Proceedings of the Annual Meeting of the Cognitive Science Society

Learning relations between concepts: classification and conceptual combination

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 26(26)

ISSN
1069-7977

Authors
Devereux, Barry
Costello, Fintan

Publication Date
2004

Peer reviewed

Learning relations between concepts: classification and conceptual combination

Barry Devereux (Barry.Devereux@ucd.ie), Fintan Costello (Fintan.Costello@ucd.ie)
Department of Computer Science, University College Dublin,
Belfield, Dublin 4, Ireland.

Abstract

People interpret noun-noun compounds like “wind power” by inferring a relational link between the compound’s two constituent concepts. Various studies have examined how people select the best relation for a compound from a set of candidate relations. However, few studies have investigated how people learn such relations in the first place. This paper describes an experiment examining how people learn which relations are possible between concepts. Participants in this experiment learned artificial, laboratory controlled relations between pairs of items and then judged how likely those relations were for new pairs of items. The results showed that people’s judgement of relation likelihood was reliably influenced by the presence of facilitating features for relations and by the diagnosticity of features for relations. A simple exemplar-based model of classification, using both diagnostic and facilitating features, was applied to people’s judgements of relation likelihood. This model accurately predicted people’s judgements of relation likelihood in the experiment, using no free parameters to fit the data.

Introduction

When, in everyday discourse, people encounter noun-noun compounds such as “mountain stream” or “lake boat”, they interpret those compounds by inferring a relation that can be used to combine the two constituent concepts (inferring that a “mountain stream” is a stream that flows down a mountain, that a “lake boat” is a boat that sails on a lake). In theoretical accounts of conceptual combination, this process involves selecting the best relation for a compound from a set of candidate relations. Some theories give a standard set of candidate relations to be used in all compounds (Gagné & Shoben, 1997), while others derive candidate relations from the internal structure of the concepts being combined (Costello & Keane, 2000; Wisniewski, 1997). Many studies have investigated how people select the best relation for a given compound (e.g. Costello & Keane, 2001, Wisniewski, 1996). However, there have been very few studies investigating how people learn and form these relations in the first place. In this paper we aim to fill this gap.

The paper describes an experiment investigating how people learn relations between two sets of novel concepts. In the experiment we designed four different relations that could hold between artificial, laboratory-generated ‘beetle’ and ‘plant’ concepts. Participants learned these relations from sets of examples, with each example showing one sort of relation holding between one type of plant and one type of beetle. After learning, participants were shown new pairs of plants and beetles, and asked to say which of the four learned relations could hold between those two items.

This experiment was designed to examine two different possible factors in people’s learning of relations between concepts: the presence of *diagnostic* features for those relations, and the presence of *facilitating* features. By diagnostic features for a relation we mean features of a constituent concept that are strongly associated with a particular relation. Diagnostic features are most familiar in the case of single concepts: for example “has four legs” and “is made of wood” are diagnostic features for the single concept *chair*: most things that are chairs have those features, and most things that are not chairs do not. Similarly, the feature “has a flat surface raised off the ground” might be diagnostic for the relation *is-sat-on-by*: most instances of the *is-sat-on-by* relation have that feature; most instances of other relations do not. In the experiment we asked whether people would use the diagnostic features for relations when selecting likely relations for beetle-plant pairs.

By facilitating features we mean the features of a pair of concepts that are necessary for a given relation to be possible, and without which that relation cannot hold. For example, while the compound “steel chair” can easily be interpreted using the *made-of* relation, the compound “kitchen chair” cannot possibly be interpreted as “a chair made of kitchens” simply because kitchens are not a type of substance. Being a substance is a necessary facilitating feature for an item to take part in the *made-of* relation. Again, in the experiment we asked how people would use such facilitating features when selecting likely relations for beetle-plant pairs.

This paper is organised as follows. In the next section we discuss the representation of relations in terms of sets of examples, as used in our experiment. We then describe the experiment in detail. To foreshadow the results, we found that both diagnostic and facilitating features had a reliable influence on people’s selection of likely relations for pairs of items. We then describe how an exemplar-based model of concept conjunction (Costello, 2000, 2001) can be applied to the results of this experiment, giving a close fit to people’s judgements of relation likelihood. Finally, we conclude by discussing the implications of our findings for theories of conceptual combination.

Learning Relations from Exemplars

Our primary assumption is that the relations selected during conceptual combination are essentially categories, just as the concepts that they link are essentially categories. We use an exemplar representation to describe these relational categories. Exemplar theories of classification, which propose that a category is represented as the set of remembered instances of that category and that new items are classified on the basis of their similarity to those instances (e.g. Medin & Schaffer, 1978; Nosofsky, 1984), have successfully accounted for a number of patterns seen in people’s learning of single categories. We extend the exemplar approach to allow both relations and the concepts that they link to be represented by sets of instances.

In a category representing a single concept, each exemplar consists simply of a single set of features. For a category representing a relation, however, each exemplar consists of two sets of features: the features of the two single-category exemplars that are being linked by that relation. For example, suppose we have two categories *A* and *B* consisting of the set of exemplars $\{a_1, a_2, a_3, a_4, a_5, a_6\}$ and $\{b_1, b_2, b_3, b_4, b_5\}$ respectively. Each category represents a single concept, and each exemplar contains features describing one example of that concept. We can compute the membership of a given item in category *A*, for example, by comparing that item to the set of stored exemplars of category *A*. Then a relation *R* linking the concepts *A* and *B* might be represented as the set of exemplars of that relation, for example $\{(a_1, b_1), (a_1, b_2), (a_3, b_3), (a_4, b_4), (a_4, b_5)\}$. Regarding *R* as a category, we can compute the membership of any pair of items (x, y) in the relation *R* by comparing that pair to the set of exemplars of that relation. If more than one relation is defined, we can compute membership in each of the relations and make assertions about which relation the given pair of items is most likely to belong to.

This representation of relations in terms of a collection of pairs of category exemplars is motivated by how mathematical relations are defined in set theory. In set theory, binary relations are defined as sets of ordered pairs. To take a well-known example, the "is equal to" relation is a set denoted by =, and is defined on the integers to be the set $\{\ldots, (-1, -1), (0, 0), (1, 1), (2, 2), \ldots\}$. Thus the "is equal to" relation holds between two integers *x* and *y* if and only if the ordered pair (x, y) is a member of the set denoted =. We extend this set-theoretic idea of relations and propose that a relation between two concepts can be represented as a set of relation exemplars, where each of these relation exemplars is an ordered pair of category exemplars. Membership of a pair of items in a relational category is then computed by comparing that pair of items to the stored exemplars of that relational category, as is precisely the case for exemplar models of simple categories.

This approach assumes that when people are selecting the correct relation for a pair of items they may be performing a classification task in which they compare that pair of items to various sets of relation exemplars. What factors would we expect people to be influenced by in such a classification task? First, we would expect people pay attention to the features in items that are diagnostic for particular relations (that is, features that are present in most of the items that take part in the relation, and absent in most items that do not take part in the relation). If a particular feature is diagnostic for a relation, people should use that feature to identify new items likely to take part in that relation. Such a result would be consistent with other results in the classification literature, which reveal that people are attentive to diagnostic features when making determinations of category membership.

Second, we would expect people to pay attention to whether or not a given item has the facilitating features required for a given relation (as in our "kitchen chair" example). If a particular feature is present in every item that takes part in a certain relation, then we can assume that that feature may be necessary for that relation to take place: the feature may facilitate that relation. When confronted with a new pair of items which do not possess that facilitating feature, we would expect people not to select that relation for that pair of items. Note that a facilitating feature for a given relation may also be a diagnostic feature for that relation (if it occurs in every item that takes part in that relation *and* in no items that do not take part in that relation). However, a facilitating feature for a relation may also be non-diagnostic for that relation (if the feature occurs in every item that takes part in that relation, but also occurs in many items that do not take part in that relation). Next we describe an experiment examining the influence that facilitating and diagnostic features have on people's selection of relations for pairs of items.

A Categorisation Experiment

This experiment aims to test three hypotheses: that people can learn relations from sets of examples of those relations; that diagnostic features are important in people's selection of relations for pairs of items; and that facilitating features are also important in relation selection. The design of the experiment is essentially the same as other experiments in the category learning literature: a preliminary training phase where participants are exposed to exemplars of different artificial, laboratory controlled categories is followed by a test phase where participants are presented with new items and are asked to make judgements of category membership. The categories in this experiment are four different relations that can hold between pairs of objects. Each of the training items consists of two objects linked by one of these relations. The test items also consist of two objects; however these objects are not linked by any relation, and participants are asked to judge the likelihood of different relations holding between these items.

To examine the role of facilitating features in relation selection, the training items were designed so that two of the learned relations had facilitating features: we called these two relations the facilitated relations. Every time one of these relations occurred in the set of training items, a particular beetle or plant feature (the facilitating feature) was also present in that item. The other two relations did not have facilitating features: we called these two relations the independent relations. Similarly, the training items were designed so that some beetle and plant features were particularly diagnostic for some relations, and some features were not. In the training items, the diagnostic features for a particular relation occurred most frequently in beetle or plants taking part in that relation, and occurred rarely in other relations. The pairs of objects used in the test phase of the experiment consisted of various combinations of facilitating and diagnostic features for different relations. By examining participants' choice of relations to link the objects in these test items, we can then assess the influence of facilitating and diagnostic features in relation selection.

Method

Participants. 16 postgraduate students or recent college graduates volunteered to take part in the experiment. All were native speakers of English.

Materials. The materials for the training phase consisted of 18 visual stimuli on an A5-sized card depicting a cartoon

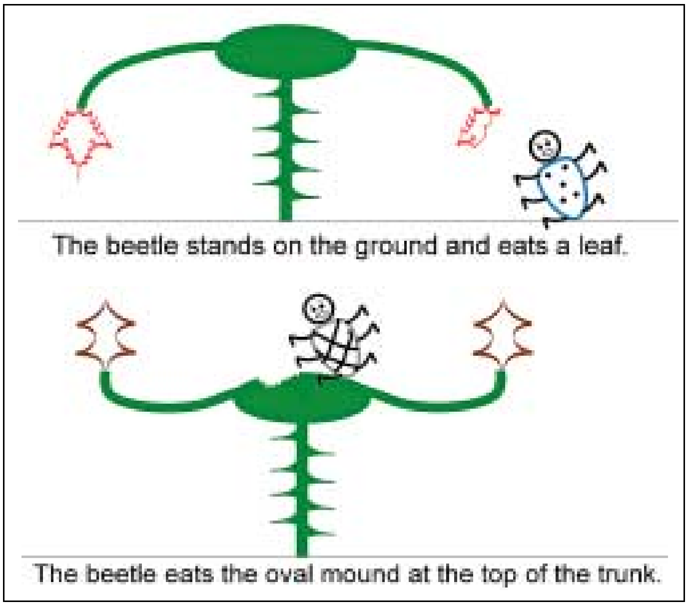

Figure 1: Two training phase stimuli.

beetle eating a plant. The beetles varied on three quaternary-valued feature dimensions: color of the shell, pattern on the shell, and facial expression. The plants varied on four feature dimensions: color of the leaves (quaternary-valued), shape of the leaves (quaternary-valued), droop of the branches (binary-valued), and whether there were buds or thorns on the trunk (binary-valued). There were four possible ways in which a beetle could eat a plant, corresponding to the four relational categories: the beetle could land on a leaf of a plant and eat the leaf (an independent relation), the beetle could eat from the top of the trunk of the plant (an independent relation), the beetle could eat from the trunk of the plant if there were buds rather than thorns on the trunk (a facilitated relation), or the beetle could stand on the ground and eat the leaf of a plant that had drooping branches (a facilitated relation). Underneath each picture was a sentence describing the eating behaviour that was talking place. Examples of the training phase's stimuli are presented in Figure 1.

These 18 items described a category structure for four relations, each relation being one of the different ways in which a beetle could eat a plant. The distribution of beetle and plant features was controlled so that some features would be facilitating for relations and so that some features would be diagnostic for relations. The distribution of beetle and plant features across the four different relation categories is shown in abstract form in Table 1. The numerical values in columns B1, B2 and B3 represent the different possible features that beetles could have; the values in columns P1, P2, P3, and P4 similarly represent the different possible features of plants. Each row in this table represents a different particular exemplar of one of the four relations R1, R2, R3 and R4.

In this experiment, we were interested in the influence which the distribution of features across items would have on people's selection of relations between items. We were not concerned with any effect which the physical properties of stimuli (e.g. the salience of different colours, the distinctiveness of different shapes) would have on people's relation selection. Thus, while each participant's set of training items had the abstract structure shown in Table 1, each participant saw a unique set of physical stimuli. The abstract-to-physical mappings for the category dimensions and values for the two independent relations and for the two facilitated relations were balanced across participants. This was done so that the physical dimensions, values and relations would not be confounded with their abstract counterparts.

Table 1: The abstract relational category structures used in the training phase.

Item	Relation	Insect Features			Plant Features			
		B1	B2	B3	P1	P2	P3	P4
1	R1	1	1	1	1	1	1	2
2	R1	4	1	1	1	1	1	1
3	R1	2	2	2	4	1	1	2
4	R1	3	4	2	1	4	2	1
5	R1	3	3	3	1	1	1	1
6	R2	1	1	2	3	2	1	2
7	R2	2	2	2	3	2	1	1
8	R2	2	3	4	1	1	1	1
9	R2	3	4	3	2	3	2	2
10	R3	1	4	1	4	2	2	1
11	R3	1	1	4	2	2	2	1
12	R3	2	2	2	2	3	2	2
13	R3	3	2	4	2	3	2	1
14	R3	2	2	3	3	2	2	1
15	R4	1	3	1	4	3	2	2
16	R4	2	3	3	2	2	1	2
17	R4	3	2	3	3	3	1	2
18	R4	3	3	3	3	3	1	2

Facilitating and Diagnostic Features. The four relations in Table 1 were designed so that two relations had facilitating features (R3 and R4) and two did not (R1 and R2), and two relations had highly diagnostic features (R1 and R3) and two had less diagnostic features (R2 and R4). Relations R1 and R2 were the independent relations, while relations R3 and R4 were the facilitated ones. P3 is the facilitating dimension for relation R3: every exemplar of R3 involves an item with a value of 2 on dimension P3. Similarly P4 is the facilitating dimension for relation R4: every exemplar of R4 involves an item with a value of 2 on dimension P4. (In the experimental materials, the facilitating features were instantiated in a causally meaningful way. For example, the physical relation "the beetle stands on the ground and eats the leaf" had the facilitating feature "drooping branches on the plant").

Of the two independent relations, R1 has highly diagnostic features while R2 has less diagnostic features. Relation R1 has two highly diagnostic features: a 1 on P1 and a 1 on P2. Relation R2, however, has no particularly diagnostic features. (R2 is therefore a very vague category, not well distinguished by either diagnosticity or facilitating features). Of the two facilitated relations, R3 has highly diagnostic features while R4 has less diagnostic features. For relation R3, a 2 on dimension P3 is a highly diagnostic feature for that relation, occurring five out of five times in examples of that relation and only three times outside it. (Note that this feature, a 2 on P3, is also the facilitating feature for relation R3). Relation R4, however, has no such highly diagnostic feature, although a 3 on B2 and a 3 on B3 are both moderately diagnostic for that relation. (R4's facilitating feature is not very diagnostic, occurring four times within the category and five times outside it.)

The materials for the test phase consisted of more visual stimuli depicting beetles and plants; however in these pictures the beetles and plants were shown separately, without any eating or any other interaction taking place. Underneath each test picture was the question "How likely are the different types of eating behavior?", followed by the four relation description sentences, each of which was accompanied by a 7-point scale ranging from -3 (labelled "not at all likely") to +3 ("extremely likely"). The order in which the four scales were presented was balanced across participants.

The test phase consisted of 29 beetle-plant pairs. Of these, nine pairs were selected from the 18 beetle-plant pairs that had been presented in the training phase, but now without the eating behaviour shown. These nine previously-seen pairs were used to assess how accurately participants learned the training items they had studied. The remaining 20 test items (beetle-plant pairs) had not been seen previously by participants. For these items the properties of interest are whether or not the facilitating feature of relation R3 or of relation R4 is present, and whether or not the item had features diagnostic for particular relations.

Procedure. The experiment consisted of two sections: a training phase where participants studied the training items (pairs of beetles and plants taking part in particular relations), and a test phase where they had to rate the likelihood of the different possible relations for a sequence of beetle-plant pairs. Participants were asked to pretend to be biologists interested in learning about imaginary plants and beetles and the relationship between them. The seven dimensions on which the beetles and plants could vary were explicitly pointed out to participants. It was pointed out to participants that they might find it useful to try to learn about the eating behaviour by looking for relationships between features and types of eating, or by learning the features of individual examples and remembering the type of eating occurring with them. Participants spent about five minutes reading the instructions, during which time the experimenter answered any questions they had. After reading the instructions participants were presented with the 18 training items at a large desk area. Participants were given 12 to 15 minutes to study the training items.

After the training phase, the 18 training items were removed and participants were given the 29 test items. Participants were first shown the nine test items corresponding to items they had studied in the first part of the experiment. Participants were told to mark an integer value on each of the four scales describing how likely they felt the four possible types of eating behaviour were. Following these nine items 20 new test items were presented to the participants. The order in which the items were presented was randomized for each participant, and participants were allowed to rate the items at their own pace.

Results

Participants' learning of the training items. For these nine items there was a "correct" relation (each item was a member of one category during the learning phase) and three "incorrect" relations (corresponding to the other three categories). The responses for each relation and each experimental item were classified as either positive (> 0) or non-positive (≤ 0), depending on how the participant responded on each scale. On average, participants gave a positive rating to correct relations 71% of the time and a positive rating to incorrect relations 33% of the time. Two participants gave a positive score to only four correct relations; these two participants were excluded from the analysis. The remaining 14 participants gave a positive rating to correct relations 75% of the time and a positive rating to incorrect relations only 25% of the time. These results indicate that participants learned to distinguish between the categories in the training phase.

Participants' sensitivity to facilitating features. One-tailed binomial tests with $\alpha = 0.05$ were used to identify whether the presence or absence of the facilitating features for a relation had an effect on how participants responded when grading the likelihood of that relation. The proportion of positive responses for each of the four relations was the statistic of interest.

First we considered the items in which the facilitating features for a given relation were absent. Of the 29 test items, 16 were items for which the facilitating feature for relation R3 was absent and 16 were items for which the facilitating feature for relation R4 was absent. For relation R3, the binomial test was significant for 13 of the 14 participants; in other words, 13 of the 14 participants were significantly more likely to produce a non-positive rather than a positive response to relation R3 when the facilitating feature for relation R3 was absent. (Indeed, 10 participants *never* produced a positive response). For relation R4, 10 of the 14 participants were significantly more likely to produce a non-positive rather than a positive response. (Here, five participants never produced a positive response).

A similar analysis was performed looking at the items where the facilitating feature was present. Of the 29 test items, 13 were items for which the facilitating feature for relation R3 was present and 13 were items for which the facilitating feature for relation R4 was present. For relation R3, 8 of the 14 participants were significantly more likely to produce a positive rather than a non-positive response. For relation R4, 7 of the 14 participants were significantly more likely to produce a positive rather than a non-positive response. These results are sensible considering that in many cases participants will rate a relation as having low likelihood for a given item, even when that relation's facilitating feature is present in the item: the facilitating feature doesn't mean that the relation must be selected for this item, only that it is a possibility. The difference in these results between items that had and items that had not the facilitating feature for a relation indicate that participants were highly sensitive to the presence and absence of those features.

Participants' sensitivity to diagnostic features. The diagnosticity of a feature for a category is a measure of how good that feature is at identifying membership of that category. If a feature appears in many items in a category and few items outside a category then that feature will have high diagnosticity. More formally, we can define the diagnosticity of a feature *f* for a category *C* to be

$$D(f,C)=\frac{\left|C\cap E_f\right|}{\left|C\cup E_f\right|} \qquad \text{(Eq. 1)}$$

where E_f denotes the set of exemplars that have feature *f*. Using this formula we calculated the average diagnosticity of the features of each of the 29 test items for each of the four relational categories and compared this to the observed data. For two of the four relations, the amount of diagnosticity for items had a high correlation with the observed membership ratings for the items (for R1, $r = 0.83$, $p < 0.01$, *%var* = 69%; for R4, $r = 0.81$, $p < 0.01$, *%var* = 65%). For the other two relations, the correlation was less strong though still significant (for R2, $r = 0.66$, $p < 0.01$, *%var* = 43%; for R3, $r = 0.70$, $p < 0.01$, *%var* = 49%). These results indicate that participants were sensitive to diagnostic features when making their category judgements.

Diagnostic and facilitating features. Relations R1 and R2 were the independent relations: membership in these relations did not depend on facilitating features. Diagnosticity was very important for identifying members of R1 but was not very important for identifying members of R2. We would therefore expect the correlation of diagnosticity to the observed memberships to be higher for relation R1 than relation R2; this is the case in the above analysis of the effect of diagnostic features.

Relations R3 and R4 were the facilitated relations: membership in these relations depended on both diagnostic features and facilitating features. R3 was designed to have highly diagnostic features, while R4 was designed to have less diagnostic features. The diagnosticity analysis above, however, shows that the correlation of the observed memberships with diagnosticity was lower for R3 than for R4. The divergence between diagnosticity and membership ratings for these relations suggests an interaction between diagnosticity and facilitating features.

As a way of examining this interaction we looked at the total number of positive and non-positive responses across all participants for relations R3 and R4; they are presented in Table 2. For cases where the facilitating feature is present, there are less positive responses for R4 than for R3 and more non-positive responses for R4 than R3. Conversely, for cases where the facilitating feature is absent, there are less positive responses for R3 than for R4 and more non-positive responses for R3 than R4. These data suggest that participants are more sensitive to the presence or absence of the facilitating feature for R3 than they are for R4. This is consistent with the fact that the facilitating feature for R3 is more diagnostic for R3 than the facilitating feature for R4 is diagnostic for R4. In other words, people seem to be using both the facilitating nature and the diagnosticity of features together in deciding relation likelihood for these relations. In the next section we investigate this interaction between diagnosticity and facilitating features in more detail by applying a model of classification to our data.

Table 2: Total number of positive and non-positive responses across all participants

	Facilitating Feature Present		Facilitating Feature Absent	
	Positive	Non-positive	Positive	Non-positive
R3	130	52	14	210
R4	123	59	30	194

Modelling Relation Selection

We are interested in how people used facilitating and diagnostic features when making judgements of relation likelihood in our experiment. As we have seen, our results indicate an interaction between diagnosticity and facilitating features. We are also interested in whether or not our view of relations in terms of exemplar-represented categories can successfully account for the results of the above experiment. To explore these issues, we examined whether an exemplar model of categorization could be used to model participants' responses of how likely each relation is for each item in the experiment. We used as our starting point Costello's (2000, 2001) Diagnostic Evidence Model (DEM) which uses diagnostic features to model classification in concept combination. This model calculates an evidence score for an item *x* in a category *C* using the diagnosticity of each feature of the item for that category according to the formula

$$E(x,C)=1-\prod_{f\in F_x}\left(1-D(f,C)\right) \qquad \text{(Eq. 2)}$$

where F_x is the set of features of *x* and $D(f, C)$ is computed as in Equation 1. (Equation 2 essentially sees category membership as a disjunction of the feature diagnosticities and is not dissimilar to simply averaging the diagnosticities as we did in the previous section). As a preliminary step we applied this model to the experimental data without using any information about facilitating features: the model is only sensitive to features' diagnosticity. In this form the model still produces a reasonable fit to the data (for R1, $r = 0.88$, $p < 0.01$, *%var* = 78%; for R2, $r = 0.63$, $p < 0.01$, *%var* = 40%; for R3, $r = 0.73$, $p < 0.01$, *%var* = 53%; for R4, $r = 0.72$, $p < 0.01$, *%var* = 52%), with no free parameters.

The model in this form uses diagnostic information alone. However, the results of our experiment indicate that people make use of both diagnosticity and facilitating features in determining the relations. One possible account of how people use both these types of information is that people are applying diagnosticity information after they have been constrained by the presence or absence of the facilitating features. Perhaps people check if an item has the necessary facilitating features for a particular relation and then, if it does, use the diagnostic evidence of the features. In modifying the model we therefore assume that participants' do not use known exemplars which depict a relation that is impossible for the test item at hand: facilitating features restrict the universe of discourse so that membership of an item in a relational category is a calculation across the subset of the learned exemplars that belong to relational categories that are possible for the current item. Our formula for diagnosticity then becomes

$$D(f,x,C)=\frac{\left|C\cap E_f\cap R_x\right|}{\left|\left(C\cup E_f\right)\cap R_x\right|} \qquad \text{(Eq. 3)}$$

where R_x is the set of known exemplars that belong to relations that are not impossible given the features of x. This modified model gives a much closer fit to the observed relation selection ratings (for R1, $r = 0.89$, $p < 0.01$, $\%var = 79\%$; for R2, $r = 0.78$, $p < 0.01$, $\%var = 61\%$; for R3, $r = 0.98$, $p < 0.01$, $\%var = 0.96$; for R4, $r = 0.96$, $p < 0.01$, $\%var = 0.92$), again with no free parameters. Clearly, both information about the diagnosticity of features and information about the presence or absence of facilitating features are required to accurately model the experimental data.

Though this modified DEM model may not be the best way of modelling the data, it does suggest that using information about both the facilitating features and diagnosticity of features of an item are important in selecting relations. The fact that an exemplar model of classification can predict how people rate the likelihood of different relations linking pairs of items is also evidence in support of the hypothesis that relations can be represented as categories. This suggests that relation selection can be thought of as a kind of classification task.

Conclusions and Discussion

Our study yields three findings. First, people can learn which relations are possible between concepts from sets of examples of those relations. Second, people pay attention to facilitating features for those relations and use those features when judging relation likelihood for new examples. Third, people also pay attention to, and use, diagnostic features for those relations. Such findings are consistent with our hypothesis that relations can be represented with an exemplar category structure, and that the selection of a relation between two constituents can be seen as a categorisation task. These findings have implications for current theories of how relational links are used in conceptual combination. In particular, these findings may be problematic for Gagné & Shoben's (1997) CARIN model, which proposes that in conceptual combination people select the correct relational link between two concepts from a fixed set of 16 relational links called *thematic relations*. First, the thematic relations used in the CARIN model have no internal structure: there is no way, in that model, in which facilitating or diagnostic features could be associated with those relations (for example, the MADE-OF relation in the CARIN model how has no way of requiring that a concept taking part in it is type of substance). Furthermore, the four different relations we used in our experiment do not occur in the CARIN model's fixed set of thematic relations: it would be hard for that model to explain how people used these relations in our experiment.

Our findings are consistent with other theories of conceptual combination (e.g. Costello and Keane, 2000; Murphy, 1988; Wisniewski, 1997) which do allow internal conceptual structure to influence relation selection. These theories use some variation of the idea that a concept representation can contain 'slots' such as MADE-OF or LOCATED and that conceptual combination involves one concept filling a slot in another concept (so that "kitchen chair", for example, would involve the "kitchen" concept filling the LOCATED slot in "chair"). The exemplar-based model of relation selection described in this paper provides an alternative to this slot-based representation of relations, showing that relations can be represented as sets of paired-item exemplars, rather than as slots in concepts. This exemplar-based model has the advantage of giving a simple account of how people learn which facilitating and diagnostic properties are associated with each relation.

As for future work: in our experiment, participants learned relational categories only, and did not learn conceptual categories (they did not learn different categories of beetle or plant, for example). A possible extension of this work would be to have participants learn, from sets of exemplars, both conceptual categories and the relational categories that link them. This experiment could reveal more both about how relations are learned and used, and about how exemplar-level and conceptual-level information interact in conceptual combination.

Acknowledgements

This research was supported by a grant from the Irish Research Council for Science, Engineering and Technology, funded by the National Development Plan.

References

Costello, F. J. (2000). An exemplar model of classification in simple and combined categories. *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, (pp 95-100). Mahwah, N. J.: Erlbaum.

Costello, F. J. (2001) A computational model of categorisation and category combination: Identifying diseases and new disease combinations, *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, (pp. 238-243). University of Edinburgh: Erlbaum.

Costello, F., & Keane, M.T. (2000). Efficient Creativity: Constraints on conceptual combination. *Cognitive Science, 24,* 299-349.

Costello, F. J., and 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,* 255-271.

Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory and Cognition, 23* (1), 71-87.

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

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, 12, 529-562.

Nosofsky, R. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* (1), 104 114.

Wisniewski, E. J. (1996). Construal and similarity in conceptual combination. *Journal of Memory and Language,* 35, 434-453.

Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review,* 4(2), 167-183.