Fuzzy sets. Information and control, 8, 338-353.
Zadeh, L. (1982) A note on prototype theory and fuzzy sets. Cognition, 12, 291-297.