European Journal of Cognitive Psychology

Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/pecp20

Interaction of knowledge-driven and data-driven processing in category learning

Andre Vandierendonck & Yves Rosseel

Published online: 21 Sep 2010.

To cite this article: Andre Vandierendonck & Yves Rosseel (2000) Interaction of knowledge-driven and data-driven processing in category learning, European Journal of Cognitive Psychology, 12:1, 37-63, DOI: 10.1080/095414400382190

To link to this article: http://dx.doi.org/10.1080/095414400382190

EUROPEAN JOURNAL OF COGNITIVE PSYCHOLOGY, 2000, *12* (1), 37–63

Interaction of knowledge-driven and data-driven processing in category learning

André Vandierendonck and Yves Rosseel
University of Ghent, Belgium

The present paper argues that category learning is both a data-driven and a knowledge-driven process. This is described in a generic model that distinguishes between categorical knowledge, conceptual knowledge, and implicit cognitive theories. The model assumes that each of these knowledge aspects may affect the process of category learning by affecting the way similarities between objects are perceived. This central assumption of the model is tested in two experiments. The first experiment shows that the presence or absence of prior categorical and conceptual knowledge affects the psychological stimulus space by changing the saliency of the stimulus dimensions. The second experiment uses these weights to predict the distribution of errors over the stimuli and the number of trials to criterion in category learning by other participants under the same knowledge conditions. We conclude that prior categorical and conceptual knowledge affect category learning by mediation of similarity perception, and discuss the implications of these results.

INTRODUCTION

In the last decade several authors have stressed the importance of prior knowledge and general background knowledge in the formation of efficient and coherent categories (see e.g., Lakoff, 1987; Medin & Wattenmaker, 1987; Murphy & Medin, 1985). The hypothesis that prior knowl-

Requests for reprints should be addressed to A. Vandierendonck, Department of Experimental Psychology, University of Ghent, Henri Dunantlaan 2, Belgium.
Email:andre.vandierendonck@rug.ac.be

Yves Rosseel is Research Assistant of the Fund for Scientific Research—Flanders (Belgium). Presently, Yves Rosseel is affiliated to the University of Leuven.

The research reported in this article was supported by grant 32000994 from the Fund for Scientific Research (Belgium) and grant 01102594 from the Research Council of the University of Ghent (Belgium) to the first author.

http://www.tandf.co.uk/journals/pp/09541446.html

edge affects categorisation has found support in several empirical studies that tested for the effects of knowledge.

Basically, three different methodologies have been used to study knowledge effects. Some investigators assume that the *evocation* or *activation* of knowledge in semantic memory will affect processes evoked during category learning or category usage. By selecting specific labels or examples, knowledge is activated that affects categorisation performance (see e.g., Barsalou, 1985; Hayes & Taplin, 1995; Lamberts, 1994; Pazzani, 1991; Wisniewski, 1995; Wisniewski & Medin, 1994). Other investigators have approached the study of knowledge effects by varying *the meaningfulness* and the *familiarity* of the stimulus materials (see e.g., Murphy & Allopenna, 1994; Murphy & Spalding, 1995; Spalding & Murphy, 1996; Vandierendonck, 1978). A few researchers directly *manipulated prior knowledge* (see e.g., Heit, 1994, 1998; Nakamura, 1985).

The present paper specifically addresses the problem of interaction between prior knowledge, similarity, and category learning. First, we will argue that knowledge is organised in structures operating at different levels. Next, a generic model is presented that not only embodies these assumptions, but also implies that effects of prior knowledge on category learning may be mediated by processes of similarity judgement. In order to test this implication, two experiments are reported. The first one shows that similarity judgements are influenced by manipulations of prior knowledge. The second one is intended to demonstrate that effects of knowledge manipulation on category learning can be predicted from changes in the similarity judgements obtained in an independent sample of subjects.

KNOWLEDGE, SIMILARITY, AND CATEGORIZATION

Knowledge is assumed to be the complete body of accessible information stored in the memory system.[1] This information includes concepts, schemata, frames, models, relationships, episodes, procedures, etc. In addition, we assume that knowledge is *structured*. This means that knowledge consists of interrelated entities that together form a structure; concepts are basic components in these structures.

In their extensional meaning, concepts are related to a system of categories, a categorisation. A central idea in many theories of categorisa-

[1]This is a psychological definition of knowledge, and it excludes information available in books, on audiotape, videotape, film, etc. It also excludes information that is accumulated in the memory system but which is temporarily or definitively not accessible.

tion is the notion of similarity, and the acquisition of a categorisation is generally thought to depend on similarity. In fact, it has been amply shown that *similarity* is one of the key factors in categorisation (for a collection of papers addressing this issue, see e.g., Thibaut & Vandierendonck, 1995). According to *exemplar models* (e.g., Medin & Schaffer, 1978; Nosofsky, 1984; Nosofsky & Kruschke, 1992), categorisation is based on the perceived similarity between a to-be-categorised instance and one or more exemplars stored in memory, such that the new instance will be placed in the category that contains the most similar exemplars. In *prototype models* (e.g., Homa, 1984; Posner & Keele, 1968; Reed, 1972; Rosch, 1975) similarity between the instance and one or more prototypes is the basis of the categorisation. These prototypes are inferred from previous experiences and they represent the central tendency and the variability of the categories. It can even be said that in some *rule models* (e.g., Anderson, Kline, & Beasley, 1979; Vandierendonck, 1995), where the categories are represented by means of production rules, the similarity or the correspondence between a to-be-categorised instance and the rule is at the basis of the categorisation.

However, Medin (Medin, 1983; Medin, Goldstone, & Gentner, 1993; Medin & Wattenmaker, 1987; Murphy & Medin, 1985) has quite convincingly argued that similarity does not suffice to explain conceptual coherence. The category of "protected buildings" gets its coherence from the fact that all buildings in the category have been found worthy of conservation. Even though the buildings may be as diverse as windmills, churches, former city halls, farms, etc., while very similar buildings, such as other windmills, farms, and so on do not belong to the category. Clearly, perceptual similarity is sometimes subordinated to categorical knowledge.

LEVELS OF KNOWLEDGE

It is clear from this exposition that categorisation and knowledge can be described at a number of levels. Figure 1 presents a schematic overview of these levels and the way they are supposed to interact with data-driven processes in category learning. This viewpoint proposes that category learning consists of the perception of similarities between objects and the grouping of objects implied by these similarities. On this basis, a category representation is constructed that contains the information necessary to assign an instance to its category. This representation corresponds to the extensional meaning of the concept (see e.g., Johnson-Laird, Herrmann, & Chaffin, 1984; Vandierendonck, 1991), and in the overview this is located at the level of *categorical knowledge*.

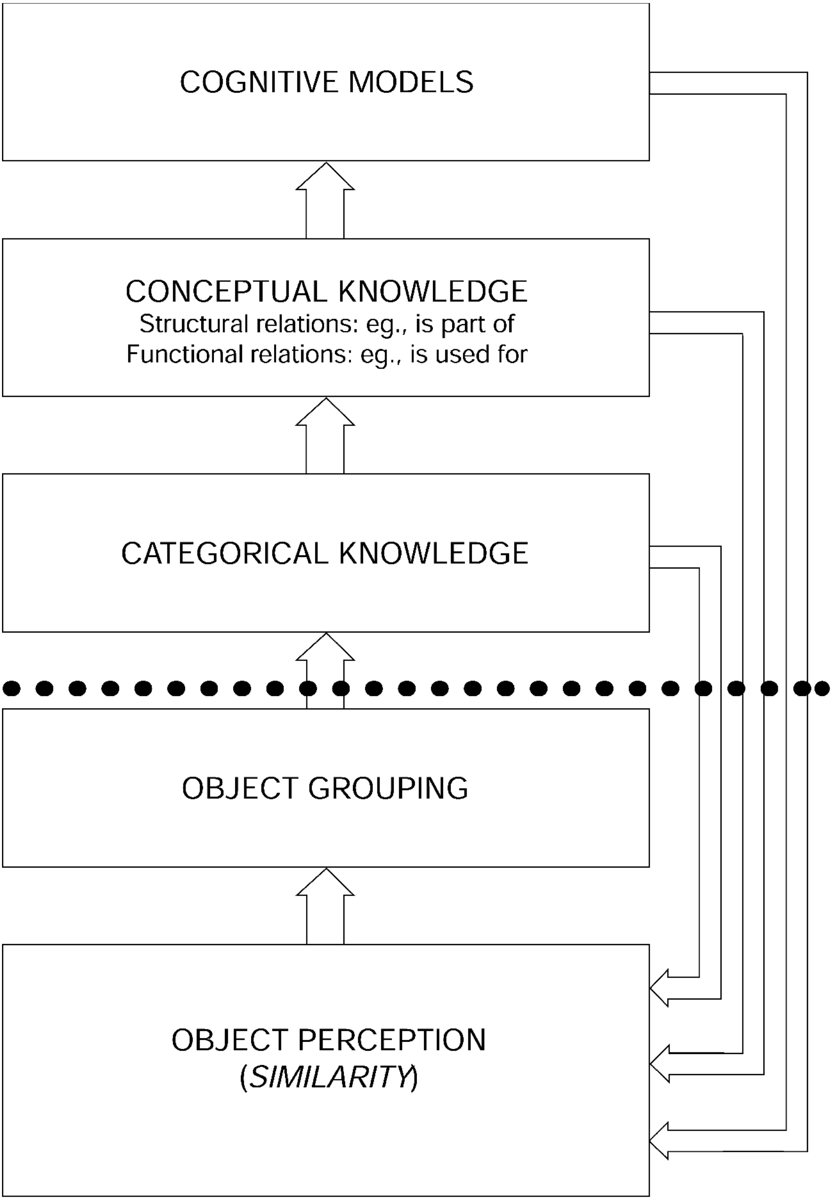

Figure 1. Schematic representation of different levels of similarity and knowledge involved in category learning and in the study of the effects of knowledge on similarity and category learning. The levels above the dotted line represent the different levels of knowledge. The top-down arrows represent the effects of knowledge on similarity perception. The bottom-up arrows show the steps involved in inductive or data-driven learning.

Categories and their representations are usually coupled to concepts, which are considered to be cognitive entities that represent the relations between different categories in terms of class inclusion relations, part-of relations, schemata, mental models, etc. This is the *conceptual level* and it refers to structural as well as functional aspects of the concept. This level of representation is assumed to be the one necessary for using categorical distinctions in reasoning, in discourse, etc. It corresponds to the intensional meaning of the concept (again see Johnson-Laird et al., 1984; Vandierendonck, 1991).

Clusters of categories that are linked by all sorts of interrelations may be part of an implicit theory about the world or an implicit cognitive model (Lakoff, 1987). Such a model forces particular interrelations between categories to become salient and particular categorisations to be more salient than other ones. When such an implicit model is activated by presenting a prime for instance, this activation is hypothesised to spread to relevant conceptual representations, and so further on to particular relations and to particular categorisations.

The generic model in Figure 1 shows that the process of category learning is *both* a *data-driven* and a *knowledge-driven process*. Considering the data-driven aspects, it is a process in which the learner must acquire knowledge allowing perceived objects to be assigned to a predefined category. Whether the categorisation is eventually performed on the basis of exemplar similarity, similarity to a prototype, or on the basis of rules does not matter at this point. In any case, the criterion used to relate the objects and the categories is the categorical knowledge obtained. However, in this process it may be necessary to overcome the perceived similarities and the tendency to group similar things together. Homa, Rhoads, and Chambliss (1979) and Livingston, Andrews, and Harnad (1998) report evidence supporting such a process; they have shown that category learning changes the multidimensional similarity space. The latter authors have also shown that within-category compression may occur under some learning conditions. These findings testify to the importance of the similarity processes in categorisation.

Category learning is also a knowledge-driven process. Knowledge originating either at the categorical level, the conceptual level, or at the level of implicit cognitive models may affect the process of similarity perception. The changed estimation of similarities then affects the data-driven construction of a categorical representation. Hence, the perceived similarity of objects is the joint result of perceiving the objects and the utilisation of relevant knowledge to select aspects or attributes to judge the likeness. Consequently, if at any level, knowledge is activated that is relevant for judging the similarity of the present object to objects or other information in memory, then the similarity judgement will be affected by that knowledge, and in its turn the altered similarity judgements will affect the process of category learning.

The model proposed here is one in which knowledge is in an interactive causal relationship with the process of object perception: Experience with new objects may result in a change of the representations at the different levels of knowledge, and any level of knowledge may affect the perception of similarities among objects. To be sure, this framework yields only a first rough and rather vague approximation to the problem of the interaction between data-driven and knowledge-driven processes in category learning.

It is the purpose of the present paper to present a test of the central thesis of the model, namely that prior knowledge operating at different representational levels affects similarity judgements and that the perceived similarities, in their turn, affect the process of category learning. A first experiment tests whether the availability of prior knowledge affects similarity judgements. In the second experiment, the effects shown in the first experiment will be used to predict the speed of learning and the errors committed during category learning.

EXPERIMENT 1

In the first experiment the participants rated the similarity of all possible pairs of a set of stimuli. In the prior knowledge conditions, they did this after going through a knowledge acquisition phase, in which they memorised a series of facts. It was expected that the perceived similarities would be systematically different as a function of the kind and of the amount of knowledge acquired in the knowledge acquisition phase. More specifically, the relative importance of the stimulus dimensions was expected to change due to the availability of prior knowledge.

This method constitutes a natural extension of the usage of multidimensional scaling in studies of concept learning. Homa et al. (1979) for example, compared the multidimensional psychological space before and after an extensive categorisation training, with the aim of clarifying the conceptual changes induced by the learning process. In a similar way, tests of the Generalised Context Model (GCM; e.g., Nosofsky, 1992a) are based on predictions of changes in the multidimensional psychological space due to attentional shifts evoked by the learning process. In all these studies, it is assumed that the similarity scaling yields data that are independent of the categorisation and typicality data used to index the learning process. The present study follows the same methodology and extends it by collecting similarity data to study the influences of (prior) knowledge on the psychological object space.

Method

Participants. Ninety first-year students in psychology and educational sciences at the University of Ghent (Belgium) participated for course requirements and credit. They volunteered to participate in this particular experiment. They were randomly and equally assigned to the 10 cells of the design (see Materials and Design section).

Materials and design. As the purpose was to test the effects of the availability of particular sources of knowledge, it was necessary to develop an artificial mini-world for the test. A stimulus set, a categorisation system, and two different cover stories were developed for this purpose.

The *stimuli* were so-called robots that varied in three five-valued dimensions: the number of pairs of arms, the size of the hat, and the orientation of the antennae. Figure 2 displays two examples of these stimuli.

Only 15 of the 125 possible combinations of the three dimensions were implemented as stimuli. Two "categories" were developed around two prototypes. With the values along the dimensions scaled from 0 to 4, one prototype had the values 1,3,2 (one pair of arms, hat size 3, and horizontal antennae), whereas the other had the values 3,1,2 (three pairs of arms, hat size 1, and horizontal antennae). Starting from these prototypes, transformations were generated by increasing or decreasing the number of arms and the hat size. Increments in the hat size and decrements in the number of arms were correlated with a more upward

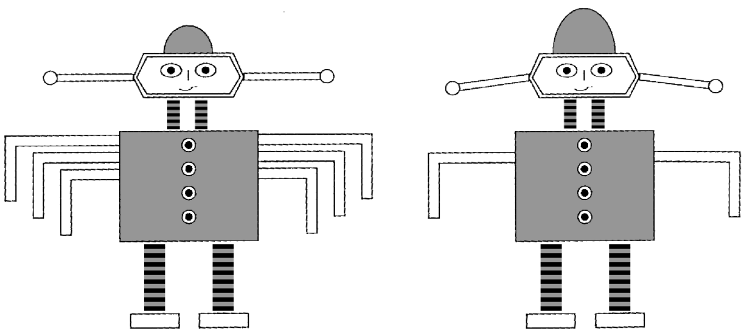

Figure 2. Two examples of stimuli used in the present study. The stimuli vary in three attributes: arms, hat, and antennae. The remainder of the stimulus layout was held constant. As used in the experiments, the stimuli were in colour (red, pink, yellow, white, and black).

position of the antennae. Conversely, decrements in hat size and increments in the number of arms resulted in a more downward position of the antennae. All stimuli used were obtained by the application of one or two such transformations.

Table 1 displays the 15 stimuli and the two categories developed around the two prototypes. As can be seen in the table, in addition to the prototype, each category contained six instances, and one category (O) also contained an exception. This exception did not follow the rules of generation. The rationale for the inclusion of the exception was to enable an analysis of category learning in ill-defined categories. Except for the mere fact that an exception is present, this aspect is not an issue of the present study.

From the two *cover stories* a number of statements were generated at the categorical and the conceptual levels of knowledge as defined in the introduction. Because the same stimuli were categorised in the same way under both cover stories, the categorical knowledge statements were the same for both stories, and only the conceptual statements differed. The Appendix displays all the statements generated from the two cover stories (translated from Dutch).

As categorical knowledge concerns features that are relevant to the categorisation, statements that specify the relevant or the typical features in a category provide categorical knowledge. This includes statements such as "Kwarks mostly have many arms", "Orkels mostly have few arms", etc. "Very often Kwarks have a small hat", "The exact number of arms varies both in Kwarks and in Orkels". Also statements about the correlation between characteristics are situated at this level, such as "The

TABLE 1
Description of the Stimulus Patterns and the Categorisation System Used in the Present Study

	Number of Pairs of Arms				
Hat Size	*0*	*1*	*2*	*3*	*4*
0				K2–	K3– –
1		K6+ +	K5+	K1	O15+ +
2			K7+ +	K4+	
		O11–	O14– –		
3		O8	O12–	O13– –	
4	O10+ +	O9+			

Stimulus patterns are indicated in the appropriate cells. The category is indicated by the letters K and O. The numbers refer to a stimulus identification, and the + and – signs refer to the position of the antennae.

fewer the number of arms, the larger the size of the hat", and "The more arms, the more downward the orientation of the antennae tends to be". The complete list of categorical statements used is presented in the Appendix.

Knowledge at the conceptual level, relates the category to other known categories. Hence, statements such "All Kwarks are workers" (cover 1) and "Orkels bore the mine galleries" (cover 2) are examples of conceptual knowledge. The same is true for statements specifying the function of the category or of certain features: "The bigger the hat, the bigger the brain for thinking" (cover 1) and "The arms are used to transport the ore" (cover 2). Explanation of why there is a correlation between certain features also exemplifies conceptual knowledge: "The happier they are, the more the antennae are oriented upwards" (cover 1) and "The deeper the robot can go into the mine, the more upward the antenna orientation" (cover 2).

The materials were varied in four respects: Participants can be presented categorical information or not; they can be given conceptual information or not; conceptual information can be based on two different cover stories; and the conceptual information does or does not include information about the exception. A complete factorial design based on these four variables is not possible, because the latter two variations are conditional on the presence of conceptual information. Table 2 shows the design used in the present experiment.

Procedure. Participants were tested individually at an IBM-compatible PC with a 14-inch colour monitor. First the experiment was explained and a few examples of the stimulus materials were shown. All participants

TABLE 2
Incomplete Factorial Design Used in Experiments 1 and 2

	$\overline{CO}$	*CO*			
		Cover 1		*Cover 2*	
		$\overline{E}$	*E*	$\overline{E}$	*E*
$\overline{CA}$	1	3	5	7	9
CA	2	4	6	8	10

The values in the cells of the table refer to a number scheme of the cells used in the data analysis.

The abbreviations *CA*, *CO*, and *E* refer to the presence, respectively of categorical information, conceptual information, and information about the exception. The same symbols overlined indicate the absence of such information.

were told that the stimuli they would see belonged in a natural way to two categories. Next, if applicable, the participants went through the knowledge acquisition phase, in which the texts specified in the Appendix were presented. The participants were instructed to memorise the information. After a study period, they were presented a quiz containing 8–20 questions, depending on the number of statements presented in the condition. If more than 15% of the answers were incorrect, the text was presented again, followed by a similar quiz. If still too many errors were committed, the text was presented for a third and last time.

After this knowledge acquisition phase, the similarity rating part of the experiment was introduced. Participants in all the conditions of the experiment were explained the procedure. It was stressed that they were to use all the values of the scale. On each trial two stimuli were displayed, one in the left and the other in the right half of the screen. Below the two stimuli a row of 2×2 cm response keys was displayed numbered 1 through 9. One of these keys could be selected by positioning the cursor over the key and then clicking the left mouse button. In this phase, all possible pairs of stimuli were presented once. Not only the sequence of the pairs but also the intra-pair order was randomised. First there was a short practice session (20 pairs) based on a random selection of pairs to give the participant the opportunity to adapt to the situation and to the response scale. The entire session, knowledge training included, lasted for about 45 minutes, and was completely presented on the computer.

Results

Knowledge training. Of the 80 participants who were in the knowledge conditions, 60 passed the test after one single presentation; 2 required three presentations of the information. The number of participants that succeeded the test was spread about equally over all conditions (eight or nine participants per condition succeeded). It seems safe to assume, therefore, that knowledge acquisition was similarly successful in all the conditions.

Similarity ratings. Because the similarity ratings collected in the 10 conditions are expected to differ due to the instructions, it does not make sense to pool these data for an overall multidimensional scaling. To show that the stimulus set was by and large perceived as intended, the data of the control condition were subjected to a multidimensional scaling. To that end, trimmed means (3 standard deviations) of the similarities of each pair of stimuli were calculated and formed the input of a multidimensional scaling. As there were three physical stimulus dimensions, a solution in three dimensions was obtained in a city-block metric by means of the KYST program (Kruskal, Young, & Seery, 1977). The solution had a stress of 0.059. The three dimensions in this solution corresponded quite well with the physical dimensions in order of importance, arms ($r = -.97$), hat ($r = -.18$), and antennae ($r = -.65$).

The experiment addressed the question whether the availability of prior knowledge about the stimuli and their categorisation would have an effect on the perceived similarity of the stimulus space. One possible way for such an effect to materialise is that the importance or the weight of particular stimulus dimensions used in the similarity judgement is changed. In order to test whether this was the case, the similarity ratings for each participant were subjected to a multiple regression analysis with the physical intra-pair distance on each dimension as the predictors.

Averaged over the 90 regression analyses, the regression coefficients for the arms, hat, and antennae dimensions amounted to 12.38, 5.10, and 2.63 respectively, and the standard deviations on these coefficients were 4.31, 2.40, and 2.24 respectively. It appeared that in 89 of the 90 analyses the regression coefficient for the arms dimension was significant at the .05 level. For the hat dimension the coefficient was significant in the data of 78 participants, whereas the antennae dimension attained significance in 53 data-sets. So it seems that the order of importance of the three dimensions was similar for all participants.

To facilitate comparisons of the importance of each of the three stimulus dimensions across participants and conditions, an estimate of the relative weight of each dimension is needed for each participant. Such relative weights were obtained for each participant by taking the ratio of the regression coefficient over the sum of all three regression coefficients in the analysis.[2] The rationale of this method is that each regression coefficient represents a unique part of the covariance between the predictor and the criterion. The three regression coefficients taken together explain all the variance due to the stimulus dimensions, so that the ratio expresses the relative importance of the stimulus dimension (predictor) in the similarity rating (criterion).

On the basis of the analysis it is predicted that the importance of the stimulus dimensions may differ as a function of the knowledge conditions. In order to test this prediction, the relative weights were subjected to a multivariate analysis with the 10 knowledge conditions as dummy coded independent variables, and the 3 dimensional weights as the dependent variables. Hypotheses were tested by means of contrasts in the independent and the dependent variables. This procedure follows a suggestion of McCall and Appelbaum (1973) for a correct analysis of repeated measures designs. The significance level was set at .05.

[2]Note that this is not the same as the calculation of normalised coefficients.

As could be expected on the basis of the regression analyses, the three weights were reliably different [$F(2, 79) = 151.28$, $p < .001$]. The average weights were respectively 0.60, 0.26, and 0.14. The average on the arms dimension was larger than the sum of the other two [$F(1, 80) = 256.91$, $p < .001$], and the average on the hat dimension was still reliably larger than the average on the antennae dimension [$F(1, 80) = 49.48$, $p < .001$].

Table 3 presents the average weight on the arms and the hat dimensions in the 10 conditions of the design. An analysis of the effects of the knowledge conditions on the weight of the arms dimension revealed an effect of the presentation of categorical information [$F(1, 80) = 4.78$, $p < .05$], such that the mean weight was smaller when no categorical knowledge was available ($M = 0.60$), than when categorical information had been presented ($M = 0.67$). The presence of conceptual information also affected the importance of the arms dimension ($M_{\overline{CO}} = 0.68$ and $M_{CO} = 0.59$) [$F(1, 80) = 5.66$, $p < .05$], but the interaction of categorical and conceptual information presence was not significant ($F < 1$). Similar analyses on the hat dimension and on the antennae dimension did not reveal any reliable effects. This is probably because the complementary effects were spread over the two dimensions: It is clear in Table 3 that increases in the arms-weights correspond to decreases in the hat-weights, but the differences in the hat-weights were smaller.

It could be the case that the effects shown are an artifact of the averaging of the regression coefficients. Therefore, the same analyses were performed on the regression coefficients as they were obtained in the

TABLE 3
Average Relative Weight of the Arms and the Hat Dimensions in Experiment 1 as a Function of the Information Presentation Conditions

$\overline{CO}$		CO			
		Cover 1		*Cover 2*	
		$\overline{E}$	E	$\overline{E}$	E
$\overline{CA}$					
Arms	0.65	0.57	0.56	0.58	0.47
Hat	0.23	0.32	0.23	0.27	0.32
CA					
Arms	0.71	0.70	0.55	0.62	0.54
Hat	0.21	0.23	0.22	0.24	0.29

The abbreviations *CA*, *CO*, and *E* refer to the presence, respectively, of categorical information, conceptual information, and information about the exception. The same symbols overlined indicate the absence of such information.

individual analyses. In general, the effects tended to be larger, but the effects found significant were the same. In addition, the effect of the presence of exception information on the regression coefficient of the antenna dimension also attained significance ($M_{\overline{E}} = 2.23$ and $M_E = 3.44$) [$F(1, 80) = 5.56$, $p < .05$].

Discussion

The results show that the availability of prior knowledge affects the importance assigned to the stimulus dimensions in judging the similarity of pairs of stimuli. Although the arms dimension, which is a salient attribute of the stimulus set, is the most important one in all conditions, it is clear that when categorical knowledge is available its importance is even larger. However, when conceptual knowledge was available the importance of this dimension was smaller than when no such knowledge was available. Both these effects appeared to be additive.

These findings are consistent with the model presented in the introduction. The model assumes that the knowledge acquired in the first phase of the experiment is stored in memory and is activated during the stimulus comparison task of the second phase. As categorical knowledge specifies the stimulus features that are characteristic of the category, it is expected that the perceptual stimulus features will result in activation of that knowledge. Salient stimulus features are more likely to attract attention and are therefore thought to be more efficient in evoking the related categorical knowledge. In the present context this means that activation of categorical knowledge about the arms dimension is more likely and this is expected to result in a larger impact of that dimension in the similarity judgements.

In contrast, conceptual knowledge represents the relations between the categories and their function. Because this knowledge is not directly tied to perceptual features, it is not dependent on saliency for its activation. Once activated on the basis of relevant cues, this knowledge provides information about several features and their relation to the categories so that the perceptual importance of the features is overruled to some extent by the retrieved conceptual information. More specifically, in the present context, it is expected that the evocation of conceptual knowledge will result in a decreased weighting of the arms dimension.

The model does not specify an interaction between the categorical and the conceptual knowledge levels. This means that the increased weighting of the arms dimension due to categorical knowledge and the decreased weighting of that dimension due to conceptual knowledge are obtained independently from each other. Consequently, no statistical interaction of the two effects is expected.

All the findings seem to be covered by the relations specified in the model, but it may be objected that the effects were significant only with respect to the arms dimension and not with respect to the other dimensions. Inspection of Table 3 shows that the differences for the hat dimension are in the predicted direction. On average, the weight for the hat dimension tended to be lower with than without categorical knowledge (*M*s respectively, 0.24 and 0.28) and it tended to be higher with than without conceptual knowledge (*M*s respectively, 0.27 and 0.22). These differences were too small, however, to attain significance.

An alternative explanation would be that the participants assumed that the knowledge presented in the first phase of the experiment was relevant to the task in the second phase. This would have led them to use this information to guide their stimulus comparisons. However, as only the categorical information is directly related to the stimulus features, it may be expected that only this information is used and there is no reason to expect that conceptual information is used and if it is used, unless the model specified in the introduction is correct, there is no reason to expect that the effects of conceptual information would be opposite to those of the categorical information.

Taken all together, Experiment 1 provides evidence in favour of the view that the manipulations of prior knowledge result in changes of the psychological stimulus space. This interpretation could gain in strength if the average relative weights per information condition could be used to predict learning performance of another group of participants in a similar experiment with category learning instead of similarity rating. Furthermore, such a study could support the hypothesis that effects of knowledge on category learning are mediated by perceived similarity, namely by changes in the psychological stimulus space. Such a test was realised in the second experiment.

EXPERIMENT 2

The purpose of Experiment 2 was to test whether measures of category learning speed and learning efficiency could be predicted on the basis of the average dimension weights per condition as observed in Experiment 1. The same design was used with category learning instead of similarity judgement as the critical phase.

Method

Participants and design. One hundred first-year students at the faculty of psychology and educational sciences of the University of Ghent participated for course requirements and credit. They all volunteered for this particular experiment, and none of them had participated in Experiment 1. They were randomly and equally assigned to the 10 conditions of the design displayed in Table 2.

Materials and procedure. The materials were the same as in Experiment 1, and the procedure was different only for the last phase of the study. Instead of rating the pair-wise similarity, the participants learned to categorise the individual stimuli. In this part of the experiment, only stimulus patterns K1, K2, K6, O8, O9, O13, and O15 (see Table 1) were presented. The block of seven stimuli was presented several times, each time in a different random order. Learning continued until no more errors were committed during the presentation of two consecutive blocks, or until 70 trials had passed, whichever came first.[3]

Results

Knowledge training. Of the 90 participants who were in the knowledge conditions, 72 passed the test after a single presentation; 10 required three presentations of the information. The numbers of successful participants were spread equally across all conditions (between 7 and 10 participants per condition succeeded).

Methodology. Before presenting the results, an explanation of the methodology used is called for. The first part of the method is concerned with model testing; the result of this part yields a rather general assessment of the usefulness of the knowledge-based predictions. The second part uses the parameters obtained in the first part to generate a number of predictions concerning some aspects of the training performance and to estimate the importance of the knowledge effects.

A first question to be considered concerns the way relative dimensional weights can be used to predict learning performance. A possible solution is suggested by applications of the Generalised Context Model (e.g., Nosofsky, 1984, 1986, 1992b). This model predicts category learning performance on the basis of three sets of parameters: (1) the relative importance assigned to the stimulus dimensions (the "attentional" weights), (2) the steepness of the similarity gradient (c), and (3) the category biases. Given stimulus objects grouped in two categories, C_1 and

[3]The learning phase was followed by a transfer phase in which the participants were requested to categorise swiftly all 15 stimulus patterns several times in a random order. Finally, there was a typicality phase in which participants rated the typicality of each stimulus with respect to the K-category. The latter two phases are of no concern for the present study, however.

C_2, this model defines the probability that an instance S_i is assigned to category R_1 as

$$P(R_1|S_i) = \frac{\sum_{j\in C_1} N_j b_1 s_{ij}}{\sum_{j\in C_1} N_j b_1 s_{ij} + \sum_{k\in C_2} N_k(1-b_1) s_{ik}} \tag{1}$$

where N_j represents the relative frequency with which instance j was presented during training, where b_1 ($0 \leqslant b_1 \leqslant 1$) is a free parameter representing the bias towards category 1, and where s_{ij} refers to the similarity between exemplars i and j. This similarity is defined as exponentially related to distance.

$$s_{ij} = e^{-c d_{ij}} \tag{2}$$

where c ($c > 0$) is a free parameter representing the steepness of the similarity gradient, and d_{ij} is the distance between stimuli i and j according to a Euclidean or a city-block metric in which the dimensions of variation are assigned different weights. In general, the distance for a n-dimensional stimulus set is given as

$$\mathrm{d_{ij}} = \left(\sum_{k=1}^{n} \mathrm{w_k} |x_{ki} - x_{kj}|^r\right)^{\frac{1}{r}} \tag{3}$$

where w_k ($0 \leqslant w_k \leqslant 1$ and $\sum_k w_k = 1$) are free parameters, x_{ki} is the kth coordinate of the ith stimulus, and r is taken to be 1 (city-block metric) or 2 (Euclidean metric). For the present application, a city-block metric is used.

Basically, two different processes operate in the GCM. One process is attentional selectivity which is reflected in the relative weights (w_k) assigned to the stimulus dimensions by the end of training. The second process is one of generalisation: The smaller the value of the parameter c, the more two stimuli tend to be considered similar; in other words, the more the behaviour associated to one stimulus tends to be generalised to the other one.

The prediction of the present study is related to the aspect of attentional selectivity, and so it seems that the GCM-framework can be used to verify the prediction. Application of GCM to the present data requires four free parameters, namely c, b_1, w_1 (arms dimension), and w_2 (hat dimension; the weight for the antennae dimension w_3 is simply $1-w_1-w_2$). The predictions obtained in this way are the best predictions that can be achieved within the framework. Taking into account the characteristics of the GCM, this means that three sources of information contribute to the prediction in the most optimal way: attentional selectivity, category bias (a decision aspect), and stimulus generalisation. By means of a maximum likelihood procedure, optimal values for the four free parameters can be estimated. The details of this procedure have been extensively described by Nosofsky (1992a).

The knowledge framework used in the present study involves only the prediction of attentional selectivity, however. Applied to the GCM, this means that the parameters w_1 and w_2 are fixed: The values obtained in Experiment 1 are thought to predict performance in categorisation learning. This results in a restricted version of the GCM with only two free parameters (c and b_1), which can be estimated by the same procedure. It is now possible to test whether this restricted model deviates significantly from the complete four-parameter model (see e.g., Nosofsky, 1992a). To that end, a statistic (G^2) is calculated

$$G^2 = -2(\ln L_r - \ln L_c) \tag{4}$$

where $\ln L_r$ is the natural logarithm of the (maximum) likelihood of the restricted model and $\ln L_c$ represents the natural logarithm of the likelihood of the complete model. Asymptotically, G^2 follows a χ^2 distribution with the difference in the number of free parameters between the two models as the degrees of freedom.

Suppose the restricted model yields a good fit. This may happen for a number of reasons. Our prediction is that the attentional weights obtained in the 10 conditions of Experiment 1 are especially suitable to explain learning performance. It may, however, also occur that the two free parameters (c and b_1) explain most of the variance. To control for the latter case, a second restricted model can be used in which $w_1 = w_2 = 0.333$. If this restricted model performs well, it would mean that the exact value of attentional weights is indeed not important.

Further control can be achieved by invoking a third restricted model, in which the attentional weights can vary freely, but in which the c and b_1 parameters are fixed. Because especially the c parameter is a very powerful one, it could be the case that deficiencies in the attentional weights can to a large extent be compensated for by an appropriate value of c. Therefore, c and b_1 were estimated over the complete data set and the same values were used in each of the 10 conditions. This way a restricted model is realised in which the c and b_1 parameters are still optimal to the data set at hand, but such that variations specifically linked to the conditions in the experiment cannot be reflected in the differences of these parameters over the conditions.

On the basis of the parameter estimations obtained, each model can be used to predict training performance. This performance can be measured in a number of ways. In the present study we shall use the proportion of errors per stimulus during training (the variable which is used to fit the models) and the trial number with the last error. Whereas the first measure is especially sensitive to difficulties related to particular stimuli, the latter measure gives an indication about the total duration of learning.

Model fitting. For each stimulus presented during acquisition, the average frequency with which the stimulus was assigned to each of the two categories was calculated per condition. These values were used to obtain a fit of each of the four models discussed in the previous section. The complete four-parameter version of the GCM yielded fits that were quite good in most of the conditions, with log-Likelihoods varying from –20.12 to –78.16 and Root Mean Squared Deviations (RMSDs) varying from 0.040 to 0.160.

The four-parameter version of the GCM was also used to obtain estimates of the c and b_1 parameters that would be valid in all 10 conditions of the experiment. The RMSD of this application amounted to 0.131, which is not excellent but quite good.

The three restricted (two-parameter) versions of the model yielded poorer fits overall than the complete model: The log-Likelihoods ranged from –28.34 to –146.26 in the knowledge-based version, from –21.45 to –24.18 in the equal weights version, and from –22.56 to –103.57 in the fixed generalisation version. The RMSDs varied between 0.105 and 0.271 in the knowledge-based version; between 0.122 and 0.292 in the equal weights version and between 0.056 and 0.217 in the fixed generalisation version.

Application of the G^2-statistic (equation 4) revealed that the fit of the knowledge-based version was significantly poorer than the complete model in each of the 10 conditions: $\chi^2(2)$ ranged from 13.98 to 166.82 (all $p < .001$). Except for the no-knowledge control condition, $\chi^2(2) = 0.20$, the same was true for the equal weights restricted model; in the other conditions $\chi^2(2)$ varied from 11.34 to 122.65 (all $p < .01$). In the fixed generalisation application, the statistic was significant in 8 of the 10 conditions with $\chi^2(2)$ ranging from 2.43 to 81.44.

All these findings show that the complete four-parameter version of the GCM yields a better fit than each of the three restricted two-parameter versions. This does not mean, however, that the more restricted versions fail to predict some important aspect of learning performance. In order to assess the value of the restricted versions, and the knowledge-based version in particular, a second series of tests was performed in which the predictive values of all four versions were compared over all conditions. Two separate tests were performed, one on the proportion of errors per stimulus (the same variable used to obtain the fits) and the other on the trial number on which the last error was committed.

Errors per stimulus. All four versions were used to generate predictions on the expected average number of errors per stimulus in each of the 10 conditions of the experiment. With seven stimuli and 10 conditions, each application yields 70 values that were used in a multiple regression analysis to predict the observed proportion of errors. As model fitting has shown that the complete model yields better fits than each of the restricted models, it is evident that the complete model is expected to yield the better predictions. Because the knowledge-based version of the model is based on information that is not included in the complete model, it is expected that this model contributes significantly to the prediction.

Table 4 displays the inter-correlations between the predictors and the observed average proportion of errors per stimulus over the 10 conditions. The correlations in this table are moderate to high, but it is interesting to note that the knowledge-based predictions do not correlate very well with the predictions of the complete GCM and of the fixed GCM. The multiple regression coefficient amounted to 0.93 (coefficient of determination 0.87) [$F(4, 65) = 106.01$, $p < .001$]. Table 5 displays the details of the analysis. It is clear from this table, that the regression coefficients of the complete GCM, the knowledge-based version, and the equal weights version were all significant. However, the predictions of the equal weights version, though significant, were in the opposite direction to those observed, so that only the complete GCM and the knowledge-based version appear to be useful predictors.

Trial of last error. Prior knowledge may also have an effect on the overall rate of learning, so that in the end fewer trials are needed to master the categorisation. The trial number on which the last error is committed was used to index this aspect of learning rate. On average, 64.0 trials were needed to attain the learning criterion or the end of the learning phase; the averages per condition varied between 56.7 and 68.7.

An exact prediction of the trial number is not possible with the predic-

TABLE 4
Correlations Between the Four Predictors and the Average Proportion of Errors

	Complete GCM	*Knowledge-based*	*Equal-weights*	*Fixed Generalisation*
Knowledge-based	0.42			
Equal-weights	0.87	0.69		
Fixed generalisation	0.96	0.37	0.81	
Errors	0.89	0.57	0.79	0.84

TABLE 5
Results of the Multiple Regression Analysis in Experiment 2 for the Number of Categorisation Errors Per Stimulus

	Regression	*t(65)*	*p-level*
Complete GCM	1.92	6.42	.001
Knowledge-based	1.41	5.97	.001
Equal-weights	–2.03	–3.81	.001
Fixed generalisation	–0.27	–1.07	–

tors used. However, relative predictions were obtained by taking in each condition the largest proportion of errors predicted, and entering this as the prediction for that condition, so that for each knowledge condition a relative prediction is made of the duration of learning. The rationale behind this choice is that the larger the proportion of errors to a stimulus, the more trials will be needed to achieve complete learning for that particular stimulus.

Table 6 displays the correlation matrix of the four predictors and the dependent variable, the trial of last error. In contrast to the proportion of errors per stimulus where the test concerned 70 data points, with the trial of last error, there are only 10 data points, 1 per condition. Nevertheless, the multiple regression yielded interesting results. The multiple correlation was 0.95 (coefficient of determination 0.90) [$F(4,5) = 11.57$, $p < .01$]. Only the regression coefficient of the knowledge-based prediction attained significance, as can be seen in Table 7.

Discussion

Four realisations of the GCM were defined and optimal values of their parameters were estimated. On average, the fits of all four variations varied over the conditions from reasonably good to very good. Occasion-

TABLE 6
Correlations Between the Four Predictors and the Average Proportion of Errors

	Complete GCM	*Knowledge-based*	*Equal-weights*	*Fixed Generalisation*
Knowledge-based	0.26			
Equal-weights	0.41	0.77		
Fixed generalisation	0.59	–0.22	–0.34	
Last error	0.67	0.46	0.16	0.70

TABLE 7
Results of the Multiple Regression Analysis in Experiment 2 for the Trial of the Last Error

	Regression	*t(5)*	*p-level*
Complete GCM	17.69	0.86	–
Knowledge-based	41.53	3.45	.05
Equal-weights	–29.47	–1.04	–
Fixed generalisation	42.53	2.12	.09

ally, the fit was excellent. As could be expected, the four-parameter version of the GCM outperformed each of the three two-parameter versions. This was shown in the large number of significant G^2 tests.

Our main concern in this experiment, however, was not to test the validity of the GCM. The usefulness of this model has been shown in numerous studies, and there were no a priori reasons to believe that the model would fail in the present data. On the contrary, the learning materials used in the present study correspond to ill-defined categories and previous research has shown that the context model (e.g., Medin & Schaffer, 1978; Medin & Smith, 1981) and its extension, the generalized context model (e.g., Nosofsky, 1984, 1986, 1988; Nosofsky, Clark, & Shin, 1989), give a very powerful description of such categorisation phenomena. On the basis of this extensive literature it was expected that the GCM would provide a good to excellent fit of the data, and this expectation was confirmed.

The four sets of predictions have in common that they are derived from the GCM. If the mechanisms described by GCM are sufficient to explain the learning data, then a multiple regression analysis is expected to reveal that one of the predictors is sufficient to explain the variance in the dependent variable. Because the complete GCM version is the most powerful one, this predictor would be the only one contributing significantly to the multiple correlation. If, on the contrary, the restricted versions contain information that is not accounted for by the commonality of the models, then these predictors would also explain some of the variance in the dependent variable.

The multiple regression analysis showed that indeed the complete GCM contributes significantly to the number of errors per stimulus during category learning, but it appeared that the regressions of the knowledge-based version and the equal weights version were also significant. However, only the regressions of the complete and the knowledge-based versions were positive, indicating that these models capture the difficulty of the individual stimuli during category learning. It is interesting to note that this result is obtained in a situation where the correlation of the knowledge-based predictor and the dependent variable is lower than the correlations of the other predictors with the dependent variable. This supports the conclusion that the variance shared by the knowledge-based prediction and the dependent variable is unique.

The same predictors were used to predict the trial of last error. In this multiple regression analysis, one predictor appeared to be sufficient to explain the variance in the dependent variable, namely the knowledge-based prediction. It was the only predictor whose regression coefficient was significant. Again the correlations of the knowledge-based prediction with the other predictors and with the dependent variable tended to be small.

These findings taken together constitute clear evidence that both the average number of errors per stimulus and the trial of last error were dependent on the dimensional weights of the psychological space induced by the knowledge manipulations. It must be stressed that this is not a circular prediction. The dimensional weights were obtained in Experiment 1 from an independent sample of participants who went through the knowledge training and then judged the pair-wise similarity of instances. They had not seen these instances before and they had not learned how to categorise the individual instances. The sample of participants in Experiment 2 was obtained independently and received the same knowledge instructions before the category learning took place. The prior exposure to categorical and conceptual knowledge notwithstanding, category learning was still needed. As the average trial of last error occurred around trial 64 and the average number of correct category assignments of the individual stimuli varied largely, the presentation and memorisation of the knowledge did not reduce the category learning part to a trivial exercise.

GENERAL DISCUSSION

The first experiment showed that prior knowledge influences similarity judgements by changing the experienced saliency of the stimulus dimensions involved. Experiment 2 demonstrated that these changes in dimensional saliency predicted the speed of learning and the distribution of errors over the stimuli in category learning. It was concluded that effects of prior knowledge on category learning are mediated by changes of the attentional weights in the psychological space.

The risk of building a circular argument was avoided in this study. First, the knowledge effects were observed both in similarity ratings and in indices of category learning. Second, manipulation of the knowledge presented was done in a controlled way. There was an internal replication by having two different cover stories, and the information presented was selected on the basis of a model that allowed a distinction between different levels of knowledge. Finally, the knowledge given did not constrain interpretation of stimulus information in such a way or to such an extent that category learning became superfluous. This is supported by the finding that the average trial of the last error was rather large in all conditions.

Even though the present data clearly show that effects of prior knowledge on category learning are mediated by effects on similarity perception, some limitations of these experiments should be kept in mind. Only one type of category structure was used in the present study and it was implemented on a set of artificial stimuli. All too often in research on categorisation results based on one type of stimuli or on a single categorical structure are generalised to the entire domain. Therefore, future studies should assess the generality of the present finding.

Nevertheless, the results obtained in the research reported here have some important implications for our views on categorisation and its relationship to knowledge. For one thing, the present paper shows that even if categorisation is considered to be a data-driven process, knowledge acquired prior to the category learning experiment is bound to affect the learning process. In other words, learning a new categorisation is at the same time data-driven and knowledge-driven, and the present study shows that the knowledge input in the process is mediated by the perceived similarity. Whether this is the only route by which knowledge influences category learning is a task for future research to find out.

In any case, the present results suggest that knowledge and similarity are closely interrelated, and this raises the question concerning the relative impact of conceptual and perceptual processes on similarity judgement, and the question of whether the process of perception itself is affected by prior knowledge. It is also clear from the present study that even if knowledge may affect the perceived similarities between exemplars, there remains room for learning the categorisation. Goldstone (1994, 1995) has demonstrated that category learning itself changes the similarity judgements (see also Homa et al., 1979). The question may be raised how category learning affects similarity perception. In terms of the present model, it may be suggested that category learning results in a categorical representation that may be activated by the context of the experiment, so that the activated categorical representation or knowledge may affect the perception and the inference of stimulus similarities. The present results are encouraging for research that intends to specify more precisely how stimulus perception is affected by prior knowledge, and whether the knowledge effects occur at a perceptual level indeed. At least some other studies suggest ways in which such an effect may occur. Schyns and Rodet (1997), for example, have collected evidence showing that the features of the stimuli are not fixed characteristics but may be constructed during category learning, such that the perception of the category would change as a consequence of this process of feature learning (see also Thibaut & Schyns, 1995).

Manuscript received July 1998
Revised manuscript received January 1999

REFERENCES

Anderson, J.R., Kline, P.J., & Beasley, C.M., Jr. (1979). A general learning theory and its application to schema abstraction. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 13, pp. 277–218). New York: Academic Press.

Barsalou, L.W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 626–654.

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

Goldstone, R.L. (1995). Mainstream and avant-garde similarity. *Psychologica Belgica*, *35*, 146–165.

Hayes, B.K., & Taplin, J.E. (1995). Similarity-based and knowledge-based processes in category learning. *European Journal of Cognitive Psychology*, *4*, 383–410.

Heit, E. (1994). Models of the effects of prior knowledge on category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1264–1282.

Heit, E. (1998). Influences of prior knowledge on selective weighting of category members. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 712–731.

Homa, D. (1984). On the nature of categories. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 18, pp. 49–94). New York: Academic Press.

Homa, D., Rhoads, D., & Chambliss, D. (1979). Evolution of conceptual structure. *Journal of Experimental Psychology: Human Learning and Memory*, *5*, 11–23.

Johnson-Laird, P.N., Herrmann, D.J., & Chaffin, R. (1984). Only connections: A critique of semantic networks. *Psychological Bulletin*, *96*, 292–315.

Kruskal, J.B., Young, F.W., & Seery, J.B. (1977). *How to use KYST-2, a very flexible program to do multidimensional scaling and unfolding* [Computer program manual]. Bell Laboratories.

Lakoff, O. (1987). Cognitive models and prototype theory. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 63–100). Cambridge, UK: Cambridge University Press.

Lamberts, K. (1994). Flexible tuning of similarity in exemplar-based categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1003–1021.

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.

McCall, R.B., & Appelbaum, M.I. (1973). Bias in the analysis of repeated measure designs: Some alternative approaches. *Child Development*, *44*, 401–415.

Medin, D.L. (1983). Structural principles in categorization. In T.J. Tighe & B.E. Shepp

(Eds.), *Perception, cognition, and development: Interactional analyses* (pp. 203–230). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Medin, D.L., Goldstone, R.L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, *100*, 254–278.

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

Medin, D.L., & Smith, E.E. (1981). Strategies and classification learning. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 241–253.

Medin, D.L., & Wattenmaker, W.D. (1987). Category cohesiveness, theories, and cognitive archeology. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 25–62). Cambridge, UK: Cambridge University Press.

Murphy, G.L., & Allopenna, P.D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 909–919.

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

Murphy, G.L., & Spalding, T.L. (1995). Knowledge, similarity, and concept formation. *Psychologica Belgica*, *35*, 127–144.

Nakamura, G.V. (1985). Knowledge-based classification of ill-defined categories. *Memory and Cognition*, *13*, 377–384.

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

Nosofsky, R. (1988). Similarity, frequency, and category representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 54–65.

Nosofsky, R.M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.

Nosofsky, R.M. (1992a). Exemplar-based approach to relating categorization, identification, and recognition. In F.G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363–393). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Nosofsky, R.M. (1992b). Exemplars, prototypes, and similarity rules. In A. Healy, S. Kosslyn, & R. Shiffrin (Eds.), *Essays in honor of William K. Estes* (Vol. 1, pp. 149–167). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Nosofsky, R.M., Clark, S.E., & Shin, H.J. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 282–304.

Nosofsky, R.M., & Kruschke, J.K. (1992). Investigations of an exemplar-based connectionist model of category learning. In D.L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 28, pp. 207–250). New York: Academic Press.

Pazzani, M.J. (1991). Influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 416–432.

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

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

Rosch E.H (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*, 192–233.

Schyns, P.G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 681–696.

Spalding, T.L., & Murphy, G.L. (1996). Effects of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 525–538.

Thibaut, J.-P., & Schyns, P.O. (1995). The development of feature spaces for similarity and categorization. *Psychologica Belgica*, *35*, 167–187.

Thibaut, J.-P., & Vandierendonck, A. (Eds.). (1995). Similarity and categorization. [Special Issue]. *Psychologica Belgica*, *35*.

Vandierendonck, A. (1978). Interaction of cue saliency, stimulus unitization, and familiarization in concept identification: Attention or memory phenomenon? *Psychologica Belgica*, *18*, 224–237.

Vandierendonck, A. (1991). Are category membership decisions based on concept gradedness? *European Journal of Cognitive Psychology*, *3*, 343–362.

Vandierendonck, A. (1995). A parallel rule activation and rule synthesis model for generalization in category learning. *Psychonomic Bulletin & Review*, *2*, 442–459.

Wisniewski, F.J. (1995). Prior knowledge and functionally relevant features in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 449–468.

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

APPENDIX

This Appendix contains the paragraphs used in the knowledge acquisition phase of the two experiments. The sentences are shown in the order they were presented to the participants. Before each sentence an indication is given concerning the cover story on which it is based (C1 or C2) or whether it is a categorical statement (CA) or conceptual information (CO), and the other numbers refer to the cells of the design in which the sentence was used.

General information

C1: 3,4,5,6. All Kwarks are workers; their arms are for working. All Orkels are thinkers; the bigger their hat, the bigger their brain for thinking.

C2: 7,8,9,10. Kwarks transport the ore in the mine; they use their arms to transport the ore. Orkels bore the mine galleries; the bigger the hat, the deeper they can drill.

CA: 2,4,6,8,10. Kwarks mostly (but not always) have many arms and a small hat. Orkels mostly have few arms and a big hat.

C1: 3,4,5,6. Kwarks and Orkels are not all equally happy, but, in general, the bigger the hat and the fewer the number of arms, the happier the robot is. Robots with a smaller hat and with more arms are less happy.

C2: 7,8,9,10. Kwarks and Orkels are not able to work at the same levels in the mine, but in general it can be said that the bigger the hat and the smaller the number of arms the better the robot is equipped for working at deep levels. It can also be said that robots with smaller hats and more arms are not equipped for working at deep levels.

CA: 2,4,6,8,10. The exact number of arms and the size of the hat varies as well in Kwarks as in Orkels. In general it can be said that the more arms a robot has, the smaller the hat, and the more downward the orientation of the antennae will be. It is also true that the fewer the number of arms on a robot, the larger the hat and the more upward the antennae orientation is.

C1, 3,4,5,6. The happier the robots are, the more their antennae are oriented upwards; the more unhappy they are, the more the antennae are oriented downwards.

C2: 7,8,9,10. The deeper the robots can go in the mine, the more upwards the orientation of their antennae; the less deep the robots can go, the more the orientation of the antennae is downwards.

Information about the exception

C1: 5,6. The previous text contained a description of characteristic features of Kwarks and Orkels. These characteristic features are not always present: Kwarks are trained by a special Orkel, which is able to explain and to demonstrate the simple tasks the Kwarks have to perform. As a consequence, this Orkel with a small hat and many arms is happy.

C2: 9,10. Via the mine shaft all robots are provisioned by an Orkel that transports all supplies. Because this Orkel is not required to drill, it has a small hat, and it still can reach the deeper mine galleries.

CA, CO: 5,6,9,10. So there is an Orkel with many arms and a small hat.