

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

An Exemplar Model of Classification in Single and Combined Categories

#### **Permalink**

<https://escholarship.org/uc/item/99d6x64v>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 22(22)

#### **Author**

Costello, Fintan

#### **Publication Date**

2000

Peer reviewed

# An Exemplar Model of Classification in Single and Combined Categories.

Fintan Costello ([fintan@compapp.dcu.ie](mailto:fintan@compapp.dcu.ie))

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

## Abstract

This paper describes an exemplar-based model of people's classification and typicality judgements in both single and combined categories. This model, called the diagnostic evidence model, explains the observed family resemblance structure of single categories; the productive nature of category combination; the observed overextension of typicality judgments in some combined categories; and the situations in which that overextension occurs. The model also gives a close fit to quantitative results from a representative single-category classification data-set.

Models of categorisation need to explain two basic aspects of human cognition: our ability to classify items as members of single categories such as *fish* or *cat*, and our ability to classify items as members of combinations of categories such as *wild cat* or *pet fish*. A successful model should account for the graded structure of classification in single categories: the fact that people's judgements of membership typicality for items in categories are proportional to the items' family resemblance to members of those categories (Rosch, 1978; Rosch & Mervis, 1975). A successful model should also account for the productivity of category combination: the fact that people are able to understand and judge membership in new combinations of categories, even if no already-existing examples of those combinations are known. This combinatorial productivity is important because it underlies our ability to think new thoughts and understand new expressions. In many currently popular models of categorisation (e.g. the context theory; Medin & Schaffer, 1978), an item's membership in a category is proportional to its similarity to the stored exemplars of that category. While this approach gives a good account for the graded structure of single categories, it has difficulty explaining the productivity of category combination, which involves classification in combinations for which no stored exemplars are available (Rips, 1995). This paper describes an exemplar-based model of classification in single and combined categories which explains the family resemblance structure of single categories, the productivity of category combination, and other specific results in both domains. The model, called the diagnostic evidence model, extends a successful earlier theory (Costello & Keane, 1997, in-pressA, in-pressB).

The first part of the paper presents the diagnostic evidence model of categorisation in single and combined categories, and gives its account for family resemblance and productivity in combination. The second part demonstrates the model by showing how it explains the observed overextension of typicality in some combined categories.

Overextension occurs when people rate an item as a poor member of both constituents of a combination, but as a good member of the combination as a whole; for example, when goldfish are rated as untypical members of the categories *pet* and *fish*, but as typical members of the combination *pet fish* (Hampton, 1988). Overextension has posed problems for a number of theories of category combination. The diagnostic evidence model accounts for results on overextension, and explains why overextension occurs in some combinations but not in others. The third part of the paper demonstrates this model further by showing how it gives a good fit to quantitative results from a representative classification data-set (Nosofsky, Palmeri, & McKinley, 1994); a fit as close as that given by exemplar-similarity models such as the context theory.

## The diagnostic evidence model

The diagnostic evidence model extends an earlier theory of the interpretation of noun-noun combined phrases, called the constraint theory (Costello & Keane, 1997, in-press-A). That theory set out to explain the diversity of interpretations which people produce for noun-noun combinations: the fact that people sometimes interpret combinations by forming conjunctions between the combining categories (as in the interpretation "*pet bird*: a parrot or some other bird which is also a pet"), sometimes by asserting relations between the categories (as in "*jungle bird*: a bird that lives in jungles"), and sometimes by transferring properties from one concept to the other (as in "*skunk bird*: a bird that smells bad"). Constraint theory explains this diversity by describing a combination process that forms mental representations satisfying three constraints of diagnosticity, plausibility and informativeness. Each interpretation type represents a different way of satisfying these constraints. The theory has been tested in a computer program which simulates the interpretation of noun-noun combinations, producing each interpretation type and generating results that agreed with people's interpretations of those combinations (Costello & Keane, in-press-A). Further, Costello & Keane (in-press-B) have provided direct experimental evidence for diagnosticity's role in the formation of combined categories.

Where the Constraint theory gave a qualitative account of noun-noun interpretation, the diagnostic evidence model aims to give a quantitative account of people's classification of items in single and combined categories. The model focuses on the diagnosticity constraint. The model assumes that people represent categories by storing sets of category exemplars in memory. From these sets, diagnostic attributes for categories are computed: these attributes serve to identify category members. An item's membership

typicality 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 high membership typicality in a category if it has attributes that are highly diagnostic for that category. An item has high typicality in a combination if it has some attributes highly diagnostic for one constituent of the combination, and other attributes highly diagnostic for the other. Two novelties in this model are its method for computing attribute diagnosticity, and its logic for combining the diagnosticity of multiple attributes to compute membership in single or combined categories. I describe these below.

### Attribute Diagnosticity

Diagnostic attributes are attributes which occur frequently in stored instances of a category, but rarely in that category’s contrast set (the set of stored instances which are not members of the category). These attributes serve to identify members of a category: a new item possessing a attribute which is highly diagnostic for a given category is likely to be a member of that category. The diagnosticity of attribute  $x$  for category  $C$  is defined in Equation 1. Let  $K$  be the contrast set for  $C$ . Let  $j_x$  be 1 if instance  $j$  possesses attribute  $x$ , and 0 otherwise.  $D(x/C/K)$ , the diagnosticity of  $x$  for  $C$  relative to  $K$ , is equal to the number of instances in  $C$  that possess  $x$ , divided by the total number of instances in  $C$  plus the number of instances in  $K$  that possess  $x$ :

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

If the attribute  $x$  occurs in all instances in category  $C$ , but no instances 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  $C$ : every instance possessing  $x$  is a member of  $C$ , every instance not possessing  $x$  is not a member. 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  $C$ : not every instance possessing  $x$  will be a member of  $C$ , not every instance not possessing  $x$  will be a non-member.

An important novelty in the diagnostic evidence model is that the diagnosticity of an attribute for a category can change depending on whether the category occurs singly or as part of a category combination. This change in diagnosticity arises because the contrast set used for computing diagnosticity is different in single and combined categories. For single categories, the contrast set consists of all instances which are not members of the category in question. For combined categories, however, the contrast set consists of instances which are not members of any of the constituents of the combination. The contrast set for a combination is thus a subset of the contrast sets for the single categories which make it up. This change in contrast set means that some attributes which are not diagnostic for a category when it occurs singly (because they occur frequently in that category’s contrast set), will be diagnostic

**Table 1.** An illustrative array of exemplars

Exemplar category labels		Attributes			
		FOUND	KEPT-IN	COLOR	HAS-PART
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 dog spaniel	house	basket	brown	tail
8	pet dog doberman	house	kennel	black	tail
9	pet dog bulldog	house	basket	brown	-----
10	pet terrapin	house	tank	green	skin

for that category when it occurs in a combination (if they occur only rarely in that combination’s contrast set).

The computation of attribute diagnosticity can be demonstrated using an illustrative set of stored exemplars of categories such as *pet*, *fish*, *dog* and *lobster*, shown in Table 1. These exemplars are described in attribute-value pairs on four dimensions: FOUND, KEPT-IN, COLOUR, and HAS-PART. Consider the diagnosticity of the attribute FOUND:HOUSE for the single category *fish*, which has 4 stored exemplars (exemplars 3, 4, 5, 6).  $K_{fish}$ , the contrast set for the category *fish*, contains exemplars 1, 2, 7, 8, 9, 10. FOUND:HOUSE occurs in 2 of the 4 *fish* exemplars in Table 1, and in 4 exemplars in the contrast set  $K_{fish}$ . The diagnosticity of FOUND:HOUSE for the *fish* is thus

$$D(\text{FOUND} : \text{HOUSE}/fish/K_{fish}) = \frac{2}{4 + 4} = 0.25 \quad (2)$$

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 exemplars that are members neither of *pet* nor of *fish* (exemplars 1 and 2). FOUND:HOUSE does not occur in any exemplars in the contrast set  $K_{petfish}$ . The diagnosticity of FOUND:HOUSE for *fish* relative to the contrast set  $K_{petfish}$  is thus

$$D(\text{FOUND} : \text{HOUSE}/fish/K_{petfish}) = \frac{2}{4 + 0} = 0.5 \quad (3)$$

The attribute thus gives a greater degree of diagnostic evidence for membership in the *fish* constituent of *pet fish*; in other words, the attribute FOUND:HOUSE is more diagnostic in identifying an item as a pet fish than it is identifying an item as a fish. This effect of contrast set on diagnosticity is central to the diagnostic evidence model’s account for overextension in combined categories, and is discussed in the section on overextension, below.

### A Continuous-valued Logic for Evidence

Diagnostic attributes, then, give evidence for an instance’s classification in a category. Instances 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. How is the diagnostic evidence from an item's attributes combined to produce an overall measure of evidence for category membership? The diagnostic evidence model uses a continuous-valued logic to combine diagnostic evidence from multiple attributes. This logic assumes continuous variables with values between 0 and 1, and uses the following 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 derive from standard probability theory, and can be justified by considering the operations *AND*, *OR*, and *NOT* for samples of independently distributed variables. Suppose variables *A* and *B* have 0.75 and 0.5 probability of being true, respectively. 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 ( $0.75 \times 0.5$ ): of the 75% of cases in which *A* is true, 50% of those are cases in which *B* is also true. Finally, the probability of *A OR B* being true is 0.875 ( $1 - (1 - 0.75) \times (1 - 0.5)$ ): of the 25% of cases in which *A* is false, 50% of those are cases in which *B* is also false; thus *A OR B* is true in 87.5% of cases. Similar (though often more complex) logics have been used in various areas (e.g. in models of evidence-based reasoning; Shafer, 1976). The current model is unique in using this approach to compute the contribution which different attributes make in people's classification of items in single or combined categories.

### Combining Diagnosticity of Multiple Attributes.

To combine the diagnostic evidence from multiple attributes for membership in a category, the diagnostic evidence model uses the equation for *OR*. An instance *i* with a set of attributes  $x_1, x_2, x_3$ , will be a member of category *C* if  $x_1$  or  $x_2$  or  $x_3$  serves to identify the instance as a member of *C* (if  $x_1 OR x_2 OR x_3$  is diagnostic for *C*). This is formalised in Equation 7, which has the form of the equation for *OR* (Equation 6, above). Let *A* be the set of attributes of instance *i* and  $D(x|C/K)$  be the diagnosticity of attribute *x* for *C*. Then  $E(i|C/K)$ , the overall evidence for classifying instance *i* as a member of *C*, is

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

This equation accounts for people's classification in both strictly defined and "family resemblance" categories. If an attribute *x* strictly defines a category *C* (occurs in all instances 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 instance *i* will definitely be a member of *C*. In categories which have no single perfectly diagnostic attribute but rather have 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 instance has, the higher its degree of membership will be. In other words, the more of a family resemblance an instance has to the members of a

category, the higher its membership typicality will be. This relationship between diagnostic attributes and membership has specific support in Rosch & Mervis' (1975) finding that people's judgements of an instance's typicality in a single category rose reliably with the number of diagnostic attributes for that category which the instance possessed.

The combination of diagnostic evidence can be illustrated using the exemplars in Table 1. For example, consider the evidence for exemplar 5 (*salmon*) as a member of the category *fish*. This exemplar has attributes LIVES:SEA, COLOUR:SILVER, and HAS-PART:SCALES. The diagnosticities of these attributes for *fish* are relatively high (0.4, 0.75 and 0.5 respectively, as computed from Equation 1). From Equation 7, these diagnostic evidence values are combined to obtain an overall measure of evidence for exemplar *salmon*'s typicality in category *fish* as follows:

$$E(\text{salmon} | \text{fish} | K_{\text{fish}}) = 1 - (1 - 0.4)(1 - 0.75)(1 - 0.5) = 0.925 \quad (8)$$

The exemplar *salmon* has good evidence for membership in the category *fish* because it possesses highly diagnostic attributes for that category: in other words, *salmon* is a highly typical *fish*. Other exemplars have less diagnostic attributes for the category *fish*, and thus have lesser degrees of evidence and are less typical category members. For example, the exemplar *shark* has the less diagnostic attribute HAS-PART:SKIN and is a less typical member of the category *fish* ( $E(\text{shark} | \text{fish} | K_{\text{fish}}) = 0.91$ , computed as above); the exemplar *goldfish* has two less diagnostic attributes LIVES:HOUSE and COLOUR:GOLD and is less typical again ( $E(\text{goldfish} | \text{fish} | K_{\text{fish}}) = 0.813$ ); the exemplar *spaniel* has no diagnostic attributes and is a poor member of the category ( $E(\text{spaniel} | \text{fish} | K_{\text{fish}}) = 0.25$ ).

### Diagnostic Evidence for Combined Categories

The diagnostic evidence model of classification, then, is consistent with observed patterns of typicality in single categories. The model extends easily to account for classification in category combinations: an item will be a member of a combined category if it gives diagnostic evidence for membership in each constituent in that combination. In computing an item's membership in a combined category, the model uses the continuous-valued *AND* described above (Equation 5) to combine the item's evidence for membership in each constituent of the combination. An instance *i* will be classified as a 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. More formally,  $E(i|C_1 \dots C_N | K_{1 \dots N})$ , the evidence for classifying *i* as a member of combination  $C_1 \dots C_N$ , is

$$E(i|C_1 \dots C_N | K_{1 \dots N}) = \prod_{n=1}^N E(i|C_n | K_{1 \dots N}) \quad (9)$$

where the contrast set  $K_{1 \dots N}$  is the set of instances not in any of the categories  $C_1 \dots C_N$ . Note that an instance *i* will give evidence for membership in each constituent of a combination if it has some attributes diagnostic for each constituent: some attributes diagnostic for one constituent, other attributes diagnostic for others.

Because the diagnostic evidence model computes evidence for membership in a combination by combining evidence for membership in its constituent categories, it can explain people’s ability to classify items in new combinations, even if they have no stored exemplars of those combinations. An item is classified as a member of a combination, even one with no stored exemplars, if the item has diagnostic attributes for each constituent category in the combination. For example, in Table 1, there are no stored exemplars of the combination *pet lobster*. However, an item could be classified as a good member of the combination *pet lobster* if it possessed the attribute HAS-PART:CLAWS (perfectly diagnostic for *lobster* in Table 1) and the attribute FOUND:HOUSE (highly diagnostic for *pet*).

In accounting in this way for the productivity of category combination, the model goes beyond theories such as the context theory, in which classification is based on similarity to stored exemplars of a category. Such theories cannot account for classification in new combinations for which there are no stored exemplars. For example, in an exemplar-similarity based model, people would judge membership in *pet lobster* by computing an item’s similarity to stored exemplars of that combination (by comparing the item to previously seen examples of pet lobsters). Since *pet lobster* has no stored exemplars, this computation would be meaningless (see Rips, 1995).

### Accounting for Overextension

Various studies have examined overextension of classification in combined categories. Overextension occurs when people rate an item as a poor member of both constituents of a combination, but as a good member of the combination as a whole. For example, people might rate goldfish as typical members of the combination *pet fish*, but as untypical members of the single categories *pet* and *fish*. Hampton (1988) found that overextension was more likely for some combinations than for others: the lower the degree of overlap between combining categories (the fewer exemplars the categories had in common) the more likely the combinations were to be overextended. For example, the constituents of *pet fish* have low overlap (many fish are not pets; many pets are not fish), and that combination was often overextended. By contrast, combinations of categories with many common members were usually not overextended. For example, the constituents of *pet dog* have high overlap (most dogs are also pets), and that combination was usually not overextended.

Overextension poses a challenge for theories of category combination (Osherson & Smith, 1981). In the diagnostic evidence model, overextension arises because of changes in attribute diagnosticity: because some attributes may have low diagnosticity for a category when it occurs singly, but high diagnosticity for that category when it occurs as part of a combination. As we saw earlier, the attribute FOUND:HOUSE was less diagnostic for the single category *fish*, but was more diagnostic for the category in the context of the combination *pet fish* (because the attribute occurred often in the contrast set for the category *fish*, but not in the contrast set for the combination *pet fish*). This change in

**Table 2.** Overextension of exemplar *goldfish* in *pet fish*

Evidence for membership in	Exemplar	Attribute Diagnosticity			
		FOUND	KEPT-IN	COLOR	HAS-PART
	<i>goldfish</i> :	house	tank	golden	scales
<i>pet</i> singly :	0.714	0.67	0.14	0	0
<i>fish</i> singly:	0.813	0.25	0.33	0.25	0.5
<i>pet fish</i> :	0.89				
<i>pet</i>	1	1.0	0.2	0	0
<i>fish</i>	0.89	0.5	0.4	0.25	0.5

diagnosticity means that an item with that attribute could give good evidence for membership in the combination *pet fish* (and therefore high typicality in that combination), but poor evidence for membership in the single categories *fish* and *pet* (low typicality in those single categories).

Table 2 illustrates this account of overextension, showing computed evidence for the exemplar *goldfish* as a member of the single categories *pet* and *fish*, and the combination *pet fish*. Note that *goldfish* gives higher evidence for membership in *pet fish* (0.89) than in either *pet* (0.714) or *fish* singly (0.813). *Goldfish* would thus be judged a highly typical *pet fish* but a less typical *pet* or *fish*. This is because the exemplar’s attributes have higher diagnosticity for the combination than for the single categories. For example, FOUND:HOUSE has a diagnosticity of 0.67 for the single category *pet* and of 0.25 for the single category *fish*. In the context of *pet fish*, however, FOUND:HOUSE has a higher diagnosticity both for the constituent *pet* (1.0) and the constituent *fish* (0.5). (In Table 2, evidence for membership in the single categories is computed by combining attribute diagnosticity as in Equation 7. Evidence for membership in the combination is obtained by computing evidence for membership in each constituent category as in Equation 7, and combining that evidence as in Equation 9).

In this account, overextension arises from a difference between the contrast sets for single categories and those for a combination, which leads to a difference in diagnostic evidence for membership in the single categories and the combination. If there is little difference between these contrast sets, overextension won’t occur. Table 3 illustrates this for the combination *pet dog*. *Pet dog* is not overextended: the exemplar *spaniel* gives more evidence for membership in the single categories *pet* (0.96) and *dog* (0.98) than in the combination *pet dog* (0.95). Because the categories *pet* and *dog* have a high overlap (in Table 1, all

**Table 3.** Non-overextension of exemplar *spaniel* in *pet dog*

Evidence for membership in	Exemplar	Attribute Diagnosticity			
		FOUND	KEPT-IN	COLOR	HAS-PART
	<i>spaniel</i> :	house	basket	brown	tail
<i>pet</i> singly :	0.96	0.67	0.5	0.5	0.5
<i>dog</i> singly:	0.98	0.5	0.67	0.67	0.67
<i>pet dog</i> :	0.95				
<i>pet</i>	0.96	0.67	0.5	0.5	0.5
<i>dog</i>	0.98	0.6	0.67	0.67	0.67

**Table 4.** Predicted and observed probability of classification of exemplars in Nosofsky, et al., (1994) Experiment 1.

Exemplar labels	Exemplars	Diagnostic evidence ( $C(i/A), W = 8$ )	Predicted classification probability (linear transform of $C(i/A)$ )	Classification probability observed in Experiment
A1	1 1 1 2	0.69	0.77	0.77
A2	1 2 1 2	0.65	0.74	0.78
A3	1 2 1 1	0.75	0.83	0.83
A4	1 1 2 1	0.52	0.6	0.64
A5	2 1 1 1	0.52	0.6	0.61
B1	1 1 2 2	0.37	0.46	0.39
B2	2 1 1 2	0.37	0.46	0.41
B3	2 2 2 1	0.13	0.23	0.21
B4	2 2 2 2	0.07	0.17	0.15
T1	1 2 2 1	0.45	0.54	0.56
T2	1 2 2 2	0.31	0.4	0.41
T3	1 1 1 1	0.78	0.86	0.82
T4	2 2 1 2	0.31	0.4	0.4
T5	2 1 2 1	0.16	0.26	0.32
T6	2 2 1 1	0.45	0.54	0.53
T7	2 1 2 2	0.07	0.17	0.2

dogs are also pets) there is little difference between the contrast sets for the single categories *pet* and *dog* and the contrast set for the combination *pet dog*. There is thus little difference in the diagnosticity of attributes for the single categories and for the combination; hence, there is no overextension.

This account explains Hampton’s (1988) finding that overextension is rare for combinations whose constituent categories have a high degree of overlap. The greater the overlap between the constituent categories of a combination, the less of a difference there is between the contrast sets for those categories occurring singly, and the contrast set for that combination. The less of a difference between contrast sets, the less of a difference between diagnostic evidence for membership in the single categories and in the combination; the less chance of overextension.

### Fitting Classification Data-sets

As described above, the diagnostic evidence model can explain various results in natural-language categorisation and category combination. In this section I fit the model to results obtained in a study of classification in artificial laboratory-learned categories: Nosofsky, Palmeri, & McKinley’s (1994) replication of Medin & Schaffer’s (1978) study. In Nosofsky, Palmeri, & McKinley’s experiment, participants learned to classify drawings of rocketships as coming from planet A (category A) or planet B (category B). The rocketships varied on four dimensions (shape of tail, wings, nose, and porthole) each with two values, represented by 1 and 2. Rockets from planet A had values of 1 on most dimensions, while rockets from planet B had values of 2. An abstract representation of this category structure is shown in Table 2. In an initial training phase, participants learned 9 training items: A1...A5 from category A and B1...B4 from category B. In the test phase

participants categorised the 9 training items and 7 new test items T1...T7. Test item T3 was the prototype for category A (having a value 1 on all dimensions).

In this experiment participants classified items into one of only two possible categories (A or B). Classification in this two-category task is different from classification in natural-language categories: when only two categories are available, an item’s membership in a category depends both on evidence that the item is a member of the category, and on evidence that the item is not a member of the other category. In applying the diagnostic evidence model to this two-category task, the model was extended (using the continuous-valued logic described above) to take account of both sources of evidence: an item was classified in category A if it gave evidence for membership in A, OR did NOT give evidence for membership in B. Formally,  $C(i/A)$ , the classification score for  $i$  as a member of category A, is

$$C(i/A) = E(i/A/K_A) \text{ OR } (\text{NOT } E(i/B/K_B)) \quad (10)$$

$$= 1 - (1 - E(i/A/K_A))^W (1 - (1 - E(i/B/K_B)))$$

where  $E(i/A/K_A)$  and  $E(i/B/K_B)$  give measures of evidence for membership in A and B respectively (computed according to Equation 7), and where parameter  $W$  represents the relative importance of evidence for membership in A versus evidence for membership in B in classification.

The diagnostic evidence model was applied to the data-set using only the training stimuli (exemplars A1...A5 and B1...B4). These training exemplars were used to compute the diagnosticity of the values 1 and 2 on each dimension for the categories A and B. These diagnosticities were then used to compute the diagnostic evidence score  $C(i/A)$  for both training and test exemplars as members of category A. These scores are shown in the “diagnostic evidence” column in Table 4. These scores are those for the value of  $W$  for which the correlation between predicted and observed

classification was highest ( $W = 8$ ;  $r = .99$ ,  $\%var = 98\%$ ). The model's predicted classification probabilities (shown in the next column in Table 4) were obtained by a linear transformation of the diagnostic evidence scores, mapping the mean diagnostic evidence score onto the mean observed classification probability, and the standard deviation of the diagnostic evidence score onto the standard deviation of observed classification probabilities. (This transformation introduces no extra degrees of freedom into the model's fit to the data; it simply allows direct comparison between computed evidence for classification and the classification probabilities observed in the experiment). The diagnostic evidence model's computed classification scores for items closely follow people's classifications of those items, as comparison of the predicted and observed classification probability columns in Table 4 shows. The model accounts for the qualitative finding that the test exemplar T3 (the prototype for category A) gets a higher classification probability than all other test exemplars. The percentage of variance explained by the diagnostic evidence model (98%) is in the same range as that produced by other models (the context model explains 96% of variance in these results; the Rulx model explains 98%; see Nosofsky, Palmeri, & McKinley, 1994). However, those models used four free parameters to fit the data (varying the selective attention paid to the 4 dimensions on which exemplars were described), as opposed to the single parameter used by the diagnostic evidence model.

### Conclusions and Future work

The diagnostic evidence model of classification described here goes beyond other theories of classification in giving an account for both single and combined categories. The model explains the family resemblance structure of single categories, the productivity of category combination, and the occurrence of overextension in some combined categories. That the model is exemplar-based is significant: a number of results have shown that exemplars are important both for simple and combined categories (e.g. Gray & Smith, 1995). Some argue that exemplar-based models cannot account for the productivity of combination (Rips, 1995); the current model provides evidence against this argument. The model fits a representative classification data set as closely as Medin & Schaffer's (1978) context theory, while using fewer free parameters.

There are, however, various classification results which the diagnostic evidence model cannot currently explain. Because the model does not provide a mechanism for learning, it cannot address the role of learned attribute correlations in classification. A number of studies (e.g. Medin, Altom, Edelson, & Freko, 1982) show that people learn to associate correlated pairs of attributes with categories, and to use those correlated attributes in classification. The diagnostic evidence model as it currently stands cannot account for this result because it treats all attributes independently. Extending the model with a learning mechanism which can recognise attribute correlations and use those correlations to form new "composite" attributes may allow the model to account for

the role of correlation in classification. In an initial test of this approach, in which composite attributes were hard-coded into the representation used, the diagnostic evidence model was able to give a good fit to Medin et al.'s results. In future work I aim to develop the diagnostic evidence model in this direction.

### References

- Costello, F. J., & Keane, M. T. (in press A). Efficient creativity: Constraint guided conceptual combination. *Cognitive Science*.
- Costello, F. J., & Keane, M. T. (in press B). Testing two theories of conceptual combination: Alignment versus diagnosticity in the comprehension and production of combined concepts. *Journal of Experimental Psychology: Learning, Memory & Cognition*.
- Costello, F. J., & Keane, M. T. (1997). Polysemy in conceptual combination: Testing the constraint theory of combination. In *Nineteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Gray, K. C. & Smith, E. E. (1995). The role of instance retrieval in understanding complex concepts. *Memory & Cognition*, 23(6), 665-674.
- 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. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101(1), 53-79.
- Osherson, D. N. & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9, 35-58.
- Rips, L. J. (1995). The current status of research on concept combination. *Mind & Language*, 10(1/2), 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.
- Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press.