

Testing a computational model of categorisation and category combination: Identifying disease categories and new disease combinations

Fintan Costello (fintan@compapp.dcu.ie)

School of Computer Applications, Dublin City University,
Glasnevin, Dublin 9, Ireland.

Abstract

The diagnostic evidence model gives a computational account of how people classify items in single categories and in category combinations (complex categories formed by combining two or more single categories). This model sets out to explain generativity in category combination (the fact that people can classify items in new category combinations even if they have never seen any examples of those combinations). The model also aims to explain context effects such as overextension in category combination. In an experiment people learned to identify imaginary diseases from artificially-constructed patient descriptions, and then classified new patient descriptions into combinations of those disease categories. The model accurately predicted people's classification scores for patient descriptions in these disease combinations, requiring no free parameters to fit the experimental data. The experiment showed that both generativity and overextension can occur in combinations of artificially-constructed disease categories, and confirmed the model's predictions about when overextension and generativity will occur.

Introduction

The ability to combine mental representations is a basic part of human cognition. For example, to understand a combined phrase such as *pet fish* we must somehow combine our representations of the constituent categories *pet* and *fish*. Category combination is generative: we can understand a new combined phrase such as *pet lobster*, even though we may never have seen an example of a *pet lobster* before. Generativity is important because it allows us to think new thoughts, understand new expressions, and respond to new situations. However, generativity poses a problem for theories of classification in which an item's membership in a category is proportional to its similarity to previously seen exemplars of that category (e.g. the Context theory; Medin & Schaffer, 1978). Since no previously-seen exemplar is available for combined categories such as *pet lobster*, membership in such a category cannot be computed by exemplar similarity (Rips, 1995). While such theories give a good account for classification in single categories, they do not extend well to category combination.

Context effects such as overextension occur reliably in category combination. These effects also pose a problem for current theories of classification.

Overextension occurs when people classify an item as a poor member of a constituent category of a combination, but as a good member of the combination as a whole; for example, when people rate goldfish as poor members of the single categories *pet* and *fish*, but as highly typical members of the conjunction *pet fish* (Hampton, 1988). Overextension shows that an item's category membership can change depending on the context in which classification occurs: being poor if the category occurs singly, but good if it occurs as part of a combination. For theories in which classification is based on fixed rules for category membership (e.g. Nosofsky, Palmeri, & McKinley, 1994), these changes in membership are difficult to explain. While such theories apply well to classification in single categories, they do not extend to category combination.

This paper describes a computational model which explains classification in both single and combined categories. This model, called the diagnostic evidence model, also explains generativity and overextension in category combination. This paper also describes an experiment investigating classification, category combination, generativity and overextension in artificial laboratory-learned categories. In this experiment people learned to identify imaginary diseases from artificially-constructed patient descriptions, and then classified new patient descriptions into combinations of those disease categories. Both generativity and overextension occurred reliably in these combinations of artificial categories. The model accurately predicted people's classification scores for patient descriptions in these disease combinations, requiring no free parameters to fit the data. The patterns of overextension and generativity in the experiment closely matched those predicted by the model.

The Diagnostic Evidence Model

The diagnostic evidence model is an extension of a model originally developed to explain how people interpret novel noun-noun phrases (Costello & Keane, 1997, 2000, 2001). The model aims to explain classification in both single categories (see Costello, 2000) and category combinations. The model assumes that people represent categories by storing sets of category members in memory. From these sets, diagnostic attributes for categories are computed: these attributes serve to identify new category members. An item's classification in a single or combined category is a function of the diagnosticity of its attributes for that

category or for the constituent categories of that combination. An item has a high classification score in a category if it has diagnostic attributes of that category. An item has a high score in a combination if it has some attributes diagnostic for one constituent of the combination, and others diagnostic for the other.

Attribute Diagnosticity

Diagnostic attributes are attributes which occur frequently in members of a category, but rarely in that category's contrast set (the set of non-members of that category). These attributes serve to identify members of a category: a new item having an attribute which is diagnostic for a category is likely to be a member of that category. Equation 1 defines the diagnosticity of an attribute x for a category C . Let K be C 's contrast set. Let j_x be 1 if an item j has attribute x , and 0 otherwise. $D(x|C|K)$, the diagnosticity of x for C relative to K , is equal to the number of members in C which have attribute x , divided by the total size of C plus the number of items in K which have x :

$$D(x|C|K) = \frac{\sum_{j \in C} j_x}{|C| + \sum_{j \in K} j_x} \quad (1)$$

If the attribute x occurs in all items in C , but no items in C 's contrast set, then x is fully diagnostic for C ($D(x|C|K) = 1$). Such an attribute is a perfect guide to membership of C : a new item having that attribute is most likely a member of C . An attribute which does not occur in all members of C , or which occurs in some members of C 's contrast set, will be less diagnostic for the category. Such an attribute will be a poorer guide to membership of C : a new item with that attribute is less certain to be a category member.

Diagnosticity changes in combination

The contrast set is important in computation of attribute diagnosticity: the fewer occurrences of an attribute in the contrast set for a category or combination, higher its diagnosticity will be. The contrast set for a single category consists of all items which are not members of that category. The contrast set for a combined category, however, consists of all items that are not members of any constituent of the combination. This change in contrast set means some attributes that are not diagnostic for a category occurring singly can be diagnostic for that category in a combination. This change in diagnosticity is the basis for overextension in category combination.

Table 1 shows 10 stored members of categories such as *pet* and *fish*, described on 4 dimensions. Computation of attribute diagnosticity can be illustrated using this set of stored category members. Consider the diagnosticity of the attribute *<found:house>* for the category *fish*. *<found:house>* occurs in 2 of the 4 members of *fish* in Table 1, and occurs 4 times in the

Table 1. Example items: 10 stored category members.

| Item | Categories | Item Attributes | | | |
|------|---------------|-----------------|--------|--------|--------|
| | | FOUND | KEPT | COLOR | PARTS |
| 1 | lobster | sea | ----- | pink | claws |
| 2 | lobster | aquarium | tank | pink | claws |
| 3 | fish goldfish | house | tank | gold | scales |
| 4 | fish guppy | house | tank | silver | skin |
| 5 | fish salmon | sea | ----- | silver | scales |
| 6 | fish shark | sea | ----- | silver | skin |
| 7 | pet spaniel | house | basket | brown | tail |
| 8 | pet pitbull | house | kennel | black | tail |
| 9 | pet bulldog | house | basket | brown | ----- |
| 10 | pet terrapin | house | tank | green | skin |

contrast set K_{fish} (the set of items which are not members of the category *fish*). The diagnosticity of *<found:house>* for *fish* is

$$D(\text{found} : \text{house} | fish | K_{fish}) = \frac{2}{4+4} = 0.25 \quad (2)$$

This attribute has a low diagnosticity for the single category *fish*: *<found:house>* does not identify members of category *fish* well. In the context of the combination *pet fish*, however, the attribute has a higher degree of diagnosticity for *fish*. $K_{petfish}$, the contrast set for the combination *pet fish*, consists of items that are members neither of *pet* nor of *fish* (items 1 and 2 only). *<found:house>* does not occur in any items in $K_{petfish}$. The diagnosticity of *<found:house>* for *fish* relative to the contrast set $K_{petfish}$ is thus

$$D(\text{found} : \text{house} | fish | K_{petfish}) = \frac{2}{4+0} = 0.5 \quad (3)$$

Attribute *<found:house>* is thus more diagnostic for *pet fish* than for *fish* alone. Given this, the diagnostic evidence model would predict overextension for the combination *pet fish*: an item such as *goldfish*, which possessed the attribute *<found:house>*, could be classified as an untypical *fish*, but as a typical *pet fish*.

A logic for evidence

Diagnostic attributes give evidence for an item's classification in a category. Items usually contain a number of different attributes, however, which may be more or less diagnostic for the category in question, or diagnostic for other categories. The diagnostic evidence model uses a continuous-valued logic to combine the diagnosticity of multiple attributes. This logic assumes continuous variables with values between 0 and 1, and uses the logical operations

$$NOT A = 1 - A \quad (4)$$

$$A AND B = AB \quad (5)$$

$$A OR B = 1 - (1 - A)(1 - B) \quad (6)$$

These equations can be justified by considering the operations *AND*, *OR*, and *NOT* for samples of independent variables. Suppose A is true in 75% of samples, and B

is true in 50% of samples. Then the probability of *NOT A* being true is 0.25 (1-0.75). The probability of *A AND B* being true is 0.375 (0.75 x 0.5): *A* is true in 75% of samples, *B* is true in 50% of those. Finally, the probability of *A OR B* being true is 0.875 (1-(1-0.75) X (1-0.5)): *A* is false in 25% of samples, *B* is false in 50% of those, and thus *A OR B* is true in 87.5% of samples.

Combining attribute diagnosticities

To compute an item's overall evidence for membership in a category, the diagnosticity of the item's attribute are combined using the equation for *OR* (Equation 6). An item *i* with a set of attributes x_1, x_2, x_3 will be a member of category *C* if $x_1 \text{ OR } x_2 \text{ OR } x_3$ is diagnostic for *C*. This is formalised in Equation 7. Let *A* be the set of attributes of item *i* and let $D(x|C|K)$ be the diagnosticity of attribute *x* for *C*. Then $E(i|C|K)$, the overall evidence for classifying item *i* as a member of *C*, is

$$E(i|C|K) = 1 - \prod_{x \in A} (1 - D(x|C|K)) \quad (7)$$

If an attribute *x* strictly defines a category *C* (occurs in all members of *C* and never occurs outside *C*), then *x* is perfectly diagnostic of *C* ($D(x|C|K) = 1$). If any item *i* possesses attribute *x*, then by Equation 7 $E(i|C|K)$ will be 1, and the item *i* will definitely be a member of category *C*. In categories which have no single defining attribute but rather a range of attributes of medium diagnosticity, Equation 7 combines evidence from different attributes in computing evidence for category membership: the more diagnostic attributes the item has, the higher its degree of membership will be. This fits with the observed family resemblance structure of natural categories (Rosch, 1978). The relationship between diagnosticity and membership is supported by Rosch & Mervis' (1975) finding that people's judgements of an item's typicality in a category rises with the number of the item's diagnostic attributes.

Diagnostic evidence in combinations

In the diagnostic evidence model, an item will be a member of a combined category if it gives evidence for membership in each constituent category in that combination: if it has some attributes diagnostic for one constituent of the category, and other attributes diagnostic for the other. In computing an item's membership in a combined category, the model uses the equation for *AND* to combine the item's evidence for membership in each constituent. An item *i* will be classified as member of a combined category $C_1 \dots C_N$ if it gives evidence for membership in C_1 *AND* evidence for membership in C_2 *AND* evidence for membership in C_3 and so on. Formally, $E(i|C_1 \dots C_N|K_1 \dots N)$, the evidence for classifying *i* as a member of $C_1 \dots C_N$, is

$$E(i|C_1 \dots C_N|K_1 \dots N) = \prod_{n=1 \dots N} E(i|C_n|K_1 \dots N) \quad (8)$$

Table 2. Classification of the item *goldfish* in single categories *pet* and *fish* and combination *pet fish*.

| Evidence for membership in | Attribute Diagnosticity | | | |
|----------------------------|-------------------------|-----------|--------------|-------------|
| | FOUND house | KEPT tank | COLOR golden | PART scales |
| <i>pet</i> singly : | 0.7 | 0.7 | 0.1 | 0.0 |
| <i>fish</i> singly: | 0.8 | 0.2 | 0.3 | 0.2 |
| <i>pet fish:</i> | | | | |
| constituent <i>pet</i> | 1 | 1.0 | 0.2 | 0.0 |
| constituent <i>fish</i> | 0.9 | 0.5 | 0.4 | 0.2 |
| <i>Pet fish</i> overall: | 0.9 | | | |

where the contrast set $K_{1 \dots N}$ is the set of items not in any category $C_1 \dots C_N$. In this equation an item *i* gives evidence for membership in each constituent of a combination if it has attributes diagnostic for each. Note that, in computing the evidence for membership in each constituent category (r.h.s. in Equation 8), the contrast set for the combination as a whole is used. In computing membership in those categories occurring singly, their single contrast sets would be used.

Table 2 illustrates the diagnostic evidence model by showing the computed membership for the item *goldfish*, which has attributes `<found:house>`, `<kept:tank>` `<color:golden>` and `<part:scales>`, in the single categories *pet* and *fish* and in the combination *pet fish*. The diagnosticity of the item's attributes for single categories and for constituents of the combination are listed in columns under those attributes. The item's membership scores in the single categories and the constituents of the combination are computed from those diagnosticities (shown in bold, to the left of those diagnosticities). At the bottom of Table 2 is the item's overall membership score in the combination (obtained by multiplying its constituent membership scores).

Explaining overextension and generativity

In the diagnostic evidence model, overextension arises if some attributes have low diagnosticity for a single category but high diagnosticity for that category in a combination. Table 2 illustrates this overextension. The item in Table 2 has a higher overall membership score in the combination *pet fish* than in the categories *pet* or *fish* presented singly, because the item's attributes are more diagnostic for the combination than for the single categories. For example, `<found:house>` has lower diagnosticity for the single category *fish*, but higher diagnosticity in context of the combination *pet fish* (it occurred often in the contrast set for the single category *fish*, but not in the contrast set for *pet fish*). The model thus predicts overextension for that item.

The diagnostic evidence model gives a generative account of category combination, in which an item can be classified in a new combination even if no previous examples of that combination have been seen. An item

Table 3. Training materials for learning diseases.

| Training Item | Item features | | | Member of Category or Combination |
|---------------|---------------|----|----|-----------------------------------|
| | D1 | D2 | D3 | |
| 1 | A | X | C | A |
| 2 | A | Y | Y | A |
| 3 | A | A | X | A |
| 4 | Y | A | Y | A |
| 5 | X | A | B | A&B |
| 6 | A | B | X | A&B |
| 7 | Z | B | B | B |
| 8 | X | B | B | B |
| 9 | Y | X | B | B |
| 10 | Z | Y | B | B |
| 11 | C | A | Y | C |
| 12 | C | X | B | C |
| 13 | C | Y | C | C |
| 14 | C | A | C | C |
| 15 | C | X | C | C |
| 16 | X | Y | C | C |

is classified in such a combination if it has diagnostic attributes for each constituent category in the combination: some attributes diagnostic for one constituent, other attributes diagnostic for the other. For example, in Table 1, there are no stored members of the combination *pet lobster*. However, an item could be classified as a good member of the combination *pet lobster* if it possessed the attributes *<has-part:claws>* (diagnostic for *lobster* in Table 1) and *<found:house>* (diagnostic for *pet*). The model thus predicts generativity for that combination. The next section describes an experiment testing the diagnostic evidence model’s predictions about classification, overextension, and generativity in category combination. This experiment uses artificial categories in the domain of disease diagnosis.

Disease Diagnosis: An Experiment

Most experiments investigating category combination examine how natural-language categories are combined. The current experiment examines category combination with artificial, laboratory-generated categories representing imaginary diseases. In this experiment, every subject was given a set of 16 patient descriptions (16 training items), each with 3 symptoms and each having a given disease or disease combination. The abstract distribution of symptoms in training items was identical for all subjects, and is shown in Table 3. The training materials used ill-defined categories: no symptom perfectly indicated any disease. In the training phase of the experiment, subjects used these training items to learn to identify diseases. In the transfer phase subjects were given new patient descriptions (transfer items) and asked, for each disease and each possible disease combination, to indicate

whether the new item was a member of that disease category or disease combination.

Method

Subjects. 19 Dublin City University undergraduates.

Materials. Each subject received a set of 16 patient-description cards (training items) with the abstract structure shown in Table 3. In these, abstract attribute A (on any dimension) is most diagnostic for category A, attribute B for category B, and C for category C. Each subject received a different set of patient descriptions, generated via a unique mapping from abstract attributes to concrete symptoms. For example, for one subject, attribute *<A>* on dimension D1 became symptom *eyes:puffy*; *<A>* on D2 became *skin:flaking*, and *<A>* on D3 became *mucles:taut*. For other subjects the attributes were mapped to different symptoms.

Procedure. In the training phase subjects spent 15 minutes learning to identify diseases by studying their 16 patient-description cards. Subjects were then shown, in random order, patient descriptions with the same symptoms as those they had learned, but sometimes with incorrect diagnoses. Subjects indicated whether diagnoses were correct and incorrect. If a subject got a diagnosis wrong, they were shown the correct answer. The transfer phase of the experiment examined subjects classification of 5 new patient descriptions (the transfer items). Table 4 gives abstract representations for these items. Each subject’s transfer items were formed by applying their attribute-to-symptom mapping to this representation. Each item was presented 6 times, each time with a different single or combined category. Subjects rated the given item as a member or non-member of the given category or combination, using a -10 to +10 rating scale, with a positive rating indicating membership and a negative rating non-membership.

Results

Analysis of subject’s performance in the training phase showed that most had no problem in classifying items. One subject got most of the training-phase test classifications wrong and was excluded from analysis. The 2nd-last column in Table 4 (‘classification probability: observed’) shows the observed probability (proportion) of subjects rating each transfer item as a member of the given combination. (For space reasons the corresponding data for single categories are not shown.) For example, the observed probability of transfer item *<ABY>* being classified in combination *A&B* was 0.5: 50% of subjects rated that item as a member of that combination.

The data in Table 4 shows that subjects responded consistently to items. For example, there were some items which had high observed classification probabilities for particular combinations, indicating that many subjects agreed that those items did belong in

Table 4. Observed and predicted classification of the 5 transfer items in 3 different category combinations.

| Item | Combination | Classification probability | | Item | Combination | Constituent | Overextension | |
|------|-------------|----------------------------|-----------|------|-------------|-------------|---------------|-----------|
| | | Observed | Predicted | | | | Observed | Predicted |
| D1 | D2 | D3 | | | | | | |
| A | B | Y | A&B | 0.50 | 0.47 | A | Yes (0.50) | Yes |
| A | B | Y | A&C | 0.11 | 0.13 | C | Yes (0.56) | Yes |
| A | B | Y | B&C | 0.11 | 0.07 | C | Yes (0.67) | Yes |
| C | Y | B | A&B | 0.06 | 0.19 | Y | Yes (0.56) | No |
| C | Y | B | A&C | 0.28 | 0.21 | X | No (0.41) | Yes |
| C | Y | B | B&C | 0.72 | 0.77 | X | Yes (0.61) | Yes |
| Y | A | C | A&B | 0.22 | 0.14 | | | |
| Y | A | C | A&C | 0.50 | 0.50 | | | |
| Y | A | C | B&C | 0.17 | 0.17 | | | |
| X | B | C | A&B | 0.28 | 0.23 | | | |
| X | B | C | A&C | 0.17 | 0.27 | | | |
| X | B | C | B&C | 0.39 | 0.42 | | | |
| X | X | B | A&B | 0.28 | 0.29 | | | |
| X | X | B | A&C | 0.11 | 0.15 | | | |
| X | X | B | B&C | 0.39 | 0.45 | | | |

those combinations. For example, 72% (0.72 observed classification probability) of participants identified item <CYB> as a member of combination *B&C*: most subjects agreed in classifying that item as a member of that combination. Conversely, a number of items had very low observed classification probabilities for particular combinations, indicating that a large proportion of subjects agreed that those items did not belong in those combinations. For example, only 6% (0.06 observed classification probability) of participants classified item <CYB> as a member of combination *A&B*. The remaining 94% of participants indicated that the item did not belong in that combination. Because each subject's patient descriptions used a unique mapping from abstract attributes to symptoms, this consistency depended only on the distribution of those symptoms in the learned categories.

Model fit. To apply the diagnostic evidence model to the experimental materials, the equations described earlier were used to compute the classification score for each of the transfer items in Table 4 in every possible single and combined disease category. The diagnosticity of each item's attributes for each category was computed from the distribution of those attributes in the training items shown in Table 3. The last column in Table 4 ('classification probability: predicted') shows computed classification scores for each item in each combined category. These computed scores were compared with the observed probability with which people classified the items as members of each combination. There was a strong correlation between the predicted and observed classification scores for the combined categories ($r=.95$, $p <.01$, $\%var=.9$). Comparing predicted and observed classification scores for items across all single and all combined categories

Table 5. The 6 item-combination-constituent triplets for which overextension occurred or was predicted.

| Item | Combination | Constituent | Overextension | |
|-------|-------------|-------------|---------------|-----------|
| | | | Observed | Predicted |
| A B Y | A&C | C | Yes (0.50) | Yes |
| C Y B | A&C | A | Yes (0.56) | Yes |
| C Y B | B&C | B | Yes (0.67) | Yes |
| Y A C | A&C | C | Yes (0.56) | No |
| XX B | A&B | A | No (0.41) | Yes |
| XX B | B&C | C | Yes (0.61) | Yes |

also showed a significant correlation ($r=.85$, $p <.01$, $\%var=.73$). No free parameters were used to fit the model's classification scores to the experimental data.

Generativity. The generativity of category combination was examined by comparing the classification of transfer items in the combinations *A&B*, *A&C*, and *B&C*. In the training phase, subjects saw examples of the combination *A&B* but not of the other combinations. If combination is not generative, participants will only be able to identify items as members of the previously-seen combination *A&B*, but not as members of the other two combinations. Table 3 shows that 72% of participants classified the item <CYB> as a member of the previously-unseen combination *B&C*. More participants classified the item in that new combination than classified the item <ABX> in the previously-seen combination *A&B*. There was no significant difference between the number of items classified in the previously-seen combination *A&B* and in the other, new combinations. This supports a generative view of category combination.

Overextension. The occurrence of overextension in the experimental data was analysed at the individual subject level. Overextension was taken to have occurred every time an individual subject gave a particular transfer item a higher score as a member of a combined category than they gave that item as a member of one of the constituent categories of that combination. For example, if a given subject gave the transfer item <ABY> a high classification score as a member of *A&C*, but the same subject gave <ABY> a lower classification score as a member of *C*, that would be taken as a case of overextension (an overextension response). In the experiment there were 5 different transfer items, each of which was classified in 3 different category combinations (*A&B*, *A&C*, *B&C*), where each combination had two different constituent categories. There were thus 30 possible cases in which overextension could arise. Out of those 30 cases, there were 5 in which 50% or more subjects produced overextension responses. Table 5 shows all item-combination-constituent triplets for which overextension

either occurred in the experiment or was predicted by the model. The 5 cases with at least 50% overextension responses in the experiment are indicated in Table 5 by a "Yes" in the 2nd-last column (with the proportion of overextension responses for each case). For these cases, at least 50% of subjects rated the given item as a better member of the given combination than of the given constituent category presented singly. In a significant number of cases subjects rated the given item as a member of the combination, but as a non-member of the constituent category. For example, 44% of subjects rated the item <CYB> as a member of the combination *B&C*, but the same subjects rated that item as a non-member of the constituent category *B*. These cases show that overextension occurs reliably even for artificial categories. (Order of presentation had no reliable influence on overextension in these cases.)

To analyse the model's predictions about overextension, the model's computed classification scores for all transfer items in all category combinations were compared with the scores for those items in the constituents presented singly. If an item had a higher classification score in a combination than in one of its constituent categories presented singly, the model predicted overextension for that item. Again, there were 30 cases in which overextension could happen: out of those 30 cases, the model predicted overextension in only 5 cases. These are indicated by a "Yes" in the last column in Table 5. Of the 5 cases in which overextension occurred in the experiment, 4 were cases in which overextension was predicted by the model. Of the remaining 25 cases in which overextension was not observed, 24 were cases in which the model predicted overextension would *not* occur. The model accurately predicted the occurrence and non-occurrence of overextension in these materials.

Discussion and Conclusions

The results obtained in the above experiment are important for a number of reasons. They show that both overextension and generativity occur even for combinations of artificial laboratory-learned categories. Previous research has investigated these factors in natural-language category combinations alone. They show that the patterns of overextension and generativity seen in the experiment have a close quantitative match with those predicted by the diagnostic evidence model. Other models give a looser qualitative account of overextension and generativity. Finally, these results show that the diagnostic evidence model accurately predicts people's classification of items in both single categories and category combinations, needing no free parameters to fit the data. Other models typically apply either to single categories or to category combination, but not to both. These models typically require a number of free parameters to fit the relevant data.

The diagnostic evidence model, in accounting accurately for the results of the above experiment,

represents an advance on other current theories of classification. However, there are results which the model cannot currently explain. For example, studies show that people can learn correlations between pairs of attributes and use those correlations in classification (Medin, Altom, Edelson, & Freko, 1982). The diagnostic evidence model, because it has no mechanism for learning correlations between attributes, cannot account for these results. In future work the model will be extended to learn attribute correlations by forming new "composite" attributes, and to use those attributes in classification. This may allow the model account for these findings.

References

Costello, F. J. (2000). An exemplar model of classification in single and combined categories. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*.

Costello, F. J., & Keane, M. T. (1997). Polysemy in conceptual combination: Testing the constraint theory of combination. In *Proceedings of the nineteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint guided conceptual combination. *Cognitive Science*, 24(2).

Costello, F. J., & Keane, M. T. (2001). Testing two theories of conceptual combination: Alignment versus diagnosticity in the comprehension and production of combined concepts. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 27(1), 255-271.

Hampton, J. A. (1988). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 55-71.

Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 8, 37-50.

Medin, D. L. & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207-238.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. K. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53-79.

Rips, L. J. (1995). The current status of research on concept combination. *Mind & Language*, 10, 72-104.

Rosch, E. & Mervis, C. D. (1975). Family resemblance studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.) *Cognition and Categorization*. Hillsdale, NJ: Erlbaum.