Classification in combined phrases: noun-noun conjunctions and intersective, subsective and privative adjective-noun combinations.

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

Abstract

How is the meaning of an adjective-noun phrase related to the meanings of its adjective and noun constituents? Different types of adjective produce different relations between a combination and in its constituent adjectives and nouns. Current theories give separate mechanisms for each adjective type. This paper shows how a computational model originally developed to explain classification in noun-noun conjunctions also gives a single explanation for combinations involving standard intersective adjectives, subsective adjectives such as "skilful", and privative adjectives such as "fake". This model may lead to a unified account of natural-language word combination.

1 Introduction

The ability to put words together in new ways to produce meaningful word combinations is a defining feature of language. But how are the meanings of multiword phrases such as "pet fish", "skilful violinist", "fake surgeon", or "red-headed butcher" produced from the meaning of their constituent words? Meaning depends on categorisation: to understand a single word such as "fish" is to be able to correctly classify items in the single category *fish*; to understand a phrase such as "pet fish" is to be able to classify items in the combined category *pet fish*. To explain how the meaning of a combination like "pet fish" is produced from its constituent words, we must first explain how people classify items in single categories like *pet* and *fish* and in combined categories like *pet fish*.

The first part of this paper describes a computational model of how people classify items in single and combined categories. This model, called the diagnostic evidence model, accounts for various psychological results: it fits data on classification in both natural-language and artificial categories and combinations; explains graded typicality as observed in single categories; and accounts for the *overextension* seen in some noun-noun conjunctions (Costello, 2000a, 2000b; Costello & Keane, 2000, 2001). The second part of the paper shows how the model explains classification in combinations involving three different types of adjective: intersective adjectives ("red-headed butcher"), subsective adjectives ("skilful vioinist"), and privative adjectives ("fake surgeon"). For these different adjective types, different relationships hold between a combination and its adjective and noun constituents. Most current theories of adjective-noun combination give separate mechanisms for each adjective type. No current theory gives a single account for both noun-noun and adjective-noun combinations. The diagnostic evidence model sets out to give a unified account for intersective, subsective and privative adjective-noun combinations, and for noun-noun conjunctions.

2 The Diagnostic Evidence Model

The diagnostic evidence model was developed as a computational model of classification in noun-noun conjunctions (Costello, 2000a, 2000b). 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 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.

As well as addressing classification in noun-noun conjunctions, the diagnostic evidence model sets out to explain the overextension effects observed in such conjunctions. Overextension occurs when people classify an item as a poor member of both constituents of a conjunction, but as a good member of the conjunction 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 is an example of the influence of context on classification. The model's account for overextension in noun-noun conjunctions also explains context effects observed in adjective-noun combinations such as "skilful violinist" or "fake surgeon". The model explains these effects in terms of how diagnosticity changes in combined categories.

2.1 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:

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

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.

Table 1. Ten stored members of *pet* and *fish.*

Item	Categories		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

2.2 Diagnosticity changes in combination

The idea of the contrast set is important in computing attribute diagnosticity: the fewer occurences of an attribute in the contrast set for a given 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. Because of this, some attributes which are not diagnostic for a category occuring singly can be diagnostic for that category as part of a combination. This change in diagnosticity allows the model to explain overextension in noun-noun conjunctions.

Table 1 shows 10 stored members of categories *pet* and *fish*, described on 4 dimensions. Computation of attribute can be illustrated using this set of stored category members. Consider the diagnosticity of 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 contrast set K_{fish} (the set of items which are not members of the category *fish*). The diagnosticity of <found:house> for *fish* is thus

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

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). <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

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

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*.

2.3 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

$$\text{NOT } A = 1 - A \quad (4)$$

$$A \text{ AND } B = AB \quad (5)$$

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

These equations can be justified by considering the operations *AND*, *OR*, and *NOT* for samples of independent variables. Suppose a variable *A* has a *0.75* probability of being true (true in 75% of samples), and a variable *B* has a *0.5* probability of being true (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)*: of the 75% of samples 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) X (1-0.5))*: of the 25% of samples 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.

2.4 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*. An item *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 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-language categories (Rosch, 1978). This 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.

2.5 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 $C_1 ... C_N$ 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, $E(i|C_1...C_N|K_{1...N})$, the evidence for classifying *i* as a member of combination $C_1...C_N$, is

$$E(i \mid C_1 ... C_N \mid K_{1...N}) = \prod_{n=1}^{N} E(i \mid C_n \mid K_{1...N}) \quad (8)$$

where the contrast set $K_{1...N}$ is the set of items not in any category $C_1...C_N$. In this equation an item *i* gives evidence for membership in each constituent of a combination if it has attributes diagnostic for each.

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

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

Table 2 illustrates the diagnostic evidence model by showing the computed membership for an item with 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 in the constituents of the combination are computed from those diagnosticities, and are shown in bold to the left of those diagnosticities. At the bottom of Table 2 is the overall membership score for the item in the combination (computed by multiplying its constituent membership scores).

Notice that the item in Table 2 has a higher overall membership score in the combination *pet fish* than in either of the categories *pet* or *fish* presented singly. This illustrates the model's account of overextension, which arises if some attributes have low diagnosticity for a single category but high diagnosticity for that category as part of a combination. As Table 2 shows, the attribute <found:house> has lower diagnosticity for the single category *fish*, but higher diagnosticity in the context of the combination *pet fish* (the attribute occurred often in the contrast set for the single category *fish*, but not in the contrast set for *pet fish*). The model would thus predict overextension for that combination.

2.6 Empirical support for the model

Empirical support for the model comes from a number of sources. First, in an experiment investigating how people classify items in artificial laboratory-learned categories and in combinations of those categories (Costello, 2000b), the model computed classification scores which agreed closely with people's classification scores for items in both the combined categories (r=.99 %var=98, p <.01) and single categories (r=.85 %var = 0.65, p<0.01). The pattern of overextension observed in the experiment also agreed with that predicted by the model. No free parameters were needed to fit the data. Second, the model gave a close fit (r=.99, %var=.98, p<0.01) to experimental data on the classification of items in single categories only (Costello, 2000a). General support for the model comes from studies showing the importance of diagnosticity in noun-noun combination (Costello & Keane, 2001), and from other studies showing that overextension is more likely in a combination whose categories do not overlap (Hampton, 1988). This matches the diagnostic evidence model's predictions: such combinations produce the greatest change in contrast set, and so are expected to cause overextension.

3 Adjective Combinations

This paper focuses on intersective, subsective, and privative adjectives, which behave in quite different ways in adjective-noun combinations (Kamp, 1975). In general an item will be a member of an adjective-noun combination if the item has some attributes identifying it as a member of the adjective category, and other attributes identifying it as a member of the noun category. Intersective, subsective, and privative adjectives differ in how the attributes identifying adjective membership are selected. For intersective adjectives such as *red-headed*, the same attribute always identifies members of the adjective category: in classifying someone as a *red-headed butcher*, a *red-headed surgeon*, or a *red-headed violinist*, the same feature (<hair colour:red>) indicates membership in the adjective. For subsective adjectives such as *skilful*, however, different attributes identify members of the adjective category depending on the noun with which the adjective is combined. For example, in identifying a *skilful violinist* the attribute <expressive> may be important, while for *skilful surgeon*, <dextrous> would be important and <expressive> would be less so.

For privative adjectives such as *fake*, the pattern is different again. For a privative combination such as *fake surgeon*, the attribute which identifies an item as *fake* is the reverse of

Table 3. An illustrative array of stored items for adjective and noun categories.

Item	Category membership			Attributes					
	surgeon	violinist	adjective	MOVEMENT	ATTITUDE	TRAINING		CLOTHES	HAIR
						MEDICAL	MUSICAL		
1	**Yes**	No		dextrous	normal	surgery	none	white coat	brown
2	**Yes**	No		dextrous	reserved	surgery	none	white coat	black
3	**Yes**	No		normal	reserved	surgery	none	white coat	blond
4	No	**Yes**		normal	expressive	none	violin	black-tie	black
5	No	**Yes**		rythmic	expressive	none	violin	black-tie	brown
6	No	**Yes**		rythmic	normal	none	violin	black-tie	blond
7	No	No		rythmic	normal	none	none	white coat	brown
8	No	No		normal	reserved	none	none	black-tie	black
9	No	No		rythmic	combative	none	none	white coat	blond
10	No	No		normal	reserved	none	none	black-tie	brown
11	No	No		clumsy	normal	none	none	casual	black
12	No	No		clumsy	combative	none	none	casual	blond
13	No	No	red-headed	clumsy	combative	none	none	casual	red
14	No	No	skilful	dextrous	expressive	none	none	casual	brown
15	No	No	skilful	dextrous	expressive	none	none	casual	black
16	No	No	skilful	dextrous	expressive	none	none	casual	blond

an attribute which would be important in identifying a member of the noun category *surgeon*. If having surgical training is an attribute which is important in identifying *surgeons*, then having no training would be important in identifying a *fake surgeon*. Other attributes which are typical but not defining for *surgeons* (i.e. being reserved, wearing a white coat) would also be important for membership in a privative combination such as *fake surgeon*. In other words, someone is a *fake surgeon* if they seem to be a surgeon but are not really a surgeon; that is, if they have a high membership in both the category *surgeon* and in the negated category *non-surgeon*.

Current theories (e.g. Franks, 1995) typically provide separate mechanisms to explain intersective, subsective, and privative adjectives. The diagnostic evidence model provides a single explanation for all three combination types. The model's account for subsective and privative combinations is closely linked to its account for overextension in noun-noun conjunctions, and is best illustrated by some examples. Table 3 shows an set of items marked as members of categories *surgeon*, *violinist*, *red-headed*, *skilful*, *non-surgeon* and *non-violinist* (the categories *non-surgeon* and *non-violinist* are marked by a 'No' in the *surgeon* and *violinist* columns). These items have attributes on dimensions movement style, emotional attitude, professional training, clothes worn, and hair colour. These items represent a number of facts. First, all members of the category *surgeon* have surgical training; all non-surgeons do not; all violinists have violin training; all non-violinists do not. Second, wearing a white coat and being reserved is typical of surgeons, but some non-surgeons also have those attributes. Similarly, wearing black tie and being expressive is typical of violinists, but some non-violinists also have those attributes. Third, all and only members of the category *red-headed* have red hair. Finally, the 3 items marked as *skilful* in the example set (items 14, 15, 16) represent knowledge about the possible ways in which someone can be skilful. These items all contain the attributes <dextrous> and <expressive>, indicating that someone can be *skilful* if they have either the attribute <dextrous> or the attribute <expressive>. These items are intended to stand for a larger number of items describing people who are skilful in various different ways.

3.1 Intersective adjectives

For intersective adjectives such as red-headed the same attribute always identifies members of the adjective category. The diagnostic evidence model's account for intersective adjective-noun combinations is straightforward. In the model an item will be classified as a member of an adjective-noun combination if it has some attributes diagnostic for the adjective and others diagnostic for the noun. For example, in Table 3 attribute <hair:red> is perfectly diagnostic for *red-headed* and <medical-training:surgery> is perfectly diagnostic for *surgeon* (both attributes occur in all members of their categories and in

Table 4. Computed membership score for two items in the combination *skilful surgeon*.

Category		Attribute Diagnosticity			
		MOVEMENT	ATTITUDE	TRAINING	CLOTHES
		dextrous	normal	surgery	white coat
skilful	1	1.0	0.0	0.0	0.0
surgeon	1	0.7	0.2	1.0	0.6
overall	**1**				
		normal	expressive	surgery	white coat
skilful	0.6	0.0	0.6	0.0	0.0
surgeon	1	0.2	0.2	1.0	0.6
overall	**0.6**				

Table 5. Computed membership score for two items in the combination *skilful violinist*.

Category		Attribute Diagnosticity			
		MOVEMENT	ATTITUDE	TRAINING	CLOTHES
		dextrous	normal	violin	black-tie
skilful	0.6	0.6	0.0	0.0	0.0
violinist	1	0.2	0.2	1.0	0.6
overall	**0.6**				
		normal	expressive	violin	black-tie
skilful	1	0.0	1.0	0.0	0.0
violinist	1	0.4	0.7	1.0	0.0
overall	**1**				

no non- members, and so have a diagnosticity score of 1). Any items which possess both these attributes would thus be classifed as members of the combination *red-headed surgeon*. Since <hair:red> has maximum diagnosticity for *red-headed*, its diagnosticity will not improve no matter what noun that adjective combines with. While for intersective adjectives, diagnostic attributes do not change in combination with different nouns, for subsective and privative adjectives, different attributes will be diagnostic in different subsective or privative combinations, depending on the noun with which the adjective is combined.

3.2 Subsective adjectives

For subsective adjectives such as *skilful*, different attributes will be diagnostic for the adjective category in different adjective-noun combinations. In the diagnostic evidence model these changes in diagnosticity arise automatically because of differences between the contrast sets for different adjective-noun combinations. For example, Table 4 shows the overall computed evidence for membership for two different items in the combination *skilful surgeon*. The diagnosticity of each items' attributes for the constituent categories *skilful* and *surgeon* are listed in columns under those attributes. The two items' membership scores in the two constituent categories (computed from those diagnosticities) is shown in bold, to the left of those diagnosticities. Below those scores is the overall membership of the item in the combination (computed by multiplying those constituent membership score). Both items shown in Table 4 have the attributes of a typical surgeon (surgical training, white coat, etc.), and are good members of the *surgeon* category. The two items differ in the attributes which identify them as members of the adjective category *skilful*: the first item has the attribute <movement:dextrous>, while the second item has the attribute <attitude: expressive>. The first item's attribute <movement:dextrous> is perfectly diagnostic for the adjective *skilful* in the context of *skilful surgeon*. This is because the contrast set for *skilful surgeon* consists of all items in Table 3 that are neither members of the category *skilful* nor the category *surgeon*. <movement:dextrous>, while it occurs in all members of the category *skilful*, occurs in no items in that contrast set (see Table 3).

By contrast, the second item in Table 4, with the attribute <attitude:expressive>, is not a good member of the combination *skilful surgeon*. This is because the attribute <attitude:expressive> does not have high diagnosticity for *skilful* in the context of *skilful surgeon*, since the attribute occurs in two items in the contrast set for *skilful surgeon* (see Table 3). Thus, according to the diagnostic evidence model, the attribute <movement:dextrous> identifies people who are skilful in the way that is relevant to being a surgeon, and the attribute <attitude: expressive> does not.

Table 5 gives the computed evidence for membership for two items in the combination *skilful violinist*. A similar, though reversed, pattern occurs for this combination. Both items in Table 5 have the attributes of a typical violinist (violin training, black-tie, etc.). The first item has attribute <movement:dextrous>, while the second item has attribute <attitude: expressive>. The item with the attribute <movement:dextrous> has poor membership in the combination *skilful violinist*, while the item with the attribute <attitude:expressive> is a perfect member of that combination. Again, this is because the contrast set for *skilful violinist*

consists of all items which are not members of either the category *skilful* nor the category *violinist*. The attribute <movement:dextrous> occurs in two members of that contrast set, and so is less diagnostic. <attitude:expressive> on the other hand, occurs in no members of that contrast set, and so has high diagnosticity.

Tables 4 and 5 shows how the diagnostic attributes for a subsective adjective like *skilful* can change depending on the noun with which the adjective combines: for *skilful surgeon* the attribute <dextrous> is most diagnostic, while for *skilful violinist* the attribute <expressive) is most diagnostic. This change in diagnostic attribute arises because of differences in the contrast sets for the two combinations. The next section shows how the model can similarly account for classification and changes in diagnosticity in combinations involving privative adjectives such as *fake*.

3.3 Privative adjectives

Privative adjective-noun combinations such as "fake surgeon" represent a combination between a noun category and the negation of that category. For example, a *fake surgeon* is "someone who appears to be a surgeon, but is actually not a surgeon". A given item will be a member of such a privative combination if it has some attributes diagnostic for the noun category *surgeon* and other attributes diagnostic for the negated category *non-surgeon*.

In the diagnostic evidence model, the degree of membership of an item in a privative combination is obtained by computing the item's evidence for membership in a combination containing the noun category and the negation of that noun category. For example, an item's membership in the combination *fake surgeon* would be computed by obtaining its evidence for membership in the combination "*surgeon* and *non-surgeon*". As in all other combinations, evidence for membership in the constituent categories of a combination is computed relative to the contrast set for the combination: the set of items which are not members of either constituent category of that combination. The contrast set for fake surgeon (that is, for the combination "*surgeon* and *not a surgeon*") is the set of items which are neither members of the category *surgeon* nor of the category *non-surgeon*. This set necessarily contains no members, and so there will be a significant change in the diagnosticity of attributes for the privative adjective-noun combination.

Table 6. Computed classification score for two items in the combination *fake surgeon*.

Category	Attribute Diagnosticity				
	MOVEMENT		ATTITUDE	TRAINING	CLOTHES
		dextrous	reserved	surgery	white coat
surgeon	1	0.7	0.7	1.0	1.0
non-surgeon	0.5	0.2	0.2	0.0	0.2
overall	**0.5**				
		dextrous	reserved	none	white coat
surgeon	1	0.7	0.7	0.0	1.0
non-surgeon	1	0.2	0.2	1.0	0.2
overall	**1**				

Table 6 shows the computed evidence for membership of two different items in the combination *fake surgeon*.(that is, the combination *surgeon and non-surgeon*). For each item evidence for membership in the constituent categories *surgeon* and *non-surgeon* are computed relative to the contrast set for the combination. Because the first item in Table 6 has the attribute <medical-training:surgery>, it has perfect evidence for membership in the constituent category *surgeon*. However, the item has poor evidence for membership in the negated category *non-surgeon*: the item has no attributes which are diagnostic for *non-surgeon*. The item's overall membership in the combination *fake surgeon* is therefore low.

The second item in Table 6 has the attribute <medical-training: none> rather than <medical-training:surgery>. Despite not having the attribute than <medical-training:surgery>, that item also has perfect evidence for membership in *surgeon*. This is because the item possesses the attribute <clothes:white-coat>, which is perfectly diagnostic for the category *surgeon* in context of the combination *fake surgeon*: that attribute occurs in all members of the category *surgeon* and in no members of the contrast set for *fake surgeon*. That item also has perfect evidence for membership in the category *non-surgeon*, because the attribute <medical-training:none> is perfectly diagnostic for that category: all items which are members of the category *non-surgeon* have that attribute. The overall evidence for membership of this second item in the combination *fake surgeon* is thus high. Note that this second item, despite having a high degree of membership in the combination

fake surgeon, would nonetheless be a poor member of the category *surgeon* occuring on its own. For *surgeon* occuring alone the attribute <clothes:white coat> has low diagnosticity; an item having that attribute but not the attribute <medial-training:surgery> would be a poor member of the category. Only in context of *fake surgeon* does <clothes:white-coat> become diagnostic, allowing the item to be a good member of that combination.

The above examples show that the diagnostic evidence model can explain classification in privative, intersective and subsective adjective-noun combinations, as well as in noun-noun conjunctions. The model goes beyond other current theories in giving a single explanation for the various different patterns seen in these different combination types.

4 Conclusions

The diagnostic evidence model gives a unified account for classification in noun-noun conjunctions and in intersective, subsective and privative adjective-noun combinations. An important part of this account is an explanation for how changes in the contrast set for combinations can influence classification in those combinations. For intersective adjectives the changes in contrast set when an adjective and a noun are combined are minor, and have little influence on classification. For subsective adjectives, on the other hand, the contrast set for an adjective-noun combination is significantly different from the contrast set for the adjective category, and so has a significant influence on the classification. In combination with different nouns (combinations with different contrast sets) the diagnosticity of attributes for a subsective adjective category will change: in one combination one set of attributes will be diagnostic for the adjective category; in another combination other attributes will be diagnostic. For privative adjectives, finally, the change in contrast set for an adjective-noun combination is greatest. This change can reverse the normal pattern of membership in the noun constituent of the combination, so that an item which would be a good member of the noun category occuring singly would be a poor member of the privative adjective-noun combination.

There are a number of other adjective types which the model does not currently address. Because the model's representations use discrete-valued attributes (which are either present or absent in a given item), it cannot address combinations involving graded or continuously-valued adjectives such as tall or old. The model is also unable to address combinations involving adjectives such as former or future, which are based on relative location in time. Again, the model's representations cannot apply to such adjectives. An extended version of the model using continuous rather than discrete attribute representations is currently under development. This extension may provide a single account for all different types of combined phrase.

References

Costello, F. J. (2000a). An exemplar model of classification in single and combined categories. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, 95-101.

Costello, F. J. (2000b). Testing a computational model of categorisation and category combination: Identifying diseases and new disease combinations. In *Proceedings of AICS 2000, 11th Irish Conference on Artificial Intelligence & Cognitive Science*. Galway, Ireland: World Scientific.

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.

Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint guided conceptual combination. *Cognitive Science, 24(2):* 299-349.

Franks, B. (1995) Sense Generation: A Quasi-Classical Approach to Concepts and Concept Combination. *Cognitive Science 19(4)*: 441-505.

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.

Kamp, H. (1975). Two theories about adjectives. In *Formal Semantics of Natural Language*. Cambridge University Press.

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.