British Journal of Psychology (2004), 95, 371–386

www.bps.org.uk

Information about the logical structure of a category affects generalization

Emmanuel M. Pothos[1]*, Nick Chater[2] and Andrew J. Stewart[3]
[1]Department of Psychology, University of Crete, Greece
[2]Department of Psychology, University of Warwick, UK
[3]Department of Psychology, University of Manchester, UK

This paper considers whether information about the logical structure of a category affects how people generalize. We carried out three experiments with the following structure: participants were first presented with a set of training items, and were subsequently asked to decide whether new items belonged to the same category as the training items. Each experiment had two conditions that differed only in terms of the category label provided for the training items; different category labels conveyed different information about the logical structure of the category to which the training items were supposed to belong. In all cases, participants' generalization was greatly affected by such information. Our results suggest that people make the default assumption that category labels correspond to groupings of highly similar objects.

A person presented with a set of objects that are assumed to be members of the same category will have some idea as to how these objects should be generalized to other novel objects. Now, suppose that, additionally, the person is provided with some information about the logical structure of the category. We are interested in whether such information can affect the spontaneous expectations the person will have as to how the objects should be generalized.

There have been some previous studies indicating that this may be the case. For example, Goldstone (1996) used items that were labelled in ways that conveyed different information about the logical structure of the items' category; for example, As and not-As as opposed to As and Bs. Goldstone found that how the items were labelled affected whether people perceived them as belonging to interrelated categories (so that the emphasis on objects' perceptions would be in terms of diagnostic features) as opposed to isolated categories (where objects would generally be perceived in terms of a larger number of non-diagnostic features). Note that it is useful to distinguish such work from research that investigates how the logical form of a category (e.g. A OR B vs. A AND B) influences the rate of learning of the category (e.g. Dennis, Hampton, & Lea,

*Correspondence should be addressed to Dr Emmanuel Pothos, Department of Psychology, University of Crete, Rethymnon 74100, Greece (e-mail: pothos@psy.soc.uoc.gr).

1973; Shepard, Hovland, & Jenkins, 1961), as well as how negating one component of a composite concept affects perception of the concept (Hampton, 1997).

If describing a set of objects with a logical combination of default labels influences the way in which the objects are generalized, it is possible that describing objects with a single default label influences their generalization as well. Flannagan, Fried, and Holyoak (1986) showed that adults generally assume that the (properties of the) exemplars of a category would conform to a unimodal and possibly normal distribution, when the category is labelled in a default way. The corresponding research with children is more extensive; for example, Markman and Hutchinson (1984) observed that the presence of a category label induces children to generalize a set of objects in a taxonomic instead of a thematic way. Markman (1987, 1991) also noted that children avoid applying two labels to the same set of referents and that a novel noun in conjunction with a newly encountered object will be taken to correspond to the whole object.

It appears that the linguistic labelling of a category may affect assumptions about the category structure and therefore how the category would be extended to novel items. Related to this appears to be the research tradition on how teaching a particular category structure for a set of objects can affect how the objects will be perceived. For example, after learning to classify a set of objects in some particular way, the similarity of objects in the same category might increase and the similarity of objects in different categories decrease (e.g. Goldstone, 1994a). Introducing a category boundary in a dimension of stimulus variation will generally make the stimuli on either side of the boundary less similar (e.g. Livingston, Andrews, & Harnad, 1998; Özgen & Davies, 2002). Categorizing an object in a category might lead to a representation of some of the objects' features as more salient or the creation of new object features (e.g. Goldstone, 1995, 2000; Schyns, Goldstone, & Thibaut, 1997). More generally, differences in linguistic labelling of the same objects (in different cultures) sometimes implicate differences in how the objects/stimuli are represented (e.g. Roberson, Davies, & Davidoff, 2000).

This research is only partly relevant to the present project: each object can, in principle, be represented in several ways (e.g. Goodman, 1954). Therefore, when the cognitive system is required to learn a particular categorization for a set of objects, it presumably has mechanisms for changing the object representations so as to make it easier to classify the objects in the required way. We are interested in *spontaneous* changes in how a set of objects will be generalized, depending on whether the objects are labelled as 'As' or 'not-As' or 'As and Bs' or 'As or Bs' and so forth, in a way analogous to that of Goldstone (1996). In other words, does information about the logical structure of some objects' category spontaneously affect how the objects will be generalized to other novel ones? As we shall see, our experiments suggest that this might be the case.

We examined this issue by using objects whose distributional characteristics are intuitively consistent with different kinds of category structures. For example, consider the schematic representation of Experiment 1 in Fig. 1: each dot represents a training instance in the experiment, so that dots closer together imply that the corresponding objects are more similar to each other (Shepard, 1987). Such a set of objects would broadly appear to be compatible with a single convex category (as in Fig. 1a) or a bimodal category with a concave structure (as in Fig. 1b; such a category is of course formally equivalent to two convex categories, cf. McKinley & Nosofsky, 1995). Would it make a difference whether the training instances are described as 'As or Bs' as opposed to 'As'? One possibility is that participants will generalize the training instances only on the basis of their distributional characteristics. Thus, participants would decide whether the objects belong to a single convex or two convex groups, and generalize accordingly.

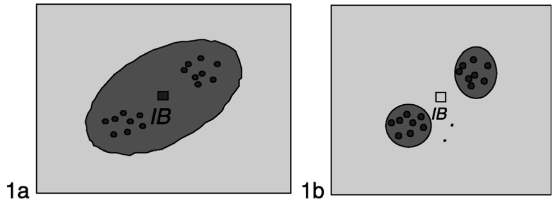

Figure 1. Schematic representation of different possible generalizations in Experiment 1. Generalization assuming that the category corresponds to a single convex region is shown in 1a and generalization assuming that the category corresponds to two convex regions in 1b.

The alternative possibility is that information about the logical structure of the category to which the objects belong would be taken into account in participants' expectations and assumptions about the form of the category, so that generalization would partly depend on how the objects are labelled. These two possibilities can be distinguished in terms of whether or not the test instance IB is classified in the same category as the training items.

In this study, all experiments had the above structure: participants were presented with a set of training items and were subsequently shown a set of test items. They were asked to decide which of the test items belonged in the same category as the training items. There were two conditions in each experiment that differed only with respect to how the training items were labelled. The labels in the two conditions of each experiment were differentiated in terms of the information conveyed about the logical structure of the category of the training items. By examining whether participants generalized differently in the two conditions of each experiment, we can clarify whether the assumptions about the category to which the training items belong are affected by labelling information.

Similarity

In this study the collected empirical data relate to how the same set of training objects is generalized to new instances, depending on how the training objects are described. Generalization is assumed to be a function of how similar the training objects are to novel objects, and possibly the information conveyed by the linguistic labelling of the objects. It is clearly important to control as carefully as possible the ways in which participants are likely to determine the compatibility of the training instances to the test ones. Goodman (1972) pointed out that there are so many ways in which similarity can be measured that in effect it is a scientifically useless concept. Goldstone (1994b) discussed the explanatory adequacy of similarity to suggest that perceptual (physical) similarity is a well-defined notion, but if one starts introducing abstract properties then this may not be the case (for further discussion, see Goldstone, 1998; Goldstone & Barsalou, 1998; Pothos, in press).

We chose stimuli for which the basis for their generalization is either physical similarity (e.g. circles that vary in diameter – Experiment 1) or an equally obvious abstract property (e.g. patterns that are obviously nearly symmetric – Experiment 2; whether a car is sports or non-sports – Experiment 3). In this way, we hoped to minimize the degree to which participants encoded the stimuli in terms of unintended abstract properties or spurious, idiosyncratic influences of general knowledge

(Murphy & Medin, 1985), so that we could assume that our participants represent the stimuli in very similar, if not identical, ways (cf. Goldstone, 1994b; Goodman, 1972). Nevertheless, general knowledge influences are an important aspect of the categorization process. There is extensive evidence that, at least in some cases, people's assumptions about the type of category to which a set of items belongs will depend on their naïve understanding of the world (Gelman & Wellman, 1991; Murphy & Medin, 1985). Thus, in Experiment 3 we tried to extend the design of Experiments 1 and 2 using realistic materials. A complication of such a design is that there may be cultural differences in people's ontological knowledge that could correspondingly affect the categorization process (Sera, Gathje, & Pintado, 1999; but see Lopez, Atran, Coley, Medin, & Smith, 1997). For Experiment 3, our only assumption is that the influence of general knowledge on how participants represented the stimuli would be analogous, because all our participants (in that experiment) came from a relatively homogenous population: they were all undergraduate students at the same university.

As we have discussed, the objective of the study was to establish whether information about the logical structure of the category to which a set of items is assumed to belong affects the way the items are generalized. If we find that such information affects generalization, can we conclude that the perceived similarity of the objects changes as well? In this study, we do not examine this issue and no conclusion is forthcoming about how information about the logical structure of a category interacts with similarity information to affect generalization. Of course, we mentioned above the extensive literature on how categorizing a set of objects in different ways frequently alters the similarity structure for these objects. If we find that labelling the training items in different ways alters the way they are generalized, would we be looking at effects like those reported by Goldstone (1994a)? The answer is most likely no, because such research concerns changes in the similarity of stimuli that have been specifically associated with fixed (experimenter defined) categories. In our case, we are interested in whether information about the logical structure of a category to which a set of items belongs would spontaneously alter the way the items would be generalized to other new ones (cf. Goldstone, 1996).

Implications from formal models of categorization

Research on formal models of categorization involves primarily two traditions, exemplar models and prototype models (for reviews, see Hahn & Chater, 1998; Komatsu, 1992). According to a common type of exemplar theories, a new instance will be more likely to be classified as an X as opposed to a Y if the average similarity of this new instance to all the Xs is higher than the average similarity to all the Ys (e.g. Hintzman, 1986; Medin & Schaffer, 1978; Nosofsky, 1989, 1991). On the basis of prototype models, a new instance will be more likely to be considered an X instead of a Y if it is more similar to the prototype of Xs than to the prototype of Ys. The prototype of a category is the central tendency of all the members of the category (e.g. Homa & Chambliss, 1975; Posner & Keele, 1968; Reed, 1972; Rosch & Mervis, 1975).

Exemplar and prototype theories represent our most valid, in terms of mathematical rigour and experimental verification, understanding of categorization processes (e.g. see Ashby & Alfonso-Reese, 1995; Ashby & Perrin, 1988; Nosofsky, 1990). It is, therefore, important to examine how the present investigation can be informed by our knowledge of formal categorization models. There are two issues to consider. First, note that all our experiments concern how participants generalize a set of training items to novel exemplars. In this sense, it might appear that formal categorization models do not apply, as such models are typically specified in terms of at least two contrasting categories. However, formal models can be trivially extended in cases in which there is a single category of training instances: the test instances can be ordered with respect to their similarity to the training instances (or to a prototype) and compatibility selections can then be guided by a response bias (e.g. participants might be told that half the test instances are compatible with the training instances). Indeed, Pothos and Bailey (2000) have successfully applied a modification of the Generalized Context Model (Nosofsky, 1989) along these lines in Artificial Grammar Learning.

The second issue applies not just to the present project, but also to any research of how categorization may affect object representations: formal categorization models are specified in terms of similarity, either to prototypes or to previous instances. However, it is rarely the case that similarity itself is specified. If the similarity between a set of objects changes between training and test then the predictions of how the training objects will be generalized change as well. In this study, our objective is simply to establish whether information about the logical structure of a category is taken into account in determining how a set of objects will be generalized. However, if we find that such information is taken into account, our research does not allow us to establish the nature of the changes that lead to differences in generalization. Hence, this research is neutral with respect to formal models of categorization.

EXPERIMENT 1

In this experiment, we used stimuli that were circles varying in diameter. The training stimuli comprised two distinct subgroups, one with circles having fairly small diameters and one with circles having fairly large diameters. In this way, the training stimuli could be seen as belonging to a single concave category or two separate convex ones (a set of X points is convex if for any pair of points x, y in the set, the line $kx + (1 - k)y$ is also in X, for every real number k in the range [0, 1]). The test items included a critical set of instances that were in between the two subgroups of training items (the 'intermediate' test items). The training items were labelled either in a way to cue a single convex category or in a way to cue two convex categories (equivalently a bimodal concave category). Examination of how participants classified the intermediate test items would reveal whether this information about the logical structure of the training items' category influenced participants' categorization. A schematic representation of Experiment 1 is shown in Fig. 1.

Design and participants

The experiment had a between-participants design, with two conditions (25 participants in the single label condition and 26 in the chomps or blibs condition). Participants were either Oxford University or University of Wales Bangor undergraduates and took part in the study for a small payment or course credit.

Materials

The stimuli were circles that varied in diameter and were filled with a neutral pattern. They were chosen to fall naturally into two groups, one group of circles had small diameters (3, 3.5, 4 and 4.5 cm) and another group of circles had large diameters (17.5, 18, 18.5 and 19 cm). In the test phase, participants were asked to decide which of a new set of items were also in this category. The crucial test items, the 'intermediate' items, had diameters roughly half-way between the training items (10, 10.5, 11, 11.5 and 12 cm). Filler test items were also included, chosen to be near in diameter to the training items (2, 2.5, 5.5, 6, 16.5, 17, 19.5 and 20 cm). All the items and the structure of the experiment are shown in Fig. 2.

Procedure

In the training phase of the experiment, participants received a set of items, which were all described using an invented linguistic label. In one condition participants were told that the training items were either 'chomps or blibs' (as if there were two categories); in the other condition, the items were described using a single label ('chomps' for half the participants, and 'blibs' for the other half). Participants were simply asked to examine the items. In the test phase, participants were presented with another set of items and were asked to identify the ones they thought were chomps or blibs/chomps/blibs, depending on the condition.

In all three experiments, in both training and test, participants were allowed to look at the items in that phase in any order and manner they wished. Typically, participants would look at each item once and would spend a few minutes going through the items. Also, in the test phase, participants were told that they were allowed to amend responses. The instructions and stimuli were printed on A4 sheets of paper (one sheet per stimulus).

Results

We classified participants as 'generalizing to the intermediate patterns' or 'not generalizing to the intermediate patterns' on the basis of a simple majority rule. Given that there were five intermediate patterns overall, if participants selected one or two intermediate patterns as compatible with the training stimuli then we considered these participants as not having generalized to the intermediate patterns. If participants selected four or five intermediate patterns as compatible with the training stimuli then the participants were classified as having generalized to the intermediate patterns. Participants who classified three intermediate patterns as compatible with the training items were considered ambiguous on the basis of the majority rule and hence were not included in the analyses. Equivalent versions of this majority rule were used in Experiments 2 and 3.

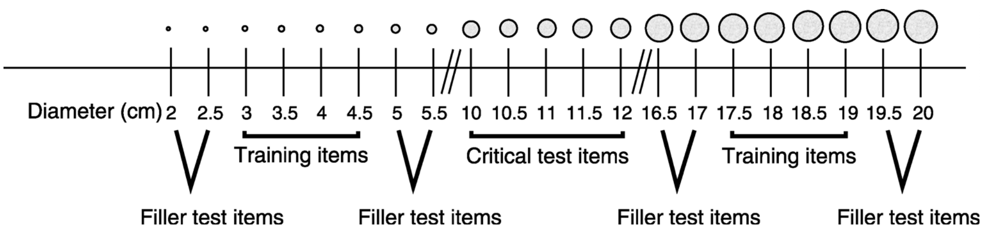

Figure 2. The structure of Experiment 1 and all the stimuli, reduced. Note that the diameter scale is not continuous.

There were only three ambiguous responses in the single label condition and no such cases in the 'chomps or blibs' condition.[1] The results are summarized in Table 1. Overall, we see that participants were more likely to generalize to the intermediate items when the training items were labelled as 'chomps' or 'blibs' than when they were labelled as 'chomps or blibs', showing that information about category logical structure is taken into account in our participants' generalizations. A significant one-tailed Fisher's exact probability test $(p = .046)$ confirmed this intuition (in this and the other experiments we computed one-tailed p-values because the only meaningful differences in the patterns of generalization would be in a specific direction). Note that the significance appears to be driven by the 'chomps or blibs' condition. However, it is unclear how we could interpret the results of any one condition independently of the other.

Table 1. Number of participants generalizing to intermediate patterns, or not, as a function of category labeling, in Experiment 1

Instructions	Selected intermediate patterns	Not selected intermediate patterns
Training items: Chomps or blibs	6	19
Training items: Single label	11	10

EXPERIMENT 2

In Experiment 1, test items would presumably be related to training items in terms of physical similarity (and, as our results showed, information about the logical structure of the category as well). In Experiment 2 we wished to examine whether reversals in generalization analogous to those observed in Experiment 1 would take place when the test items could be related to the training items in terms of an abstract property (cf. Hahn & Chater, 1998; Tversky, 1977). Accordingly, in this experiment we used training stimuli that were nearly symmetrical checkerboard patterns. In this way, the training stimuli could be seen as members of either a category of symmetrical patterns or a category of random patterns. The test stimuli were either perfectly symmetrical (across a centred vertical axis) or random. The training items were described in one condition with a positive category label ('chomps') and in the other with a negative category label ('non-chomps'). If information about the logical structure of the training items' category is taken into account, we would expect participants to associate the positive category label with the more cohesive category of the symmetrical patterns and correspondingly the negative label with the random patterns. This is because, intuitively, symmetry is a well-defined, easy to perceive (Corbalis & Roldan, 1974; Pothos & Ward, 2000) property of items, so that it would seem reasonable to define a category in terms of such a property.

[1] *We also had to reject a participant in the single label condition as this partcipant had done a similar experiment before and one participant in the 'chomps or blibs' condition who failed to understand the instructions.*

Design and participants

The experiment had a between-participant design, with two conditions. Ten participants were tested in the chomps condition and 11 in the non-chomps condition (see below). Participants were Oxford University undergraduates and took part in the study for a small payment.

Materials

The stimuli used were arrangements of 'checkerboard' patterns in a 10×10 invisible grid, with roughly equal numbers of black and white squares. A line dividing the grid into two was an axis of symmetry for the symmetrical stimuli (Fig. 3). The items used in the training phase, referred to as near-symmetrical patterns, were all nearly symmetrical apart from one randomly positioned defect (that is, moving one black square in one of its nearest neighbor positions would make the pattern perfectly symmetric about its axis of symmetry). The training items were created so that it would be fairly obvious that they were almost, but not quite, symmetrical. In the test phase, there were three kinds of stimuli: symmetrical patterns that corresponded to the near-symmetrical patterns of the training phase (i.e. where the single violation of symmetry was removed), additional symmetrical patterns that were not related to the training patterns, and also asymmetrical (random) patterns, which were created by randomizing the distribution of black and white squares. The stimuli were printed individually on A4 paper in black ink. The size of the grid was about 6.4×6.4 cm and the size of the checkerboard squares 0.65×0.65 cm.

Procedure

In the training phase, participants were given eight near-symmetrical patterns, seven of which were different (one training pattern was accidentally repeated). In the test phase, they were given 24 patterns, 8 symmetrical patterns corresponding to the training patterns, another 8 symmetrical patterns and 8 asymmetrical patterns.

In one condition, participants were told that all training items were chomps, and, in the other condition participants were told that all training items were non-chomps. In both conditions, in the test phase, participants were asked to identify the items they thought were chomps and the ones they thought were non-chomps.

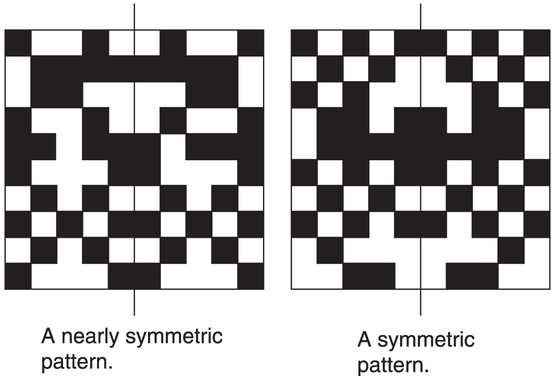

Figure 3. An example of the nearly symmetric training stimuli and the symmetric test ones.

Results

A majority rule analogous to that of Experiment 1 was used to classify participants as 'generalizing to symmetric patterns' if they classified more symmetrical patterns than (twice the) asymmetric patterns as compatible with the training ones or 'generalizing to asymmetric patterns' if they classified more asymmetric patterns as compatible with the training items than (half the) symmetrical patterns (recall that in test there would be twice more symmetrical patterns than asymmetric patterns). How participants generalized from the training items to the test items would be evident in terms of which test items they described with the same label as the training ones. Thus, when, for example, the training items were labelled as non-chomps, participants' generalization would correspond to which test items were classified as non-chomps (see Fig. 4).

The results of the experiment are shown in Table 2. In practice, in the chomps condition the eight participants who were classified as generalizing to the symmetrical patterns selected only symmetrical patterns as compatible with the training stimuli and the two participants who were classified as generalizing to the asymmetrical patterns selected only asymmetrical patterns as compatible with the training stimuli. Likewise, in the non-chomps condition, except for a single participant who selected both symmetrical and random patterns as compatible with the training items. Thus, overall participants consistently generalized from the training stimuli to new ones on the basis of symmetry.

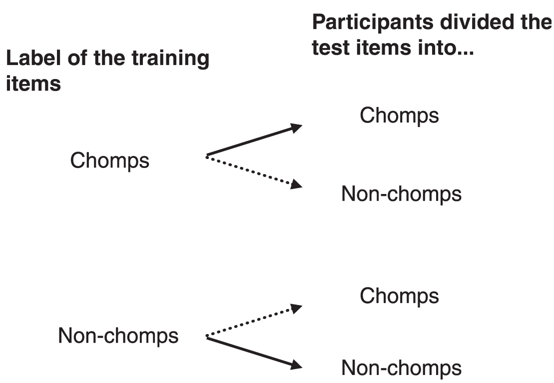

Figure 4. Characterizing performance in Experiment 2. Generalization (solid arrows) is indicated by which items in test are called in the same way as the training items.

Table 2. Number of participants generalizing to asymmetrical or symmetric patterns as a function of training instructions, in Experiment 2

Instructions	Generalization to asymmetrical patterns	Generalization to symmetrical patterns
Training items:		
Chomps	2	8
Training items:		
Non-chomps	10	1

As seen in Table 2, participants were more likely to generalize to symmetrical patterns when the training items were described with a positive label and more likely to generalize to the asymmetrical patterns when the training items were described with a negative label. A significant one-tailed Fisher's exact probability test confirmed this intuition ($p = .0017$). Thus, as in Experiment 1, information about the logical structure of a category appears to influence how participants generalize a set of objects to novel ones.

EXPERIMENT 3

In Experiments 1 and 2 we used well-controlled but highly schematic stimuli on which influences of general knowledge are likely to be small. In Experiment 3, we generalized this examination by using more realistic stimuli and a categorization situation. In Experiment 2, we used training stimuli that were nearly symmetrical patterns so that they could be equally perceived as members of the category of symmetrical patterns or the category of the asymmetrical patterns. In Experiment 3, the training stimuli were pictures of cars that were in between sports and non-sports cars, while some of the test stimuli were clearly sports cars and the rest clearly non-sports cars. The classification task in Experiment 3 was motivated in a pragmatic context (a car insurance operation). In one condition the training stimuli were described as members of a category and in the other as non-members. The property 'sports car' in Experiment 3 was exactly equivalent to the property 'symmetry' in Experiment 2. In this way, we were able to investigate whether information about the logical structure of a category will affect generalization, in a situation where presumably generalization will also be influenced by expectations and assumptions about the category, induced from general knowledge.

Design and participants

The experiment had a between-participants design, in which in one condition the training instances were labelled as members of a category (16 participants tested) and in the other they were labelled as non-members (15 participants tested). Participants were University of Wales, Bangor undergraduates and took part in the study for a small payment or course credit.

Materials

We used high-quality colour images of 24 cars, printed on A4 paper. These images were downloaded from the Internet and were of similar styles (see the Appendix for a full list of the images). Of these images, eight were clearly sports cars, eight were clearly non-sports cars, and eight cars could not be unambiguously classified as sports or non-sports cars. Appropriate images were selected on the basis of independent judgments from three observers familiar with different types of cars (EMP, AJS and one other[2]). Of course, there is little control on how people would represent realistic stimuli; thus, our participants may have encoded the cars in terms of value, size, luxury, colour, make etc. Nevertheless, consistency of responses would indicate whether participants represented the stimuli in an experiment-relevant way.

[2] *The cars included in the experiment were the ones independently characterized in the same way (as sports, non-sports, or intermediate) by the three observers.*

Although the sports cars appeared very similar to each other, this was not the case for the non-sports cars, which were comprised of luxury, family, economy, small-city and other types of cars. This situation is analogous to that in Experiment 2, in which the symmetric patterns formed a well-defined group, but the asymmetric patterns could not be characterized in any obvious way (other than being asymmetric).

Procedure

In the training phase, each participant was told that he/she would be employed to work in a foreign country as an assistant to a car insurance broker. Customers would come to him/her trying to find out how much insurance they would have to pay for their cars; the cost of insurance would depend on many factors, such as the driving history of the owner, where the owner lives, how powerful the car is, and how expensive it is. Participants were shown the eight images of cars that were neither clearly sports nor clearly non-sports cars. In one condition, participants were then told that all training items belonged to insurance group S, whereas in the other condition that none of the training items belonged to insurance group S.

In the test phase, participants were given the eight images of cars that were clearly sports cars and the eight images of cars that were clearly not sports cars. They were told that their job was to sort the car owners' application forms into two categories: those belonging to insurance group S and those not belonging to insurance group S. Overall, the Experiment 3 procedure was directly analogous to that in Experiment 2: participants were never explicitly told of the critical property distinguishing the test stimuli, but the stimuli were created in such a way that the property would be fairly obvious.

Results

As in Experiments 1 and 2, we classified participant responses in terms of whether they were predominantly sports or predominantly non-sports, on the basis of the majority rule: if a participant selected more sports cars than non-sports cars as belonging to insurance class S, then this participant would be included in the analyses as 'having selected sports cars as belonging to insurance class S', and vice versa. However, contrary to Experiment 2 in which the majority of participants generalized to either symmetric or asymmetric patterns (but not both), in this experiment most participants generalized to a mixture of sports and non-sports cars in test.

Table 3 summarizes the results of Experiment 3. When the training instances were described as members of a category (as all belonging to insurance group S), then

Table 3. Number of participants generalizing to non-sports cars or sports cars as a function of training instructions, in Experiment 3

Instructions	Generalization to sports cars	Generalization to non-sports cars
Training items: All belong to S	14	2
Training items: None belong to S	8	7

participants were more likely to generalize to the well-defined category of sports cars, than when the training instances were described as non-members of a category (not belonging to insurance group S). Although the difference in generalization performance with the two linguistic labels was not as obvious as in Experiment 2, there was still an interaction (Fisher's exact probability test one-tailed $p = .044$). Thus, information about the logical structure of a category would affect generalization even in a situation where expectations and assumptions about the category might be induced by general knowledge.

GENERAL DISCUSSION

We have identified situations in which participants would generalize differently a set of objects depending on whether they were told the objects belonged to a single category, two categories, or no particular category. The rejected null hypothesis is that how participants generalize the training objects to novel ones depends only on the abstract or physical properties on the basis of which the training objects were initially represented.

One ultimate objective of such research is to provide an explanatory framework for how it is that categorization alters object perception and generalization. When it comes to effects associated with the presence of a linguistic label, a possibility would be to consider a linguistic label as an additional feature or dimension of object variation that influences similarity computations between objects in a way analogous to physical features or dimensions. However, our research shows that it is not just the presence of a linguistic label that matters, but also its meaning, the particular information conveyed about the logical structure of the training objects' category. Does information about the logical structure of a category affect generalization in a way analogous to how general knowledge affects categorization (e.g. Wattenmaker, Dewey, Murphy, & Medin, 1986)? For example, Wisniewski and Medin (1994) asked their participants to categorize children's drawings. These investigators designed their stimuli so that there would be several possible rules for categorizing them. They found that participants' performance varied greatly depending on whether the stimuli were labelled in a meaningful way (e.g. drawn by creative children) or not. The problem is that effects of general knowledge hinge so tightly on people's specific knowledge background, cultural biases, experiences and so forth that they have been notoriously difficult to model (Heit, 1997; Pickering & Chater, 1995). Right now, it does not appear that a principled account of general knowledge effects can be forthcoming (Fodor, 1983).

An alternative explanation for how information about logical structure affects generalization might be motivated from the observation that in all our experiments generic, positive category labels appeared to be associated with groups of objects that had high within-group similarity, either in terms of physical similarity (Experiment 1) or in terms of an abstract property (in Experiment 2, symmetry; in Experiment 3, whether a car was a sports car or not). It is possible, therefore, that people generally associate generic, positive labels with groups of highly similar objects (so that logical functions of generic labels, such as negation or conjunction, would lead to corresponding expectations about the structure of a category). This is the basis for Rosch and Mervis's (1975) explanation for why basic level category terms are preferentially used in object recognition as well as Pothos and Chater's (2002) simplicity model for how people spontaneously divide a set of objects into groups (see also Chater, 1996, 1999).

With future work we hope to clarify whether people's default assumptions about category structure are consistent with the above possibility.

In conclusion, our results show that information about the logical structure of a category affects how the category members will be generalized. The nature of the interaction between logical structure information and generalization appears to be reduced to a possibility that people associate generic category labels with an expectation that category members would be highly similar to each other.

Acknowledgements

This research was supported by ERSC grant ref. R000222655 to Emmanuel Pothos and Nick Chater. Nick Chater was also partially supported by European Commission grant RTN-HPRN-CT-1999-00065. We would like to thank Ulrike Hahn and Koen Lamberts for helpful discussions.

References

Ashby, G. F., & Alfonso-Reese, A. L. (1995). Categorization as probability density estimation. *Journal of Mathematical Psychology*, *39*, 216–233.

Ashby, G. F., & Perrin, N. A. (1988). Towards a unified theory of similarity and recognition. *Psychological Review*, *95*, 124–150.

Chater, N. (1996). Reconciling simplicity and likelihood principles in perceptual organization. *Psychological Review*, *103*, 566–581.

Chater, N. (1999). The search for simplicity: A fundamental cognitive principle? *Quarterly Journal of Experimental Psychology*, *52A*, 273–302.

Corballis, M. C., & Roldan, C. E. (1974). On the perception of symmetrical and repeated patterns. *Perception and Psychophysics*, *16*, 136–142.

Dennis, I., Hampton, J., & Lea, S. (1973). New problem in concept formation. *Nature*, *243*, 101–102.

Flannagan, M. J., Fried, L. S., & Holyoak, K. J. (1986). Distributional expectations and the induction of category structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 241–256.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understandings of the nonobvious. *Cognition*, *38*, 213–244.

Goldstone, R. L. (1994a). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*, 178–200.

Goldstone, R. L. (1994b). The role of similarity in categorization: Providing a groundwork. *Cognition*, *52*, 125–157.

Goldstone, R. L. (1995). Effects of categorization on color perception. *Psychological Science*, *6*, 298–304.

Goldstone, R. L. (1996). Isolated and interrelated concepts. *Memory and Cognition*, *24*, 608–628.

Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, *49*, 585–612.

Goldstone, R. L. (2000). Unitization during category learning. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 86–112.

Goldstone, R. L., & Barsalou, L. W. (1998). Reuniting perception and conception. *Cognition*, *65*, 231–262.

Goodman, N. (1954). *Fact, fiction and forecast*. London: Athlone Press.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman, *Problems and projects* (pp. 437–447). Indianapolis: Bobbs-Merrill.

Hahn, U., & Chater, N. (1998). Similarity and rules: Distinct? Exhaustive? Empirically distinguishable? *Cognition*, *65*, 197–230.

Hampton, J. A. (1997). Conceptual combination: Conjunction and negation of natural concepts. *Memory & Cognition*, *25*, 888–909.

Heit, E. (1997). Knowledge and concept learning. In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts, and categories* (pp. 7–41). London: Psychology Press.

Hintzman, D. L. (1986). 'Schema-abstraction' in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.

Homa, D., & Chambliss, D. (1975). The relative contribution of common and distinctive information on the abstraction from ill-defined categories. *Journal of Experimental Psychology: Human Learning and Memory*, *1*, 351–359.

Komatsu, L. K. (1992). Recent views of conceptual structure. *Psychological Bulletin*, *112*, 500–526.

Livingston, K. R., Andrews, J. K., & Harnad, S. (1998). Categorical perception effects induced by category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 732–753.

López, A., Atran, S., Coley, J. D., Medin, D. L., & Smith, E. E. (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive Psychology*, *32*, 251–295.

Markman, E. M. (1987). How children constrain the possible meanings of words. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 255–287). Cambridge: Cambridge University Press.

Markman, E. N. (1991). The whole object, taxonomic, and mutual exclusivity assumptions as initial constraints on word meanings. In J. P. Byrnes & S. A. Gelman (Eds.), *Perspectives on language and cognition: Interrelations in development* (pp. xx–xx). Cambridge: Cambridge University Press.

Markman, E. M., & Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology*, *16*, 1–27.

McKinley, S. C., & Nosofsky, R. M. (1995). Investigations of exemplar and decision bound models in large, ill-defined category structures. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 128–148.

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

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.

Nosofsky, R. M. (1989). Further tests of an exemplar-similarity approach to relating identification and categorization. *Journal of Experimental Psychology: Perception and Psychophysics*, *45*, 279–290.

Nosofsky, R. M. (1990). Relations between exemplar-similarity and likelihood models of classification. *Journal of Mathematical Psychology*, *34*, 393–418.

Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 3–27.

Özgen, E., & Davies, I. R. L. (2002). Acquisition of categorical color perception: A perceptual learning approach to the linguistic relativity hypothesis. *Journal of Experimental Psychology: General*, *131*, 477–493.

Pickering, M., & Chater, N. (1995). Why cognitive science is not formalized folk Psychology. *Minds and Machines*, *5*, 309–337.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.

Pothos, E. M. (in press). The rules vs. similarity distinction. *Behavioral and Brain Sciences*.

Pothos, E. M., & Bailey, T. M. (2000). The importance of similarity in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 847–862.

Pothos, E. M., & Chater, N. (1997). Classification and prior assumptions about category 'shape': New evidence concerning prototype and exemplar theories of categorization. *Proceedings of the nineteenth annual conference of the Cognitive Science Society* (pp. 620–625). Mahwah, NJ: Erlbaum.

Pothos, E. M., & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science*, *26*, 303–343.

Pothos, E. M., & Ward, R. (2000). Symmetry, repetition, and figural goodness: An investigation of the weight of evidence theory. *Cognition*, *75*, B65–B78.

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, *3*, 382–407.

Roberson, D., Davies, I., & Davidoff, J. (2000). Color categories are not universal: Replications and new evidence from a stone-age culture. *Journal of Experimental Psychology: General*, *3*, 369–398.

Rosch, E., and Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573–605.

Schyns, P. G., Goldstone, R. L., & Thibaut, J. (1997). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1–54.

Sera, M. D., Gathje, J., & Pintado, J. (1999). Language and ontological knowledge: The contrast between objects and events made by Spanish and English speakers. *Journal of Memory and Language*, *41*, 303–326.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.

Shepard, R. N., Hovland, C., & Jenkins, H. (1961). Learning and memorization of classifications. *Psychological Monographs*, *75*, 1–42.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352.

Wattenmaker, W. D., Dewey, G. I., Murphy, T. D., & Medin, D. L. (1986). Linear separability and concept learning: Context, relational properties, and concept naturalness. *Cognitive Psychology*, *18*, 158–194.

Wisniewski, E., & Medin, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, *18*, 221–281.

Received 12 December 2002; revised version received 13 January 2004

Appendix

The types of cars used in Experiment 3

Training items
BMW 3.16 convertible
BMW
Alfa Romeo GTV
Audi TT
Honda CRX
Hyundai S-Coupe
Mazda MX5
MG Convertible

Test items

Type of Car	*Classification*
Austin	Non-sports
Rolls	Non-sports
Rolls	Non-sports
Fiat	Non-sports
Ford	Non-sports
VW	Non-sports
Volvo	Non-sports
Fiat	Non-sports
Ferrari	Sports
Porsche	Sports
Lotus	Sports
Lamborghini	Sports
Lamborghini	Sports
Lotus	Sports
McLaren	Sports
McLaren	Sports