The Effect of Category Learning on the Representation of Shape: Dimensions Can Be Biased but Not Differentiated

Hans Op de Beeck, Johan Wagemans, and Rufin Vogels
University of Leuven

Recent studies have suggested a profound influence of category learning on visual perception, resulting in independent processing of previously integral dimensions. The authors reinvestigate this issue for shape dimensions. They first extend previous findings that some shape dimensions (aspect ratio and curvature) are processed in a separable way, whereas others (radial frequency components) are not. They then show that a category-learning phase improved the discrimination of a relevant with respect to an irrelevant dimension, but only for separable dimensions. No similar effect was found on the relative sensitivity for integral shape dimensions. Thus, category learning is capable of biasing separable shape dimensions but does not alter the status of dimensions in the visual system as either separable or integral.

When we perceive an object, we automatically classify it as belonging to separate categories (dangerous or harmless, pleasant or nasty, dog or cat, etc.). Most models of categorization describe this behavior with two steps (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Ashby & Perrin, 1988; Kruschke, 1992; Nosofsky, 1984, 1986; Smith & Medin, 1981). First, they assume a psychological similarity space, with stimuli being represented by points or regions in a space in such a way that perceived similarity is reflected by spatial proximity. In the second step, stimuli within this space become linked with specific categories.

The validity of assuming a psychological similarity space was investigated previously in a series of articles that led to the development of techniques to analyze measures of pairwise similarity (for a review, see Shepard, 1980). However, until recently it was often ignored that these studies did not reveal the mechanisms that underlie the construction of such a similarity space, as these are the sort of questions tackled in studies of object perception. Recent models of categorization have incorporated the idea that a complete stimulus representation is not available at the time that a categorization response is made (Lamberts, 2000). These models assume that this representation is constructed gradually by sampling stimulus features in a serial manner. Nevertheless, visual processing is still regarded as some kind of black box that provides the feature and dimension values needed for computing similarity while neglecting how these values are obtained.

As a consequence, categorization research and research on perception have emerged as distinct domains with separate theoretical frameworks (Cutzu & Edelman, 1998). This is most regrettable, because empirical evidence suggests that these processes interact. That is, not only does the relative distance and clustering of stimuli in psychological space determine the ease of learning specific category rules (e.g., Ashby et al., 1998; Nosofsky, 1987; Op de Beeck, Wagemans, & Vogels, 2001), but categorization also influences perceived similarities (Goldstone, 1994; Goldstone, Lippa, & Shiffrin, 2001; Livingston, Andrews, & Harnad, 1998). The aim of the present article is to bridge the theoretical gap between research on categorization and perception by investigating the effect of categorization on the perception of stimuli varying along shape dimensions.

The Effect of Categorization on Visual Perception

Recently, an increasing amount of evidence has confirmed the hypothesis that the perception of differences in stimuli is influenced by their relevance for a specific behavior. Several decades ago, the concept of *selective attention* was introduced in learning theories to explain this effect of relevance. Selective attention to dimensions that are relevant for behavioral choices results in a stimulus representation that is expanded along these dimensions relative to irrelevant dimensions (for a review, see Hall, 1991).

Shepard and Chang (1963) suggested that the usefulness of incorporating selective attention into a categorization model depends on the sort of dimensions along which stimuli vary. For some combinations of dimensions (e.g., size and brightness), it is quite easy to attend to stimulus variations along one dimension while ignoring the other dimension (Foard & Kemler Nelson, 1984; Garner, 1976; Melara, Marks, & Potts, 1993). For these so-called "separable" dimensions, the fit between categorization models and behavioral data is increased by incorporating a parameter for selective attention in the models. By contrast, for other combinations of dimensions (e.g., saturation and brightness), it is difficult, if not impossible, to ignore variations along one dimension while attending to the other dimension. For these "integral" dimensions, the incorporation of selective attention into a catego-

rization model did not always result in increased fit values (see also Nosofsky, 1986, 1987).

A major problem with the applied models is that they do not differentiate between perceptual and decisional processes. As a consequence, they do not allow determination of whether the benefit of including a selective attention parameter in a model is related to a change in the perception of stimuli, or rather to a change in decision strategies. Recent work by Maddox (2001, 2002) has investigated the same problem within the framework of general recognition theory (Ashby & Townsend, 1986), a theory that differentiates between perceptual and decisional processes. Subjects learned to categorize stimuli according to two sorts of rules: *Decisional selective-attention categorization problems* could be solved by attending to one of two separable stimulus dimensions, whereas the other problems, *decisional integration categorization problems*, required subjects to combine information from the two stimulus dimensions (the category rule was oriented diagonally in the two-dimensional stimulus space). Selective attention was apparent only in the first sort of rules, and no change in the perceptual stimulus representation could be found when the category rule did not correspond with a separable stimulus dimension.

So, the effect of categorization on stimulus representation seems to interact with the way in which stimuli are encoded—that is, whether the relevant and irrelevant stimulus variations are separable. Goldstone (1994) tested this hypothesis explicitly by measuring the ability of participants to discriminate stimuli after categorization training. One group of subjects learned to categorize stimuli along one of two dimensions, and their performance in a same–different task was compared with that of control participants receiving no prior training. As expected, Goldstone found a difference between separable and integral dimensions. Learning that only one of two separable dimensions is relevant resulted in a discrimination improvement that was specific for that dimension. However, with integral dimensions, learning generalized to the irrelevant dimension. This result agrees with the hypothesis that a selective weighting of dimensions by attention is more prominent with separable compared with integral dimensions.

However, several studies have suggested that the effect of learning is more profound and that learning can influence whether dimensions are integral or separable. First, Foard and Kemler Nelson (1984) suggested that learning is one of the factors that determine whether dimensions are integral or separable (other factors include the actual dimensions used, instructions, and the available processing time).

Second, Goldstone (1994) found more general learning for integral dimensions compared with separable dimensions, as described above, but a small part of the learning effect for integral dimensions was specific for the relevant dimension. Goldstone hypothesized that learning could have influenced the processing of the integral dimensions in such a way that they were transformed into more separable dimensions.

Third, Goldstone and Steyvers (2001) proposed a new paradigm to determine whether integral dimensions are transformed into separable dimensions during a training phase. The authors referred to this transformation as *dimension differentiation*. The experiments revealed that a short period of category learning facilitates the learning of a new category rule with the same stimulus set as long as this new rule can be described with the same set of dimensions as the first rule. For example, facilitation between two category rules occurs when the irrelevant dimension of the first rule becomes relevant and the relevant dimension irrelevant. In a two-dimensional stimulus space, such a reversal of dimensional relevance corresponds to orthogonal category rules (the orientation difference between the two rules is 90°). No facilitation occurs, however, when the new category rule is defined by a novel combination of the dimensions used in the old rule (e.g., when the orientation difference between the two category rules is 45° in stimulus space).

This demonstration of dimension differentiation could be taken as evidence for the hypothesis that a relatively short learning phase can change the basic dimensions used to analyze stimuli in the visual system. This provocative hypothesis was discussed by Schyns, Goldstone, and Thibaut (1998), and it suggests that experience can induce the creation of entirely new dimensions. However, there are several difficulties with this proposal. First, it is always difficult to prove that behavioral learning effects are related to changes in visual representations, as these effects could reflect more "postperceptual" effects. Second, many of the findings that are suggestive for the creation of new functional features (e.g., Oliva & Schyns, 1997; Schyns & Murphy, 1994) can also be interpreted as evidence for a selective attention process that operates on already existing features. Third, even if the learning effects are related to changes in visual processing beyond selective attention processes (in particular, dimension differentiation), it is still not known how the underlying mechanisms could be implemented in the visual system. Such questions can be answered only by considering how the dimensions under study are processed in the visual system. These issues were neglected in the categorization literature, partly owing to the use of combinations of dimensions for which such knowledge is not available. Indeed, previous studies often involved comparisons between quite diverse dimensions (saturation, brightness, size, face dimensions such as nose length and eye separation, etc.). For example, comparing integral with separable dimensions has been operationalized as a comparison between saturation/brightness and face dimensions (e.g., Nosofsky, 1986, 1987). To avoid these drawbacks of previous studies, we decided to investigate the interaction between perception and categorization within the context of shape perception. As will be described below, we have detailed knowledge about how shapes are analyzed in the visual system through psychophysics, neurophysiology, and computational modeling.

Independent Shape Dimensions

Starting with the seminal work of Hubel and Wiesel (1962), evidence from neurophysiological studies has accumulated to show that early vision analyzes shapes in terms of simple features such as the orientation of line segments (for a review, see Hubel & Wiesel, 1977; Treisman, Cavanagh, Fischer, Ramachandran, & von der Heydt, 1990). In higher order stages of the visual system, neurons are tuned for increasingly complex combinations of these simple features (i.e., more complex shape dimensions; see Hegdé & Van Essen, 2000; Pasupathy & Connor, 1999; Tanaka, 1996; Tanaka, Saito, Fukada, & Moriya, 1991). Although it is not yet completely clear which dimensions are coded at the higher stages, the gradual increase in the complexity of tuning properties is incorporated in all major models of shape recognition (e.g.,

Biederman, 1987; Edelman & Intrator, 2000; Riesenhuber & Poggio, 1999).

Two recent studies examined behaviorally whether shapes are encoded with a set of independent shape dimensions (Arguin & Saumier, 2000; Stankiewicz, 2002). These studies showed that shape dimensions such as curvature, aspect ratio, and tapering have a special status as being processed independently by the human visual system, mirroring previous findings with simpler dimensions (e.g., orientation and spatial frequency of gratings; see Regan, 2000; Vogels, Eeckhout, & Orban, 1988). As shown by Stankiewicz (2002), discrimination of curvature is not affected by changes in orientation (and vice versa). The visual search experiments of Arguin and Saumier (2000) revealed lower search rates when targets and distractors could be differentiated by looking at only one independent dimension compared with a situation in which two dimensions had to be taken into account. Global similarity was controlled for and could not explain this result. None of these effects were found with dependent dimensions, being linear combinations of the independent dimensions in these studies.

Neurophysiological studies on the (in)dependent coding of shape dimensions have mainly concentrated on simple properties such as the orientation and spatial frequency of gratings that are coded in primary visual cortex (e.g., Jones, Stepnoski, & Palmer, 1987; Mazer, Vinje, McDermott, Schiller, & Gallant, 2002). The orientation and spatial frequency tuning curves of neurons in macaque and cat V1 were shown to be relatively invariant across a range of spatial frequencies and orientations, respectively. Relative invariance means that the relative strength of responses (i.e., the ranking) to different values of one dimension is the same across variations in the other dimension. Relative tuning invariances have been shown for other combinations of dimensions not restricted to the shape domain (e.g., shape and color, Komatsu & Ideura, 1993; size and position, Op de Beeck & Vogels, 2000). However, no study has investigated how combinations of complex shape dimensions are implemented. Possibly, neurons at the higher levels of the visual system encode shapes using a limited set of more complex independent dimensions, with tuning for one dimension being relatively invariant with respect to the value of other dimensions.

The notion of dimension differentiation suggests a mechanism that can change this tuning of units in the visual system. A pair of integral dimensions could be "differentiated" into separable dimensions by making the tuning of neurons for one dimension relatively invariant across different values of the other dimension. To date, no neurophysiological studies have addressed this issue.

The Present Investigation

In this series of behavioral experiments, we determined to what extent the processing of well-defined shape dimensions can be modified by category learning. Experiment 1 extended the visual search paradigm of Arguin and Saumier (2000) by including a more diverse set of shape dimensions. It provided direct evidence that some pairs of shape dimensions are processed in a separable way and others in an integral way. The same sets of shape dimensions were used in Experiments 2 and 3, which investigated the effect of category learning on the sensitivity for relevant and irrelevant dimensions. For separable

dimensions, we expected that selective attention would result in an increased sensitivity for a relevant dimension but not for an irrelevant dimension (i.e., *dimension-specific learning*, as found by Goldstone, 1994). If dimension differentiation occurred during category learning, then we would predict dimension-specific learning for integral dimensions. However, if dimension differentiation did not occur, then we would expect the learning effects to generalize to irrelevant integral dimensions. The results showed general learning with integral shape dimensions and specific learning with separable dimensions. This pattern of results indicates that integral shape dimensions were not differentiated by category learning.

Stimuli

We included pairs of shape dimensions for which we had at least indirect evidence that they are processed in a separable or an integral way. Arguin and Saumier (2000) showed that attention can be allocated to either the curvature or the aspect ratio ("thickness") of ellipses, while the other feature is ignored (similar results were obtained with aspect ratio and tapering). On the basis of this evidence, we used aspect ratio and curvature as separable dimensions in our experiments.

Arguin and Saumier (2000) argued that dimensions are processed in an integral way if the dimensions along which stimuli are manipulated are poorly correlated with the actual dimensions used by the visual system to analyze the stimuli. A classic example of mathematically defined shape dimensions that are not used by the visual system are Fourier descriptors, or radial frequency components (RFCs), which describe the contour of a shape as a sum of sinusoidal modulations (Zahn & Roskies, 1972). Manipulating the amplitude or phase of several RFCs causes perceptually salient shape variations (e.g., Cortese & Dyre, 1996; Wilkinson, Wilson, & Habak, 1998), but there is no behavioral evidence that shapes are decomposed into RFCs by the visual system. Neurophysiological studies showed a clear contrast between the coding of RFC dimensions and the coding of separable dimensions such as shape and color. We have already noted that the tuning of single neurons is relatively invariant for these separable dimensions, but no relative invariance was found for RFC dimensions (Albright & Gross, 1990): The preferences of single neurons for particular values of one RFC are not preserved when another RFC is changed. Nevertheless, some combinations of RFC dimensions (in particular, amplitude variations of two different RFCs) are represented in an ordinally faithful psychological space (same dimensionality and stimulus order as in parametric space) by humans and macaque monkeys (Cortese & Dyre, 1996; Op de Beeck et al., 2001). As discussed by Edelman (1999; see also Cutzu & Edelman, 1998), the dimensionality of the perceptual shape space can fit that of the parameter space, even when the actual dimensions used by the visual system to analyze the stimuli differ from the parametric dimensions. So, we included RFC variations in our experiments, because they have the properties of integral dimensions as suggested by Arguin and Saumier (2000).

Methods for Stimulus Construction

Stimuli Defined by RFC Dimensions

We randomly constructed several groups of shapes with their contour defined by seven RFCs. In our RFC algorithm, a line was

modulated by seven sinusoidal functions (RFCs, each with three parameters: frequency, phase, and amplitude), after which the line was bent to create closed contours.

All shapes in each group shared five RFCs with a particular frequency, phase, and amplitude. We varied the amplitude of two additional RFCs in steps between .03 and .46 radians, again with a fixed frequency and phase. To calibrate both dimensions, we compared the number of pixels in the shape images that changed their figure–ground status by increasing the amplitude of each RFC from minimal to maximal amplitude. The amplitude modulations of the two RFC dimensions were equated to get a maximal difference of 5% between the dimensions on this pixel-based measure.

We selected one group, referred to as *complex RFC stimuli*, out of the large number of possibilities (Figure 1A). A second stimulus set was constructed in which the only RFC components with nonzero amplitude were the two manipulated dimensions, resulting in far simpler shapes (*simple RFC stimuli*; Figure 1B). The Appendix describes the results of a pilot study, which replicated the findings of Cortese and Dyre (1996) that the perceptual similarity space of each of these shape groups was two-dimensional with the same stimulus order as in the parametric space (similar results with other shape groups were published previously; see Op de Beeck et al., 2001). As stated before, there is no evidence that humans encode these shapes using RFC dimensions, but the intershape similarities can nevertheless be captured in a two-dimensional geometric space; that is, the RFC dimensions define the psychological similarity space.

Stimuli Varying on Aspect Ratio and Curvature

A third stimulus set, referred to as *deformed ellipses*, was similar to the stimuli used in the first experiment of Arguin and Saumier (2000). These shapes corresponded to deformed ellipses that varied in their aspect ratio (i.e., ratio of lengths of minor over major axis) and curvature of the main axis (see Figure 1C). Both dimensions were calibrated in pixel space using the same procedure as for the RFC sets.

Experiment 1: Validating Separability–Integrality of Shape Dimensions

For the purpose of the present experiments, we needed pairs of shape dimensions that were undoubtedly separable or integral. Our selection of stimulus sets was based on rather indirect evidence, because none of these dimensions had been studied with classic paradigms to distinguish between integral and separable dimensions. These classic paradigms include triad sorting, Garner interference, and various single-dimension and multiple-dimension categorization and matching tasks (e.g., Foard & Kemler Nelson, 1984; Garner, 1976; Maddox, 1992; Nosofsky, 1987).

A notable exception is the study of Stankiewicz (2002), which applied a variant of Garner interference to demonstrate the independent processing of aspect ratio and curvature. Subjects in a Garner interference paradigm are required to categorize stimuli along one of two dimensions. Dimensions are labeled as separable if variation on the irrelevant dimension does not interfere with performance (the result found by Stankiewicz with aspect ratio and curvature). One cannot draw the opposite conclusion in the other
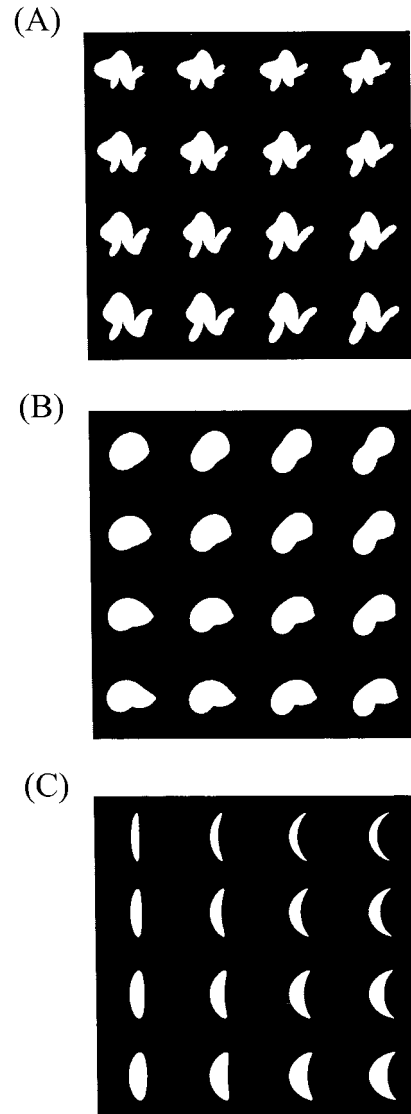

Figure 1. Sixteen stimuli from each stimulus set used in Experiments 1–3: (A) complex radial frequency component (RFC) stimuli, (B) simple RFC stimuli, and (C) deformed ellipses. In (A) and (B), the amplitude of one RFC increases from the top to the bottom row (vertical dimension) and from the left to the right column (horizontal dimension). In (C), aspect ratio increases from the top to the bottom row and curvature from the left to the right column.

case: Degradation of performance by the irrelevant variation does not necessarily imply integrality. It is possible that subjects are uncertain about the task they have to perform, in particular about the rule or stimulus features they have to use to categorize stimuli. If this kind of uncertainty exists at the decision level, then it will be impossible to find evidence for separable processing even if dimensions are separable. One can reduce decisional uncertainty by giving explicit verbal instructions to the subjects about the (ir)relevance of dimensions or by including a sufficiently long training phase (as in Stankiewicz, 2002).

Even with carefully designed studies, it is impossible to exclude completely the effect of decisional variables on performance in

any of the classic paradigms. The most systematic investigation of this issue was performed by Ashby and colleagues in the context of the general recognition theory (GRT; Ashby & Townsend, 1986; Maddox, 1992; Thomas, 1996). It was noted that performance of subjects in all of the classic paradigms reflects not only how the visual system encodes the stimulus dimensions but also task-related processes (e.g., strategies and decision rules). So, it is not surprising that different paradigms may reveal contrasting results, and it is widely acknowledged that convergent evidence needs to be found across different methods. Ashby and colleagues proposed several tasks that need to be combined in order to get more unambiguous evidence for separability–integrality. In these tasks, subjects are requested to use the values on one dimension of a pair, on the other dimension, or on both dimensions.

It is not straightforward to apply either the classic paradigms or the GRT framework to RFC dimensions. These dimensions cannot be identified by a verbal label, and a preview of stimuli of the different categories is necessary to instruct subjects in each new task (actually, the nonexistence of such a label is already suggestive for the fact that these dimensions have a different status compared with, for example, curvature). We noted before that a correct understanding of the task instructions by the subjects is a prerequisite for interpretable findings in the Garner interference paradigm, but it is also of great importance in the GRT framework. The only way to exclude these ambiguities would be to include a long training phase before each task, but the latter could modify the processing of these dimensions. It is clear that we need to avoid this problem, because the effect of experience on the processing of dimensions is exactly what is studied in our experiments.

The ambiguities with interpreting the interaction between dimensions could be solved with a task that does not require the definition of a new task rule for each different stimulus set. So, we decided to use a completely different paradigm: visual search. This paradigm has become a standard way to distinguish between differences in elementary visual features and stimulus differences that are better described as being a conjunction of elementary features (Treisman & Gelade, 1980). The visual search paradigm was applied by Arguin and Saumier (2000) to study the interaction between shape dimensions, and it provides an elegant solution to the problems mentioned previously. The instructions given to participants are always exactly the same ("Is the target present or not?"), no matter on which dimensions the stimuli vary. Only a short preview of the target is needed, and subjects can readily apply the task to a new stimulus set. Not only does this paradigm solve the problems raised above, the results of the present study as a whole become more impressive because it involves a comparison (and convergence) between very diverse paradigms (visual search and categorization). Thus, we applied the visual search paradigm to all of our stimulus sets to investigate whether some of the dimensions can indeed be considered as elementary/independent dimensions and others not. In words that are more standard in the categorization literature, it can establish whether some of the pairs of dimensions are separable (independent) and others integral (dependent).

The critical comparison in Experiment 1 is between a condition in which attention to one stimulus dimension is sufficient to discriminate the target from the distractors (condition "1D") and a condition in which two dimensions must be attended to in order to select the target (condition "CONJ"). These two conditions are analogous to feature and conjunction search, respectively, in standard applications of the visual search paradigm. If the relevant stimulus dimension in condition 1D is an elementary dimension of the visual system, then one would expect relatively shallow search slopes (little effect of display size), as in typical feature search (Treisman & Gelade, 1980). In contrast, the target in condition CONJ can be discriminated from all distractors only by processing both stimulus dimensions, and so one would expect relatively steep search slopes, as in typical conjunction search. So, if the manipulated stimulus dimensions map onto elementary perceptual dimensions (as suggested before for aspect ratio and curvature, but not for RFC dimensions), then we would expect steeper search slopes in condition CONJ compared with condition 1D.

Many experiments have shown that this distinction between feature and conjunction search is an important factor in determining search slopes, but it is not the only factor. Several studies have pointed to effects of the similarity between target and distractors and the similarity among distractors (e.g., Duncan & Humphreys, 1989; Treisman, 1991). The more the target resembles the distractors, the stronger the effect of distractor number. Similarity among distractors has the opposite effect: Searching for a target among a homogeneous set of distractors is easier compared with search among a heterogeneous distractor set. Given these known influences of similarity, we needed to control for target–distractor and distractor–distractor similarity to exclude this factor as a potential confound. In addition to the actual search task (referred to as Experiment 1A) we therefore included a control experiment to measure perceptual similarity (Experiment 1B). Experiment 1B demonstrates that the differences between stimulus sets in Experiment 1A cannot be attributed to pairwise target–distractor and distractor–distractor similarity.

Method

Participants

Twelve naive volunteers from diverse undergraduate and graduate programs at the University of Leuven participated in this study. Six participants were involved in Experiment 1A, and the other 6 participants in Experiment 1B. They were paid for their participation.

Stimuli

We included the three stimulus sets described before: complex RFC stimuli, simple RFC stimuli, and deformed ellipses. Four stimulus triads (all including one target and two distractors; see Figure 2) were constructed for each set, following closely the design of Arguin and Saumier (2000). In the first two triads (condition 1D), the target differed from both distractors on one parametric dimension only (amplitude of one RFC or aspect ratio/curvature), whereas the distractors differed only on the orthogonal dimension. In these target–distractor combinations, one and only one dimension is relevant when searching for the target. In the two other triads (condition CONJ), the target differed from one distractor on one parametric dimension and from the other distractor on the other parametric dimension.

The exact location of all stimuli in the stimulus space was chosen before the experiments started and was based on the performance of 4 pilot subjects (all authors and a naive subject). These subjects performed a pilot version of Experiment 1B with preliminary stimuli. This performance provided a crude estimation of differences in target–distractor and distractor–distractor similarity between conditions and between stimulus sets. When large differences were noted in this pilot experiment, stimuli
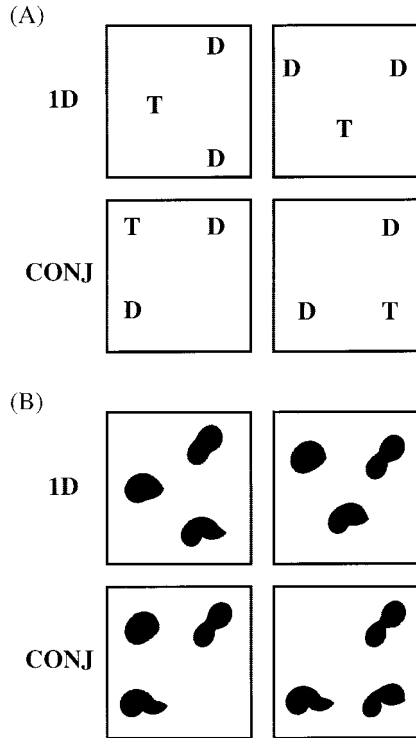
(A)

(B)

Figure 2. Orientation of the four stimulus triads in Experiment 1 in a two-dimensional parametric stimulus space. (A) Each triad contains one target image (T) and two distractors (D). The relative arrangement of these targets and distractors can be applied to any of the two-dimensional configurations of stimulus shapes in Figure 1 to obtain a specific stimulus triad, as illustrated in (B) for the simple radial frequency component stimuli. 1D = condition in which attention to one parametric dimension is sufficient to discriminate the target from the distractors; CONJ = condition in which two dimensions must be attended to in order to select the target.

were changed for the actual Experiments 1A and 1B (e.g., to increase target–distractor similarity in a particular triad).

In this experiment and all of the following ones, stimuli were presented on a black background and filled with noisy textures consisting of black and white pixels (in which the number of black and white dots was the same for 2×2 squares in the texture, resulting in highly uniform textures). Using this method, we constructed three different bitmaps for each original stimulus, each containing the same shape filled with another noise texture. The bitmap to be used was selected randomly from these three alternatives for each stimulus presentation.

Materials

The stimuli (maximum size, $2.4° \times 2.4°$) were presented on a 17-in. monitor in a darkened room. The distance between the participants and the screen was about 60 cm. In this experiment and all of the other ones, participants had to indicate their responses by means of a QWERTY keyboard on which two keys ("Z" and "/") were marked.

Procedure

Experiment 1A. Six subjects participated in Experiment 1A, which involved three sessions, one for each stimulus set (complex RFC stimuli, simple RFC stimuli, and deformed ellipses). Each session consisted of two blocks of 60 trials for each of the four triads (for a total of eight blocks).

Two factors were crossed within each block: display size (three, five, or seven stimuli) and target presence (present or absent). The search array of a target-present trial consisted of an equal number of each distractor (e.g., one target image and three images of each distractor, with a display size of seven), whereas almost equal numbers of the two distractors were present in a target-absent array (e.g., four images of one distractor and three of the other, with a display size of seven).

Each block began with a simultaneous presentation of all stimuli from a triad with the target in the top half of the screen and the two distractor images side by side in the bottom half. The first trial was started when the subject pressed a key. A fixation point was present throughout the trial, and the appearance of the target image was preceded by a warning sound. The target image was visible for 300 ms, followed by an interstimulus interval of 500 ms and the actual search array. This search array disappeared when the subject responded. The target and distractors were randomly presented at 12 positions equally spaced on an imaginary circle 12° in diameter.

The procedure was explained to each participant. Subjects were instructed to respond "present" or "absent" as accurately and quickly as possible. The key that was associated with each response was varied across subjects, as was the order of the three stimulus sets. Each subject completed six practice trials of each triad that would appear in a session before the start of the actual experiment.

Experiment 1B. Six other subjects participated in Experiment 1B, which was designed to measure target–distractor and distractor–distractor similarity in Experiment 1A. The critical difference with Experiment 1A was that targets in Experiment 1B were searched for among homogeneous distractors (i.e., all distractors in a given block of trials were the same). Each subject completed three sessions—again, one for each stimulus set. Each session consisted of two blocks of 60 trials for each of the possible pairwise stimulus permutations (4 triads × 3 permutations for each triad), for a total of 24 blocks for each session. Two of three permutations (condition T-D) for each triad involved the original target as "target" in each trial, and one of the two original distractors as "distractor" (*original* refers to the status of stimuli in Experiment 1A). The third permutation (condition D-D) involved one original distractor as target and the other one as distractor. As in Experiment 1A, two factors were crossed in each block: display size (three, five, or seven stimuli) and target presence (present or absent). The search array of target-present trials consisted of the target image and two, four, or six images of one distractor, whereas all images were identical in target-absent arrays. This procedure for controlling stimulus similarity is identical to the one of Arguin and Saumier (2000, Experiments 1B and 2B), except for the fact that we also controlled for distractor–distractor similarity.

Each block in Experiment 1B began with a presentation of the target image in the top half of the screen and the distractor image underneath it. The trial procedure and the instructions were the same as in Experiment 1A.

Results

Experiment 1A

The main dependent variable of interest in search tasks is response time (RT) in correct trials. Error rates were low overall (5% on average) and showed no speed–accuracy trade-off. On the contrary, subjects made more errors in those conditions where RT was slow (e.g., error rate of 10% in target-absent trials with a large display size). Thus, search slopes cannot be caused by a speed–accuracy trade-off.

All analyses were performed on the logarithm of RT after the removal of outliers (RT more than 3 standard deviations from a subject's mean for a condition, which occurred in less than 1% of the trials). Analyses without this data trimming revealed similar

results. The correct RT for each stimulus set was analyzed with an analysis of variance (ANOVA) including three within-subject factors: condition (1D or CONJ), target presence (present or absent), and display size (three, five, or seven items). The complete list of F values and associated probabilities in each ANOVA is provided in Table 1, and so here we describe only the most important results. The results are displayed in Figure 3.

Effects that are typically found in applications of the visual search paradigm were present for all stimulus sets. RT in target-present trials was always faster compared with target-absent trials. Trials with larger display sizes were associated with slower RTs, an effect that interacted significantly with target presence in all analyses.

The main purpose of the experiment was a comparison of the effect of condition between the stimulus sets. We expected to find more efficient search in condition 1D compared with condition CONJ for aspect ratio and curvature, but not for RFC dimensions. This prediction was confirmed by a significant interaction between condition and display size (larger effect of display size in condition CONJ) that was found only with deformed ellipses (see Figure 3C). Several other effects, including the factor "condition," were found with the deformed ellipses: a main effect of condition (1D faster than CONJ), an interaction between condition and target presence (larger effect of condition in target-absent trials), and a

Table 1
Results of Analysis of Variance on Reaction Time Values in Experiment 1

Effect	df	F	MSE	p	η^2
Complex RFC stimuli					
1D vs. CONJ (1)	1, 5	0.25	.0060	.64	
Target presence (2)	1, 5	146.80	.0016	<.01	.26
Display size (3)	2, 10	223.62	.0010	<.01	.49
(1) × (2)	1, 5	11.10	.0004	.02	.01
(1) × (3)	2, 10	0.01	.0006	.98	
(2) × (3)	2, 10	28.07	.0006	<.01	.04
(1) × (2) × (3)	2, 10	0.52	.0002	.61	
Simple RFC stimuli					
1D vs. CONJ (1)	1, 5	0.19	.0019	.68	
Target presence (2)	1, 5	41.71	.0045	<.01	.14
Display size (3)	2, 10	110.58	.0015	<.01	.24
(1) × (2)	1, 5	6.10	.0006	.06	
(1) × (3)	2, 10	1.27	.0004	.32	
(2) × (3)	2, 10	15.22	.0008	<.01	.01
(1) × (2) × (3)	2, 10	1.94	.0002	.19	
Deformed ellipses					
1D vs. CONJ (1)	1, 5	73.71	.0010	<.01	.05
Target presence (2)	1, 5	26.24	.0060	<.01	.12
Display size (3)	2, 10	104.93	.0015	<.01	.23
(1) × (2)	1, 5	15.94	.0009	.01	.01
(1) × (3)	2, 10	8.00	.0004	<.01	.01
(2) × (3)	2, 10	25.12	.0003	<.01	.01
(1) × (2) × (3)	2, 10	8.96	.0002	<.01	.00

Note. RFC = radial frequency component; 1D = condition in which attention to one dimension is sufficient to discriminate target; CONJ = condition in which attention to two dimensions is needed to discriminate target.
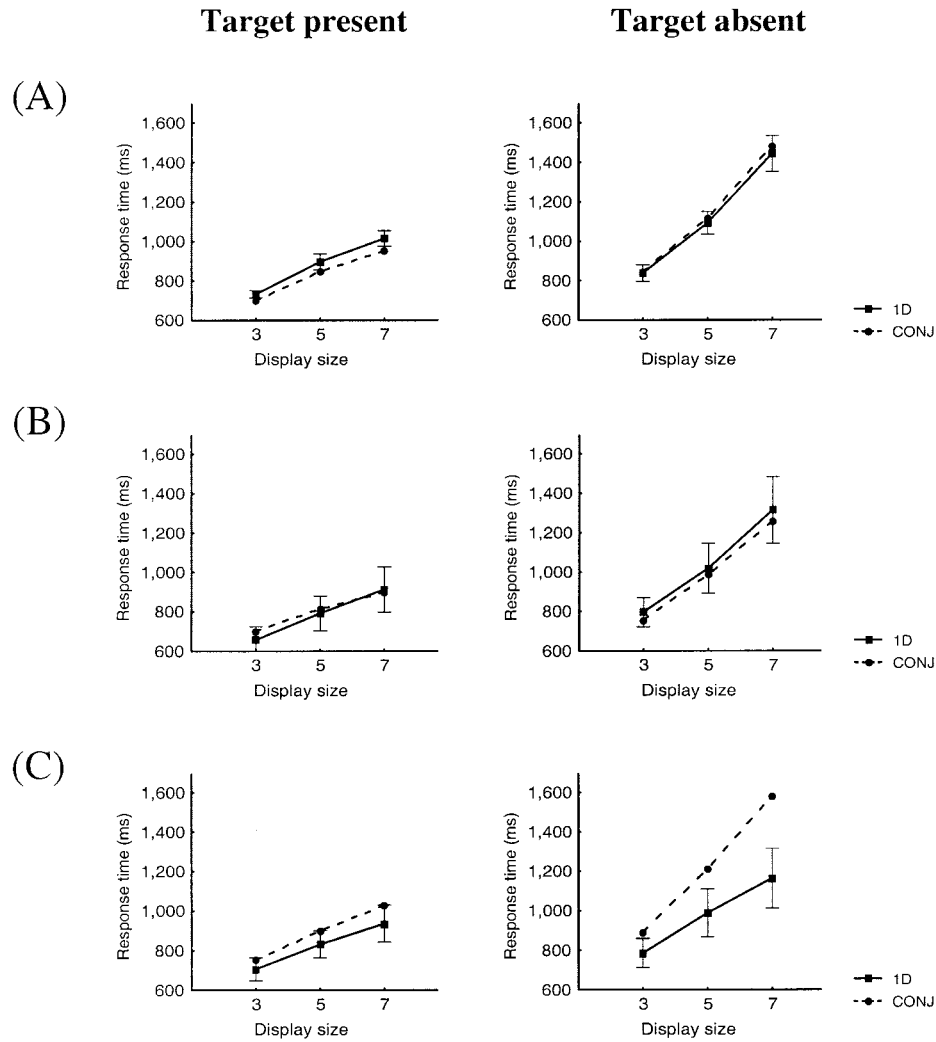
three-way interaction between condition, target presence, and display size.

The results from a visual search paradigm are often summarized by search slopes obtained from linear regression. These slopes indicate the extra time (in ms) it takes subjects to respond when display size is increased with one item. It is clear from Figure 3 that the effect of display size averaged across all subjects is linear, and we also found a good linear fit from the application of linear regression on the data of individual subjects for each combination of the variables stimulus set, condition, and target presence (mean R^2 = .95 and .99 for target-present and target-absent trials, respectively). The mean slopes for target-absent trials are shown in Figure 4 (open symbols). Statistical analyses for each stimulus set on search slopes by two-way ANOVAs with factors condition and target presence revealed similar effects as analyses on RTs—that is, only significant effects of condition with deformed ellipses. An overall three-way ANOVA including stimulus set as a third factor revealed a significant interaction between condition and stimulus set, $F(2, 10) = 7.95$, $MSE = 614$, $p < .01$, $\eta^2 = .12$, caused by the fact that the effect of condition was restricted to one stimulus set, deformed ellipses. Search was more efficient in condition 1D compared with condition CONJ for the deformed ellipses, but not for the two sets of RFC stimuli.

Experiment 1B

Experiment 1B was designed to verify that the effects obtained in Experiment 1A (especially, the effects of condition being restricted to deformed ellipses) were not due to artifactual differences across conditions and stimulus sets in the discriminability of the targets from individual distractors, or in the discriminability among distractors.

The effect of display size was small in the control task compared with the actual search task, with a mean slope of 12 in target-present trials ($R^2 = .69$) and 36 in target-absent trials ($R^2 = .86$). Figure 4 shows the slope in target-absent trials for T-D and D-D pairings as a function of condition (1D and CONJ) and stimulus set (consistent, but smaller effects were noted in target-present trials). There was no significant difference between conditions for T-D pairings in any of the stimulus sets, although the higher slope for condition CONJ compared with condition 1D for complex RFC stimuli, slope 50 versus 39 ms/item, respectively, was almost significant, $t(5) = 2.25$, $SEM = 4.6$, $p = .07$. For D-D pairings there was a trend toward higher slopes for condition 1D compared with condition CONJ in all stimulus sets, but this effect was significant only with complex RFC stimuli, slope 38 versus 21 ms/item, respectively, $t(5) = 2.62$, $SEM = 6.5$, $p < .05$.

It is clear that none of these differences in search slopes between conditions 1D and CONJ was comparable with the effect of condition with deformed ellipses in Experiment 1A (shown with open symbols in Figure 4). The largest effects in Experiment 1B were found for complex RFC stimuli in both T-D and D-D pairings. If pairwise stimulus discriminability as found in Experiment 1B would influence the results in Experiment 1A, then we would expect to find the largest influence with complex RFC stimuli. This is especially so because the effects in T-D and D-D pairwise discriminability for these stimuli predict the same result, namely, a steeper slope in condition CONJ. That is, the somewhat higher target–distractor discriminability and lower distractor heterogene-

Target present   **Target absent**

Figure 3. Response time in Experiment 1A for (A) complex radial frequency component (RFC) stimuli, (B) simple RFC stimuli, and (C) deformed ellipses, as a function of condition (1D vs. CONJ), target presence, and display size. Error bars indicate the standard error of the mean (representing the variability of each data point across subjects). 1D = condition in which attention to one stimulus dimension is sufficient to discriminate the target from the distractors; CONJ = condition in which two dimensions must be attended to in order to select the target.

ity in condition 1D would make search more efficient. This prediction was not confirmed. In Experiment 1A, search in condition 1D was not more efficient than in condition CONJ with RFC stimuli. As a consequence, we can exclude the possibility that the large difference between 1D and CONJ for deformed ellipses in Experiment 1A is related to differences in pairwise stimulus discriminability.

Discussion

The crucial result of Experiment 1A is the more efficient search in condition 1D compared with condition CONJ with ellipses varying in aspect ratio and curvature. This effect of condition was not found with the other two stimulus sets, in which stimuli varied on RFC dimensions. This pattern of results suggests that aspect ratio and curvature are separable/independent dimensions in the

visual system and that RFC dimensions are integral/dependent. The RFC stimuli do not appear to be encoded in terms of independent RFC dimensions, as it makes no difference whether two stimuli differ on two separate RFC dimensions or only on one (overall similarity being equal).

Experiment 1B enabled us to falsify alternative hypotheses in terms of target–distractor discriminability or distractor heterogeneity. We found small differences in target–distractor and distractor–distractor discriminability between conditions 1D and CONJ, but there was no indication that these differences were larger for deformed ellipses than for the other two stimulus sets. In addition, the results of Experiment 1A cannot be explained by stimulus complexity instead of dimensional separability, because the same pattern of results was obtained with the complex and simple RFC stimuli. Indeed, the latter stimulus set is similar to the
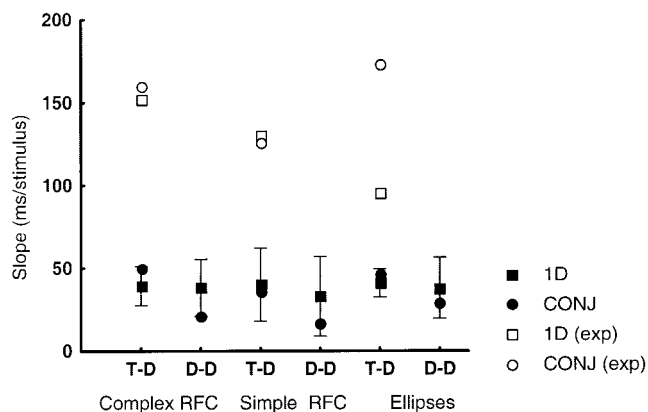
Figure 4. Search slopes in target-absent trials in Experiment 1B as a function of stimulus pair (target–distractor [T-D] or distractor–distractor [D-D]) and condition. Error bars indicate 95% confidence intervals based on within-subject variance. The search slopes in Experiment 1A (exp) are shown with open symbols. 1D = condition in which attention to one stimulus dimension is sufficient to discriminate the target from the distractors; CONJ = condition in which two dimensions must be attended to in order to select the target.

deformed ellipses with respect to stimulus complexity (e.g., the number of protrusions).

The results of Experiment 1, together with the data described in the Appendix, inform us about the way in which RFC dimensions are processed by naive observers. First, the perception of shape similarity is faithful, indicated by the fact that the stimulus configurations in psychological space fit the stimulus configurations in parameter space (both in dimensionality and stimulus order). Second, human observers do not characterize the differences between these shapes in terms of RFC dimensions that are processed independently. It makes no difference whether two shapes differ on two separate RFC dimensions or only on one (overall similarity being equal). In other words, the low-dimensional psychological space is faithful, but the orientation of stimulus configurations in this space is arbitrary. A different result was obtained with the dimensions aspect ratio and curvature. In the latter case, the orientation of a stimulus configuration in psychological space is not arbitrary. This was shown by the finding that searching a target is most efficient when only one of these dimensions is relevant to differentiate the target from all distractors. This result is comparable to the established finding that a search involving color *or* orientation (two classic separable dimensions) is more efficient than a search involving color *and* orientation (Treisman & Gelade, 1980).

The remaining experiments investigate whether a categorization phase during which one RFC dimension is relevant and another one irrelevant can prompt subjects to process these integral dimensions in a separable way—that is, whether categorization learning can induce dimension differentiation.

Experiment 2: Integral–Separable Dimensions and Category Learning

The design of Experiment 2 is related to the study of Goldstone (1994), but using RFC variations as integral dimensions and aspect

ratio/curvature as separable dimensions. Goldstone compared the discrimination performance of experimental participants who had learned one of several category rules with the performance of naive control participants. Goldstone found that both irrelevant and relevant dimensions tended to be discriminated better as a consequence of training with integral dimensions, whereas learning was more specific to the relevant dimension with separable dimensions. In addition, improvements in the discrimination task were always larger for the relevant dimension than for the irrelevant dimension, even with integral dimensions. It is the latter aspect of the data that Goldstone suggested reflects the process of *dimension differentiation*. So, the main question in our study is whether categorization training will always result in a higher sensitivity for a relevant compared with an irrelevant dimension, even with the integral RFC dimensions.

An additional finding in Goldstone (1994) was a main effect of training: Performance always improved, either for the relevant dimension only or for both the relevant and the irrelevant dimension. This main effect of practice was accompanied by a response bias of the control participants to respond as not seeing any difference between paired shapes. It is possible, however, that this learning effect reflects strategic knowledge that is not stimulus specific. Indeed, Doane, Alderton, Sohn, and Pellegrino (1996) showed that exposure to difficult shape discriminations leads to superior transfer performance with novel shapes. Applied to the category-learning design, it is possible that the category-learning phase suggested to the participants that they had to attend to quite small stimulus differences, regardless of the stimuli. This alternative hypothesis could be ruled out by using a control condition that is more stringent than the one with naive participants with no prior experience in the experimental context. So, the control participants in Experiment 2 were also involved in the category-training phase, but with different stimuli than those used later in the discrimination phase. Given this other operationalization of the control condition, it was an open question whether we would find a main effect of training causing better discrimination.

Method

Participants

One hundred eight students from the undergraduate program of psychology at the University of Leuven participated in this study for course credits. Each participant was randomly assigned to one of six groups of equal size ($n = 18$).

Stimuli

The experiment involved the complex RFC stimuli for three of the subject groups and the deformed ellipses for the other three groups. We constructed 16 (4×4) stimuli for each stimulus set (Figures 1A and 1C) for use in the test phase and more extended sets of 64 (8×8) intermediate stimuli for use in the training phase (see Figure 5 for a two-dimensional representation of the values of these 64 stimuli on the two parametric dimensions).

Materials

The shapes (maximum size $4° \times 4°$) were presented on a 17-in. monitor with a black background in a darkened room. The distance between the participants and the screen was about 100 cm.
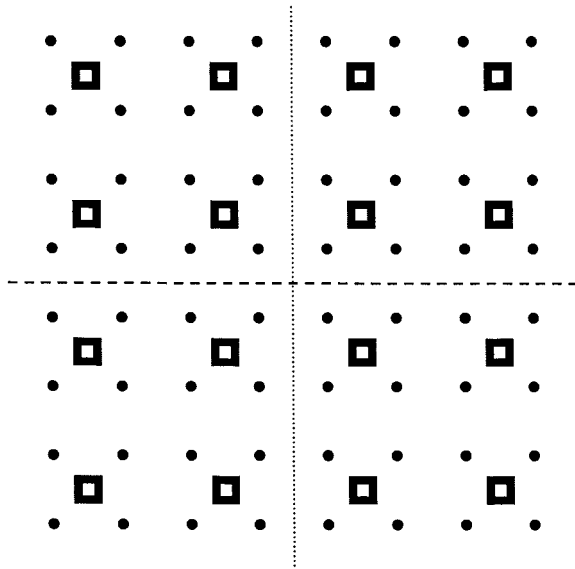
Figure 5. Relative position in parameter space of the stimuli in Experiment 2. The full circles represent the positions of the 64 stimuli used in the training phase, and the open squares represent the positions of the 16 stimuli used in the test phase. These 16 stimuli are shown in Figure 1 for each stimulus set. The dashed and dotted lines represent the category rule when the vertical and horizontal dimensions are relevant, respectively.

Procedure

The experiment consisted of a training phase and a test phase. During the training phase, participants had to categorize 64 similar shapes into two categories separated by a linear category rule (each category was associated with one of the two marked keys). The orientation of the category border determined the (ir)relevance of a particular dimension for the categorization task. We used the complex RFC stimuli for two participant groups. One group had to learn the horizontal category border drawn as a dashed line in Figure 5 (i.e., the vertical dimension was relevant), whereas the other group had to learn a vertical category border (drawn as a dotted line in Figure 5). Two other participant groups were trained with the deformed ellipses (with aspect ratio relevant for one group and curvature for the other). The remaining two groups of participants (control groups) had to learn a horizontal or vertical category border with stimuli from another set of complex RFC stimuli (described by seven RFCs with different parameter values).

The training phase consisted of 15 blocks of 64 trials and lasted for about 1 hr. All stimuli were presented once in a random order in each block. A fixation spot remained on the screen for an entire block. A trial began with a sound of 200 ms, followed by the presentation of a stimulus for 150 ms in one of four positions on the screen. The four positions were the corners of a square centered on the fixation spot, at an eccentricity of about 4°. As a consequence of this position randomization, subjects cannot predict the screen position at which a shape will be shown, and so selectively attending to spatial locations will not improve task performance. Stimulus presentation was followed by the presentation of a mask for 40 ms at the same position on the screen. The mask consisted of 36 (6 × 6) squares each containing a fragment of a shape constructed with the same algorithms as the stimuli used in this experiment. Finally, the computer waited another 1,650 ms to allow participants to make a response, after which feedback was given. If a correct response was made (i.e., the key was pressed that was associated with the category to which the stimulus belonged), the trial was ended by two high-frequency sounds. If an incorrect response was

made (either the wrong key was pressed or no response was made at all), participants received two low-frequency sounds. The intertrial interval was 500 ms.

The procedure of the training phase was explained to each participant. It was made clear that they would have to learn to categorize a set of similar stimuli into two groups without any prior knowledge (i.e., the only thing they could do in the first trial of the first block was guessing). They were told explicitly to look carefully at the shape of the stimuli and to ignore their position or the mask.

The test phase was completed on the day following the training phase and was preceded by three additional training blocks of 64 trials. Three groups (the two participant groups trained with the RFC stimuli of Figure 1A and one control group) discriminated the complex RFC stimuli during the test phase, and the three other groups (the two participant groups trained with the deformed ellipses and the other control group) performed the test phase with deformed ellipses. During the test phase, two stimuli were presented in one trial, and participants had to indicate whether their shape was the same or different using the same two keys as in the training phase (now labeled differently). Two variables were counterbalanced across participants: the key that was associated in the training phase with the category containing the lower amplitude values on the relevant dimension, and the key that had to be used to indicate "same."

The test phase consisted of eight blocks of 64 trials. One fourth of the trials were "same" trials, in which the same shape was presented twice. Only stimuli that are vertically or horizontally adjacent in parametric space were paired in "different" trials, so only the smallest differences were presented, resulting in 24 possible combinations of stimuli. Each of these combinations was presented for 16 trials, while each stimulus was presented in 8 "same" trials. The progression of the different intervals in each trial was as follows (the fixation spot was always present): warning sound (200 ms), first stimulus (150 ms) in the center of the screen, mask (40 ms), interstimulus interval (1,000 ms), and second stimulus (150 ms) at one of four eccentric positions (see training phase). The allowed maximal response latency was again 1,650 ms.

The procedure of the same–different task was explained to all participants. Participants completed 10 practice trials to allow them to get used to the task in general and the response–key association in particular. It was stressed that the probability of "same" trials could deviate from .5, but the direction of the deviation was left unspecified. Again, participants were told explicitly to look carefully at the shape of the stimuli and to ignore their position (which was always different for first and second stimulus) or the mask.

Results

Training Phase

All six groups of participants showed progressive learning in the training phase. Figure 6 shows the performance of the four experimental groups at different time points of the training phase as a function of the position of the stimuli along the relevant dimension. Three main observations can be made. First, performance got progressively better toward the last training blocks. It clearly took some time to learn the category rule, and even final performance was not perfect.

Second, performance was not distributed evenly across all stimuli, but it depended on the distance between a stimulus and the category border. In the end, performance was quite good for stimuli far away from this border, but it was worse for stimuli close to it.

Third, there was a strong difference in performance level between the experimental groups. Performance was superior for participants categorizing deformed ellipses. These participants per-
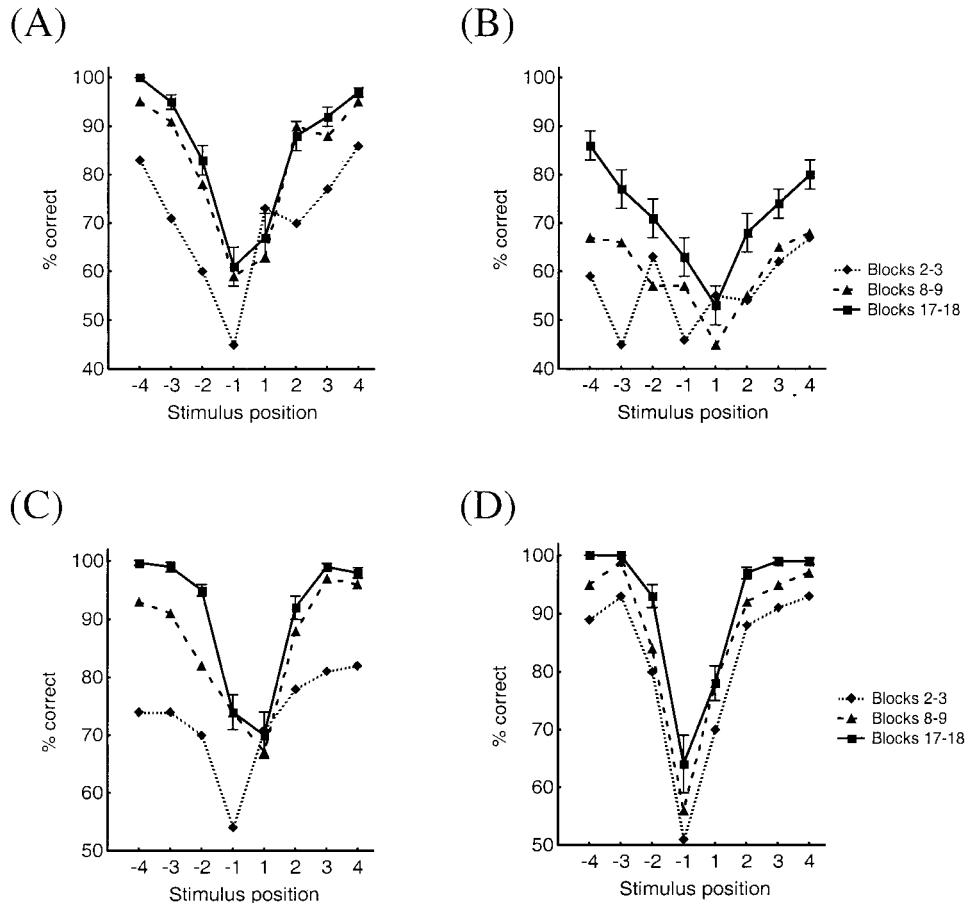
(A)

(B)

(C)

(D)

Figure 6. Performance in the training phase of Experiment 2. Top row: Performance with complex radial frequency component stimuli when the relevant dimension is the (A) vertical or (B) horizontal dimension in parametric space. Bottom row: Performance with deformed ellipses when the relevant dimension is (C) aspect ratio or (D) curvature. Stimulus position is expressed relative to the position of the category border on the relevant dimension, with positive numbers for positions with higher amplitude values. Error bars indicate the standard error of the mean.

formed at 91% correct (*SEM* = 1%) in the last two blocks of the training phase, compared with 76% correct (*SEM* = 2%) for participants categorizing the complex RFC stimuli. Additional differences were present between participants learning different category rules with the same stimuli. With complex RFC stimuli, participants learning to categorize along the vertical dimension (Figure 6A) showed a relatively good final performance level. They performed at 83% correct (*SEM* = 2%) in the last two blocks of the training phase, compared with 68% correct (*SEM* = 3%) for participants categorizing along the horizontal dimension (Figure 6B). With deformed ellipses (Figure 6C and 6D), the two groups performed at the same level in the last two blocks of the training phase (91% correct, *SEM* = 1%), although performance in the first blocks was better with curvature as the relevant dimension (82% correct, *SEM* = 2%, compared with 73%, *SEM* = 3%). The two control groups did not differ significantly from each other (70% and 72%, *SEM* = 3%). As participants were assigned randomly to the six groups, the differences between groups were probably related to a differential perceptual salience of the relevant stimulus differences.

Test Phase

To analyze the data from the same–different task, we converted the hit and false alarm rates for each stimulus combination into the bias-free sensitivity measure d' (Macmillan & Creelman, 1991). There was a large variability between participants in their bias to respond "different," with a mean percentage of "different" answers of only 47%. Only 6 participants had a bias close to the actual proportion of "different" trials in the test session (responding "different" in at least 70% of the trials); 10 participants even had a robust bias in the other direction (responding "different" in less than 30% of the trials). Most important, the bias did not differ significantly between the groups, $F(5, 102) = 1.15$, $p > .30$.

We first performed separate analyses for the two stimulus sets. For each set, we analyzed the d' values in an ANOVA with three variables: group, a between-subjects variable with three levels referring to the three groups of participants; dimension, a within-subject variable referring to the dimension along which the stimuli in a trial varied (either the vertical or the horizontal dimension in Figure 5); and position, a within-subject variable referring to the

position of the stimulus difference along a dimension (between the smallest, middle, or largest amplitudes).

The results of the ANOVA for complex RFC stimuli are summarized in Table 2. There was no significant main effect of subject group. Participants in the experimental groups tended to have slightly higher d' values, but even an a priori contrast based on the prediction that both experimental groups (mean $d' = 1.4$) would perform better than the control group (mean $d' = 1.2$) did not reach significance, $F(1, 51) = 2.84$, $p = .10$. There was a strong effect of dimension, with a higher sensitivity for the vertical dimension (mean $d' = 1.5$) compared with the horizontal dimension (mean $d' = 1.1$). This effect of dimension interacted significantly with position (largest difference between both dimensions for the middle amplitudes). No other effects reached significance. Most important, there was no interaction between the group variable and the dimension, $F(2, 51) < 1$ (see Figure 7A). This indicates that the relevance of a dimension for categorization did not affect the sensitivity for differences along that dimension.

The results of the ANOVA for deformed ellipses are summarized in Table 2B. There was no main effect of subject group. The performance level of the experimental and control groups was similar (mean $d' = 2.5$ for the experimental compared with 2.4 for the control group). Participants were better in discriminating curvature (mean $d' = 2.8$) relative to aspect ratio (mean $d' = 2.1$). This effect of dimension interacted significantly with position (largest difference between both dimensions for the lower dimension values). Most important and in contrast to the results for the complex RFC stimuli, there was a significant interaction between the group variable and the dimension (see Figure 7B). Although all subject groups performed better with differences in curvature relative to aspect ratio, this difference became smaller when aspect ratio had been relevant during categorization and larger when curvature had been relevant.

We tested whether the difference in the interaction between group and dimension between stimulus sets was significant by
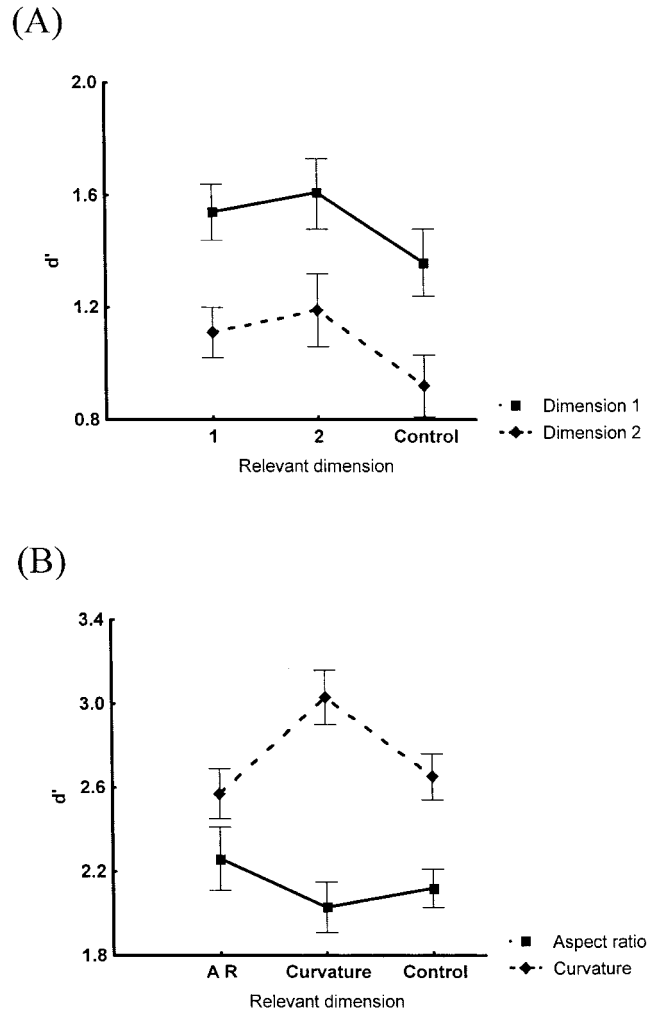

Figure 7. Sensitivity in Experiment 2 for (A) complex radial frequency component stimuli and (B) deformed ellipses as a function of participant group and dimension. Error bars indicate the standard error of the mean. AR = aspect ratio.

performing a four-way ANOVA on all data including the same three variables as before with an additional between-subjects "stimulus" variable. This analysis revealed a significant higher order interaction between the variables stimulus, group, and dimension, $F(2, 102) = 5.54$, $MSE = .28$, $p < .01$. This confirms that the absence of an interaction between group and dimension for complex RFC stimuli is a result that differs significantly from the clear interaction between group and dimension for deformed ellipses.

Discussion

The results during the categorization phase are generally consistent with existing models of categorization. These models describe categorization as a process of linking the representation of a stimulus to the probability of this stimulus belonging to particular categories. The closer a stimulus is located to a category border (i.e., the more similar it is to exemplars from another

Table 2
Results of Analysis of Variance on d′ Values in Experiment 2

Effect	df	F	SEM	p	η^2
		Complex RFC stimuli			
Subject group (1)	2, 51	1.54	1.26	.22	
Dimension (2)	1, 51	119.03	0.12	<.01	.11
Position (3)	2, 102	2.87	0.18	.06	
(1) × (2)	2, 51	0.03	0.12	.97	
(1) × (3)	4, 102	0.71	0.18	.59	
(2) × (3)	2, 102	28.42	0.19	<.01	.08
(1) × (2) × (3)	4, 102	0.20	0.19	.94	
		Deformed ellipses			
Subject group (1)	2, 51	0.60	1.14	.55	
Dimension (2)	1, 51	71.04	0.44	<.01	.08
Position (3)	2, 102	270.09	0.26	<.01	.37
(1) × (2)	2, 51	7.57	0.44	<.01	.02
(1) × (3)	4, 102	2.07	0.26	.09	
(2) × (3)	2, 102	130.26	0.23	<.01	.16
(1) × (2) × (3)	4, 102	0.63	0.23	.64	

Note. RFC = radial frequency component.

category), the more the latter probability will deviate from 1 (certainty). This property of the categorization process is clearly present in our results: Performance was modulated by the distance between stimulus and category border. In addition, the difficulty of category learning along a dimension was clearly related to the a priori discriminability of the relevant dimension. For control participants tested with complex RFC stimuli, we found that differences along the vertical dimension were easier to discriminate compared with differences along the horizontal dimension (indicating, not surprisingly, that calibrations in pixel space do not necessarily correlate with perceived similarities). The higher sensitivity for the vertical dimension agrees with the finding that category learning proceeded fastest when the vertical dimension was relevant. Control participants tested with deformed ellipses showed superior performance in the discrimination of differences in curvature compared with differences in aspect ratio. In line with this result, we initially found better performance in the category-learning phase when curvature was the relevant dimension. In sum, the data of the training phase reveal a strong effect of shape similarity on the difficulty of learning specific category rules.

The main subject of the present study was the reverse, namely, the effect of categorization on similarity. Attentional learning theories (e.g., Goldstone, 1994; Hall, 1991; Nosofsky, 1986) predict that the learning of specific stimulus–category linkages could bias the processing of stimulus differences that are relevant for differentiating members of different categories. We predicted that such a biasing mechanism would work only in the case of dimensions that can be processed independently. In Experiment 2, a bias would occur if the status of a dimension as relevant or irrelevant would affect discrimination performance. For integral RFC dimensions, we did not observe such modulation. In contrast, it was clearly present for curvature and aspect ratio. So, the separable processing of dimensions as operationalized in Experiment 1 predicts the effect of dimensional relevance during category learning on discrimination performance. The absence of an effect of dimensional relevance with integral dimensions suggests that categorization training did not induce dimension differentiation in this experiment.

When comparing our results with those described by Goldstone (1994), two differences are worth noting. First, the effect of the integral–separable distinction is larger than in Goldstone's study. Dimensional relevance had no effect at all with integral RFC dimensions, whereas Goldstone still found some selective learning with saturation and brightness (although smaller than with size and brightness). Possibly, this difference is due to the use of other stimulus dimensions (see General Discussion). Alternatively, differences between our and Goldstone's procedures might explain the different results. Goldstone allowed his participants to look for a long time at each stimulus in each trial, whereas we restricted the exposure duration to 150 ms. Foard and Kemler Nelson (1984) suggested that separable processing of dimensions is more likely to occur with longer exposure durations. Longer stimulus durations allow participants to perform additional cognitive explorations (e.g., covert and overt shifts of attention) beyond the initial, more bottom-up processing of stimuli, which could help in differentiating stimulus dimensions.

Second, the control group in the Goldstone (1994) study performed markedly worse compared with the experimental group. In contrast, we found no or only a minor tendency toward better performance by participants in the experimental groups compared with controls. The main difference between the two studies is the procedure followed for the control participants. Our control group received the same amount of experience in category learning as our experimental participants. They learned to categorize other stimuli than the ones that were used in the test phase. Apparently, this more stringent control was sufficient to erase the general learning effects shown by Goldstone. This suggests that a substantial part of the behavioral improvements found previously were due to strategic knowledge that is not stimulus specific (Doane et al., 1996).

Experiment 3: Integral–Separable Dimensions and Long Category Learning

Experiment 2 confirmed the importance of the integral–separable distinction for predicting the effect of category learning on shape similarity. These results suggest that shapes are encoded in terms of a set of elementary independent dimensions. However, there are several reasons why we need further evidence for substantiating this claim. First, there is a rapidly growing literature providing suggestive evidence that several aspects of shape representation are quite malleable (e.g., Schyns et al., 1998; Schyns & Murphy, 1994; Sigala & Logothetis, 2002). Concerning the effect of categorization, Goldstone and Steyvers (2001) suggested that category learning could possibly result in the differentiation of dimensions. In an attempt to establish even more persuasively that the basic dimensions used to represent shapes do not change during categorization learning, we included a prolonged learning phase. We reasoned that the longer the duration of a category-learning phase, the more pronounced the resulting differentiation could be. If an extended period of categorization would indeed induce dimension differentiation, then integral dimensions would become separable whereas separable dimensions would not change their status. So, we would expect that the effect of categorization on shape similarity would always follow the predictions for separable dimensions. At least, the difference between originally integral and separable dimensions should decrease substantially.

Furthermore, we need to falsify two possible alternatives for the explanation in terms of the integral–separable distinction in Experiment 2. First, there was a difference in performance between complex RFC stimuli and deformed ellipses. The category rules were learned fastest with deformed ellipses. The low performance in the training phase with complex RFC stimuli could explain at least partially the absence of an effect of dimensional relevance. In the extreme case, if participants would not be able to learn a category rule at all (performance at chance during the learning phase), then they would have no clue regarding the (ir)relevance of the dimensions. Although participants performed above chance in all conditions, it is still possible that the absence of any effect of dimensional relevance with RFC stimuli is partly caused by lower performance in the training phase. Including a longer training phase offers a direct way to solve this problem. Even those participants with the disadvantage of being confronted with the most difficult category rule would finally end up with reasonable performance. Given a sufficiently long training period, all participants would have learned which stimulus differences are (ir)relevant in the final training sessions. To avoid this alternative explanation for differential effects with our different types of

dimensions, we also reduced the differences between the deformed ellipses to half of their size in Experiment 3.

A second alternative explanation for the difference between RFC dimensions and aspect ratio and curvature involves differences in stimulus complexity. Not only do the complex RFC dimensions differ from the deformed ellipses in the dimensions that are varied, but in addition, they are far more complex stimuli (this could also be the cause of the worse performance in the training phase). We controlled for stimulus complexity by including both the simple and the complex RFC stimuli (see Figure 1A and 1B) in Experiment 3. If the specificity of learning effects in the test phase is as small for simple RFC stimuli as it is for complex stimuli, then stimulus complexity is highly unlikely to be the explanation for the difference between RFC dimensions and aspect ratio and curvature.

Finally, to make the absence of any effect of dimensional relevance with integral dimensions even more compelling, we included an additional manipulation to increase the odds of getting an effect. Until now, all stimulus differences were intermixed in the training trials. At each trial, a stimulus could be different from previous stimuli on both the relevant and the irrelevant dimension. We reasoned that this would make it difficult for participants to attend to one dimension and ignore the other one. So, we changed this aspect of the training phase for the complex RFC stimuli by including a blocked stimulus presentation: Within each block of training trials, stimuli differed only on the relevant dimension; their value on the irrelevant dimension was altered only between blocks. For all other stimulus conditions, we included the same training procedure (but now five times as long) as used in Experiment 2.

Method

Participants

Thirty volunteers participated in this study. They were students from diverse undergraduate and graduate programs at the University of Leuven. They were not familiar with the topic of research, and they were paid for their participation.

Stimuli

For complex RFC stimuli, stimuli and their presentation were the same as in Experiment 2. Experiment 3 also included the simple RFC stimuli shown in Figure 1B. The third stimulus set was based on the deformed ellipses used before, but in Experiment 3 we restricted the stimulus space to the middle values of aspect ratio and the higher values of curvature (only two thirds of the interval used before).

Procedure

For 18 participants (experimental groups), the experiment consisted of seven sessions comprising five training sessions (TR) and two test sessions (TEST) in the following order: TEST1–TR1–TR2–TR3–TR4–TR5–TEST2. The experiment included three experimental groups with different stimulus sets (complex RFC, simple RFC, and deformed ellipses). The time interval between first and second test session was 13 days on average. Two control groups with 6 participants each performed only the test sessions (with a time interval of 13 days on average), with either complex RFC stimuli or deformed ellipses.

For each session of the training phase with complex RFC stimuli, the main difference with the procedure in Experiment 2 was a blocked trial order. Within each block (64 trials), only stimuli with one particular value on the irrelevant dimension were presented (eight repetitions of eight stimuli). Each value of the irrelevant dimension was used in two blocks, resulting in a total number of 16 blocks. One particular dimension was relevant for 3 experimental participants. Each session of the training phase with simple RFC stimuli and deformed ellipses followed exactly the same procedure as used in Experiment 2 (no blocked trial order). At the beginning of the first training session, all participants received the same instructions as in Experiment 2.

The test phase was the same for all participants, except that other stimuli were presented to different groups. Task and trial procedure were not changed relative to Experiment 2. However, this phase included only 12 of the 16 shapes used in Experiment 2. In a given trial, we only paired different stimuli that were horizontally or vertically adjacent in Figure 5 *and* that crossed the two possible category borders (one border for each dimension). In this way we reduced the number of possible stimulus combinations from 24 to 8, making the d' computation more reliable for each of them. Note that Experiment 2 showed that the effect of categorization on a particular dimension did not depend on the particular value on this dimension. Each of the combinations was presented for 54 trials in each session, and each stimulus was presented in 12 "same" trials (a total of 576 trials divided into six blocks).

The procedure of the same–different task was explained to each participant in the same way as in Experiment 2. We did this before each test session, and participants always went through 10 tryout trials. It was stressed that the probability of "same" trials could deviate from .5.

Results

Training Phase

On the first training session, participants showed the same pattern of results as noticed previously in Experiment 2, albeit performance was generally better. Figure 8 shows the evolution in performance across the five sessions as a function of the distance between a stimulus and the category border that had to be learned (averaged across all experimental groups). No clear difference in performance was noted between subject groups (88%, 88%, and 87% correct overall with complex and simple RFC stimuli and deformed ellipses, respectively). Performance was worst for stim-
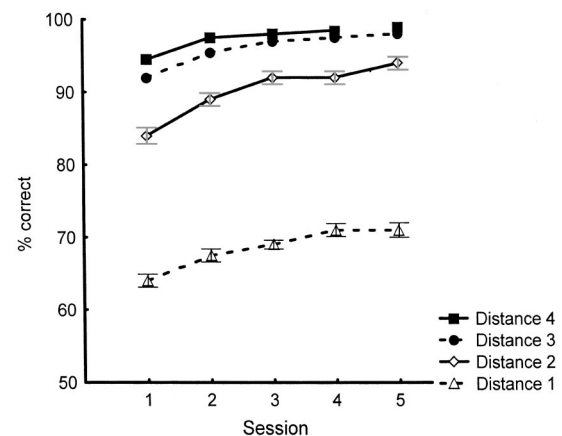

Figure 8. Performance in the different sessions of the training phase in Experiment 3 as a function of the distance between stimuli and the category border (e.g., Distance 1 is the grouping of positions −1 and 1 on the *x*-axis of Figure 6). Error bars indicate the standard error of the mean.

uli closest to the category border. There was a gradual improvement from the first to the last session at all distances.

Test Phase

For each participant, we computed the mean d' value for each dimension and each session separately. Figure 9 shows the difference between the two sessions for relevant and irrelevant dimensions for the three groups of experimental participants, and averaged across dimensions and stimulus sets for control participants (the mean d' in the first test session was 1.4). Performance was generally better in the second compared with the first session, an effect that was relatively small for control participants and for the irrelevant dimension with ellipses.

To assess the significance of the difference between sessions, we analyzed the data for each group of experimental participants with an ANOVA including two within-subject variables (pre- vs. posttraining and relevant vs. irrelevant dimension). The performance of experimental participants improved in the second session, as indicated by a main effect of training with complex RFC stimuli, $F(1, 5) = 23.54$, $MSE = .17$, $p < .01$, $\eta^2 = .25$; with simple RFC stimuli, $F(1, 5) = 13.16$, $MSE = .36$, $p = .02$, $\eta^2 = .40$; and with deformed ellipses, $F(1, 5) = 22.17$, $MSE = .07$, $p < .01$, $\eta^2 = .18$. There was no significant interaction between training and dimensional relevance for RFC stimuli for complex RFC stimuli, $F(1, 5) = 1.45$, $MSE = .04$, $p > .25$, or for simple RFC stimuli ($F < 1$). In contrast, there was a significant interaction between training and relevance for deformed ellipses, $F(1, 5) = 57.01$, $MSE = .02$, $p < .01$, $\eta^2 = .12$.

Control participants showed a small improvement in the second test session relative to the first, but this effect was not significantly different from zero for RFC stimuli, $F(1, 5) = 1.65$, $MSE = .10$, $p > .25$, or for deformed ellipses ($F < 1$).
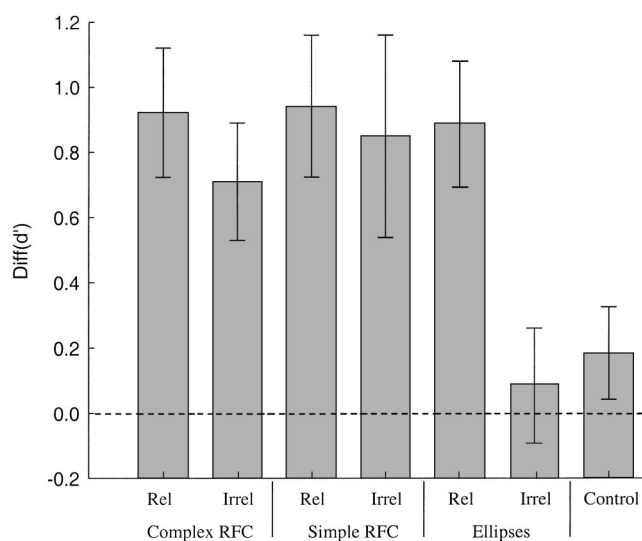
Until now, statistical analyses have been performed separately for the different stimulus sets. We tested whether the interaction of the training and relevance factors differed for the RFC and aspect ratio/curvature dimensions by performing a three-way ANOVA including the between-subjects variable stimuli (RFC stimuli vs. deformed ellipses) in addition to two within-subject variables (training and relevance). This analysis revealed a significant higher order interaction of the variables stimuli, training, and relevance for complex RFC stimuli, $F(1, 10) = 8.50$, $MSE = .03$, $p = .02$, $\eta^2 = .01$, and for simple RFC stimuli, $F(1, 10) = 7.71$, $MSE = .05$, $p = .02$, $\eta^2 = .01$. This confirms that the absence of an interaction between training and relevance for RFC stimuli is a result that differs significantly from the clear interaction between training and relevance for deformed ellipses.

Finally, we checked whether an analysis with more subjects would reveal a significant difference between relevant and irrelevant stimulus differences for RFC dimensions. Such an effect is suggested by the fact that the improvements on the irrelevant dimension were slightly smaller than those on the relevant dimension in all conditions with mainly general learning (simple and complex RFC stimuli). We analyzed all these data in one ANOVA with three factors (stimuli, training, and relevance), but this analysis still revealed no significant interaction between training and relevance, $F(1, 15) = 1.09$, $MSE = .06$, $p > .30$.

Discussion

The long training phase in Experiment 3 improved the sensitivity of participants to differences between the shapes during the test phase. This improvement generalized to an irrelevant dimension for RFC dimensions, but it was specific for the relevant dimension in the case of aspect ratio and curvature. Even a long-duration categorization-learning phase revealed discrimination improvements to be modulated strongly by the encoding of dimensions as either integral or separable. So, our results falsify the hypothesis that category learning is able to transform originally integral dimensions into separable dimensions.

In addition, the design of Experiment 3 helps us to unconfound the integral–separable distinction from several alternative variables. First, there was no longer a substantial difference in categorization performance between the integral and separable dimensions, making it implausible as an explanation for the effect of the integral–separable distinction in each experiment. Participants knew which stimulus differences were relevant during categorization and which differences were not, even in the first session (more than 80% of responses were correct).

Second, we included the simple RFC stimuli to investigate the effect of stimulus complexity. If the higher complexity of the complex RFC stimuli was the reason no effects of dimensional relevance were found with these stimuli, then one would expect specific learning with simple RFC stimuli. This prediction was falsified by our data.

Experiment 3 revealed improvements in experimental compared with control participants as a consequence of the training phase (as found by Goldstone, 1994), whereas Experiment 2 did not. Several explanations can be offered for this difference. First, the training phase was almost five times as long. Categorization during these sessions gradually improved from day to day, at least for stimuli close to the category border. It is probable that its effect on

Figure 9. Changes in sensitivity, Diff (d'), in the second relative to the first test phase in Experiment 3 for the relevant (Rel) and irrelevant (Irrel) dimension (experimental groups of participants) and for both dimensions in the case of the control group. Error bars indicate the standard error of the mean. RFC = radial frequency component.

discrimination also proceeds gradually, being larger after five sessions than after one.

Second, control participants in Experiment 3 were balanced only with respect to their initial exposure to the stimuli during the first test phase and did not perform any categorization training (in contrast to the control group in Experiment 2). As noted previously, it is possible that experimental participants acquire skills or get some strategic knowledge during the training phase that are not stimulus specific and that can be useful in the second test phase. However, this explanation for the improvements confined to experimental participants is less compelling than it would have been for control participants without any training in Experiment 2. Most important, it would have to be skills and knowledge that are acquired in a rather slow manner, as all participants performed the second test phase after the first session, during which they became familiar with all aspects of the test task. So, the improvements in experimental participants were found relative to control participants who also had some initial experience with the stimuli and as much experience with the test task. Whatever the explanation for these improvements may be, their existence suggests that category learning has gradual effects on discrimination performance that are not tied to the first encounter with stimuli.

General Discussion

Experiment 1 established that shape dimensions such as aspect ratio and curvature are elementary dimensions in the encoding of visual shape, whereas RFC dimensions are not. Experiments 2 and 3 addressed the question of whether the processing of the integral RFC dimensions could be changed by category learning. Experiment 2 focused on the effect of 1 hr of category learning on the discrimination of RFC dimensions and the dimensions of aspect ratio and curvature. Participants did not show relatively better discrimination performance on relevant compared with irrelevant dimensions in the case of RFC dimensions. For aspect ratio and curvature, however, learning that only one dimension is relevant during categorization biases the processing of that dimension relative to an irrelevant dimension. The same basic dissociation between RFC dimensions and aspect ratio/curvature was found in Experiment 3 after 5 hr of category learning. We conclude that even long training phases do not allow the integral RFC dimensions to become processed in a separable way. Thus, our results indicate a rather limited effect of category learning on shape perception. Category learning is able to improve the processing of one out of a set of separable dimensions. With integral dimensions, however, all dimensions are affected by category learning in the same manner, whether they were relevant or irrelevant during category learning. The effect of dimension differentiation was very small, if present, even if the training phase included five sessions of 1 hr.

Differences From Previous Studies on Dimension Differentiation

Why have previous studies by Goldstone and his colleagues found substantial evidence for dimension differentiation, whereas we have not? The first study (Goldstone, 1994) mainly focused on the effect of category learning in biasing the processing of relevant

stimulus dimensions. The study included both separable and integral dimensions. Integral dimensions did reveal less biasing compared with separable dimensions, but the difference between the two types of dimensions was much less clear-cut compared with our results. Even with the traditional integral dimensions of brightness and saturation, a substantial part of the learning effect was confined to the relevant dimension (after only 1 hr of category learning). If one assumes that brightness and saturation are indeed a good example of integral dimensions, then this result suggests that these dimensions have been partially differentiated during category learning.

However, as noted by Goldstone and Steyvers (2001), there is evidence that brightness and saturation are not integral in an absolute sense (Melara et al., 1993). Even these single color dimensions are psychologically privileged compared with their combinations. Apparently, there is no other way to process brightness and saturation differences than to look at brightness and saturation (even if it is hard). This result provides a clue to understand why we found much more compelling effects for integral dimensions within the domain of shape perception. Several studies and models of shape perception have suggested that shapes are encoded using a restricted set of elementary dimensions that include aspect ratio and curvature but not RFC variations (e.g., Arguin & Saumier, 2000; Biederman, 1987; Stankiewicz, 2002; Zusne, 1970). (Note that the separable–integral distinction is likely to be a reflection of an underlying continuum; likewise, the distinction between elementary and other shape dimensions is probably a relative distinction that can be made clearly only for some dimensions, such as the ones we used.) Of course, whatever the kind of dimension one is manipulating, the shape variations will always yield differences on one or several of the elementary dimensions. As a consequence, it is not surprising that human perceivers encode all of these shapes using these elementary dimensions, even if the parametric dimensions do not match them. Arguin and Saumier (2000) already stated that evidence for dependent, integral dimensions is found when the dimensions along which stimuli are manipulated are poorly correlated with the actual dimensions used by the visual system during encoding. As suggested in the experiments presented here, evidence for selective attention effects due to category learning is found only if the parametric dimensions correspond to these elementary dimensions.

A later study by Goldstone and Steyvers (2001) provided more compelling evidence for dimension differentiation. As integral dimensions, they used face-morphing dimensions in most experiments and saturation and brightness in another. They suggested that a relatively short category-learning phase can turn integral dimensions into more separable dimensions as measured by transfer category learning. So, why did a much longer training phase fail to result in dimension differentiation in our study? Again, a possible route for explaining this discrepancy would be to point to the different stimuli and dimensions being used. We used stimuli (shapes) that are preferentially encoded in terms of some elementary dimensions, whereas this type of encoding has not been proposed for faces or brightness/saturation. For example, Biederman and Kalocsai (1997) suggested a completely different type of representation for faces. In addition, many studies have indicated

more "holistic" processing of faces compared with other objects (Maurer, LeGrand, & Mondloch, 2002).

An alternative hypothesis for the apparent discrepancy between our results and those obtained by Goldstone and Steyvers (2001) points to the different experimental designs as the critical factor. We looked for dimension differentiation by studying the transfer of learning effects from a category-learning phase to another task (same–different judgment). Goldstone and Steyvers (2001) investigated the transfer between different category-learning tasks. From this viewpoint, the conflicting results suggest that dimension differentiation does generalize among category-learning tasks but not to different tasks such as a same–different judgment. What could be special about shape processing during category learning? A suggestion came from the introspective reports of participants after the training and the test phase, when we questioned them about the strategy they followed to solve the tasks. With complex RFC dimensions, several participants said they interpreted the stimuli in a complex way during category learning. They reported using category rules such as "looks like an angel or not," "resembles a rabbit or not," and so forth. No one used similar words to verbalize the stimuli during the same–different judgment; instead, everyone reported using cues such as global shape, relative orientation of shape parts, and so on. So, insofar as one can trust verbal reports, participants seemed to process the stimuli in a more "interpretative" way during the category task.

The Effect of Category Learning: Encoding Versus Interpretation

A relevant study reporting differences between shape perception and interpretation was published decades ago (Shepard & Cermak, 1973). These authors also manipulated RFC-like dimensions. If participants were asked to group shapes according to their similarity, then a faithful representation of the shape dimensions was found: Stimuli that were close in parametric space were perceived as being very similar. Interestingly, the authors reported another experiment asking participants to group shapes that could be interpreted as depicting the same real-life objects. Participants had no problem solving this Rorschach-like task, but the data revealed surprising results: Shapes that were proximal in parametric space were regularly interpreted as representing different objects (and vice versa). So, the task people are engaged in has a tremendous effect on how stimulus similarity is actually judged. The apparent contradiction between our results and the study of Goldstone and Steyvers (2001) could be solved by assuming that the interpretation of shapes during category learning helps human perceivers to differentiate dimensions, but this differentiation is not found in a task that does not prompt perceivers to interpret stimuli in the same way (same–different judgments). From this viewpoint, dimension differentiation is linked to "postperceptual" stimulus interpretation and not to stimulus encoding.

This distinction between two levels is inherent in several models of categorization and recognition (e.g., Edelman, 1999; Kruschke, 1992; Riesenhuber & Poggio, 1999). The representation of complex objects in all of these models is based on units that implement basis functions that compute the distance between a given input and the optimal input for each unit. Op de Beeck et al. (2001) confirmed this assumption by showing that neurons in inferotemporal cortex are activated according to the similarity between a

stimulus and the most preferred stimulus of each neuron. At the next level, the output of units from the first level is selectively combined to construct units that differentiate maximally between stimuli from different categories.

In this framework, selective attention and dimension differentiation operate at different levels. Convergent evidence from neurophysiological studies suggests that selective attention directly influences the representation of stimulus dimensions whereas processes such as dimension differentiation are deferred to a later level. Indeed, many neurophysiological studies have revealed effects of selective attention in visual cortex, in both monkeys and humans (for reviews, see Kanwisher & Wojciulik, 2000; Treue, 2001). Generally, selective attention enhances (or depresses) the responses of neurons that are tuned for the attended (or inattended) stimulus features and dimensions.

Dimension differentiation, however, requires a more drastic change of neuronal response properties. At least for simple stimulus dimensions, separability found at the behavioral level (e.g., orientation and spatial frequency; Regan, 2000; Vogels et al., 1988) might be reflected by a tuning of single neurons for one dimension that is relatively invariant across changes in another separable dimension (see introduction). A possible neuronal correlate of dimension differentiation would be that a dependent coding of dimensions by single neurons (as found for RFC dimensions; Albright & Gross, 1990) could be changed into independent tuning properties. There is no evidence for such a dramatic effect of relatively short (or even long) training procedures on tuning properties of visual neurons. Our proposal that dimension differentiation is linked to stimulus interpretation suggests that dimension differentiation is implemented in areas such as prefrontal cortex and basal ganglia that reorganize the information from high-level visual areas such as inferotemporal cortex (Ashby & Ell, 2001; Freedman, Riesenhuber, Poggio, & Miller, 2001; Miller & Cohen, 2001; Thorpe & Fabre-Thorpe, 2001). According to this hypothesis, the visual system itself analyzes shapes by evaluating their value on a fairly fixed set of independent dimensions. The differentiation of dimensions as cued by the context (e.g., an experiment in which some dependent dimensions have a different relevance) does not change the status of dimensions in the visual system but only in later postperceptual processing stages.

This framework provides an explanation for the apparent discrepancy between the findings of two recent studies investigating the effect of categorization learning on the tuning properties of single neurons in macaque inferotemporal cortex. One study (Op de Beeck et al., 2001), using integral RFC dimensions, found no effect of categorization on the selectivity of single neurons, whereas the other study (Sigala & Logothetis, 2002), using separable face dimensions (Nosofsky, 1986), suggested an increased neuronal sensitivity specific for relevant dimensions. These results mirror our behavioral findings in that specific learning effects are found only with separable dimensions.

Conclusion

The present investigation allows us to draw the following conclusions. First, category learning can bias the processing of separable shape dimensions in the visual system. If two shape dimensions are processed independently, then the processing of one dimension can be relatively enhanced if it was relevant during

categorization. Second, category learning cannot bias the processing of integral shape dimensions in the visual system. Even if two integral shape dimensions differ in their relevance during an extended period of categorization, they remain processed in a dependent way. Third, a comparison of our results with previous behavioral and neurophysiological findings suggests that dimension differentiation could be associated with shape interpretation, as such involving those parts of the brain that are able to reorganize input from the visual system in a flexible way.

References

Albright, T. D., & Gross, C. G. (1990). Do inferior temporal cortex neurons encode shape by acting as Fourier descriptor filters? In *Proceedings of the International Conference on Fuzzy Logic & Neural Networks* (pp. 375–378). Izuka, Japan.

Arguin, M., & Saumier, D. (2000). Conjunction and linear non-separability effects in shape encoding. *Vision Research, 40,* 3099–3115.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review, 105,* 442–481.

Ashby, F. G., & Ell, S. W. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences, 5,* 204–210.

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

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review, 93,* 154–179.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115–147.

Biederman, I., & Kalocsai, P. (1997). Neurocomputational bases of object and face recognition. *Philosophical Transactions of the Royal Society of London, 352B,* 1203–1219.

Borg, I., & Leutner, D. (1985). Measuring the similarity of MDS configurations. *Multivariate Behavioral Research, 20,* 325–334.

Cortese, J. M., & Dyre, B. P. (1996). Perceptual similarity of shapes generated from Fourier descriptors. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 133–143.

Cutzu, F., & Edelman, S. (1998). Representation of object similarity in human vision: Psychophysics and a computational model. *Vision Research, 38,* 2229–2257.

Doane, S. M., Alderton, D. L., Sohn, Y. W., & Pellegrino, J. W. (1996). Acquisition and transfer of skilled performance: Are visual discrimination skills stimulus specific? *Journal of Experimental Psychology: Human Perception and Performance, 22,* 1218–1248.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433–458.

Edelman, S. (1999). *Representation and recognition in vision.* Cambridge, MA: MIT Press.

Edelman, S., & Intrator, N. (2000). Coarse coding of shape fragments plus retinotopy approximate representation of structure. *Spatial Vision, 13,* 225–264.

Foard, C. F., & Kemler Nelson, D. G. (1984). Holistic and analytic modes of processing: The multiple determinants of perceptual analysis. *Journal of Experimental Psychology: General, 113,* 94–111.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2001, January 12). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science, 291,* 312–316.

Garner, W. R. (1976). Interaction of stimulus dimensions in concept and choice processes. *Cognitive Psychology, 8,* 98–123.

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

Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition, 78,* 27–43.

Goldstone, R. L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General, 130,* 116–139.

Hall, G. (1991). *Perceptual and associative learning.* Oxford, England: Clarendon Press.

Hegdé, J., & Van Essen, D. C. (2000). Selectivity for complex shapes in primate visual area V2. *Journal of Neuroscience, 10,* RC61.

Hubel, D., & Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology, 160,* 106–154.

Hubel, D. H., & Wiesel, T. N. (1977). Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London, 198B,* 1–59.

Jones, J. P., Stepnoski, A., & Palmer, L. A. (1987). The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology, 58,* 1212–1232.

Kanwisher, N., & Wojciulik, E. (2000). Visual attention: Insights from brain imaging. *Nature Reviews Neuroscience, 1,* 91–100.

Komatsu, H., & Ideura, Y. (1993). Relationships between color, shape, and pattern selectivities of neurons in the inferior temporal cortex of the monkey. *Journal of Neurophysiology, 70,* 677–694.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22–44.

Lamberts, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review, 107,* 227–260.

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.

Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide.* Cambridge, England: Cambridge University Press.

Maddox, W. T. (1992). Perceptual and decisional separability. In F. G. Ashby (Ed.), *Multidimensional probabilistic models of perception and cognition* (pp. 147–180). Hillsdale, NJ: Erlbaum.

Maddox, W. T. (2001). Separating perceptual processes from decisional processes in identification and categorization. *Perception & Psychophysics, 63,* 1183–1200.

Maddox, W. T. (2002). Learning and attention in multidimensional identification and categorization: Separating low-level perceptual processes and high-level decisional processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 99–115.

Maurer, D., LeGrand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences, 6,* 255–260.

Mazer, J. A., Vinje, W. E., McDermott, J., Schiller, P. H., & Gallant, J. L. (2002). Spatial frequency and orientation tuning dynamics in Area V1. *Proceedings of the National Academy of Sciences, USA, 99,* 1645–1650.

Melara, R. D., Marks, L. E., & Potts, B. C. (1993). Primacy of dimensions in color perception. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 1082–1104.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24,* 167–202.

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

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

Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 87–108.

Oliva, A., & Schyns, P. G. (1997). Coarse blobs or fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology, 34,* 72–107.

Op de Beeck, H., & Vogels, R. (2000). Spatial sensitivity of macaque

inferior temporal neurons. *Journal of Comparative Neurology, 426,* 505–518.

Op de Beeck, H., Wagemans, J., & Vogels, R. (2001). Macaque infero-temporal neurons represent low-dimensional configurations of parameterized shapes. *Nature Neuroscience, 4,* 1244–1252.

Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque Area V4. *Journal of Neurophysiology, 82,* 2490–2502.

Regan, D. (2000). *Human perception of objects.* Sunderland, MA: Sinauer Associates.

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience, 2,* 1019–1025.

Schyns, P. G., Goldstone, R., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences, 21,* 1–54.

Schyns, P. G., & Murphy, G. L. (1994). The ontogeny of part representation in object concepts. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 31, pp. 305–354). San Diego, CA: Academic Press.

Shepard, R. N. (1980, October 24). Multidimensional scaling, tree-fitting, and clustering. *Science, 210,* 390–398.

Shepard, R. N., & Cermak, G. W. (1973). Perceptual-cognitive explorations of a toroidal set of free-form stimuli. *Cognitive Psychology, 4,* 351–377.

Shepard, R. N., & Chang, J. J. (1963). Stimulus generalization in the learning of classifications. *Journal of Experimental Psychology, 65,* 94–102.

Sigala, N., & Logothetis, N. K. (2002, January 17). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature, 415,* 318–320.

Smith, E. E., & Medin, D. L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Stankiewicz, B. J. (2002). Empirical evidence for independent dimensions in the visual representation of three-dimensional shape. *Journal of*

Experimental Psychology: Human Perception and Performance, 28, 913–932.

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience, 19,* 109–139.

Tanaka, K., Saito, H., Fukada, Y., & Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology, 66,* 170–189.

Thomas, R. D. (1996). Separability and independence of dimensions within the same-different judgment task. *Journal of Mathematical Psychology, 40,* 318–341.

Thorpe, S. J., & Fabre-Thorpe, M. (2001, January 12). Seeking categories in the brain. *Science, 291,* 260–263.

Treisman, A. (1991). Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 652–676.

Treisman, A., Cavanagh, P., Fischer, B., Ramachandran, V. S., & von der Heydt, R. (1990). Form perception and attention: Striate cortex and beyond. In L. Spillmann & J. S. Werner (Eds.), *Visual perception: The neurophysiological foundations* (pp. 273–316). San Diego, CA: Academic Press.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136.

Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences, 24,* 295–300.

Vogels, R., Eeckhout, H., & Orban, G. A. (1988). The effect of feature uncertainty on spatial discriminations. *Perception, 17,* 565–577.

Wilkinson, F., Wilson, H. R., & Habak, C. (1998). Detection and recognition of radial frequency patterns. *Vision Research, 38,* 3555–3568.

Zahn, C. T., & Roskies, R. Z. (1972). Fourier descriptors for plane closed curves. *IEEE Transactions on Computing, 21,* 269–281.

Zusne, L. (1970). *Visual perception of form.* New York: Academic Press.

(*Appendix follows*)

Appendix

Perceptual Space of Radial Frequency Component Stimuli

A critical assumption throughout this article is that the amplitude variations of two radial frequency components (RFCs) are perceived faithfully—that is, that these stimulus differences are represented in a two-dimensional psychological space. The validity of this assumption has been confirmed in several studies (Cortese & Dyre, 1996; Op de Beeck et al., 2001). As described in this appendix, we replicated these findings with the RFC stimuli that were used in the present article.

The experiment included two groups of RFC stimuli (complex and simple), each containing 16 exemplars (see Figures 1A and 1B). Twelve naive participants (6 for each stimulus group) were asked to rate the similarity between pairs of shapes on a scale of 1 (*most similar*) to 9 (*least similar*). Participants were instructed to use the whole scale, and it was stressed that all shapes were somewhat similar. Stimulus presentation was exactly the same as in the test phase of Experiments 2 and 3, but now subjects were asked to press a number from 1 to 9 after the presentation of the second stimulus. Each combination of two stimuli was presented twice to each participant.

We constructed a symmetrical 16×16 similarity matrix containing the mean similarity rating that each individual gave to each combination of stimuli. These individual matrices (as well as the averaged matrix for each group of 6 subjects) were analyzed with nonmetric multidimensional scaling (MDS). The results of exploratory MDS are traditionally evaluated by inspecting the fit of MDS-derived configurations as a function of the number of dimensions (a *scree plot*). When MDS was applied to the averaged data of all 6 subjects, the scree plot of each stimulus group revealed a knee at two dimensions (see Figure A1). The analyses on individual data showed a knee only for half of the subjects (4 out of 6 subjects for simple RFC stimuli and 2 subjects for complex RFC stimuli). This result is difficult to interpret, as it could be caused either by noise variables due to the lower number of trials or by more systematic factors (e.g., subjects with a psychological space that deviates from the expected two-dimensional parametric configuration).

Although the information contained in a scree plot is almost the only statistical criterion to evaluate the results of exploratory MDS, we can solve this problem by using indices that are typically applied in confirmatory MDS. Indeed, we know the parametric stimulus configurations, and as a result we can assess the fit between the MDS-derived and these expected parametric configurations. A visual inspection of the geometric configurations found with two dimensions already suggests a good fit. Figure A2 shows these configurations for the averaged data and for two typical subjects. Stimuli that are adjacent in parameter space are always adjacent in the MDS-derived configurations. This correspondence was found for all subjects except one.
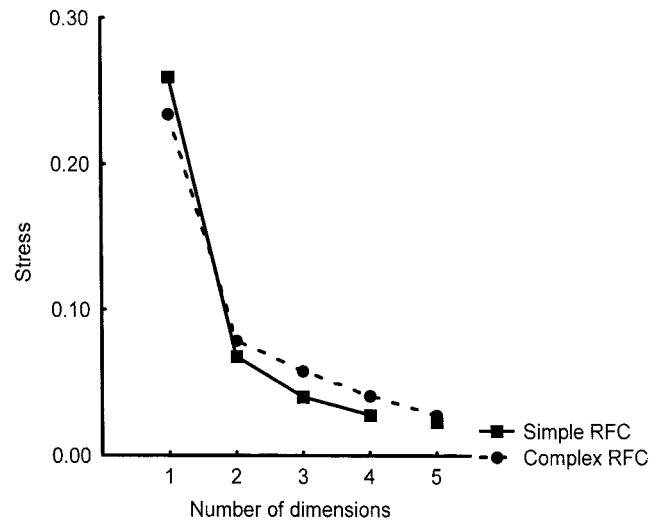

Figure A1. Scree plot showing stress as a function of number of dimensions. RFC = radial frequency component.

We can quantify this correspondence by calculating the congruence coefficient (Borg & Leutner, 1985; Op de Beeck et al., 2001) between the similarities as expected from the parametric configurations and the empirical similarity matrices. We found a significantly high coefficient ($p <$ .001) for each individual participant (minimum .937) and also for the averaged matrices (.985 and .980 for simple and complex RFC stimuli, respectively). Interestingly, similar results were found when the distances between stimuli in the MDS-derived configurations were taken instead of the original similarity ratings. As long as the number of dimensions was reduced to no less than two, the congruence coefficients were not affected. However, the coefficients decreased dramatically (to .89 and .91) with one-dimensional solutions (randomly permuted data sets show coefficients around .88, $SD =$ .013). So, the application of confirmatory MDS reveals that shape similarities are represented in a two-dimensional psychological space that corresponds closely to the parametric configurations.

Of course, these fits are not perfect for any of the stimulus groups. For each stimulus group, the most obvious discrepancy is a higher sensitivity for one of the two parametric dimensions.
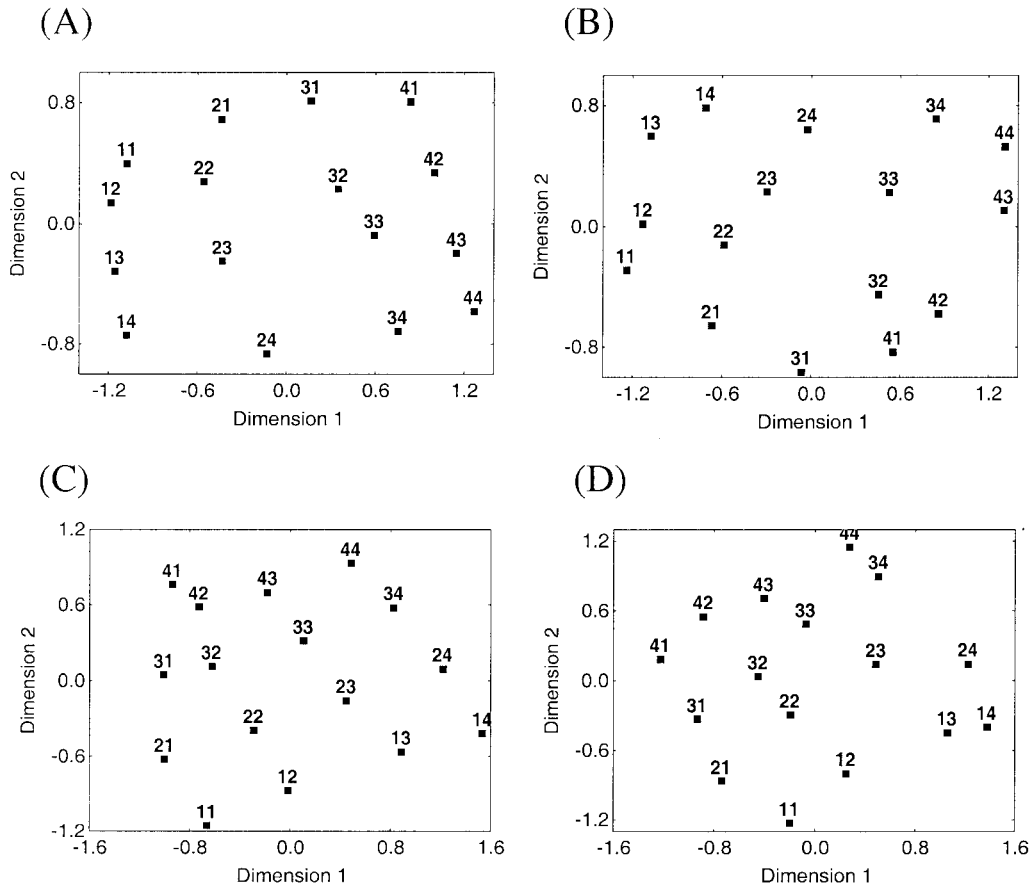
(A)

(B)

(C)

(D)

Figure A2. Two-dimensional solutions as a result of applying multidimensional scaling on the following data sets: (A) similarity ratings of complex radial frequency component (RFC) stimuli averaged across all subjects; (B) similarity ratings of complex RFC stimuli of an individual subject; (C) similarity ratings of simple RFC stimuli averaged across all subjects; and (D) similarity ratings of simple RFC stimuli of an individual subject.