

Generalization of Visual Shapes by Flexible and Simple Rules

Bart Ons and Johan Wagemans*

Laboratory of Experimental Psychology, University of Leuven (K.U. Leuven), Tiensestraat 102,
P.O. Box 3711, B-3000 Leuven, Belgium

Received 2 November 2010; accepted 25 March 2011

Abstract

Rules and similarity are at the heart of our understanding of human categorization. However, it is difficult to distinguish their role as both determinants of categorization are confounded in many real situations. Rules are based on a number of identical properties between objects but these correspondences also make objects appearing more similar. Here, we introduced a stimulus set where rules and similarity were unconfounded and we let participants generalize category examples towards new instances. We also introduced a method based on the frequency distribution of the formed partitions in the stimulus sets, which allowed us to verify the role of rules and similarity in categorization. Our evaluation favoured the rule-based account. The most preferred rules were the simplest ones and they consisted of recurrent visual properties (regularities) in the stimulus set. Additionally, we created different variants of the same stimulus set and tested the moderating influence of small changes in appearance of the stimulus material. A conceptual manipulation (Experiment 1) had no influence but all visual manipulations (Experiment 2 and 3) had strong influences in participants' reliance on particular rules, indicating that prior beliefs of category defining rules are rather flexible.

© Koninklijke Brill NV, Leiden, 2011

Keywords

Categorization, perceptual organization, simplicity, regularity, part configuration, part saliency, shape transformations

1. Introduction

Generalization of a category is the extension of a category to new instances. In a variety of categorization models, similarity is given a central place to explain the process by which different objects are grouped together in one class and considered as equivalent (for seminal work, see Ashby and Perrin, 1988; Edelman,

* To whom correspondence should be addressed. E-mail: johan.wagemans@psy.kuleuven.be

This article is part of the Shape Perception collection, guest edited by Z. Pizlo and S. Dickinson.

1999; Nosofsky, 1984; Shepard, 1987; Tversky, 1977). Generally, these models explain generalization and categorization in two steps. First, the similarity between categories and the to-be-categorized object is mapped geometrically to a psychological space or a feature representation. Secondly, a distance measure corresponding to similarity underlies object identification and categorization. In contrast to similarity-based models, rule-based models describe categories by means of causal rules and constraints of the form, “If properties x , then category C ” (e.g., Barsalou, 1985; Feldman, 1997, 2000, 2006; Rips, 1989, 2001; Smith and Medin, 1981; Smith *et al.*, 1998) in which properties can be regarded as premises or predicates.

There is a common intuition that coherent categories group things together that are similar. A set of new objects is expected to be grouped into categories such that between-category similarity is as low as possible, while within-category similarity is maximized (Rosch and Mervis, 1975). Many similarity-based models are building on this notion but differ considerably in the formal implementation of it. For instance, there are many different ways to define similarity between two objects and different definitions lead to different predictions. Practically, it is very hard to compare two intuitive notions like similarity and rules as the influence of both determinants can be expressed by a variety of different models each leading to a different prediction. A way to steer away from this problem is by evaluating the partition of clusters obtained in a categorization task and determining the glue between objects that might have led to the observed distribution of partitions (see also Pothos and Chater, 2001). In the present study, we evaluated the observed distribution of partitions following a name generalization task collected from a large group of participants. Before we discuss the results of our experimental study in more detail, we will substantiate our design choices and develop a method to establish the link between the expected frequency distribution of formed partitions and the pairwise relation between objects.

1.1. *The Name Generalization Task*

A generalization task has been used before in different forms to study participants' category formations (e.g., Abecassis *et al.*, 2001; Feldman, 1997; Landau and Leyton, 1999; Mash, 2006). A simple version of a generalization task is a triad categorization task where a participant receives three objects, one reference object and two indicative objects. The participant indicates which of the two objects goes best with the reference. A triad categorization task is one of the first tasks that have been used to map similarity in multidimensional spaces by the method of multidimensional scaling (Tongerson, 1952).

In the task which is used in all three experiments below, nine shapes were presented simultaneously and randomly positioned on a sheet of paper. For three distinctive shapes, category pseudo-names were provided that consist of non-existing words similar in form to existing Dutch nouns. For the remaining six shapes having no names yet, the participants were instructed to indicate the most appropriate name between the three alternatives. However, they were informed that each kind

consisted exactly of three objects constraining the task to the formation of three categories each consisting of three objects. Clustering methods have great difficulty to provide an ‘objective’ answer as to how many clusters can be postulated for a set of items and how many items should be grouped together in each cluster (e.g., Corter and Gluck, 1992). Even the K-means clustering method requires input about the number of clusters that need to be identified. However, by constraining the task to the formation of three categories with three objects each, the number of categories and the number of elements in each category are fixed, and the number of possible partitions becomes easier to manage.

In all experiments, we used a similar property structure adopted from experiments in Lee and Navarro (2002; see also Aitkin and Feldman, 2006). In Fig. 1, we present an abstract, schematic overview of the property structure in the stimulus sets. Each cell represents a visual object with common properties shared among all objects and particular properties shared among members in the same row (letters) or in the same column (numbers). The cells on the diagonal are named examples of the three clusters provided to the participants. The imposed constraints lead to three clusters containing three objects each: one provided example and two selected items. When participants generalize the pseudo-names to the row members, they generalize according to the letter properties. In contrast, when participants infer new name members according to the columns, they generalize according to the number properties. Many other outcomes are still possible and each outcome can be thought of as a different partition of the set of nine objects.

The structure of the stimulus set

common a 1 name 1	common a 2	common a 3	dimension Y (marked by letter)
common b 1	common b 2 name 2	common b 3	
common c 1	common c 2	common c 3 name 3	
dimension X (marked by number)			

Figure 1. Each cell represents one stimulus consisting of shared properties. All stimuli share some common properties (common) and all stimuli on the same row and the same column share some specific properties: a, b and c for the rows (dimension Y) and 1, 2 and 3 for the columns (dimension X). Three stimuli that are not forming part of the same row or column are labeled (name 1, name 2 and name 3) and serve as examples for the pseudo-categories.

Actually, the number of possible partitions is much larger than one expects intuitively. Out of nine objects, three stimuli are provided with labels and two new members out of the remaining six need to be inferred for the first category, two new members out of the remaining four need to be selected for the second category, and, finally, two stimuli are left over and become part of the third category. The number of possibilities is therefore $\binom{6}{2}\binom{4}{2}\binom{2}{2} = \frac{6!}{4!2!} \cdot \frac{4!}{2!2!} \cdot \frac{2!}{0!2!} = 90$, so the chance level for each partition outcome including the row-wise and column-wise partition is only about 1.1%. In other words, despite the imposed constraints in the name generalization task, there are still 90 different ways to classify the set of nine objects in three groups of three objects each. The row-wise and the column-wise partitions in Fig. 1 only mark two possibilities in a pool of 88 others.

1.2. Similarity-Based Generalization

The almost universal presumption of researchers in distance models of categorization is the premise that categories correspond to topological neighbourhoods in a representation space. An example of a neighbourhood is, for instance, the ‘consequential region’ in Shepard’s (1987) universal law of generalization. In these models, similarity is defined geometrically in the form of vector distances and is believed to constitute the key feature that drives identification, categorization and generalization (see also Jäkel *et al.*, 2008; Vanpaemel *et al.*, 2005). When one category exemplar is provided to the participant, generalization will occur less likely for new objects if accommodating vectors in the space are more distant from the vector of the category exemplar.

Similarity information has also been represented in set-theoretic terms where shared features are mapped as branches between items in a tree-like structure (Tversky, 1977). Similarity can be represented in different ways predicting different categorizations. Here, we will not collect similarity data for the experimental sets of objects as similarity data is only useful in the light of a particular similarity representation, a choice that we just try to circumvent. Alternatively, we will adopt a virtual representation of similarity as general as possible, so that our experimental evaluation will be flexible with respect to how similarity is defined. More in particular, we assume that the most frequent chosen partition in a group of participants is reflecting the average similarity information of the ‘true model’ (if indeed a true model exists) and, as we will explain further on, we do not need the ‘true’ definition of similarity, once the most favored partition is put forward as a reference. In light of the notion of similarity, we can assume that highly similar object pairs are more likely to be grouped together. Different partitions can be preferred but they should share many highly similar pairs in the same cluster. In other words, different partitions could emerge, but all observed partitions should correspond in grouping together highly similar pairs of objects and this limit should restrain and shape the frequency distribution of partitions in a particular manner. In the following paragraph, we will work out this idea more technically.

1.3. Rand Index and Correspondence

Assuming that the within-category similarity and between-category dissimilarity are maximized, the outcome of a generalization task should lead to a frequency distribution where the most preferred partitions share a lot of correspondences. There are different options to define correspondence between partitions in terms of pairwise relations between objects. A common measure is the Rand index (Rand, 1971). The Rand index between two different partitions X and Y is denoted by the following relation between three numbers: $(n_1 + n_2)/n$, where n_1 is the number of item pairs that are in the same cluster in X and in the same clusters in Y , where n_2 is the number of item pairs that are in two different clusters in X and in two different clusters in Y and where n is a normalization constant consisting of the total number of pairs in the set. For nine objects, n is $9 \times 8/2 = 36$. n_1 corresponds to the notion that similar items should always be grouped together and n_2 corresponds to the notion that dissimilar items should always be grouped in different clusters. In Fig. 2, the frequency distribution of the Rand index is plotted counting all possible partitions just one time, where the Rand index indicates the correspondence between all partitions and any arbitrary partition chosen as a reference. For instance, when the column-wise partition is chosen as a reference, then only the column-wise partition shares full correspondence with itself, leading to one count in the bar with Rand index equal to one. The row-wise partition is then counted in the bar with Rand index equal to 0.5 together with seven other partition possibilities. The total number of partitions in all bars sums to 90. It is possible to make a similar frequency distribution for the selected partitions from a group of participants. For instance, when 90 participants would arbitrarily select one of the possible partitions, at chance level, we can expect that each partition occurs just once and the same distribution plotted in Fig. 2 can then be expected.




Figure 2. The frequency of partitions with a particular Rand index in descending order.

The Rand index measure allows us to group partitions together with the same Rand index in relation to a reference partition. When we choose the most frequently selected partition in a categorization experiment as a reference, then similarity-driven generalization should lead to high counts for partitions with a high Rand index and low counts for partitions with a low Rand index. To establish the last step in our procedure, we need to develop a measure to evaluate how high and how low these counts should be. At chance level, the same relative frequencies that can be derived from Fig. 2 can be expected. Generally, similarity driven categorization should lead to more partitions with a high Rand index (and fewer with a lower one) than relative frequencies expected from chance level. The hypothesis verifying whether the obtained distribution in the experiments deviates from a similarity-based distribution is based on a measure that we label ' Φ ', representing the ratio between the frequency of partitions that have Rand index 0.5 against the frequency of partitions that have a Rand index between 0.5 and 1, 0.5 and 1 not included (1 is not included in order to create a measure independent from the most selected partition). At chance level, i.e., when categorization is not driven by any underlying determinant, Φ is expected to be 0.099 ($= 8/81$ in Fig. 2). Similarity driven categorization is at least expected to reach a Φ measure smaller than 0.099, that is, more partitions with high correspondence. In all the following experiments, this hypothesis will be put to the test.

1.4. Attentional Weights in Similarity-Based Representations

Naturally, participants infer new exemplars from idiosyncratic representations that underlie their choices and therefore similarity-based representations between participants are not necessarily identical and choices can differ between participants. However, a certain degree of coherence between individuals can still be expected, meaning that individuals can differ from each other to the extent that the distribution of the underlying determinants show a large spread but never to the extent that the distribution becomes significantly bimodal or multimodal. Bimodality would suggest that there are two kinds of people, in the sense that the underlying determinants of their generalization behaviour are categorically different.

However, in many original models of categorization, the codominance of the column-wise and row-wise partitions ($\Phi \gg 0.099$) can still be predicted by a modulation of attentional weights. For instance, by using the exemplar-based general context model (Nosofsky, 1984) on the stimulus set in Fig. 1, the letter and number properties might form two separate dimensions and one group of participants might attend to the letter dimension while the other group of participants might put more attentional weight on the number dimension. In the absence of a group of participants that attend to both dimensions more or less equally, the row-wise and column-wise partition might prevail, leading to high counts for the partition with Rand index 1 and high counts for the partitions with Rand index 0.5. Also in the feature-based contrast model of Tversky (1977), a similar prediction is possible when two coherent groups of participants perceive the saliency of the features

in a different way. Two different groups of coherent participants could weigh the set of letter dimensions or the set of number dimensions more heavily, inducing an advantage for the row- and column-wise partition, respectively.

In sum, although many similarity-based models of categorization can easily predict a Φ -value above 0.099, it will be at the cost of the rather unnatural assumption of bimodality in (attention) weighting space. In other words, it would violate a notion that we call coherence between participants.

1.5. Rules Based on Simplicity

The Rand index is a fruitful concept when any pairwise relation like pairwise similarity is supposed to play a key role in the formation of different partitions. However, rule-based models are not building on pairwise relations between objects. One of the most explicitly formalized rule-based categorization principles is the minimal complexity principle of Feldman (1997, 2000, 2006; but see also Leeuwenberg and van der Helm, 1991). In Feldman's view, a category can be defined by a set of rules of different complexity describing the regularities among the members of the same category. The minimal complexity principle is a formalization for the notion that people tend to define categories by a minimal set of simple rules. The coexistence of features among the known objects of a category allows us to define a category based on rules of the form “when property *a* is present then property *b* is not allowed”. The rule indicates that objects with property *a* and *b* are not part of the category.

Complexity depends on the number of rules and the degree of complexity for each individual rule. In Feldman's approach, the degree of complexity, labeled by ‘*K*’, depends on the number of premises in a rule. In the upper row of Fig. 3, four possible partitions are depicted by way of example. The dots represent the cells from Fig. 1 and the clusters are symbolized by the contours encircling the dots. The filled dots represent the provided name-shape examples and the unfilled




Figure 3. Four possible outcomes of the name generalization task of the objects represented on two dimensions are shown on top. The charts below represent the power spectra of the algebraic complexity in Feldman's logic model.

dots represent the to-be-categorized objects. In the column-wise partition in Fig. 3, the first column can be described by two rules with no premises, and therefore, of complexity degree $K = 0$, namely, ‘not 2’ and ‘not 3’. Similarly, the other columns can be described by two rules of zero complexity. In the first lower panel of Fig. 3, a bar graph is depicted presenting a total of 6 rules of zero complexity. It should be noted that this is the summed complexity for the three categories individually but not for the entire partition as such. The row-wise partition is congruent with the column-wise partition and the same bar graph is depicted in the last lower panel. The first mixed partition (second upper panel) has one entire column cluster and analogous to the previous case, two rules of complexity degree zero are needed. However, the other two categories are more complex. The category $\{b1, c1, c3\}$ has two rules of degree zero: ‘not a ’ and ‘not 2’ and one rule of degree $K = 1$: ‘if b , then not 3’. The complexity degree of the last rule is 1 because there is one premise in the rule. The complexity of the category $\{a1, a3, b3\}$ is identical. The total sum for the three categories counts six rules of degree zero and two rules of degree one. The third partition (mixed partition 2) has an identical structure and the bar graphs below are identical.

Feldman’s principle of minimal complexity is one of the rule-based models that can account straightforwardly for a high prevalence of both the column-wise and the row-wise partition and a low occurrence for all in-between partitions. Thus, instead of plotting all possible partitions on a scale of correspondence in function of a Rand index, all partitions should rather be plotted on a scale of complexity in order to verify the notion that logical complexity of a rule set is at the core of categorization. We can then expect a frequency distribution of partitions with high occurrence of low complex rule sets and low occurrence of high complex rule sets. In the current context, the row-wise and column-wise partitions are the simplest possibilities and actually constitute the same location on such a hypothetical scale of complexity.

Evaluating a Rand index makes sense only when some relations between elements are supposed to drive cluster formation. Rules are able to tie objects together with a different mechanism. They are depending on regularities which can be shared by multiple objects. At the level of the single object, a regularity is a structure that cannot be expected to arise coincidentally. For instance, when 1000 points are plotted randomly on a sheet of paper, it is not expected to see them appearing collinearly next to each other. Such an appearance forces the interpretation of an appearing line or a contour segment. Many bits and pieces of this kind constitute the representations of an object. Similarly, at the level of a group of objects, it is hard to believe that many independent objects share a property coincidentally. Once a property is shared by multiple objects, it can become a regularity determining the formation of clusters. The question addressed in rule-based accounts is when a property becomes a regularity. As explained above, when the mechanism behind this process is complexity, then row-wise and column-wise properties can both become regularities determining category formations.

1.6. The Present Study

We created three different kinds of stimulus sets and we implemented in each set the same structure as depicted in Fig. 1. However, different from former studies, we did not use simple features, but rather features that are difficult to extract (e.g., instead of using three distinct colors, we used rather complex shape differences between objects).

We tested Φ in each dataset against the limit of 0.099. The overall goal of the present study is to gather evidence from different sorts of stimulus sets to verify how well-obtained frequency distributions on the correspondence scale comply with pairwise similarity relations.

Additionally, we investigated the degree of invariance of the precedence of particular properties on which participants relied when different groups of participants were submitted to slightly different conditions of a moderating variable. The moderating variables involved were the instructions referring to different superordinate categories (Experiment 1), stimulus orientations and saliency of stimulus properties (Experiment 2), and finally, shape complexity (Experiment 3). These manipulations allow us to verify how solid the underlying representations were, no matter whether they are based on similarity, rules or both.

2. Experiment 1

In Fig. 4, we created a stimulus set with the property structure depicted in Fig. 1. In the columns, all stimuli have similar contours. We refer to the column properties by ‘sharing basic shape’. From the shapes in the middle row, the other shapes

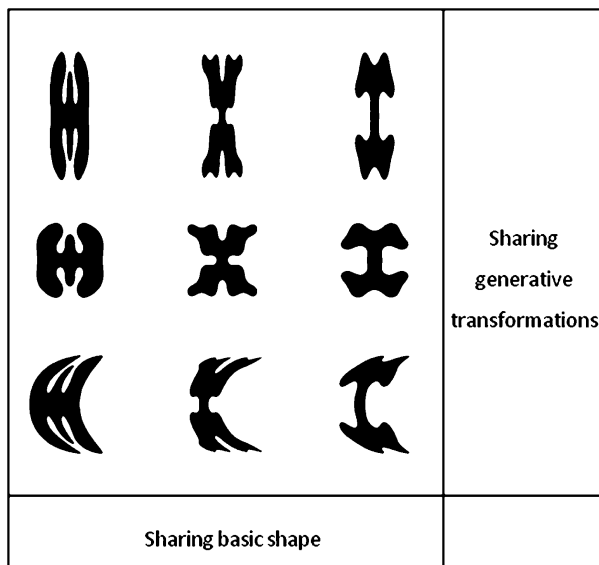


Figure 4. The structure of the stimulus set in Experiment 1.

in the upper and lower row were derived by applying deformations like stretching the image plane vertically and thus changing the aspect ratio (upper row) and bending the image plane and thus changing curvature too (bottom row) (see Ons *et al.*, 2011). We refer to the row properties by ‘sharing generative transformations’. We will evaluate to what extent participants rely on the introduced properties while we do not explicitly probe for these properties. The primary goal is to compare the obtained distribution of partitions (see Introduction) with the expected one for similarity-based generalization.

Category formation and generalization has been associated with visual similarity in adults and children (Abecassis *et al.*, 2001; Diesendruck *et al.*, 2003; Ons and Wagemans, 2011; Quinn and Eimas, 1996; Rakinson and Butterworth, 1998). Although the direction of influence has been described mainly from visual similarity to concepts, in Experiment 1 the reverse influence from conceptual knowledge to the reliance on the kind of visual properties for generalization is investigated too. For instance, knowing that the shapes are supposed to be tools might lead to participants believing that the shapes are rather rigid, and transformations like curvature might become more salient as category-discriminating features (see Landau and Leyton, 1999). Therefore, a converse influence of category ownership on the visual properties used in a generalization task can also be expected. In this experiment, we will examine the extent to which generalization by visual properties depends on instructed rules of superordinate category ownership in the form of ‘all shapes are tools’ (amulets, insects, or vegetables in the other conditions).

2.1. Method

2.1.1. Participants

200 pupils between 17 and 20 years of age (mean 18.21, SD 0.47) were randomly assigned to one of the four experimental conditions (see below). They were recruited from a secondary school, the Kardinaal van Roey Instituut in Vorselaar, Belgium. Participation was voluntary and the study was conducted according to the APA Guidelines (including informed consent procedure). There were 63 male participants (mean 18.17, SD 0.38) and 137 female participants (mean 18.23, SD 0.51).

2.1.2. Stimuli

Contours were created by means of Fourier Descriptors (Zahn and Roskie, 1972). The three basic stimuli resulting from this procedure are depicted in the middle row of Fig. 4. When contour fragments between inflection points (points with zero curvature) would be described categorically by the labels ‘concave’ or ‘convex’, the objects in the columns of Fig. 4 would share the same sequence of labels (see De Winter and Wagemans, 2006). Generative shape transformations are created by the procedures explained in Ons *et al.* (2011).

2.1.3. Procedure

All stimuli in the set were positioned randomly on a sheet of paper (with some constraints on the distance to the border and between neighbouring shapes to prevent overlap between stimuli) and the pseudo-name was provided for only three randomly selected stimuli in the set (constrained to include only one stimulus for each row and column in Fig. 4). The pseudo-names were ‘gnoel’, ‘braal’ and ‘rees’. The task has been explained previously in Section 2, The Name Generalization Task. The instructions were written at the top of the sheet of paper. Participants were instructed to write down one of the three provided pseudo-names on the dashed lines below the six stimuli and they were instructed to use each pseudo-name only two times. Completions with stimulus groups of more or less than three instances were labeled as invalid and excluded from the analysis.

There were four experimental conditions and all participants were assigned randomly to one of these four conditions, resulting in fifty participants per condition. In each condition participants were instructed to think of the objects as belonging to a particular superordinate category, namely, tools, amulets, insects or vegetables.

2.2. Results and Discussion

We counted the frequencies for the row-wise and column-wise partitions and we used the log-ratio of both counts as dependent measure in a log-linear regression model with the instructed superordinate category as explanatory variable. In the upper panel of Fig. 5, the relative frequencies are depicted for each superordinate condition. There were only very few invalid completions (i.e., 2, 1, 1, and 1 for the four experimental groups, tools, amulets, insects, and vegetables, resp.). For all experimental groups there was a higher reliance on basic shape in generalization (see Fig. 5). The odds of relying on basic shape properties against generative transformations was significantly different from 1 for tools (4.76 , $\chi^2_1 = 10.18$, $p < 0.005$), vegetables (3.25 , $\chi^2_1 = 8.5$, $p < 0.005$) and insects (3.11 , $\chi^2_1 = 8.77$, $p < 0.005$), but not for amulets (1.64 , $\chi^2_1 = 2.14$, $p = 0.14$). However, we did not find any significant difference between any pair of experimental conditions. The superordinate level that we imposed on the participants did not bias the category inferences in favour of one kind of properties. A possible explanation for the absence of such an effect might be the poor induction of superordinate category ownership. Instructing participants to think of the stimuli as tools, amulets, insects or vegetables might be too weak as a manipulation to perceive the artificial shapes as members belonging to one of the proposed superordinate categories. In essence, we cannot conclude from the present findings that superordinate category ownership does influence the dominance of particular visual properties in generalization.

We also compared the absolute frequencies of the column-wise and the row-wise partitions of the aggregated data against chance level ($1/90$) and we found that basic shape (column-wise) was 46 times above chance level ($\chi^2_1 = 820$, $p < 0.0001$) and generative shape transformations (row-wise) was 16 times above chance level on

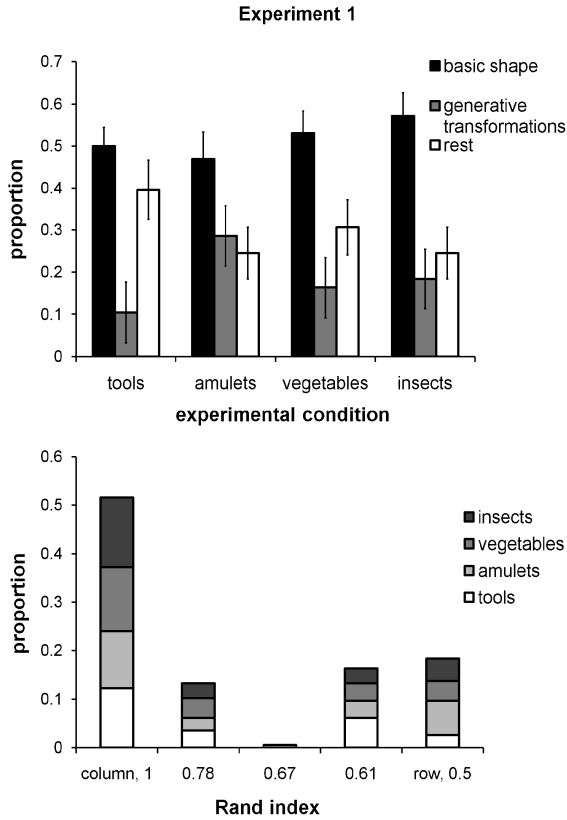


Figure 5. In the upper panel, the proportions of column-wise (basic shape) and row-wise (generative transformations) generalizations are depicted per condition. Error bars are standard errors. In the lower panel, the proportions of partitions for each correspondence group are aggregated over all conditions with the column-wise partition as reference.

average ($\chi^2_1 = 229$, $p < 0.0001$). Although stimuli were randomly positioned on a sheet of paper, participants were still able to detect the regularities in the set and to relate the kind of regularities with each other.

Finally, we tested whether the frequency distribution for all groups of correspondences is compatible with similarity-based generalization. In all conditions, the column-wise partitions were the most preferred ones. We did not find any difference of importance between the four conditions, and therefore, we aggregated the data over all conditions and we found 36 partitions with Rand index equal to 0.5 (all of them were row-wise partitions) and 58 partitions with Rand index between 0.5 and 1 ($\Phi = 36/58$). We found a strong deviation between the expected and the observed frequency distribution of partitions relative to a Φ measure of 0.099 (see lower panel of Fig. 5). The observed Φ value was 6.27 times larger than the limit of 0.099 introduced previously ($\chi^2_1 = 75.04$, $p < 0.0001$).

3. Experiment 2

One particular view on visual representation is the part-based view (e.g., for the importance of parts in children, see Bhatt *et al.*, 2006; Rakinson and Butterworth, 1998). In a series of part-based approaches (e.g., Biederman, 1987; De Winter and Wagemans, 2006; Hoffman and Richards, 1984), it is assumed that the visual system tracks the shape boundary and searches for concave dips and discontinuities along the edges of objects that are important to signal the connecting points between parts. Two concepts are important in nearly all part-based views: the shape of the parts as separate units of analysis and the overall spatial configurations of the parts. In Fig. 6, we implemented these two key concepts as row and column features in the structural grid of Fig. 1. All shapes had four peripheral parts attached to the central part, two pointy triangles and two rounded semi-ellipses. In the upper left panel of Fig. 6, the shape of the central part varied from a circle into a square-like and triangle-like shape and constituted the column properties. Part configuration refers to the relative positions of the attached parts on the central part and constituted the row properties. The primary goal in Experiment 1 was to compare the



















		central part shape			
		regular		salient	
orientation	horizontal				Sharing part configuration
					
					
	Sharing part shape		Sharing part shape		
orientation	varied				Sharing part configuration
					
					
	Sharing part shape		Sharing part shape		

Figure 6. Stimulus sets and design of Experiment 2.

generalization frequency pattern of the participants against the expected pattern of a similarity-based generalization account, and likewise, the same test is carried out in Experiment 2.

Similarly to Experiment 1, we manipulated the context for these two part-related shape properties. In one manipulation, we emphasized the distinction between central and peripheral parts by an additional salient shape property (aspect ratio and curvature in the right panels of Fig. 6) (for part saliency, see also De Winter and Wagemans, 2006, 2008). We crossed this manipulation with a second dichotomous manipulation where we kept standardized horizontal orientations against a second condition where the same stimuli were arbitrary oriented (compare the upper two panels to the lower two panels in Fig. 6). We presumed that part saliency would draw attention to the shape of the central part. Secondly, we presumed that the reluctance of regular orientations would make it harder to detect the overall configuration properties. Although all four stimulus sets in Fig. 6 had identical property structures, we hypothesized that the two moderating variables would influence the pattern of generalization.

3.1. Method

3.1.1. Participants

Subsequently to the first experiment, the same 200 pupils participated in the second experiment. To minimize the sequential effects from the first onto the second experiment, different kinds of stimuli were created, and to balance out any potential carry-over effect all pupils were randomly assigned for a second time to one of the four experimental conditions in Experiment 2.

3.1.2. Stimuli

Four sets of nine shapes were created. The basic set consisted of three types of central parts combined with three types of configurations defined by the relative positions of the smaller shape protrusions. This combination of part configuration and part shape was then also crossed with two moderating variables: orientation of the stimuli (horizontal *versus* varied) and central part shape (regular *versus* salient) (see Fig. 6).

It is worthwhile to mention that the central part shape can be conceived differently in view of regularities: The square central shape possesses four axes of symmetry (left panels), while the rectangular shape has only two axes of symmetry (right panels), the compact triangular shape has three axes of symmetry (left panel), while the more salient and less compact triangular shape does not have any axis of symmetry (right panels). The manipulation of shape can therefore also be conceived as a manipulation in degree of mirror symmetry, a factor which is known to influence regularity detection in dot patterns (e.g., Wagemans *et al.*, 1991) and shape encoding (e.g., Kayaert and Wagemans, 2009).

3.1.3. Procedure

The procedure was identical to Experiment 1. Contrary to Experiment 1, no contextual superordinate category was imposed and identical instructions were provided to all participants. The stimuli were all referred to as ‘figures’. The three pseudo-names were ‘vek’, ‘saf’ and ‘mig’.

A fully-factorial manipulation of the two perceptual variables of interest was induced (see Fig. 6) over four groups of participants. Each group of fifty participants was faced with one of the stimulus sets depicted in Fig. 6.

3.2. Results and Discussion

We used the log-ratio of the column-wise frequencies to the row-wise frequency of partitions as dependent measure in a log-linear regression model with three explanatory variables: central part shape (regular *versus* salient), orientation (horizontal *versus* varied) and the interaction between them. In the upper panel of Fig. 7, the

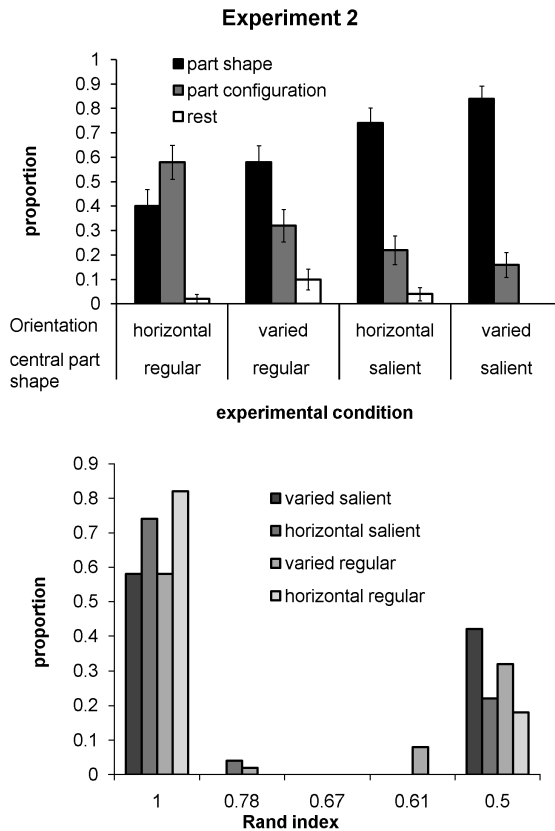


Figure 7. In the upper panel, the proportions of column-wise (part shape) and row-wise (part configuration) generalizations are depicted per condition. Error bars are standard errors. In the lower panel, the proportions of partitions for each correspondence group with the column-wise partition as reference in the first condition and the row-wise partition as reference in the last three conditions.

relative frequencies are depicted for each combination of conditions. There were no invalid completions.

Whereas the relative use of the row-wise and column-wise features was not influenced by the superordinate category labels in Experiment 1, this was now strongly influenced by the visual manipulations of shape. We found a significant effect for orientation ($\chi^2_1 = 4.45$, $p < 0.05$) and a significant effect for central part shape ($\chi^2_1 = 15.65$, $p < 0.0001$) but the interaction between orientation and central part shape was not significant ($\chi^2_1 = 0.61$, $p = 0.43$). The odds for inferring category members by part shape properties against inferring members by part configuration was 2.96 ($\chi^2_1 = 21.10$, $p < 0.0001$) on average for the groups with differently oriented stimuli, and 1.43 ($\chi^2_1 = 2.95$, $p = 0.086$) for the groups with the same oriented stimuli. Clearly, participants had more difficulties in detecting the overall part configuration when the stimuli were oriented differently. Participants also relied more on the shape of the central part when this part was more salient. The odds for inferring category members by part shape properties against inferring members by part configuration was 4.16 ($\chi^2_1 = 31.10$, $p < 0.0001$) on average for the groups that received salient central parts, and 1.31 ($\chi^2_1 = 1.87$, $p = 0.17$) for the groups with regular shaped central parts. Contrary to Experiment 1, where aspect ratio and curvature were implemented globally and seemed to have a moderate importance, aspect ratio and curvature seemed to induce strong effects in Experiment 2 when it was applied locally on just one part.

We also compared the absolute frequencies of the column-wise and the row-wise partitions against chance level (1/90) and we found that part shape (column-wise) was 36, 52, 66 and 75 times larger than chance level for the depicted conditions in the upper panel of Fig. 7, going from left to right, respectively (all $\chi^2_1 > 180$, $p < 0.0001$). We found that part configuration (row-wise) was 52, 28, 20 and 14 times above chance level, respectively (all $\chi^2_1 > 50$, $p < 0.0001$). Although stimuli were randomly positioned on a sheet of paper, participants were still able to detect the regularities and the kinds of regularities in the set.

Finally, we tested whether the frequency distribution for all groups of correspondences is compatible with similarity-based generalization. Unlike Experiment 1, we did not aggregate the data because the preferred partitions between conditions differed greatly. The odds Φ were compared against the limit 0.099 and were much larger: all Φ were larger than 3.2. In the horizontal regular and varied salient condition, there were only partitions with Rand index 0.5 and 1. To perform statistical tests we added one count for the Φ denominator in these two conditions. All Φ measures were significantly larger than 0.099: $\chi^2_1 = 17.17$, $p < 0.0001$; $\chi^2_1 = 46.08$, $p < 0.0001$; $\chi^2_1 = 27.34$, $p < 0.0001$ and $\chi^2_1 = 27.48$, $p < 0.0001$ for each condition going from left to right on the horizontal axis of the upper graph in Fig. 7, respectively, and thus, we found evidence against a frequency pattern that can be expected from a similarity-based generalization.

4. Experiment 3

The shape complexity of planar shapes can be quantified by a dimensionless length-to-area measure, namely, the ratio between the squared contour length and its enclosed area (Zusne, 1970). In Experiment 3, we implemented a twofold manipulation of stimulus complexity by increasing the number of equivalent parts, and thus, increasing the contour length relative to its enclosing area. Column-wise, stimuli had the same number of convex parts; row-wise, we implemented generative transformations similar to the ones that we used in Experiment 1 (see Fig. 8). The stimulus structure from Fig. 1 was adopted and, similarly to Experiments 1 and 2, we evaluated how well similarity based models can account for the frequency distribution of partitions. To give similarity-based generalization somewhat of an advantage, we used stimuli that were looking qualitatively similar: all stimuli are sinusoidal waves bent around a circle.

A second manipulation on complexity was added between the different stimulus sets which we called ‘complexity by the number of parts’. It involved a larger difference in complexity between different stimulus sets administered to different groups of participants. The manipulation within and between stimulus sets was carefully balanced, respecting Weber’s fraction. We added one part per column in the first set (left panel in Fig. 8) while we added two parts in the third set (third panel), so both additions were equivalent with respect to the average number of parts in the sets ($1/4 = 2/8$). Similarly, we added one and two parts in the second and fourth set, respectively (panel two and the last panel), so both variations were equivalent too with respect to the average number of parts ($1/6 = 2/12$). By controlling for discriminability, we wanted to investigate whether participants would still rely on an assumedly equally discriminable but more complex property in a generalization task. Based on Feldman’s principle of minimal logical complexity in category rules, the column-wise and the row-wise partitions are equally likely. However, besides logical complexity, complexity can also be approached from a visual point of view. Although column-wise and row-wise partitions have the same rule-based logical complexity, they might differ in the visual complexity of the shape properties embedded in those rules. Therefore, we hypothesized that participants would rely less

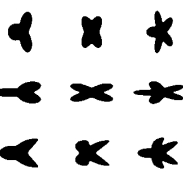
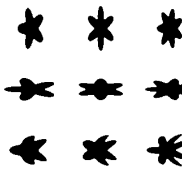
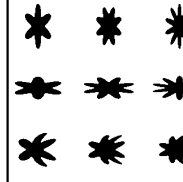
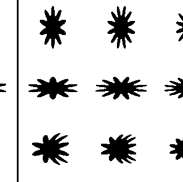
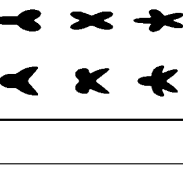
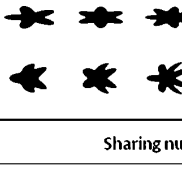
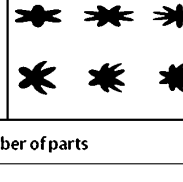
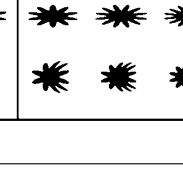
Complexity of stimulus set (number of parts)				Sharing generative transformations
				
				
Sharing number of parts				

Figure 8. Experimental conditions in Experiment 3.

on the part number property (column-based generalization) when the average number of parts in the set increases. In other words, by incorporating visual complexity as a variable between stimulus sets, we sought evidence that visual complexity is also important in generalization besides logical complexity of rules.

Similarly to Experiments 1 and 2, we evaluated how well similarity based models can account for the distribution of the frequencies. To give similarity based generalization an advantage, we used stimuli that were looking qualitatively similar: all stimuli are sinusoidal waves bent around a circle.

4.1. Method

4.1.1. Participants

Subsequent to the second experiment, the same 200 pupils participated in Experiment 3. To minimize the sequential effects from the previous experiments on this one, completely different sets of stimuli were created and all pupils were randomly assigned for a third time to one of the experimental conditions in Experiment 3.

4.1.2. Stimuli

All stimuli were created by means of a singular radial frequency component added on a circle with a constant radius (for radial frequency components, see Shepard and Chipman, 1970). The complexity of the stimuli was manipulated by increasing the frequency of the radial frequency component. The lowest frequency consisted of three cycles on the perimeter of the complete contour, resulting in three convex parts (see Fig. 8, leftmost column). The highest frequency consisted of 14 cycles along the contour, resulting in 14 convex parts (see Fig. 8, rightmost column). The stimuli on the left of Fig. 5 have the lowest complexity while the stimuli on the right have the highest complexity. Similar to Experiment 1, global generative transformations like aspect ratio (compare the first row of Fig. 8 with the second row of Figure 8) and curvature (see bottom row in Fig. 8) were applied differently between the rows.

4.1.3. Procedure

The task was identical to Experiments 1 and 2. The three pseudo-names used here were ‘vlonk’, ‘krins’ and ‘lo’. There were four experimental conditions and in each condition a different stimulus set was provided to the participant, corresponding to one of the four sets depicted in Fig. 5. In the first group, the presented stimuli consisted of 3, 4 and 5 convex parts; in the second experimental group, the local part variation consisted of 5, 6 and 7 convex parts. In the third and the fourth experimental groups, the numbers of convex parts were 6, 8 and 10; and 10, 12 and 14, respectively.

4.2. Results and Discussion

We used the log-ratio of the column-wise frequencies to the row-wise frequency of partitions as dependent measure in a log-linear regression model with complexity of the stimulus set as explanatory variable (4, 6, 8, 12 parts on average). In the upper panel of Fig. 9, the relative frequencies are depicted for each condition. The

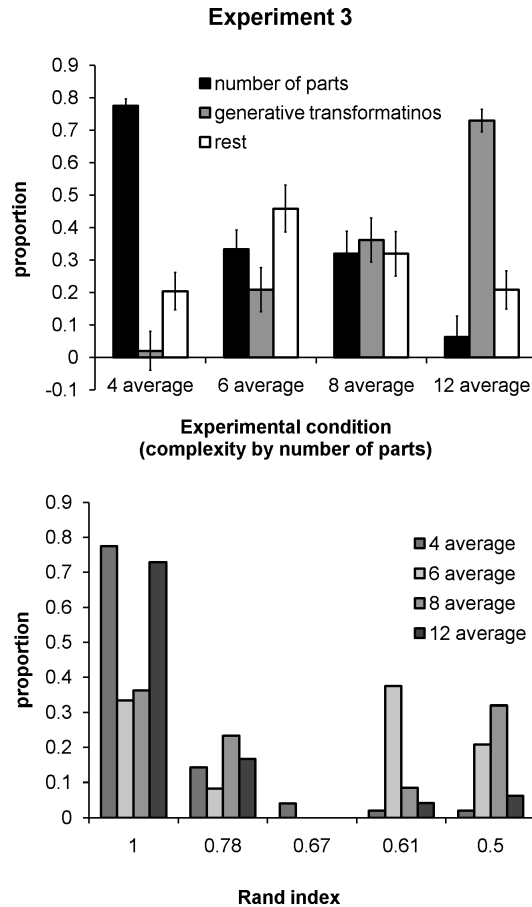


Figure 9. In the upper panel, the proportions of column-wise (number of parts) and row-wise (generative transformations) generalizations are depicted per condition. Error bars are standard errors. In the lower panel, the proportions of partitions for each correspondence group with the column-wise partition as reference in the first two conditions and the row-wise partition as reference in the last two conditions.

numbers of invalid completions were 1, 2, 3 and 2 for the four experimental groups, from left to right in Fig. 8, respectively. The invalid completions were omitted from the analysis.

By increasing the complexity of the stimuli, the number of parts seemed to become less important and the global generative transformations seemed to gain importance (see upper panel of Fig. 9). Complexity by the number of parts induced a significant effect on the column-wise *versus* row-wise dominance ($\chi^2_1 = 50.68$, $p < 0.0001$). The odds for generalization based on generative shape transformations against the number of parts were 0.03 ($\chi^2_1 = 12.89$, $p < 0.0005$), 0.63 ($\chi^2_1 = 1.36$, $p = 0.24$), 1.13 ($\chi^2_1 = 0.12$, $p = 0.72$), 11.67 ($\chi^2_1 = 16.68$, $p < 0.0001$) from the left stimulus set to the right stimulus set in Fig. 8, respectively. Participants

seemed to prefer less complex visual properties. The trend is unlikely to involve discriminability because the Weber fraction of part number was similar for the first and third conditions, as well as for the second and fourth conditions. For instance, differentiating two figures with 11 and 12 parts is more difficult than for two figures with 3 and 4 parts. To prevent the situation that participants would rely less on the number of parts property because they did not perceive this feature adequately for the more complex stimuli, we added two parts differences between the columns for the more complex stimulus sets instead of one part difference for the more simple sets.

In a second series of analyses, we compared the absolute frequencies of the column-wise and the row-wise partitions aggregated over all conditions against chance level and we found that the addition of parts (column-wise) was 44.4 times above chance level ($\chi^2_1 = 42.11$, $p < 0.0001$) and overall part structure (row-wise) is 32.3 times above chance level on average ($\chi^2_1 = 32.17$, $p < 0.0001$). Participants were able to detect the consistent structure of regularities in the set.

We also tested whether the frequency distribution of partitions on the correspondence scale is compatible with generalization based on similarity between objects. The ratio Φ was 0.1, 0.45, 1 and 0.3, for the conditions going from left to right in Fig. 9, respectively. We compared them to the limit 0.099 and we found a significant difference for conditions with six parts on average: $\chi^2_1 = 16.02$, $p = 0.0001$, and for the condition with eight parts on average: $\chi^2_1 = 40.19$, $p < 0.0001$. Additionally, we found a marginally significant difference for the condition with twelve parts on average: $\chi^2_1 = 2.85$, $p < 0.1$. In the conditions with four parts on average, the column-wise partition was too dominant and therefore, we did not find evidence against a similarity-based frequency distribution. It should be noted that the Φ ratio being equal to 0.1 cannot count as evidence in favor of similarity driven categorization. The Φ ratio should be significantly smaller than 0.099, which is clearly not the case.

5. General Discussion

5.1. Rule-Based Versus Similarity-Based Generalization

The main issue is whether subjects are more inclined to extend category membership to objects that are similar to the provided category examples, or that obey salient rules common to the provided category examples. Obviously, in many real situations and experiments, these two proposals would be confounded, because objects that obey common rules are also more similar. However, by constructing a 3×3 grid of objects in each stimulus set, in which rows or columns corresponded to common rules, but 2D regions of various shapes are more likely to correspond to similarity pools, we tried to unconfound both underlying determinants of categorization. In brief, participants overwhelmingly tend to organize categories by rows or columns and they seem to favor the use of simple rules. Actually, when looking at the stimulus sets in Figs 4, 6 and 8, the results may appear quite intuitive and

perhaps even obvious for most readers, but this is somewhat misleading given that participants were not informed about the rule structure and that the objects were presented in quasi-random locations to the participants. Moreover, by adopting the notion of coherence as explained in the Introduction, we show that these more or less intuitive results were not easily predicted distribution-wise under standard similarity models. Pothos and Chater (2001) proposed that the underlying determinants of human categorization can be investigated by looking at the partitions that humans make instead of comparing two families of models building on these determinants. That is exactly what we intended to do here: instead of comparing the plausibility of different formalizations (models) built on some notion, we compared the notions rather directly — the one that predicates similarity-based categorization and the one that predicates logical relations between rules and categories. Overall, participants seemed to use rules in the current experimental context.

A principle that can account naturally for the observed distributions of partitions in the data is the logical minimal principle formulated by Feldman (1997, 2000, 2006; see also Goodman *et al.*, 2008). The row-wise and the column-wise partitions are the only two partitions allowing simple rule-based category formulations with rules that simply state the presence or absence of a property. All other partitions have categories that are to be defined by rules with one premise in the form of ‘if A, then B’. In other words, participants favour simple rules to define categories. However, by introducing the moderating variables between the different stimulus sets, the preference for one of the two partitions with the same logical complexity changed spectacularly. The strong influence of the moderating variables shows that the logical complexity of the partition-wise rule sets are only telling one part of the story. Although row-wise and column-wise partitions have equal logical complexity, particular rules are favoured above others depending on the visual properties implemented in the rules. The visual analysis of particular properties is also depending on shape complexity (e.g., Experiment 3). To detect a visual property as a regularity in the stimulus set, a pairwise comparison between objects is not necessary. Gestalt principles like grouping by similarity could also explain how regularities become detected. Any recurrent property in a group of simultaneous presented objects can attract attention and regularities might even be detected before individual objects are analyzed.

5.2. *Detecting Regularities*

In perceptual organization, a considerable number of studies have stressed the importance of regularities or non-accidental properties (e.g., Bennett *et al.*, 1989; Feldman, 1997, 2003; Hulleman and Boselie, 1999; Kanizsa, 1979; Wagemans, 1992, 1993). Intuitively, a regularity can be thought of as a property that is appropriate for a whole set of stimuli and not just one (e.g., mirror symmetry), and as a consequence, the occurrence of such a property will be interpreted as non-accidental in the sense that an observer believes that it has some suspicious causal reason. The notion of coincidence can best be understood in the following manner.

Suppose that a contour generator machine would draw closed contours in a random fashion; then a symmetric figure would occur only occasionally. When only symmetric figures would be created from the generator, the observer would assume that the generator is constrained to create symmetric contours and they will associate the attribute symmetry to the nature of the machine (read ‘category’). An observer would therefore believe that the observed symmetry is not a coincidence. Likewise, many features presented in the stimulus sets can be perceived as regularities when multiple stimuli in the set possess the same property, like compounded parts in Experiment 2 and regular convex parts in the stimulus sets of Experiment 3; even a straight side can be considered as a regularity or a deviation from an average generic curved side. Moreover, the properties relate to properties that have some perceptual relevance according to various views of visual processing. According to Feldman’s proposal (2003, 2009), humans will construct a space of all possible interpretations built from regularities. The interpretation space can be ordered hierarchically by the degree of regularity and the interpretation with the highest degree of regularity or the interpretation by which most coincidental properties are appearing as non-accidental is the one that will be preferred. Similar views on shape processing have been proposed in the generative approach of Leyton (1987, 1988, 1989), and the descriptive minimum principle of Leeuwenberg (1971) and Leeuwenberg and van der Helm (1991).

5.3. *The Precedence of Particular Properties*

The task that we used seems to provide a powerful way to measure the relative preference for one regularity above the other. One of the most astonishing results in the experiments was the enormous shift in relative preference by some moderating variables like part saliency and complexity. In Experiment 2, participants relied more on the central part when the stimuli were oriented differently (for orientation dependency in one-shot generalization, see also Feldman, 1997) and when the saliency of the central part increased. To perceive the overall part configuration in shapes that are oriented differently, participants should take into account the relative spatial distances between the parts, which might involve more complex processing than absolute part positions. Therefore, participants might have relied less on part configuration when the stimuli were oriented differently. However, it is difficult to determine whether participants could not detect part configuration as a regularity or whether they just did not prefer the more complex visual property (see also Keane *et al.*, 2003). Therefore, we controlled for detectability in Experiment 3, and visual complexity seemed to drive preferences for particular properties in the last experiment. Participants preferred the part property when it constituted a simple visual property with a few parts, but not when it constituted a more complex property with many parts. In sum, just like participants seemed to form categories based on simple logical rules, they preferred to use the simplest visual features. Simplicity might seem to suggest that participants sought the path of least effort. However, the path of least effort would lead to a haphazard generalization instead of a column-

or row-based one, and therefore, the results seems to suggest that simplicity forms an inherent part of the nature of people's behaviour.

5.4. General Conclusion

We used a particular name generalization task in three experiments and we showed that similarity-based processing is unlikely to be involved. However, evidence should be evaluated against the specificity of the task, and it does not necessarily afford generalization to other categorization paradigms. Therefore, we believe that similarity-based processing does exist, but we claim that similarity-based processes are not always involved in generalization. A rule-based account like Feldman's minimal complexity principle provides a more natural account for the presented data. Generally, participants seem to seek the simplest category description in the extension of categories to new exemplars and they prefer visually simple and easy detectable properties.

Acknowledgements

This work was supported by research grants from the Fund for Scientific Research (FWO Flanders, G.0218.06), from the University Research Council (IDO/02/004) and from the Methusalem grant by the Flemish Government awarded to JW (METH/08/02). We would like to thank Niké Mertens for the assistance in data collection, and two anonymous reviewers for helpful comments.

References

- Abecassis, M., Sera, M. D., Yonas A. and Schwade, J. (2001). What's in a shape? Children represent shape variability differently than adults when naming objects, *J. Exper. Child Psychol.* **78**, 213–239.
- Aitkin, C. D. and Feldman, J. (2006). Subjective complexity of categories over three-valued features, in: *Proc. 28th Ann. Conf. Cognit. Sci. Soc.*, pp. 961–966. Lawrence Erlbaum, Mahwah, NJ, USA.
- Ashby, F. G. and Perrin, N. A. (1988). Toward a unified theory of similarity and recognition, *Psycholog. Rev.* **95**, 124–150.
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in category, *J. Exper. Psychol., Learning, Memory Cognition* **11**, 629–654.
- Bennett, B., Hoffman, D. and Prakash, K. (1989). *Observer Mechanics: A Formal Theory of Perception*. Academic Press, New York, USA.
- Bhatt, R. S., Hayden, A., Reed, A., Bertin, E. and Joseph, J. (2006). Infants' perception of information along object boundaries: concavities *versus* convexities, *J. Exper. Child Psychol.* **94**, 91–113.
- Biederman, I. (1987). Recognition-by-components: a theory of human image understanding, *Psycholog. Rev.* **94**, 115–147.
- Cortier, J. E. and Gluck, M. A. (1992). Explaining basic categories: future predictability and information, *Psycholog. Bull.* **2**, 291–303.
- De Winter, J. and Wagemans, J. (2006). Segmentation of object outlines into parts: a large-scale, integrative study, *Cognition* **99**, 275–325.

- De Winter, J. and Wagemans, J. (2008). Perceptual saliency of points along the contour of everyday objects: a large-scale study, *Percept. Psychophys.* **70**, 50–64.
- Diesendruck, G., Hammer, R. and Catz, O. (2003). Mapping the similarity space of children and adults' artifact categories, *Cognit. Develop.* **18**, 217–231.
- Edelman, S. (1999). *Representation and Recognition in Vision*. MIT Press/Bradford Books, Cambridge, MA, USA.
- Feldman, J. (1997). The structure of perceptual categories, *J. Mathemat. Psychol.* **41**, 145–170.
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning, *Nature* **407**, 630–633.
- Feldman, J. (2003). Perceptual grouping by selection of a logically minimal model, *Intl J. Computer Vision* **55**, 5–25.
- Feldman, J. (2006). An algebra of human concept learning, *J. Mathemat. Psychol.* **50**, 339–368.
- Feldman, J. (2009). Bayes and the simplicity principle in perception, *Psycholog. Rev.* **116**, 875–887.
- Goodman, N. D., Tenenbaum, J. B., Griffiths, T. L. and Feldman, J. (2008). Compositionality in rational analysis: grammar-based induction for concept learning, in: *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*, N. Chater and M. Oaksford (Eds), pp. 377–406. Oxford University Press, Oxford, UK.
- Hoffman, D. D. and Richards, W. A. (1984). Parts of recognition, *Cognition* **18**, 65–96.
- Hulleman, J. and Boselie, F. (1999). Perceived shape regularity does not depend on regularities along the contour, *Perception* **28**, 711–724.
- Jäkel, F., Schölkopf, B. and Wichmann, F. A. (2008). Similarity, kernels, and the triangle inequality, *J. Mathemat. Psychol.* **52**, 297–303.
- Kanizsa, G. (1979). *Organization in Vision: Essays on Gestalt Perception*. Praeger Publishers, New York, USA.
- Kayaert G. and Wagemans J. (2009). Delayed shape matching benefits from simplicity and symmetry, *Vision Research* **49**, 708–717.
- Keane, S. K., Hayward, W. G. and Burke, D. (2003). Detection of three types of changes to novel objects, *Visual Cognition* **10**, 101–127.
- Landau, B. and Leyton, M. (1999). Perception, object kind, and object naming, *Spatial Cognit. Comput.* **1**, 1–29.
- Lee, M. D. and Navarro, D. L. (2002). Extending the ALCOVE model of category learning to featural stimulus domains, *Psychonomic Bull. Rev.* **9**, 43–58.
- Leeuwenberg, E. L. J. (1971). A perceptual coding language for visual and auditory patterns, *Amer. J. Psychol.* **84**, 307–349.
- Leeuwenberg, E. and van der Helm, P. (1991). Unity and variety in visual form, *Perception* **20**, 595–622.
- Leyton, M. (1987). Nested structures of control: an intuitive view, *Comput. Vision, Graphics, Image Proc.* **37**, 20–53.
- Leyton, M. (1988). A process-grammar for shape, *Artific. Intell.* **34**, 215–247.
- Leyton, M. (1989). Inferring causal history from shape, *Cognit. Sci.* **13**, 357–387.
- Mash, C. (2006). Multiple shape similarity in the development of visual object classification, *J. Exper. Child Psychol.* **95**, 128–152.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification, *J. Exper. Psychol.: Learning, Memory, Cognit.* **10**, 104–114.
- Ons, B. and Wagemans, J. (2011). Development of differential sensitivity for shape changes resulting from linear and nonlinear planar transformations, *i-Perception* (in press).

- Ons, B., De Baene, W. and Wagemans, J. (2011). Subjectively interpreted shape dimensions as privileged and orthogonal axes in mental shape space, *J. Exper. Psychol.: Human Percept. Perform.* **37**, 422–441.
- Pothos, E. M. and Chater, N. (2001). Categorization by simplicity: a minimum description length approach to unsupervised clustering, in: *Similarity and Categorization*, U. Hahn and M. Ramscar (Eds), pp. 51–72. Oxford University Press, Oxford, UK.
- Quinn, P. C. and Eimas, P. D. (1996). Perceptual organization and categorization in young infants, in: *Advances in Infancy Research*, Vol. 10, C. Rovee-Collier and L. P. Lipsitt (Eds), pp. 1–36. Ablex, Norwood, NJ, USA.
- Rakinson, D. H. and Butterworth, G. E. (1998). Infants' attention to object structure in early categorization, *Develop. Psychol.* **34**, 1310–1325.
- Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods, *J. Amer. Statist. Assoc.* **66**, 846–850.
- Rips, L. J. (1989). Similarity, typicality, and categorization, in: *Similarity and Analogical Reasoning*, S. Vosniadou and A. Ortony (Eds), pp. 21–59. Cambridge University Press, Cambridge, UK.
- Rips, L. J. (2001). Necessity and natural categories, *Psycholog. Bull.* **127**, 827–852.
- Rosch, E. and Mervis, C. B. (1975). Family resemblances: studies in the internal structure of categories, *Cognit. Psychol.* **7**, 573–605.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science, *Science* **237**, 1317–1323.
- Shepard, R. N. and Chipman, S. (1970). Second-order isomorphism of internal representations: shapes of states, *Cognit. Psychol.* **1**, 1–17.
- Smith, E. E. and Medin, D. (1981). *Categories and Concepts*. Harvard University Press, Cambridge, MA, USA.
- Smith, E. E., Patalano, A. L. and Jonides, J. (1998). Alternative strategies of categorization, *Cognition* **65**, 167–196.
- Tongerson, W. S. (1952). Multidimensional scaling: I. Theory and method, *Psychometrika* **17**, 401–419.
- Tversky, A. (1977). Features of similarity, *Psycholog. Rev.* **84**, 327–352.
- Vanpaemel, W., Storms, G. and Ons, B. (2005). A varying abstraction model for categorization, in: *Proc. 27th Ann. Conf. Cognit. Sci. Soc.*, pp. 2277–2282. Lawrence Erlbaum, Mahwah, NJ, USA.
- Wagemans, J. (1992). Perceptual use of nonaccidental properties, *Canadian J. Psychol.* **46**, 236–279.
- Wagemans, J. (1993). Skewed symmetry: a nonaccidental property used to perceive visual forms, *J. Exper. Psychol.: Human Percept. Perform.* **19**, 364–380.
- Wagemans, J., Van Gool, L. and d'Ydewalle, G. (1991). Detection of symmetry in tachistoscopically presented dot patterns: effects of multiple axes and skewing, *Percept. Psychophys.* **50**, 413–427.
- Zahn, C. and Roskie, R. (1972). Fourier descriptors for plane closed curves, *IEEE Trans. Computers* **C-21**, 269–281.
- Zusne, L. (1970). *Visual Perception of Form*. Academic Press, New York, USA.