Conceptual combination

James Hampton City University London

Chapter to appear in: *Knowledge, Concepts and Categories* Eds: K. Lamberts and D. Shanks, UCL Press 1996

Final Revision: January 1996

Human language is compositional. When we construct a sentence in order to make an utterance, we do so by stringing together a series of words from our language in a particular order. The resulting sentence derives its meaning from three sources - the meaning of the individual words that have been selected (in particular the meanings of the "substantive" words such as nouns, verbs and adjectives), the syntactic structure of the sentence which places those words in a logical framework of syntactical roles including the subject, main verb and object, and other parts of the sentence, and finally the context and manner (including intonation and gesture) in which the sentence is uttered, which will determine what the speaker of the sentence actually intends to convey by its utterance on this particular occasion.

Assuming that a large part of our knowledge of the world is structured around the same concepts that underlie the meanings of words, then compositionality must apply equally to knowledge representation.¹ If someone "knows" something then we can describe this knowledge as a belief that some proposition is true. If I know that the sea is rough when there is a storm, then we might assume that this knowledge is represented by a proposition expressed in some symbolic "language" in the mind which maps more or less directly onto our natural spoken language. Corresponding to the words "sea", "rough", and "storm" it is proposed that we have corresponding concepts of [sea], [roughness] and [storm], so that propositional knowledge can be represented as a language-like construction in some symbolic form. This proposal is effectively the Language-Of-Thought hypothesis proposed by Fodor (1983). It is assumed that a person has a repertoire of available concepts and ways of combining those concepts into higher-order concepts and into propositions. This system allows us to represent a wide range of beliefs, and to construct elaborate representations of considerable complexity.

The process of conceptual combination is at the heart of knowledge representation, in that it asks the basic question - how is the meaning of a complex noun phrase related to the meanings of its component parts? This basic question can also be taken as asking how complex knowledge representations are constructed from simpler concepts. Interest in conceptual combination has arisen within cognitive psychology recently because of research conducted in the 1970s and 1980s into the more fundamental question of how the meanings of individual nouns, verbs and adjectives and the concepts that they represent are represented in a person's memory. One popular theory of conceptual representation, developed by Rosch (1975, 1978) and Mervis (Rosch & Mervis, 1975), suggested that the concepts which constitute word meanings for nouns like "bird" or "chair" are represented in the mind by prototypes. The Prototype Theory of concepts proposed that rather than concepts being represented by an explicit definition which could be used to clearly differentiate when an instance was or was not an example of the concept (as assumed by most linguistic theories of semantics at that time, e.g. Katz & Fodor, 1964), concepts are instead represented by an "ideal" or "average" prototype, and whether or not an instance is an example of the concept depends on the similarity of that instance to the prototype for the concept.

¹. There will of course be kinds of knowledge which do not involve language-like symbolic representations such as motor skills like a tennis serve or golf swing, such as the ability to recognize faces or musical themes, and such as episodic memory for particular events. The representation of knowledge through concepts based in language is just one form of knowledge representation, and one whose interface with other forms of memory representation is far from understood (Barsalou, 1993).

Hampton: Conceptual Combination Page 3

This "fuzzy" approach to defining the truth conditions of sentences such as "X is a bird" raised the problem of determining the semantics of more complex sentences involving logical combinations of concepts - as Osherson and Smith (1981) were first to point out. If the truth conditions of "X is a bird" can not be determined by a simple set of defining rules, then how is the truth of sentences such as "X is a pet bird" or "X is either a bird or a mammal" to be determined? Osherson and Smith argued that a major advantage of the so-called "classical" approach to concepts, by which each word had a clear-cut definition, was that the problem of assigning truth conditions to complex phrases in the language could be mapped in a direct way onto expressions in standard logic. For example the phrase "X is a pet bird" is true (by the standard classical treatment) if both "X is a pet" and "X is a bird" are true. In other words the truth of the phrase is a simple logical conjunction of the truth of the two constituent categorizations. However, if the truth of these constituent categorizations is a matter of degree (as proposed by Rosch and others) then we cannot rely on simple conjunction to provide us with the truth conditions for the complex phrase.

The reason for this difficulty comes from the assumption that the truth of categorization statements can be <u>graded</u>.² Some categorizations (for example "Dogs are animals") are more true than others (for example "Bacteria are animals"). If two of these fuzzy truth values are to be combined logically, then a new calculus is needed to determine how the truth of a logical combination is related to the truth of the constituent parts.

The problem of conceptual combination is not only a problem for the prototype theory of concepts.³ As Rips (1995) has pointed out, the problem is made no more tractable if the simple notion of a prototype is replaced by the currently favoured notion of a schema or "mini-theory" as the representational format for concepts (Murphy & Medin, 1985). The Theory View of concepts is briefly that people represent concepts through a deeper understanding of the causal connections between their observed characteristics. Rather than representing the concept of BIRD as a list of commonly observed characteristics (*flies, has wings, has two legs, builds nests in trees, has feathers*), the Theory view proposes that people's concept of BIRD is embedded in a wider set of interlocking theoretical structures corresponding to naive (and possibly fallacious) theories of animal biology, mechanics etc. Thus people may understand why birds have wings (to enable them to fly), why they fly (to escape predators and find food), why they nest in trees (because flying enables them to get into the trees in the first place), and so forth. Each characteristic is linked to others with explanatory, causal or goal-directed links.

While the Theory view undoubtedly has much to recommend it as a more powerful system for representing conceptual knowledge, when it comes to the problem of conceptual combinations it is less obvious how a set of syntactic rules for combining mini-theories could be formulated. The more powerful the representational medium employed, then the more each example appears to need its own special set of rules for

 $^{^2}$. The notion of truth here is one of psychological acceptability - how willing are people to accept the statement as true. This notion should not be confused with a metaphysical or ontological notion of truth - of what is actually the case in the world.

³. Space does not permit a full description of different theories of concepts here. For a fuller treatment see Hampton (1996c).

combination, and the more background knowledge appears to be required. Viewed from a philosophical point of view, it has been argued (Fodor, 1994) that none of the existing psychological models of concepts are adequate for giving a proper account of compositionality (and hence conceptual combination).

Apart from prototype and theory-based views of concepts, the other main view of conceptual representation is the Exemplar view. This view holds that concepts are represented in memory by their most common exemplars. Novel instances are then classified by their similarity to the remembered exemplars. This way of modelling concept representations has had best success in accounting for the results of classification learning experiments in which participants have to learn how to sort artificially constructed shapes into different categories (Medin & Schaffer, 1978). As a theory of knowledge representation, the exemplar theory would imply that we represent knowledge through particular memories of individual episodes (see e.g. Brooks, 1987). This proposal clearly has face validity -- there is little doubt that much of our knowledge is heavily dependent on a range of individual remembered experiences. However it can not work as a complete account, since an exemplar only becomes relevant to a situation when it is analysed. One can only categorize novel instances on the basis of their similarity to remembered exemplars if there is some means of determining similarity of the relevant kind. But this determination of similarity itself presupposes a deeper level of knowledge representation, since not just any kind of similarity may be used. If concepts are represented by exemplars, then conceptual combinations would also have to be represented by exemplars. "Pet birds" would be represented by remembered instances of actual pet birds. As we will see below, such a model accounts well for knowledge of familiar conceptual combinations. There are facts that we know about Pet Birds (for example that they can talk) which could not be derived from our knowledge of either Pets or Birds alone, and so must be based on knowledge of exemplars. The exemplar view of concepts would have little or nothing to say however about how we are able to construct novel unfamiliar conceptual combinations, in the case where we have no remembered instances to recall to mind.

Types of conceptual combination

Conceptual combinations can be broadly divided into three types. First there are combinations of concepts which appear to be broadly intersective in a logical sense. Some adjective noun phrases are intersective - thus a "Red apple" is both red and an apple. One can suppose (although empirical evidence might be needed to confirm this - see Hampton 1996a) that for most people the category of red apples is simply the overlap of the categories of red things and of apples. Where relative clauses are used for combining concepts the result also appears to be broadly intersective - thus "pets which are also fish" can be understood as the overlap of the categories of pets and fish.

A second type of combination involves an adjective plus noun, or a noun plus noun combination where the first word is used to <u>modify</u> the second. The first word is known as the <u>modifier</u>, and the second as the <u>head</u> noun. For example a "corporate lawyer" is a member of the category of lawyers who is concerned with the law appertaining to corporations, while a "criminal lawyer" is a member of the same category of lawyers who works on criminal cases. The difference from the first type of combination should be evident. Criminal lawyers in this sense are not the intersection of the categories of criminals and lawyers (although the intersective reading is also possible and produces an ambiguity in the meaning of the phrase).

The third type is best seen as a subtype of the modifier-head combination. When combinations become familiar and idiomatic they are known as lexical compounds and are typically marked in English by placing stress on the first (modifier) noun. Thus, to use an example from Kamp and Partee (1989), a "brick factory" is a factory that makes bricks (a lexical compound) whereas a "brick *factory*" is a factory made from bricks (a novel modifier-head combination). Many lexical compounds have become highly idiomatised so that it is no longer possible to predict their meaning on the basis of their constituent nouns. Rips (1995) cites the example of a "bull ring" which one feels ought to (but does not) mean the ring in a bull's nose. At the extreme, noun-noun compounds can be lexicalised to the point of effectively forming a single lexical unit in the mental lexicon - as in "railway" or "lipstick". Apart from these clearly lexicalised examples, it may not always be simple to distinguish modifier-head constructions from true compounds. It may in fact be better to see both types as falling on a continuum of novelty. Novel combinations such as "butter police" are clearly modifier-head constructions, whereas highly familiar combinations like "patchwork" are clearly compounds. Presumably all compounds started life as novel modifier-head forms, and then gradually became lexicalised through frequent use. A process of chaining can also occur so that the original meaning can become more and more distant. The best known example of this is Fillmore's example of "topless district" - where the object which is topless is not the district, nor the bars within it, nor the waitresses that serve in them, but the clothing that they wear.

Research on conceptual combination has focussed largely on the first two of these types. It can be safely assumed that the understanding of lexicalised compounds involves no new conceptual combination, and so the study of how their meaning relates to the meanings of their constituents is more of interest to students of historical linguistics than to psychologists. Another use of conceptual combinations which will not be explored here is one where the meaning of a conceptual combination is very heavily dependent on the pragmatic context in which the phrase is used. It was stated at the start of the chapter that when people are conversing a major source of meaning comes from the surrounding context, both situational and linguistic. Downing (1977) gives examples of novel nounnoun combinations whose interpretation would be impossible without knowing the prior context. Thus the use of "apple-juice seat" to refer to a place at a table in a restaurant where an order of apple juice had been placed is a case where the communicative context supplies the solution to interpreting the combination. It is probably the case that most novel combinations in actual speech are created in such circumstances, and their novelty may not even be noted by the hearer. Gerrig and Murphy (1992) explored this process in a series of experiments in which a novel combination such as "trumpet olive" was rendered meaningful by a prior story context such as the following:

Peter and Susan watched the skilled old woman with great awe. The woman was carving figures out of stale olives. She had been plying this craft for twenty years. Her work was remarkably detailed. Peter and Susan could see a miniature trumpet appearing out of a new one. Peter said "Would you like a trumpet olive?"

Their study illustrates very neatly how powerful an influence the discourse or story context can have on the interpretation of a novel phrase. By manipulating the story context, one could cause people to interpret "trumpet olive" to mean a wide variety of things -- the olives reserved for the trumpet players for their interval snack at a concert, a kind of olive which is stuffed into a trumpet to keep it in good working order -- there is no limit to the creative possibilities. This creativity is possible because the phrase is being used primarily as a referential phrase (Donnellan, 1966) to point to some object that has already been established within the story. The context is doing all the work of identifying the relevant object or concept, and the linguistic phrase has merely to point to it in a reasonably distinctive manner. As a result there is little conceptual work for the hearer to do.

For the most part, the study of how people combine concepts has maintained a neutral context so that the effect of manipulating conceptual content can be measured. We will see that people do in fact have the ability to select plausible interpretations for novel combinations in the absence of explicit context, and a number of studies have addressed this capability.

In the remainder of this chapter the first two types of conceptual combination -intersective combinations and modifier-head combinations -- are considered, and relevant research and theoretical models are reviewed.

Intersective combination

When Rosch and others proposed similarity to prototype as the basis for categorization, it was proposed that one could treat membership in a category as a matter of degree. A ready-made way to model the logic of such gradedness was available in a system known as fuzzy logic. Developed by Zadeh (1965), fuzzy logic is a form of set logic in which the truth value of set membership statements can take a continuous value between 1 (true) and 0 (false). For example the statement "John is a tall man" can be taken as being more or less true, depending on how close John is in height to the prototypical tall man. Zadeh's logic proposed an extension of the standard set logical operators of conjunction, disjunction and negation to include continuous truth values. For example fuzzy conjunction was defined (for simple cases) with a minimum rule, whereby the truth of the conjunction of two expressions was the minimum of the two individual truth values. Thus "John is both tall and handsome" would be only as true as the least true of the two individual statements "John is tall" and "John is handsome". The minimum rule reduces to the standard classical conjunctive operator in the case that truth values are restricted to just one or zero. (A second rule for conjunction was proposed by Zadeh for more complex statements - the rule involved multiplying the two truth values to find the truth of the conjunction, so that the truth of "John is both tall and handsome" would always be less than or equal to the truth of each individual statement. This second set of rules is a direct corollary of the axioms for combining the probabilities of independent events in probability theory.)

Although fuzzy logic had some success in accounting for intuitions about the conjunction of <u>unrelated</u> statements (Oden, 1977), it soon became clear following a key article by Osherson and Smith (1981) that not only the minimum rule, but in fact any rule that takes as input solely the truth value of the two constituent statements is doomed to failure. The reason is that the function of fuzzy conjunction appears to depend also on

the degree of relatedness of the two categories being combined. Consider first the case where the two concepts are highly related - as when one is a subset of the other.

(1) X is a kind of poultry(2) X is a kind of bird(3) X is a kind of poultry AND a kind of bird

The degree to which (3) is true of some new found specimen creature is likely to be determined largely by the degree to which (1) is true. This is because if (1) is true then (2) is also always true.

Now consider Osherson and Smith's example of two concepts that are generally contradictory

(4) X is striped(5) X is an apple(6) X is a striped apple

An actual apple that had stripes would be a better example of the conjunction "striped apple" than it would be of either of the two constituent concepts - simply because apples don't generally have stripes and striped things are generally not apples. Thus (6) should have a higher truth value than either (4) or (5).

Both of these examples run counter to the principle of the minimum rule for conjunction. A good example of poultry is likely to be a poor example of a bird (since poultry are atypical birds), but still to be a good example of the conjunction "bird which is poultry". Similarly the apple with stripes is a better example of the conjunction than of the constituent concepts - in direct contradiction of the minimum rule.

Smith, Osherson, Rips & Keane (1988) took the striped apple example from Osherson & Smith (1981) and collected empirical evidence that it is indeed true that a picture of a brown apple (for example) is considered more typical of the conjunctive concept "brown apple" than it is of the simple concept "apple". The almost trivial nature of this demonstration highlights the failing of the fuzzy logic approach to cope with predictions of typicality in complex concepts. In their Selective Modification Model, Smith et al. (1988) adopted a frame formalism (Minsky, 1975) for representing a concept such as "apple". In a frame representation, the property information which defines what apples are like (in effect one's knowledge of apples) is represented by a series of Attribute Slots such as [COLOUR], [SHAPE], [SOURCE], and so forth. Each of these slots can take different values as Slot Fillers. For example the representation of "apple" in Figure 1 has the following slot:

[COLOUR] = Red (10), Green (8), Yellow (5), Brown (2)

INSERT FIGURE 1 ABOUT HERE

The weights (w1, w2 etc.) represent the relative importance of each slot in determining similarity of any instance to the concept. The numbers in parentheses for each slot filler reflect the relative frequency (or typicality) for the different values in the population of

apples in the person's experience (they are called *votes* in the Selective Modification Model). Hence red apples are represented as being more typical than green, yellow or brown apples, simply by matching a value with a higher number of votes. (The model was developed in the USA where red apples are predominant.)

When the concept "apple" is modified with a colour such as "brown", to produce the conceptual combination "brown apple", the Selective Modification Model proposes that the representation of "apple" undergoes a <u>selective modification</u> of the [COLOUR] slot of the frame, so that whereas for "apple" [COLOUR] can take values *Red*, *Green*, *Yellow* or *Brown*, for "brown apple", all the votes for colour are transferred to the value *Brown*. In addition the model proposes that the weight of the [COLOUR] slot itself is increased in the frame. Since typicality is determined by similarity to the concept representation (in this they followed the approach adopted by Smith, Shoben & Rips (1974) and by Prototype Theory) the model proposed that typicality was determined by a weighted function of matching and non-matching features (Tversky, 1977) computed across the frame representation. Hence by increasing the weight for the [COLOUR] slot, the influence of colour on typicality is increased in the case of "brown apple" relative to "apple".

The Selective Modification Model predicts correctly that brown apples are more typical of "brown apple" than of "apple", whereas a regular apple is more typical of "apple" than of "brown apple". However it should be noted that the model is overspecified for even this relatively straightforward prediction. It incorporates two mechanisms (transfer of votes to *Brown*, and increase in the weight of the [COLOUR] slot) whereas the first alone would produce the same prediction. It is also highly problematic as a model of categorization in complex categories (as opposed to the determination of typicality).

To understand the problem, a brief digression is necessary to explain the difference between graded membership as measured by <u>typicality</u> and graded membership as it relates to <u>categorization</u>. Rosch's prototype model simply proposed that concept categories are graded, and this gradation could be seen both in variation in typicality (for example sparrows are typical examples of the category of birds, whereas ostriches are atypical), and also in the existence of borderline cases, where items are so atypical as to render their category membership questionable (as in whether carpets are furniture, or a lift is a vehicle - see Hampton & Gardiner, 1983, and McCloskey & Glucksberg, 1978 for copious examples in British and American English respectively). The prototype model proposed that both variation in typicality and the occurrence of borderline cases (whose gradedness is reflected in their probability of categorization across a sample of people or occasions), are driven by the same underlying dimension - namely similarity to the concept prototype. However they do <u>not</u> have to follow similarity according to the same function.⁴ It is clear that both sparrows and ostriches are birds (all 44 subjects agreed on a positive categorization in the Hampton & Gardiner,

⁴. It is also quite likely that similarity itself may be computed using different weights for the different dimensions in the case where Typicality is being judged as opposed to the case where Category membership is being decided. In the latter case more weight may be accorded to dimensions that are most diagnostic of category membership, whereas in the former weight may be also given to dimensions that are representative of the most central examples. Allowing weights to change between tasks in an unconstrained manner however does introduce a large number of free parameters to the prototype model which could seriously undermine its predictive value as a model.

1983 study), but whereas sparrows are universally considered highly typical in England, ostriches are considered highly atypical. So typicality can vary widely even when categorization is constant. This apparent anomaly is easily resolved by assuming that categorization is derived from similarity via a threshold criterion. Once similarity is above some level, then categorization is certain, although there will continue to be variations in rated typicality (see Hampton, 1993, 1995 for a more formal explication of the prototype model). In effect categorization probability reaches its ceiling well before similarity to the prototype reaches its maximum.

Returning to the Selective Modification Model, we can see how it is primarily a model for determining typicality rather than categorization. For example, for "red apple" the weight of the colour slot is boosted and all the votes are transferred to the colour *Red*. An actual red apple is then more similar to the concept "red apple" than it is to the concept "apple", and so should be rated as such. The weight of the colour slot needs to be boosted in the model so that the colour of an object counts as much as its "appleness" in determining its typicality. A well formed brown apple - which has all the other slots of the concept fully matching - should not be considered a better example of a "red apple" than an oddly shaped and unusually small red apple.

The effect of giving high weight to the colour slot however will be to make other red fruits like strawberries or tomatoes also highly similar to the "red apple" concept, since they will capture all the red votes on the highly weighted colour slot, as well as matching many of the fruit properties. If colour is weighted to the point where it counts as much as appleness in determining similarity, then it follows that other fruits with the right colour red but low appleness should be as similar to the "red apple" concept as apples which are not red. It is extremely unlikely therefore that the Selective Modification Model would allow the concept representation for "red apple" to correctly categorize red apples from all other fruits, and indeed Smith et al. at the end of their paper qualify their model as applying only to typicality judgements, and restate the earlier arguments (Osherson and Smith, 1981, 1982) with regard to categorization - that is that categorization must depend on a well-defined core of features, rather than on similarity to prototype.

An alternative model which does attempt to account for the range of conceptual logical functions such as conjunction, disjunction and negation while maintaining a similarity-based model of conceptual categorization is to be found in my own work (Hampton, 1987, 1988a, 1988b, 1991, 1996a). The initial aim was to examine the degree to which people's conjunctive concepts do in fact closely follow the logical function of set intersection. An earlier study (Hampton, 1982, see also Kempton, 1978, Randall, 1976) had thrown up the possibility that people's semantic categorizations were only loosely modelled by set logic. For example people were willing to accept that chairs are a type of furniture and that carseats are a type of chair, but would then deny that carseats are a type of furniture. The relation between people's use of categorization statements and the mathematician's notion of class inclusion is therefore not a direct mapping.

Hampton (1988b, Experiment 1) extended this demonstration of intransitivity to subsets of categories defined by modifier-head combinations. Consider for example "school furniture" or "office furniture". The Selective Modification Model, and most other accounts of conceptual combination, take it as given that the complex concept forms a proper subset of the head noun concept. There are a number of well known

exceptions involving *privative* modifiers like "fake", "counterfeit" and "alleged", which do not follow this rule (a "fake dollar" is not a dollar), and there are also cases such as "stone lion" or "chocolate rabbit" where the modifier has the effect of changing the reference of the head noun to objects with the shape and appearance of the original object (Franks, 1989; Wisniewski, 1996). But these examples clearly do not apply to the "school furniture" example, which could be paraphrased with a simple thematic IN relation (Levi, 1978) -- "furniture that is found in schools", or possibly a USED FOR relation -- "furniture that is used for schools".

Hampton (1988b, Experiment 1) asked subjects to classify a number of objects like desks, chairs and blackboards either as being "furniture" or as being "school furniture", and found that there were a number of cases that mirrored the earlier intransitivity result - not all objects classed as school furniture were also included in the category of furniture. When asked to classify school furniture as a type of furniture, they were happy to do so. In a follow-up test, however, the respondents were asked to choose between two alternative statements:

- (1) All school furniture is also furniture
- (2) Some school furniture is not furniture.

When faced with this choice, many people chose (2), indicating that there was no inconsistency in their minds between claiming both that school furniture is a kind of furniture, and that there are kinds of school furniture which are not furniture. Once a concept has been modified, it appears that the resulting class of objects is only an <u>approximate</u> subset of the head noun class.

In subsequent experiments, Hampton (1988b) made the modifier relation more explicitly conjunctive by using a relative clause construction. Thus the phrase "sports which are also games" certainly appears *prima facie* to involve taking the head noun class ("sports") and then identifying the subset of the head noun category that are also in the class of "games". However this is not how subjects responded. Where an activity was a very typical example of one category (for example "chess" is a typical game) then it was more likely to be classified in the conjunction ("sports which are games") than in the other constituent concept ("sports"). Across a number of pairs of overlapping concepts, the general rule seemed to apply that the likelihood of being classed in the conjunction was some average of the two constituent categorization probabilities (see Huttenlocher & Hedges, 1994, for an attempt to provide a purely statistical account of this result). Furthermore three reliable but unexpected results emerged. One was a concept dominance effect such that one of the two concepts had a greater influence on the conjunction than the other. For example membership in the conjunction "sports which are games" was more heavily determined by membership of the games category than membership of the sports category -- regardless of which was taken as head noun (see Storms, De Boeck, van Mechelen, & Geeraerts, 1993 for further studies of the dominance effect). The second result was non-commutativity -- the two forms of the conjunction "sports which are games", and "games which are sports" were not identical in either typicality or membership. The third was overextension. People were more likely to overextend the conjunction to include items that were judged not to be in one or the other constituent category than they were to underextend it by omitting items that were judged

to be in both constituent categories. None of these phenomena are readily accounted for by purely statistical effects.

Overextension has since also been found in visually defined concepts such as coloured letters and cartoon faces (Hampton, 1996a). In this study, I found additional evidence against a purely statistical account. Stimuli in one part of the experiment were cartoon faces which varied in their features so that they could be classified along two dimensions. Some looked happy while others looked sad, and orthogonal to this distinction, some looked like adults while others looked like children. Subsets of stimuli were selected in the critical condition that were clear members of one category (for example all looked like children), and borderline cases of the second category, (for example they were at the borderline in terms of looking happy or sad). Categorization in the conjunction "happy children" was then measured. It was found that variation in typicality in the first category "children" (where probability of categorization was at ceiling) was correlated with categorization in the conjunction. The link between similarity (as measured by typicality variation amongst clear category members) and categorization probability in the conjunction shows that typicality in one category can compensate for borderline status in the other. No models that propose that the conjunction is categorized through application of a rule of the kind $\{X \text{ is an } AB \text{ if and } \}$ only if X is an A and X is a B} where the two consituent categorization decisions are made independently (see for example Huttenlocher & Hedges, 1994) can account for this kind of compensation.5

An interesting corollary of the overextension effect is to be found in work by Tversky and Kahneman on probability judgments. Tversky and Kahneman (1983) showed that people will on occasion consider it more likely that an individual falls in the conjunction of two classes than in one of the individual classes alone, even although this breaks a basic axiom of probability theory. For example, they told people a story about a person Linda who had been quite radical in her views when at college. When asked to rank order a set of statements in terms of their likelihood, people tended to rank "Linda is a feminist bank teller" as more probable than "Linda is a bank teller", even though there could be no situations in which the first was true but the second was false.

Tversky and Kahneman's account of their result was based on the notion of "representativeness". If the conjunctive concept appears to be a closer description of the individual than the single constituent concept, then people judge the individual more likely to belong in the conjunction. There is clearly a close parallel between overextension of conjunctions of semantic categories and this reasoning fallacy. Both the judgment of categorization (Hampton, 1988) and of probability (Tversky & Kahneman, 1983) depend on an assessment of similarity of the instance to the representation of the conjunctive category, and a corresponding tendency to overextend the boundaries of the conjunction, or overestimate the probability that an individual belongs to it.

The model which Hampton (1987) developed to account for these results followed the same approach as the Selective Modification Model (Smith et al., 1988) in

⁵. The results were actually more complex. While typicality as an adult or a child influenced categorization in the conjunction for stimuli with borderline emotional expressions, the reverse pattern was not found. Typicality as a happy or sad face did not compensate for borderline values of the adult/child dimension. See Hampton (1996a) for details and further discussion of this result.

Hampton: Conceptual Combination Page 12

rejecting fuzzy logic formulations of conceptual combination in favour of representations involving intensional or attribute information (an approach advocated by Cohen and Murphy, 1984). Unlike the Selective Modification Model however, the Composite Prototype Model (Hampton, 1987) was intended to account for both typicality and membership in conjunctively defined categories. The model proposed that when forming a concept such as "pets that are also birds", people take their prototype representations of "pet" and "bird", and combine the prototypes into a composite to represent the conjunction. Membership and typicality in the conjunction are then determined by similarity to the Composite Prototype, with a criterion threshold being used to determine membership as in standard Prototype Theory. The way that this works can again best be represented using a frame representation with an attribute-value structure. Each attribute has certain values for one concept and certain values for the other. The composite prototype will then inherit its own attribute values from one or other constituent parent according to certain principles. To take an example, the [LOCATION] slot for "pet" has the value In the home, while the same slot for "bird" has the value In the wild. "Pet bird" inherits the value from "pet" rather than "bird" - pet birds live in the home and not in the wild. On another attribute however the composite would inherit the value from "bird" for example for the slot [COVERING] the most common "pet" value would be *Furry*, but "pet bird" would take the value Feathered which it inherits from "bird".

Having proposed that inheritance of slot values from the constituent concepts is not complete, but that conflicting values may compete to occupy the slot for the composite concept, it is then necessary to determine the principles by which the inherited value is chosen. To investigate this process, Hampton (1987) obtained attribute property listings from different groups of respondents for each constituent of a combination like "birds which are pets" and for both ways of expressing the conjunction ("birds which are pets", and "pets which are birds"). The attributes listed were then all combined into a single list, which was given to further groups of subjects who made a judgment about how important each attribute was for categorizing an instance as a member either of a constituent class ("birds", or "pets") or of the conjunctive class ("pets which are birds", or *vice versa*). This procedure was followed for six different conjunctions. The importance ratings were then analysed to see just how the process of attribute value inheritance works.

It appears that the choice of which value is inherited for any attribute is determined initially by the importance of the attribute for each constituent. Regression equations showed that a high proportion of variance in the importance of an attribute for the conjunction could be predicted from the importance for the respective constituents. (There was also a concept dominance effect which appeared to mirror closely the concept dominance effect found for typicality judgments by Hampton, 1988.) In addition to this overall effect, there was also an effect of the <u>centrality</u> or <u>mutability</u> of a property for the constituent concept. Where a property was considered a central "necessary" feature of a constituent (the top end of the "importance" scale), then it was also considered central for the conjunctive concept. Similarly where a property was considered "impossible" for a constituent (the bottom anchor point on the importance scale), it was also considered impossible for the conjunction.

In a final study, Hampton (1987) demonstrated that the degree to which an attribute value was inherited by the conjunction was also determined by the amount to

which it conflicted with all the attributes of the opposing constituent. For example one attribute of "birds" was *migrates*. However this was judged to be in conflict with many of the attributes of being a pet, and so was not judged to apply to the conjunction of "pets which are birds". Degree of conflict entered significantly into the regression equation, predicting additional variance after the two constituent importance ratings had been entered.

The result of the interaction of the two concepts in determining which attribute values are inherited is that the composite prototype representation of the conjunction will be a hybrid of the two constituents - bearing resemblance in some respects to one constituent and in other respects to the other. It is therefore easy to see how an instance class could be more similar to the composite prototype than to either of the constituent classes (the original problem identified by Osherson and Smith, 1981). Furthermore if membership of the conjunctive class is determined by how similar an instance is to the composite prototype, then it is relatively straightforward to demonstrate that one would therefore expect a pattern of overextension and/or underextension, rather than the application of a "logical" intersection rule. (The relative proportion of each would depend on where the threshold criterion for similarity to the composite prototype is placed. Empirically it appears that people place the criterion quite low, so that the conjunctive class ends up with a greater number of instances than expected on the basis of constituent class membership.)

The study of attribute inheritance (Hampton, 1987) also threw up an interesting number of attributes which were considered true of the conjunction but not true of either constituent. These were termed emergent properties and included the properties of "pet birds" that they live in cages and can talk - aspects not considered true of "pets" or "birds" alone. These properties are particularly interesting since they appear to break the basic compositionality of conceptual combination (see for example Murphy, 1988, 1990). It appears that a major source of emergent properties is simply knowledge of the world or "extensional feedback" as Hampton (1987) described it. We use the words "pet" and "bird" to identify the conjunctive class of pet birds, retrieve some familiar instances, and then proceed to describe them. We can not expect any model of conceptual combination to account directly for such effects, as they clearly relate to information that is obtained from another source - namely familiarity with the class of objects in the world (Rips, 1995, refers to this as breaking the No Peeking Principle). The exemplar approach (representing a combination through its exemplars) is the only account that will explain such emergent attributes. A study by Medin and Shoben (1988) highlights the importance of this level of exemplar knowledge. They showed that typicality of items in conjunctive categories could vary as a function of the kinds of exemplar found in those categories. For instance for the combination "small spoon", metal spoons were judged to be more typical than wooden spoons, whereas for the combination "large spoon", the reverse relation was found. Without some experiential knowledge of actual spoons, their sizes and their materials, this kind of effect would be impossible to explain. In effect then, one cannot ask theories of conceptual combination to predict such effects, nor should one take failure to predict such effects as a criticism of any theory.

A second source of emergence however may be considered fair game for models of conceptual combination. These are emergent features that appear in novel or unfamiliar concept conjunctions. A well-known example of this notion is the concept of

Hampton: Conceptual Combination Page 14

a "beach bicycle". Given the opportunity to reflect on this concept (expanding the modifier-head form into a more explicit conjunction such as "a bicycle which is used for riding on the beach"), people typically consider the problem of riding on a beach -- that the wheels would sink into the sand -- and so suppose that a beach bicycle would be equipped with particularly wide tyres. Since there is no information about wide tyres in either the bicycle or the beach concepts alone, and no knowledge of actual instances to retrieve (indeed many people may never have tried riding on a beach), then this emergent feature must have been the result of an appeal to background "theory", and in particular to naive understanding of the mechanics of bicycles and beaches. (There can be no exemplar-based account of this result if a person has never before encountered an exemplar of the concept.)

The best available source of evidence for theory-based emergent features currently lies in the domain of social stereotype categories. Kunda, Miller & Clare (1990) investigated conjunctions such as "A Harvard-educated carpenter", and found (in a relatively informal way) that subjects created quite complex and detailed stories in order to explain how an individual might end up in such a conjunction (see also Hastie, Schroeder & Weber, 1990). Hampton (1996b) reviewed the frequency of emergent attributes across a series of studies, and confirmed that theory-based emergence is more easily found in social categories than object or activity categories. Emergent attributes in novel or unfamiliar concept conjunctions are supportive evidence for Murphy & Medin's view that concepts are deeply embedded in theories of the world (see also Murphy, 1993). People have to use more abstract theoretical understanding to resolve the conflicting attribute values in constructing the composite prototype.

In another domain, Hampson (Casselden & Hampson, 1990; Hampson, 1990) investigated the combination of incongruent personality traits (following the pioneering work of Asch (1946). Although there were interactions among traits in the pattern of inheritance, there was in fact little or no evidence for emergent properties in her studies.

Evidence for the importance of background knowledge on social cognition came from another study (Hampton & Dillane, 1993) in which conjunctions of social categories were judged from different points of view. Manipulation of point of view was originally introduced by Barsalou and Sewell (described in Barsalou, 1987). Barsalou and Sewell showed that typicality in a category varied systematically when people were asked to make the judgments from the point of view of others - be they stereotypical groups like housewives versus farmers, or familiar groups like faculty members versus students. In the study by Hampton and Dillane, people were asked to generate and rate attributes about two contrasting social categories -- for example an Oxford Graduate and a Factory Worker -- and to take the point of view of either an Oxford graduate or a factory worker while completing the task. Other groups were asked to perform the same tasks for the conjunction "An Oxford graduate who is a factory worker", or its converse -- again taking one or the other point of view. Manipulation of the point of view had a huge effect on the attributes that were considered true of the conjunction. The most common pattern to emerge was one of conflict. When adopting one point of view, people judged people in the conjunction to have the same (mostly unpleasant) properties as the other category. Thus from an Oxford graduate's point of view, the graduate factory worker was mostly like other factory workers, whereas from the factory worker's point of

view he was most like other Oxford graduates (the scenarios were all about men in this study).

In a second unpublished experiment (Hampton & Oren), in addition to point of view, the sex of the person described was also manipulated. Each pair of concepts contained one male and one female stereotype (for example "A fighter pilot who is also a child minder"). Half the respondents were asked to consider males while the other half considered females. Thus for example one set of groups judged females who were either fighter pilots or child minders or both, from the point of view of either a female fighter pilot or from the point of view of a female child minder. Other groups performed the same task taking a male's point of view about males. (The sex of the subject making the judgments was not manipulated but varied randomly across groups). The results were particularly interesting. First, there were a considerable number of emergent attributes, as one would expect from the Kunda et al. (1990) study. What was especially striking was the difference between the male and female points of view. Consider for example the combination "A car mechanic who reads romantic fiction". When taking a female's point of view about a female, both points of view produced positive emergent attributes such as Ambitious, Broad-minded, Clever, Easygoing, Charming and Caring. When taking a male's point of view about a male, the emergent attributes were far more negative Dissatisfied, Elusive, Reliable, Lonely, and Soppy. Overall the male stereotypes also generated far more patterns of conflict -- whereby someone in the conjunction had none of the positive attributes of one's own group. The female stereotypes in contrast were more likely to show a pattern of complementarity. Even though the two stereotypes were strongly inconsistent, female points of view about females tended to consider someone in the conjunction as having the positive elements of both stereotypes.

The implications of these results for social psychology have yet to be explored. It is clear that social categories provide a very rich source of materials for exploring the processes of conceptual combination. Almost any two social categories can in principle overlap - there are very few ontological constraints of the kind found in artifact and natural kind categories.

Hampton (1996b) also reported a study in which a technique was developed for forcing people to come up with emergent attributes for object categories. The technique was to take two noun classes that are quite disjoint types of thing and ask people to imagine an object that falls in both classes. Imagine for example a "Fish that is also a vehicle". Some respondents failed to find a solution that truly matched the task requirements (drawing for example a car with a fish's head and tail modelled at either end), but many were able to find ingenious solutions to the problem. Successful solutions frequently involved stretching of concept categories (whales and dolphins were treated as fish - presumably because their greater intelligence could be used to good effect), and many emergent features were produced, only loosely based on existing vehicles (for example a control mechanism that involved electrodes implanted in the unfortunate animal's brain). Theory-based reasoning in conceptual combination can therefore be elicited within the domain of object categories, although many of the combinations caused considerable difficulty for the subjects in the study. (See Ward, 1994; 1995, for examples of how even when urged to come up with novel solutions to problems such as designing a new creature, people will usually stay very close to existing familiar concepts.)

Modifier-Head combinations

The study of "intersective" combinations has revealed a range of noncompositional effects, including overextension of conjunctive categories and emergent properties that are true of a conjunction but not of either constituent. These effects are however relatively marginal in the sense that people can still appreciate that intersective combinations <u>approximate</u> to the logical function of set intersection. When we turn to consider modifier-head constructions, then non-compositional effects come to the fore.

Murphy (1988) views conceptual combination as a highly creative process in which the end result of any combination may involve the introduction of deep theoretical knowledge and a wide range of facts about the world. Murphy's data depend largely on familiar modifier-head combinations. For example Murphy (1988) showed how for a range of noun-noun combinations there are properties which emerge as being true of the combination which are not true of either of the constituent parts. Casual things are not pulled over the head, and neither are shirts, but yet casual shirts are pulled over the head.

The number of examples of "emergent" properties of this kind is large, and demands an explanation in any theory of conceptual combination. The most direct explanation involves what was referred to above as "extensional feedback" (Hampton, 1988) or the No Peeking Principle (Rips, 1995). Many of Murphy's emergent features could never be predicted or derived from the constituents, because they are contingent facts about the world. The term "casual shirt" identifies a known category of objects in the world, and it is through examination of this extensional set ("peeking" at the world outside) that the emergent properties are identified. Thus it is that we know that pet birds sometimes talk (Hampton, 1987), that stop signs are hexagonal or that boiled eggs are hard whereas boiled potatos are soft.

As Rips (1995) rightly argues, such examples do not rule out the possibility of a computational account of conceptual combination in the way Murphy seems to suggest. While it is clear that these examples implicate background knowledge in the understanding of phrases such as "boiled eggs", it can be argued that there is little conceptual work being done when understanding such phrases. In effect the concept of a boiled egg already exists as a subset of the concept egg (along with scrambled, poached and fried eggs), and we simply retrieve what we know about the complex concept from our memory.

In a second experiment, Murphy (1988) considered how the meaning of adjectives changes as a function of the noun with which it is combined in a range of adjective noun combinations. For example an adjective like "long" or "new" or "open" can take a wide range of different meanings when modifying different nouns. Thus an "open year" was a flexible one, "open people" were revealing of their thoughts, an "open world" was one full of opportunity, while an "open hand" was a card hand dealt face up. This demonstration points to a second source of information that people use in interpreting conceptual combinations - the polysemy of many words (and particularly high frequency adjectives). An adjective like "open" has a highly extendable meaning, based on some very abstract schema of openness. In different domains the word's meaning has been taken and given meanings that are very specific to that domain. While this process can probably be used creatively, it is also likely that such specialised meanings are simply learnt as alternative meanings of the term by anyone learning the language. (In fact such extended uses can not always be translated with the same adjective in another language).

In order to assess the processes of conceptual combination "in the raw" it is necessary to consider novel combinations - combinations that are unfamiliar to the subject both as objects and as linguistic expressions. The influence of background knowledge and theories on the processes of conceptual combination can then be assessed. For noun-noun combinations this can be achieved by combining sets of nouns in novel ways and asking subjects for interpretations. Gagne and Shoben (1993) conducted a study of this kind. They took the corpus of naturally occurring compounds collected by Levi (1978) and selected a set of head and modifier nouns. All possible combinations were then considered and those that had a reasonably clear immediate interpretation were selected. Typically, novel noun-noun or adjective-noun combinations were interpreted using one of about 12 semantic relations, which they termed thematic relations. For example a "Mountain magazine" was interpreted with the ABOUT relation as "A magazine about mountains". In a subsequent reaction time study, Gagne and Shoben carefully selected combinations in which the common interpretation involved a thematic relation which was either (a) high frequency for both head and modifier, (b) high frequency for the head but not the modifier, or (c) high frequency for the modifier but not the head. Comprehension times for the combinations indicated that the critical variable for determining speed of comprehension was the frequency of the relation for the modifier. Provided the relation was that normally expected for the modifier, the phrase was understood equally fast, regardless of whether the head noun commonly appeared in that relation or not. On the other hand, where the relation was uncommon for the modifier, it was slower to comprehend -- again regardless of the fact that it was a common relation for the head noun.

This result is consistent with the left to right ordering of modifier and head in English. Assuming that the modifier is read first, then its common thematic relation could be retrieved and applied to the head noun. Only in the case where the resulting combination fails to make sense would the common thematic relation for the head noun be retrieved as an alternative interpretation. The result is not however a comfortable one for the Smith et al. (1988) Selective Modification model. Although that model makes no processing assumptions, it would seem to imply a computational process model in which the head noun frame is set up in working memory first, and the modifier is retrieved second. The model would therefore be more consistent with a stronger influence of the head noun relation frequency over the comprehension time than with the obtained result.

Gagne and Shoben (1993) examined the influence of "preferred" thematic relations on processing speeds. Wisniewski (1996) however took a different approach to studying modifier-head combinations. He selected a set of nouns from different taxonomic categories, such as Animals, Artifacts and Substances, and then combined them pairwise into eight different kinds of combination (he omitted Substance-Substance pairs). Subjects in his experiment were asked to paraphrase the concept combination, in order to explicate its meaning. The different paraphrases were then examined to see what kinds of relation were generated - as a function of the taxonomic categories of the head and modifier nouns. Wisniewski found that contrary to previous assumptions (Levi, 1978; Gagne & Shoben, 1993), modifier-head constructions are not always interpreted through the generation of a thematic relation whereby the modifier modifies a slot of the head noun (as in the "Mountain magazine" example). In a high proportion of cases -- and particularly in the Animal-Animal case -- the interpretation offered for the combination involved what Wisniewski termed "property mapping". In property mapping a highly salient property of the modifier was taken and mapped directly onto the head noun. To understand the difference between "slot filling" (as he termed the use of a thematic relation) and property mapping, consider a concept combination like "Tiger hound". By generating a likely thematic relation, one might paraphrase this as "Hound used to hunt tigers". In Wisniewski's terms, the representation of "Hound" has a slot [USE FOR HUNTING ---] which could be filled with values such as *Foxes, Deer, Boar* and so forth. As "Tiger" can plausibly fill this slot, so the interpretation is seen as plausible. The alternative, property mapping, strategy is to take a salient property of tigers - for example their stripes - and apply this to overwrite the corresponding property of hounds in order to come up with the paraphrase "Hound with stripes".

Property mapping requires both that the modifier noun has a highly distinctive property, and also that the alignment of the two concepts is sufficiently close that the property can easily be mapped across. This would predict that the more similar two nouns (within the same ontological category), then the more likely that property mapping will occur, and Wisniewski (1996) was able to confirm this prediction. In a further study by Davidson and Stevenson (1995), it was reported that combinations composed of two natural kinds were more likely to produce property mapping than those composed of two artifacts or a mix of an artifact and a natural kind. Similarity however did not have a major effect in this experiment, in that pairs of natural kinds and pairs of artifacts were judged equally similar. It may therefore be the case that there is an additional constraint on when property mapping is prevented by the fact that the head noun property which is to be replaced is also very central to the meaning of the head noun. For most pairs of natural kinds this constraint may be absent.

Wisniewski (1996) also noted two further strategies for interpretation of novel combinations which occurred in his data. First, when two objects were highly similar, people could form a <u>hybrid</u> of the two. Thus a "horse cow" could be some hybrid creature that was half horse and half cow. This strategy could be seen as an extension of property mapping, although it also implies a more symmetrical relation between the two concepts than is found in straight property mapping (where the combination is still considered to be in the head noun concept class). The second strategy was the construal of one or other of the two nouns to refer <u>metonymically</u> to a part or other related aspect of the concept itself (Nunberg, 1979). For example a "tiger chair" was interpreted as a chair covered with tiger skin (use of the noun to refer to a part), and an "artist collector" was interpreted as someone who collects the <u>works</u> of artists. This interpretation strategy uses a more general linguistic strategy by which complex referential phrases can be shortened. So we can speak of "a Monet" and refer not to a member of Claude Monet's family, but to one of his paintings. Subjects using the construal strategy were looking for plausible complex phrases that could have been shortened in a similar way.

A final plausible strategy for interpreting novel combinations is through analogy with already familiar combinations. Recognizing a good analogy commonly requires property mapping and alignment of features. In this case however, if the novel combination <u>reminds</u> us of a familiar combination with the same head or modifier noun,

then we might adopt the same thematic relation to provide an interpretation. For example we could interpret "tiger hound" by analogy with "fox hound", or alternatively by analogy with "tiger shark" depending on the availability of such more familiar combinations. Existing familiar combinations would then be recruited to aid in the search for an interpretation. Although linguistic intuition suggests that established combinations can be extended to novel ones in just this way, there has as yet been no direct experimental test of this notion.

Conclusions

We have seen that conceptual combination covers two rather distinct sorts of psychological process. The first involves the way in which conjunctions of prototype concepts are formed. The conceptual combination in this instance remains approximately intersective, although there appear to be systematic non-intersective effects, particularly when social stereotype categories are involved. Emergent attributes can be found arising from knowledge of the instances in the conjunctive category, and more rarely arising from the theory of the domain in which the concepts are embedded.

The second form of conceptual combination is perhaps specific to particular languages, and involves the interpretation of novel noun-noun modifier head constructions. (Dutch and German also allow noun-noun constructions to be generated in a free way, whereas French requires semantic marking of the thematic relation with a preposition such as \dot{a} , de or en.) It involves a search for an interpretation which can be likened to a problem solving or constraint satisfaction process. Two different strategies for interpretation have been identified by Wisniewski as Slot filling and Property mapping, and according to Gagne and Shoben, the speed of comprehension in slot filling depends to a large extent on how commonly the modifier enters into the thematic relation by which it is modifying the head noun.

Research on conceptual combination is still at a relatively early stage of development. However considerable progress has been made in the last fifteen years towards understanding how concepts as expressed through word meanings interact when placed together in phrases. It is clear that no simple set of compositional rules will be adequate to the task of a full description of the meaning of a complex concept. However by separating out the effects of familiarity, and by explicating the strategies that can be employed in combining concepts, we are now closer to the goal of explaining how the meaning of phrases relates to that of their constituent parts. It has become clear that the process of achieving this goal must depend heavily on an understanding of knowledge representation, and of knowledge revision. The strong conclusion that can be drawn from the work of Medin, Murphy and others is the degree to which many of the phenomena of conceptual combination can only be properly understood in the context of a fuller theory of the knowledge and naive theories of the world in which concepts are embedded.

Author notes

I would like to acknowledge the help of a number of colleagues who contributed in various ways to the development of the ideas and research described in this chapter. In particular I should like to thank Lawrence Barsalou, Nick Braisby, Herbert Clark, Margaret Dillane, Danièle Dubois, Bradley Franks, Janellen Huttenlocher, Frank Keil, Iven van Mechelen, Douglas Medin, Helen Moss, Gregory Murphy, Laura Oren, Lance Rips, Edward Shoben, Gert Storms, Jean-Pierre Thibaut, Thomas Ward and Edward Wisniewski.

References

- Asch, S.E. (1946). Forming impressions of personality. <u>Journal of Abnormal and Social</u> <u>Psychology</u>, <u>41</u>, 258-290.
- Barsalou, L.W. (1987). The instability of graded structure: implications for the nature of concepts. Chapter in U.Neisser (Ed.), <u>Concepts and conceptual development:</u> <u>Ecological and intellectual bases of categories</u>. Cambridge: Cambridge University Press.
- Barsalou, L.W. (1993). Structure, flexibility, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A.C.Collins, S.E.Gathercole, & M.A.Conway (Eds.), <u>Theories of Memory</u>. Hillsdale, NJ: Lawrence Erlbaum.
- Brooks, L.R. (1987). Nonanalytic cognition. In U. Neisser (Ed.), <u>Concepts and</u> <u>conceptual development: Ecological and intellectual factors in categorization</u>, (pp. 141-174). Cambridge: Cambridge University Press.
- Casselden, P.A., & Hampson, S.E. (1990). Forming impressions from incongruent traits. Journal of Personality and Social Psychology
- Cohen, B. & Murphy, G.L. (1984). Models of concepts. Cognitive Science, 8, 27-58.
- Davidson, O., & Stevenson, R. (1995, September). Interpretation of noun-noun compounds. Paper presented to Conference of the British Psychological Society, Cognitive Section, Bristol.
- Donellan, K. (1966) Reference and definite description. <u>Philosophical Review</u>, <u>75</u>, 281-304.
- Fodor, J.A. (1983). The Modularity of Mind. Cambridge, MA: MIT Press.
- Fodor, J.A. (1994). Concepts a pot-boiler. Cognition, 50, 95-113.
- Franks, B. (1989). Concept Combination: Towards an Account of Privatives. In
 G.Dunbar, B.Franks, & T.Myers, (eds.), Papers from the 1989 Edinburgh Round
 <u>Table on the Mental Lexicon</u>. Edinburgh: Edinburgh University, Centre for
 Cognitive Science.
- Gagne, C., & Shoben, E.J. (1993). The influence of relational information on interpreting nonpredicating combinations. <u>Paper presented at the 65th Annual Meeting of the Midwestern Psychological Association</u>, Chicago, April.
- Hampson, S.E. (1990). Reconciling inconsistent information: Impressions of personality from combinations of traits. <u>European Journal of Personality</u>, <u>4</u>, 157-172.
- Hampton, J.A. (1982). A Demonstration of Intransitivity in Natural Categories. <u>Cognition</u>, <u>12</u>, 151-164.
- Hampton, J.A. (1987). Inheritance of attributes in natural concept conjunctions. <u>Memory</u> <u>and Cognition</u>, <u>15</u>, 55-71.
- Hampton, J.A. (1988a). Disjunction of Natural Concepts. <u>Memory & Cognition</u>, <u>16</u>, 579-591.
- Hampton, J.A. (1988b). Overextension of conjunctive concepts: Evidence for a Unitary Model of Concept Typicality and Class Inclusion. <u>Journal of Experimental</u> <u>Psychology: Learning, Memory and Cognition</u>, <u>14</u>, 12-32.
- Hampton, J.A. (1991). The combination of prototype concepts. <u>In</u> P.Schwanenflugel (Ed.), <u>The psychology of word meanings</u> (pp. 91-116). Hillsdale, NJ: Erlbaum.
- Hampton, J.A. (1993). Prototype models of concept representation. In I.van Mechelen, J.A. Hampton, R.S. Michalski, & P. Theuns (Eds.), <u>Categories and concepts:</u>

<u>Theoretical views and inductive data analysis</u> (pp. 67-95). London: Academic Press.

- Hampton, J.A. (1995). Testing Prototype Theory of Concepts. Journal of Memory and Language, 34, 686-708.
- Hampton, J.A. (1996a). Conjunctions of Visually Based Categories: Overextension and Compensation. <u>Journal of Experimental Psychology: Learning, Memory and</u> <u>Cognition</u>, <u>22</u>, 378-396.
- Hampton, J.A. (1996b). Emergent attributes in conceptual combinations. In T.B.Ward, S.M.Smith, & J.Viad, (Eds.), <u>Conceptual Structures and Processes: Emergence</u> <u>Discovery and Change.</u> Washington DC: American Psychological Association.
- Hampton, J.A. (1996c). Psychological representation of concepts. <u>In</u> M.A.Conway & S.E.Gathercole (Eds.) <u>Cognitive Models of Memory</u>. London: UCL Press.
- Hampton, J.A., & Dillane, M. (1993). Taking a point of view on Attribute Inheritance. Paper presented to the 34th Annual Convention of the Psychonomic Society, Washington DC, November.
- Hampton, J.A., & Gardiner, M.M. (1983). Measures of Internal Category Structure: a correlational analysis of normative data. <u>British Journal of Psychology</u>, <u>74</u>, 491-516.
- Hastie, R., Schroeder, C., & Weber, R. (1990). Creating complex social conjunction categories from simple categories. <u>Bulletin of the Psychonomic Society</u>, <u>28</u>, 242-247.
- Huttenlocher, J., & Hedges, L.V. (1994). Combining graded categories: membership and typicality. <u>Psychological Review</u>, <u>101</u>, 157-165.
- Kamp. H. & Partee, B. (1989). Prototype theory and compositionality (Draft #2000). Unpublished manuscript.
- Katz, J.J., & Fodor, J.A. (1963). The structure of a semantic theory. Language, <u>39</u>, 170-210.
- Kempton, W. (1978). Category grading and taxonomic relations: A mug is a sort of cup. American Ethnologist, 5, 44-65.
- Kunda, Z., Miller, D.T., & Clare, T. (1990). Combining social concepts: the role of causal reasoning. <u>Cognitive Science</u>, 14, 551-578.
- Levi, J. (1978). <u>The syntax and semantics of complex nominals.</u> New York: Academic Press.
- Medin, D.L., & Schaffer, M.M. (1978). Context theory of classification learning. <u>Psychological Review</u>, <u>85</u>, 207-238.
- Medin, D.L., & Shoben, E.J. (1988). Context and Structure in Conceptual Combination. Cognitive Psychology, 20, 158-190.
- McCloskey, M., & Glucksberg, S. (1978). Natural categories: Well-defined or fuzzy sets? <u>Memory and Cognition</u>, <u>6</u>, 462-472.
- Minsky, M. (1975). A framework for representing knowledge. In P.H.Winston (Ed.), <u>The</u> psychology of computer vision. New York: McGraw-Hill.
- Murphy, G.L. (1988) Comprehending Complex Concepts. <u>Cognitive Science</u>, <u>12</u>, 529-562.
- Murphy, G.L. (1990) Noun phrase interpretation and conceptual combination. Journal of <u>Memory and Language</u>, 29, 259-288.

- Murphy, G.L. (1993). Theories and concept formation. <u>In</u> I.van Mechelen, J.A. Hampton, R.S. Michalski, & P. Theuns (Eds), <u>Categories and concepts:</u> Theoretical views and inductive data analysis. London: Academic Press.
- Murphy, G.L., & Medin, D.L. (1985) The role of theories in conceptual coherence. <u>Psychological Review</u>, 92, 289-316.
- Nunberg, G. (1979). The non-uniqueness of semantic solutions. Polysemy. Linguistics & Philosophy, 3, 143-184.
- Oden, G.C. (1977). Integration of fuzzy logical information. <u>Journal of Experimental</u> <u>Psychology: Human Perception and Performance, 3</u>, 565-575.
- Osherson, D.N., & Smith, E.E. (1981) On the adequacy of prototype theory as a theory of concepts. <u>Cognition</u>, 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. <u>American Ethnologist, 3</u>, 543-553.
- Rips, L.J. (1995). The current status of research on concept combination. <u>Mind and Language</u>, <u>10</u>, 72-104.
- Rosch, E. (1975). Cognitive representations of semantic categories. Journal of Experimental Psychology: General, 104, 192-232.
- Rosch, E. (1978). Principles of categorization. <u>In</u> E. Rosch & B. Lloyd (Eds.) <u>Cognition</u> <u>and Categorization</u>. Hillsdale NJ: Erlbaum.
- Rosch, E., & Mervis, C.B. (1975). Family resemblances: studies in the internal structure of categories. <u>Cognitive Psychology</u>, <u>7</u>, 573-605.
- Shoben, E.J. (1996). Conceptual combinations. In T.B.Ward, S.M.Smith, & J.Viad, (Eds.), <u>Conceptual Structures and Processes: Emergence Discovery and Change.</u> Washington DC: American Psychological Association.
- 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. <u>Cognitive Science</u>, 12, 485-527.
- Storms, G., De Boeck, P., van Mechelen, I., & Geeraerts, D. (1993). Dominance and Non-Commutativity Effects on Concept Conjunctions: Extensional or Intensional Basis? <u>Memory and Cognition</u>, 21, 752-762.
- Thagard, P. (1983) Conceptual combination: a frame-based theory. Paper presented at the Society for Philosophy and Psychology, Wellesley, June, 1983.
- Tversky, A. (1977) Features of similarity. <u>Psychological Review</u>, <u>84</u>, 327-352.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. <u>Psychological Review</u>, <u>90</u>, 293-315.
- Ward, T. B. (1994). Structured imagination: The role of category structure in exemplar generation. <u>Cognitive Psychology</u>, 27, 1-40.
- Ward, T. B. (1995). What's old about new ideas? In S. M. Smith, T.B. Ward, & R. A. Finke, (Eds.) <u>The Creative Cognition Approach</u> (pp. 157-178). Cambridge, MA: MIT Press.
- Wisniewski, E.J. (1996). Construal and Similarity in Conceptual Combination. In T.B.Ward, S.M.Smith, & J.Viad, (Eds.), <u>Conceptual Structures and Processes:</u>

Hampton: Conceptual Combination Page 24

Emergence Discovery and Change. Washington DC: American Psychological Association.

Zadeh, L. (1965) Fuzzy sets. Information and control, 8, 338-353.

Figure 1

Frame representation of the concept APPLE (loosely adapted from Smith et al., 1988).

Concept name: "APPLE"

Weight	Slot	Fillers
w1	COLOUR	<i>Red</i> (10), <i>Green</i> (8), <i>Yellow</i> (5), <i>Brown</i> (2)
w2	SHAPE	Round (10), Stalk at one end (5)
w3	CONTAINS	Seeds (15), Juice (10)
w4	SOURCE	On apple trees (20)
w5	USED FOR	Eating whole (10), Apple pies (5), Baking (3)

---- etc.