

The impact of context on pattern category learning and representation

Martin Jüttner

*Neuroscience Research Institute, School of Life and Health Sciences—
Psychology, Aston University, Birmingham, UK*

Berthold Langguth and Ingo Rentschler

Institut für Medizinische Psychologie, Universität München, Germany

Context traditionally has been regarded in vision research as a determinant for the interpretation of sensory information on the basis of previously acquired knowledge. Here we propose a novel, complementary perspective by showing that context also specifically affects visual category learning. In two experiments involving sets of Compound Gabor patterns we explored how context, as given by the stimulus set to be learned, affects the internal representation of pattern categories. In Experiment 1, we changed the (local) context of the individual signal classes by changing the configuration of the learning set. In Experiment 2, we varied the (global) context of a fixed class configuration by changing the degree of signal accentuation. Generalization performance was assessed in terms of the ability to recognize contrast-inverted versions of the learning patterns. Both contextual variations yielded distinct effects on learning and generalization thus indicating a change in internal category representation. Computer simulations suggest that the latter is related to changes in the set of attributes underlying the production rules of the categories. The implications of these findings for phenomena of contrast (in)variance in visual perception are discussed.

Our perception of the world is highly dependent on contextual information (see Albright, 1995). However, context has remained a relatively vague concept in vision research despite its widely acknowledged importance. Early conceptions relate contextual effects to a form of “unconscious inference”. First proposed by von Helmholtz (1896) this idea reappeared, in a more generalized form, in Bruner’s (1957) notion of “readiness” in perception. The latter includes

Please address all correspondence to: M. Jüttner, Neuroscience Research Institute, School of Life and Health Sciences—Psychology, Aston University, Aston Triangle, Birmingham B4 7ET, UK. Email: m.juttner@aston.ac.uk

This study was supported by the Deutsche Forschungsgemeinschaft Grant Ju 230/5-1 to M.J. and Re 337/10-2, Re 337/12-1 to I.R.

expectations and subjects' knowledge and has been shown to influence a variety of tasks, such as the recognition threshold of words as a function of familiarity (Soloman & Postman, 1951) and predictability (Miller, Bruner, & Postman, 1954), or the critical exposure duration for the detection of anomalous versus regular playing cards (Bruner & Postman, 1949).

Yet the idea of a perceptual–cognitive continuum has not remained unquestioned. Pylyshyn (1999) distinguishes between contextual effects as purely “within-vision” or top-down effects (i.e., visual interpretations computed by early vision that affect visual interpretation) and effects of cognitive penetration, where the latter include influences from outside the visual system that may affect the content of visual perception. In contrast, in artificial intelligence and machine vision more pragmatic views prevail. Here contextual information has been used to develop robust edge-finding schemes that allow for a reliable extraction of features for object recognition and scene analysis (Clowes, 1971; Freuder, 1986). Eventually, such context-dependent influences became subsumed under the more general label “knowledge-based”; as such their importance has been emphasized both in computer vision (Grimson, 1990) and human vision (Riseman & Hanson, 1987).

Concerning the latter, more recent research has considered the role of context mainly with regard to visual selection and object recognition. Contextual information has been seen as a factor modulating the statistical pattern of saccadic eye movements during reading or while individuals view pictorial scenes. Both in reading (Morris, 1994) and picture scanning (de Graef, Christiaens, & d'Ydewalle, 1990) fixation durations are affected by semantic properties, i.e., contextual information provided by the sentence or the picture. Although the role of context in the selection of individual saccade targets during picture scanning is still unclear (see Loftus & Mackworth, 1978; but Henderson, Weeks, & Hollingworth, 1999; Krieger, Rentschler, Hauske, Schill, & Zetzsche, 2000) such information may be incorporated in an internal saliency map subserving the deployment of spatial attention and eye movements (Findlay & Walker, 1999; Henderson, 1992). In accordance with this notion, experiments on visual search demonstrate that the presence of a target can be cued by both the configuration and the shape of the surrounding distractors, thus leading to a reduction in search time (Chun & Jiang, 1998, 1999). On a more cognitive level, many studies have shown that object identification is facilitated when an object is semantically consistent rather than inconsistent with the scene in which it appears (e.g., Biederman, 1981; Biederman, Mezzanotte, & Rabinowitz, 1982; Boyce & Pollatsek, 1992; Friedman, 1979; Palmer, 1975). However, the level of processing where this contextual modulation of perception occurs is controversial. Context information has been suggested to facilitate perceptual encoding at some relatively early stage of visual processing (Biederman, 1972), it might affect the matching between the percept and a representation stored in long-term memory (Palmer, 1975; Ullman,

1996), or it might be mediated by mechanisms of response selection (Henderson & Hollingworth, 1999; Hollingworth & Henderson, 1998).

What the aforementioned approaches have in common is that they regard context as an important determinant of how *previously acquired* knowledge guides the (current) interpretation of sensory experience. In this paper we will pursue a complementary perspective, by showing that context also specifically affects *learning*, that is the acquisition of knowledge and the way in which such knowledge is mentally represented. For visual perception, such learning involves in particular the acquisition of object categories underlying object recognition (Bruner, 1957; Rosch, 1978). Furthermore, readdressing the context problem from the perspective of category learning has the advantage that context becomes a well-defined variable since it is given by the set of stimuli to be learned.

In the present study we explored, in two classification-learning experiments and by means of computer simulation, how manipulations of context affect category learning and the ability to generalize acquired class knowledge. Our paradigm involves the classification of Compound Gabor patterns (cf. Figure 1). Such grey-level patterns result from the superposition of two sinewave gratings, a fundamental plus its third harmonic, within a Gaussian aperture. Typically these patterns are specified by only two free parameters, the amplitude and phase of the third harmonic. Thus, the corresponding Fourier feature space represents a *continuum of visual shape* where each point uniquely specifies the appearance of a pattern and clusters of points of points are used to define pattern categories, or classes, to be learned by the subject (Jüttner & Rentschler, 1996; Rentschler, Jüttner, & Caelli, 1994).

Such a paradigm allows one to study visual concept formation on a relatively early level. On the one hand, Gabor signals have been regarded as an elementary stimulus in early visual processing (e.g., Rentschler & Caelli, 1990; Watson, Barlow, & Robson, 1983; Westheimer, 1998). On the other hand, such patterns are perceptually complex enough to stimulate category learning (Jüttner & Rentschler, 1996, 2000; Rentschler et al., 1994). The acquisition of pattern categories composed of Gabor patterns is completely under experimental control as such patterns are unfamiliar for naive subjects, thus excluding the problem of prior knowledge. The learning of such categories can be characterized as a form of hypothesis testing where subjects draw much upon the perceived pattern structure (Unzicker, Jüttner, & Rentschler, 1999). Categorization performance has been successfully modelled upon the assumption that subjects construct rules that “evidence” pattern classes in terms of perceived parts and their relations (Jüttner, Caelli, & Rentschler, 1997). The evidence-based systems (EBS) approach has been shown to provide a process model of categorization that allows prediction of generalization performance with regard to segmented (grey-level transformed) versions of the original learning patterns.

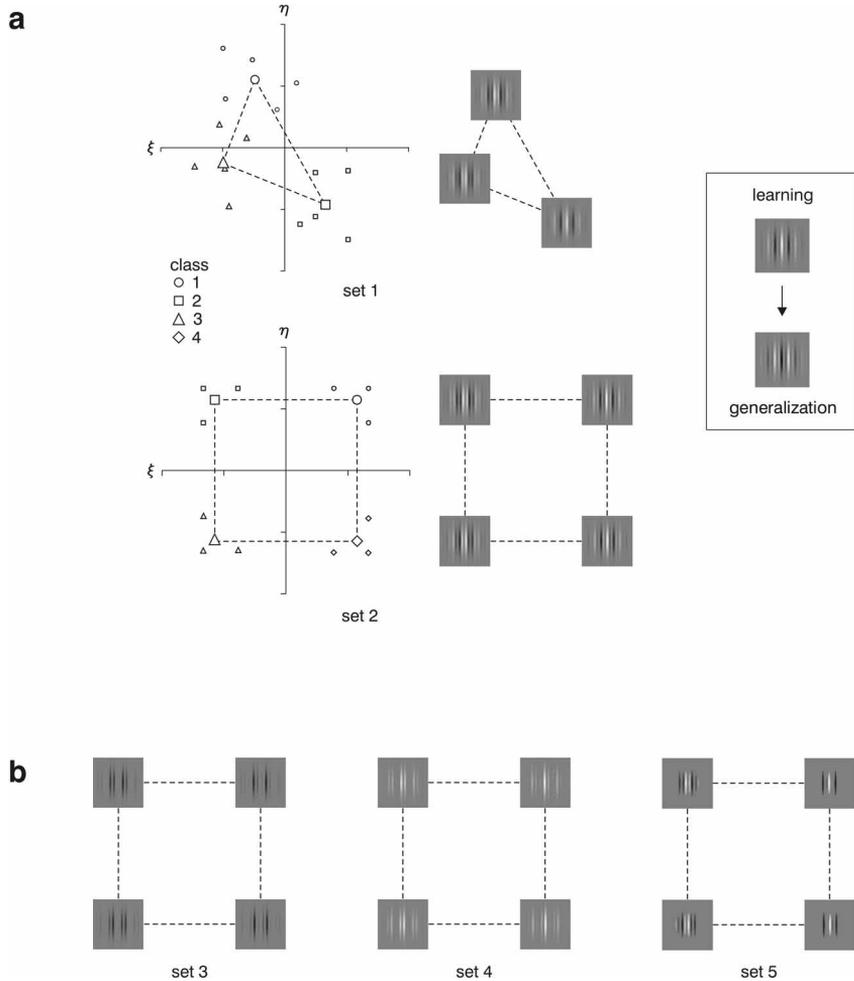


Figure 1. Definition of the five sets of grey-level patterns used in Experiments 1 and 2. As indicated in the inset for each learning set a corresponding set of test patterns was generated by inverting the contrast polarity of the patterns. (a) Sets of Compound Gabor signals used in Experiment 1. The patterns were defined in a two-dimensional (ξ, η) Fourier space. Scale: 1 unit = 20 cd/m^2 . Set 1 contained three clusters of five samples; set 2 contained four clusters of three samples. Each signal cluster defined one class to be learned by the subject. The large symbols connected by dashed lines denote the class means or prototypes. For illustration purposes these class prototypes are redrawn in their grey-level representation. (b) For Experiment 2, the signals of set 2 in Experiment 1 were degraded by replacing parts of the grey levels by mean luminance thus generating three further sets of learning signals. See text for further details.

The aim of the present study was to explore the effect of complementary manipulations of context on the internal representation of pattern categories. In Experiment 1, we changed the (local) context of individual pattern categories by changing the overall configuration of signal classes in the learning set, i.e., the number of signal classes set as defined within the generating Fourier feature space. In Experiment 2, we varied the (global) context of a fixed class configuration by changing the degree of signal accentuation. In both experiments, generalization performance was measured as the ability to transfer acquired class knowledge to contrast-inverted versions of the learning patterns. The underlying rationale was that a change in category representation induced by a change in context would affect the potential to generalize the acquired class knowledge. To corroborate our findings we then modelled the behavioural data in terms of evidence-based classification. Evidence-based classifiers allow to reconstruct combinations of nonrelational and relational attributes that provide potential solutions of a given classification problem. As these attributes can be regarded as the “signature” of the underlying conceptual representation we used the simulations to establish as to what extent these representations would change due to a change in learning context, thus providing further support for our behavioural results.

METHOD

Apparatus and materials

The experiments involved the classification of Compound Gabor patterns. Such grey-level patterns result from the superposition of two sinewave gratings, a fundamental plus its third harmonic, within a Gaussian aperture. Their intensity profile $G(x, y)$ was defined by:

$$G(x, y) = L_0 + \exp\left\{-\frac{1}{\alpha^2}(x^2 + y^2)\right\}(a \cos(2\pi f_0) + b \cos(2\pi 3 f_0 x + \phi)) \quad (1)$$

where L_0 determines the mean luminance, α the space constant of the Gaussian aperture, a the amplitude of the fundamental, b that of the third harmonic, and ϕ the phase angle of the latter. The signals were generated in a 128×128 8-bit pixel format with a linear grey-level-to-luminance function. The aperture parameter α was set to 32 pixels.

Signal variation was restricted to b and ϕ . This allowed the use of two-dimensional Fourier feature space with the Cartesian coordinates and $\xi = b \cos \phi$ and $\eta = b \sin \phi$. Thus, the (ξ, η) feature space provided a continuum of visual shape with each point uniquely specifying the appearance of a pattern. Within this continuum clusters of points were used to define the pattern classes to be learned by the subject. For Experiment 1, two configurations of learning patterns were defined within this feature space (Figure 1a): A 3-class configuration

defined by three clusters of five signals each (set 1), and a 4-class configuration composed of four clusters of three signals (set 2). For Experiment 2, three further sets of learning patterns were generated based on those of set 2 by removing image parts via a nonlinear threshold operation (Figure 1b). The samples of set 3 were obtained by removing all bright parts of the images, i.e., by setting all pixels with luminance values above L_0 to the level of the background (mean) luminance L_0 . Analogously, the patterns of set 4 were generated by removing all dark parts. The patterns of set 5 were generated by removing parts with intermediate grey-level values, i.e., by setting all pixels with luminance values in the interval $[L_0 - (a+b)/2, L_0 + (a+b)/2]$ to the level of L_0 . In addition, for each of the five sets of learning patterns a corresponding set of test patterns was generated by inverting the contrast polarity of the individual signals (see inset in Figure 1a).

The patterns were displayed on a computer-graphics display (EIZO F56) connected to a Pentium PC. Space average luminance was kept constant at 60 cd/m^2 . The stimulus patterns subtended 1.7° of visual angle at the viewing distance of 100 cm. The spatial frequency of the fundamental was set to 2.4 c/deg .

Subjects

In total, 25 paid observers participated. They gave their written informed consent to the study after the procedure had been explained to them. None of them had any prior experience with psychophysical experiments. Their ages ranged between 20 and 30 years; 13 were female, 12 were male. All had normal or corrected-to-normal vision. Each subject was randomly assigned to one of the five pattern sets, resulting in five groups of five observers each.

Procedure

The experiment was divided into two parts, a learning phase and a generalization test. The learning phase employed a supervised-learning schedule (see also Caelli, Rentschler, & Scheidler, 1987; Rentschler et al., 1994) and consisted of a variable number of learning units (Figure 2a). Each learning unit had two phases, training and (recognition) test. During the training phase, each pattern was shown three times in random order for 200 ms, followed by a number specifying the class to which the pattern belonged (as defined by the class assignments shown in Figure 1a). The class label was displayed for 1000 ms, with an interstimulus interval of 500 ms relative to the offset of the learning pattern. During the test phase of each learning unit, which served to monitor the learning status of the subject, the patterns were shown once in random order and classified by the subject by pressing the appropriate button on the computer keyboard. The series of learning units continued until the observer had achieved

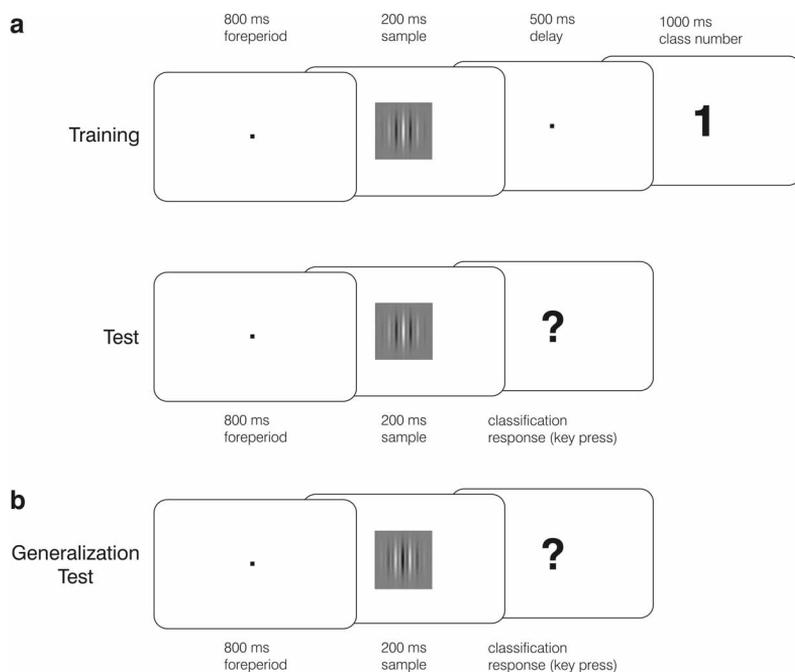


Figure 2. (a) Supervised learning schedule. The learning procedure consisted of a variable number of learning units. Each learning unit had two phases, training and (recognition) test. During the training phase, each pattern was shown three times in random order for 200 ms, followed by the corresponding class number. For the recognition test, the patterns were shown once in random order and classified by the subject. (b) During the generalization test, each test pattern was presented and classified 30 times in random order, using the same the timing parameters as in the recognition tests of the supervised learning procedure. Note that the generalization test used the contrast-inverted versions of each learning signal (cf. inset in Figure 1a).

the learning criterion of error-free classification (100% correct) in one recognition test.

Once the subjects had been trained to criterion they entered the second part of the experiment, the generalization test (Figure 2b). Here their ability was assessed to recognize contrast-inverted versions of the previously learned patterns. Each test pattern was presented and classified 30 times in random order. The timing parameters were the same as in the recognition tests during the preceding supervised learning.

Data analysis

Observer performance was assessed in terms of learning duration (i.e., the number of learning units required to reach the learning criterion) and the classification performance in the generalization test. Because the different number

of classes in set 1 (three classes) and sets 2–5 (four classes) implies different guessing rates the classification frequencies p_{gen} observed in the generalization test were transformed into a normalized generalization index κ defined by:

$$\kappa = \frac{p_{gen} - p_c}{1 - p_c} \quad (2)$$

where p_c denotes the guessing correction (being 0.25 in case of four classes and 0.33 in case of three classes).

To further explore the development of internal class concepts (in the learning phase of the experiment) and their application (in the generalization test) the confusion matrices were analysed by means of multidimensional scaling (MDS) techniques. Within the given set of n_s patterns for each pair of signals (i, j), $i, j = 1 \dots n_s$, a similarity measure $d(X_i, X_j)$ was derived using a chi-square metric:

$$d(X_i, X_j) = \sqrt{\sum_k \frac{(X_{ik} - E(X_{ik}))^2}{E(X_{ik})} + \sum_k \frac{(X_{jk} - E(X_{jk}))^2}{E(X_{jk})}} \quad (3)$$

where $E(X_{ik})$ and $E(X_{jk})$ refer to the expected values of the k -th component of the classification vectors X_i and X_j (i.e., the rows of the confusion matrix) of pattern i and j , respectively. The distance matrix was then analysed using nonmetric MDS (Kruskal, 1964; Shepard, 1962), which produced a two-dimensional configuration for the set of the n_s stimuli. As according to our paradigm the stimuli were labelled by their class labels only the class-specific means, or prototypes, reflect the similarity structure between classes. This allows a visualization of the *conceptual space* spanned by the three classes in terms of their prototype configuration.

In order to compare learning and generalization data across subjects we used weighted MDS (Carroll & Chang, 1970), which allows a parsimonious description of similarity data provided by multiple distance matrices. This procedure yields a so-called group stimulus space and a set of input-matrix specific (individual) weights. The matrix-specific stimulus space is obtained by applying the weight factors to the dimensions of the group stimulus space according to:

$$x_{ij}^{(k)} = w_{kj}^{1/2} x_{ij} \quad (4)$$

where x_{ij} denotes the coordinate of stimulus i on dimension j in group stimulus space, w_{kj} the weight of dimension j for matrix k and $x_{ij}^{(k)}$ the coordinate of x_{ij} in the stimulus space specific for matrix k . The corresponding conceptual space may be visualized by computing the configuration of the class-specific means, or prototypes, of the $x_{ij}^{(k)}$.

RESULTS

Experiment 1

In Experiment 1 we compared category learning and generalization with respect to two different class configurations (Figure 1a): A 3-class configuration defined by three clusters of five signals each (set 1), and a 4-class configuration composed of four clusters of three signals (set 2). In the first part of the experiment subjects were trained in the supervised-learning schedule (Figure 2a) to correctly categorize all signals of one of the two learning sets, i.e., assign them to their corresponding class label 1, 2, or 3 (set 1) and 1, 2, 3, or 4 (set 2), respectively. After having reached the learning criterion the observers entered the second part of the experiment. Here their ability to generalize was tested by presenting them contrast-inverted versions of the previously learned patterns, which again had to be categorized according to their class label (Figure 2b).

Figure 3 shows for the two sets of learning patterns learning duration and generalization performance. The different learning context, as expressed by the two signal configurations, had significant effects on both performance indices that were evaluated by between-subject comparisons using *signal set* as grouping variable. The 3-class configuration (set 1) was learned significantly faster than the 4-class configuration (set 2), $t(8) = 3.02$, $p < .05$. On average the subjects needed 12.1 learning units to learn set 1 but required 28.4 learning units to learn set 2. Despite the fact that the subjects of both groups are trained to 100% correct at the end of learning procedure they differed markedly in their ability to generalize the acquired categorical knowledge to contrast inversion. Generalization performance for the group assigned to set 1 is significantly

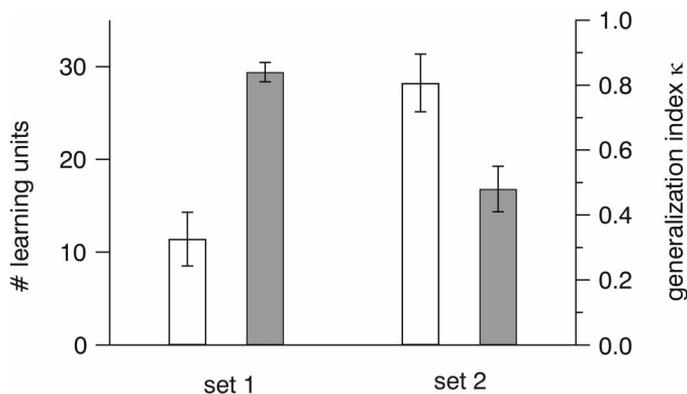


Figure 3. Mean learning duration (bright bars) and generalization performance (dark bars) in Experiment 1. Generalization was measured as the ability to classify the contrast-inverted versions of the learning patterns (see inset in Figure 1a).

higher than for the group assigned to set 2, $t(8) = 3.43$, $p < .05$. The mean of the generalization index κ reaches .81 for the 3-class configuration but drops to .45 for the 4-class configuration.

To further explore the acquisition of the pattern categories the confusion error data was analysed in terms of multidimensional scaling. When applied to the individual confusion error matrices of each learning unit this analysis yields a series of *tomograms*, which demonstrate the temporal evolution of conceptual space during learning. Figure 4a (right) shows an example of the tomogram series of one subject learning the 4-class configuration. Compared to the symmetric, square-like configuration in the defining Fourier feature space (cf. Figure 4 top right) the internal class structure during learning appears highly distorted. In general, the distances between the two pairs of classes (1,2) and (3,4) are much smaller than between the class pairings (1,4) and (2,3) indicating that the former are perceived as more similar (and therefore are confounded more frequently) than the latter. The distortion of conceptual space occasionally produces “twisted” configurations (for instance for $n = 6$ and $n = 14$ in Figure 4a right) where class pairings along the two “diagonals”—(1,3) and (2,4)—are perceived more similar than those along the “edges”—(1,2) and (3,4)—despite the fact that the latter are closer in the defining feature space.

In addition to the configurational variations during learning a more systematic distortion, or bias, is present in the data. This bias becomes most conspicuous in the conceptual space representation of the temporal average, where the confusion matrices have been accumulated across the entire sequence of learning units. Notably, an MDS analysis of the generalization data reveals that the same type of bias in the representation of the contrast-inverted stimuli (Figure 4b right). Thus confusion errors during generalization may be related to the conceptual space that was acquired during the learning phase of the experiment. For the 3-class configuration (Figure 4b left) the conceptual space does not show such pronounced bias for learning and generalization, although configurational variations during learning also do occur (Figure 4a left).

The data presented in Figure 4 refers to two individual subjects. In order to make comparisons across groups of observers weighted MDS solutions for the two groups assigned to sets 1 and 2 were derived. For each subject the confusion matrices of the cumulated learning data and the generalization test entered the analysis, and the corresponding individual conceptual-space representations were derived according to equation (4). Figure 5 (right) shows that for all subjects learning pattern set 2 the symmetry of the 4-class configuration in the defining feature space is broken in conceptual space, producing an elongated arrangement of the four class prototypes. The same type of bias is reflected in the generalization data confirming the initial observation in Figure 4b. In contrast, no such bias is present in the conceptual space representations that result from the learning and generalization of the 3-class configuration in set 1 (Figure 5 left).

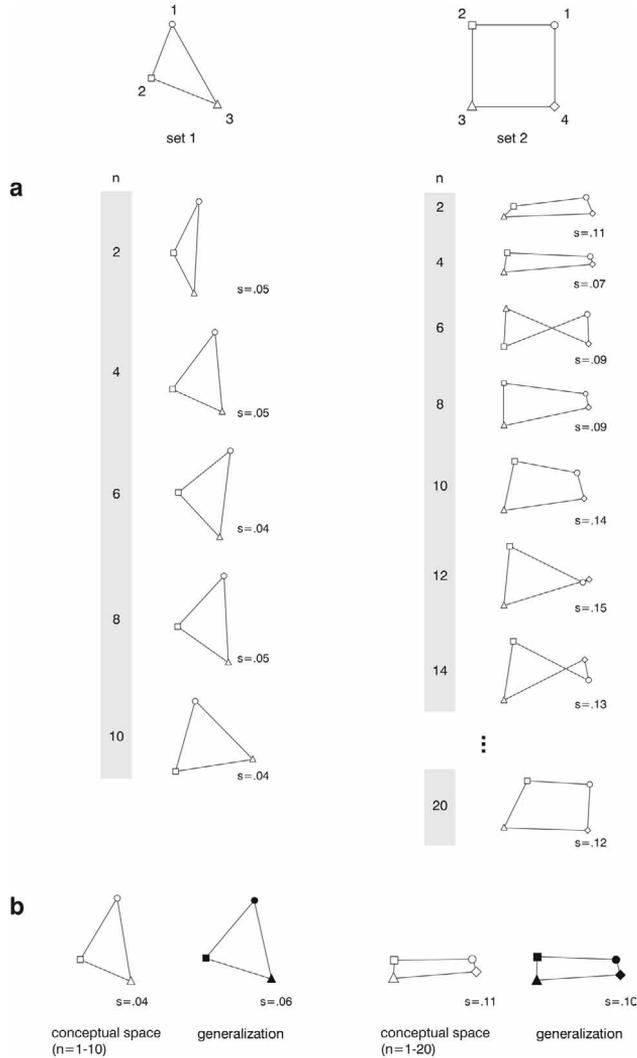


Figure 4. (a) Temporal evolution of the conceptual spaces in two subjects learning pattern set 1 (left, subject MH) and set 2 (right, subject MB). The learning tomograms are indexed by (learning-) unit number and were derived independently by multidimensional scaling (MDS) from the confusion error data. The variable s denotes the stress values of the MDS solutions. For sake of brevity only a subset of tomograms is shown. Note that for set 2 the internal class structure appears systematically distorted compared to the symmetric, square-like configuration in the defining Fourier feature space shown on top (class numbers as indicated). Despite configurational variations during learning no such distortion is present in the tomogram series of the observer learning pattern set 1. (b) Temporal average of the conceptual space representations shown in (A). Note that for set 2 the bias evident in the temporal average is reflected in the configuration derived from the generalization data. For set 1 the representation appears largely biasfree.

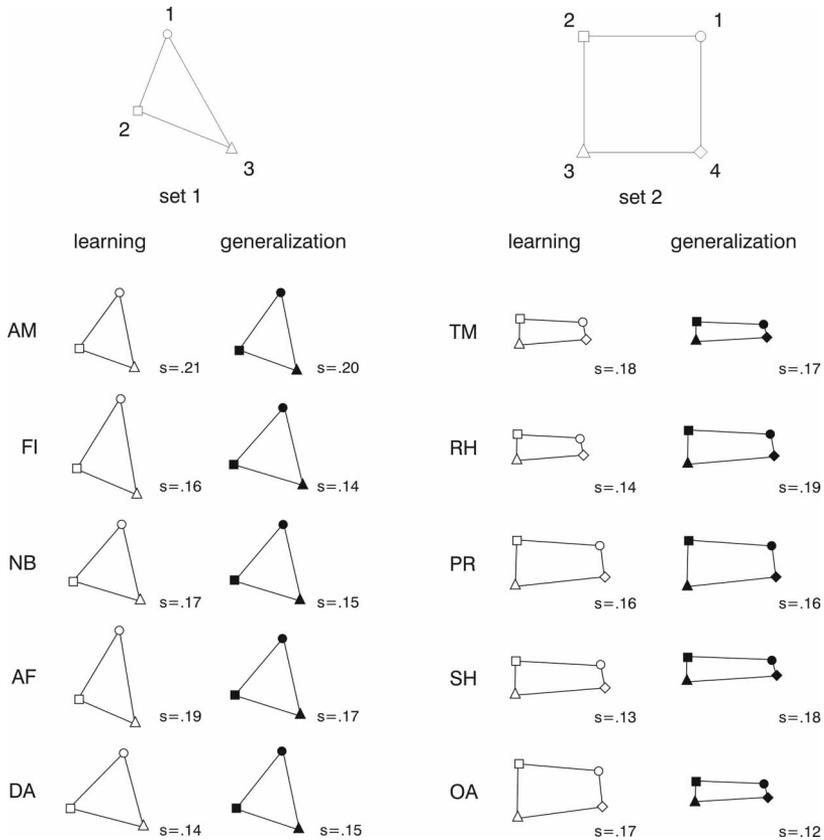


Figure 5. Conceptual spaces derived from the averaged learning and the generalization data. These representations were derived by conjointly analysing the confusion error data of all subjects assigned to sets 1 and 2, respectively. For subjects learning pattern set 2 the symmetry of the 4-class configuration appears to be systematically broken in the internal representation, both during learning and generalization. In contrast, no such bias is present in the conceptual space representations for set 1.

Experiment 2

In Experiment 1 we had changed the context *locally*, i.e., for the individual signal class, by changing the overall configuration of the learning set from three classes (set 1) to four classes (set 2). In contrast, Experiment 2 involved a *global* change of context, i.e., a manipulation of context for a complete signal configuration. The latter was achieved by modifying the degree of pattern accentuation in the same way across all signals (i.e., globally) of a fixed class configuration. Based on the 4-class configuration (set 2) of Experiment 1, three further sets of learning patterns were generated by removing systematically

image parts via a nonlinear threshold operation (Figure 1b). The thresholding led to a degraded, though more accentuated, physical stimulus representation than that of the original samples. Furthermore, for each of the three new stimulus sets a set of test patterns was generated by inverting contrast polarity. With these new sets of learning and test stimuli we trained three groups of observers to criterion, using the same supervised-learning procedure as in Experiment 1 (cf. Figure 2a) and the same class assignment as with set 2. After the subjects had reached the learning criterion they were tested for their ability to generalize to contrast inversion (cf. Figure 2b).

Figure 6 summarizes learning duration and generalization performance in Experiment 2. For illustration purposes the corresponding data for the original set 2 (Experiment 1) is also replotted. Between-group comparisons in terms of a one-way analysis of variance with *signal set* as factor variable showed that the accentuation brought about by the thresholding had a significant effect on learning duration, $F(3, 16) < 14.11$, $p < .001$, and generalization, $F(3, 16) = 14.68$, $p < .001$. Specifically, learning duration appears to decrease from set 2 to set 5. Post hoc Bonferroni multiple-comparisons yielded significant ($p < .05$) differences between the means of set 2 and those of sets 3, 4, and 5. Of the remaining pairings, only the difference between set 3 and set 5 was approaching significance ($p < .1$). Conversely, generalization improved from about .45 (set 4) to .9 (set 5). Post-hoc comparisons revealed that the improvement is mainly due to a significant difference ($p < .05$) between the mean of set 2 relative to the means of sets 3, 4, 5, and a difference between sets 3 and 5 that is approaching significance ($p < .1$).

To explore the effect of the global change of context on the conceptual space weighted MDS solutions were computed for the combined data of the stimulus sets 2–5. For each subject the confusion error matrices of the cumulated learning data and the generalization data entered that analysis. For each subject individual conceptual-space representations were derived. The group averages of those representations are depicted in Figure 7. The plot demonstrates that the

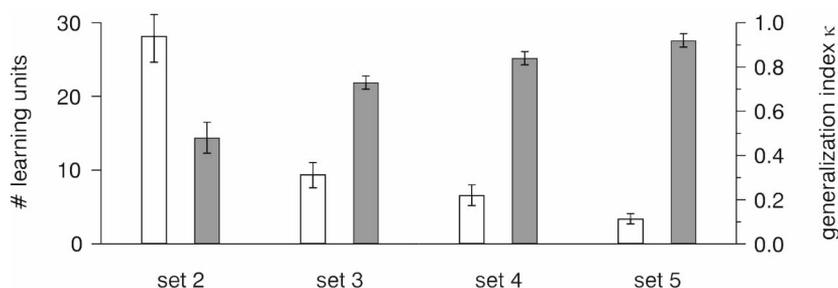


Figure 6. Mean learning duration (bright bars) and generalization performance (dark bars) for pattern sets 3–5 in Experiment 2. For comparison the data of set 2 (Experiment 1) is also included.

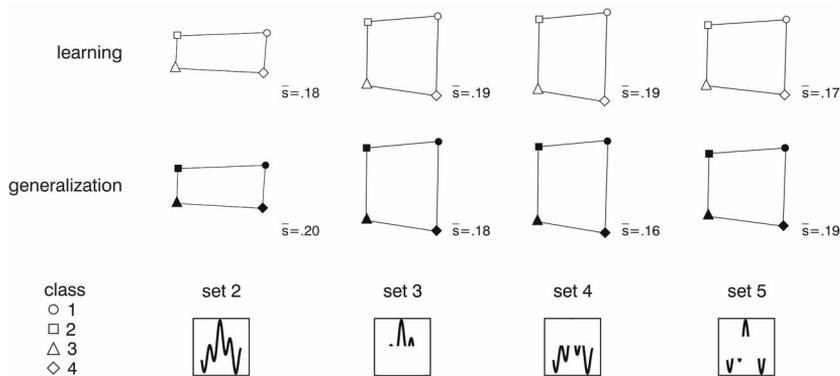


Figure 7. Group averages of the conceptual space representations for the learning and generalization of pattern sets 3–5 in Experiment 2. For comparison the corresponding data for set 2 (Experiment 1) is also included. Note that the broken symmetry apparent in the conceptual space for set 2 is restored by the context change induced in sets 3–5.

increased learning speed induced by the global context change is associated with a distinct change in the underlying conceptual space. The two pairs of classes (1,2) and (3,4) are now perceived as equally similar as the class pairing (1,3) and (3,4). By restoring the symmetry of the 4-class configuration the conceptual space becomes almost biasfree. The bias-free representation is preserved in the generalization phase. Since the prototype distances relate to the frequency of confusion errors this also accounts for the much higher generalization performance observed for the modified pattern sets.

Simulations results

The variations of learning context introduced in Experiment 1 and Experiment 2 yielded distinct effects on both learning and generalization performance. To gain further insight into the nature of the underlying mental representations we modelled human performance in terms of evidence-based pattern classification. According to the EBS approach to pattern recognition complex objects are encoded in terms of parts and their relations. Originally developed in the area of machine learning and computer vision (Caelli & Dreier, 1994; Jain & Hoffman, 1988), we have previously demonstrated that EBS also provides the framework for a process model of human perceptual categorization and generalization (Jüttner et al., 1997).

In an evidence-based classification system a given pattern is first segmented into its component parts (Figure 8). Each part is characterized by a set of part-specific, or unary, attributes (e.g., size, luminance, area), and each pair of parts is described by a set of relational, or binary, attributes (e.g., distance, angles, contrast). Thus, each part may be formally represented as a vector in a feature

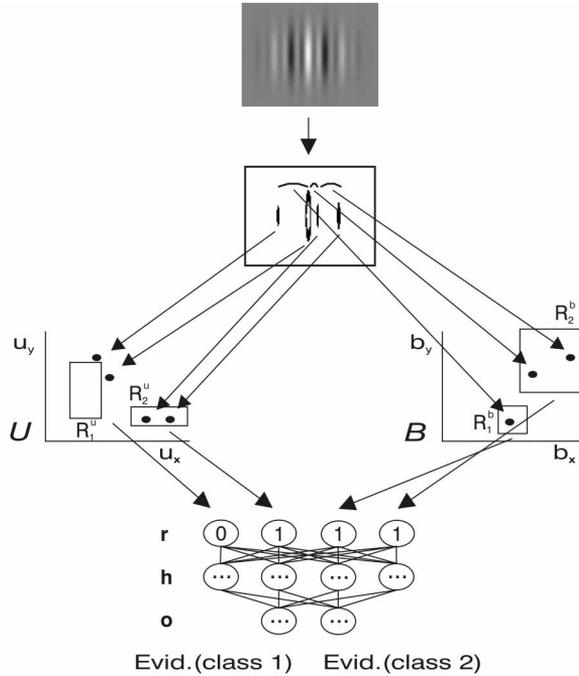


Figure 8. Sample illustration of the representational levels involved in an evidence-based classifier (after Jüttner, Caelli, & Rentschler, 1997). For convenience, only two unary and binary attributes, four rules and two classes are considered. The input image is first decomposed into its components. Each part and each pair of adjacent parts then is described by a vector of unary (part-specific) and binary (relational) attributes, respectively. Within the corresponding feature spaces the attribute vectors trigger rules providing class-specific evidences. The evidence values are implicitly represented by the synaptic weights of a three-layer neural net that is trained during the learning stage of the system. The activity of the output nodes provides a measure of the total evidence concerning the class alternatives. See main text for further details.

space spanned by the various unary attributes, and each pair of parts by a vector in a binary feature space. Within each feature space regions are defined that act as activation regions for rules. The rule regions are derived by clustering the feature spaces during the training phase of the system. During runtime, an attribute vector falling inside such a region will activate the corresponding rule (state 1); otherwise the rule remains inactive (state 0). A given object is represented by a rule activation vector, the components of which are assigned to the activation states of the individual rules. The activation of a given rule provides a certain amount of evidence for the class membership of the input object. The assignment of the evidence weights to the rules and their combinations is achieved within a neural network. Here each input node corresponds to a rule, each output node to a class, and there is one hidden layer. The

relative activity of an output node provides a measure of the accumulated class-specific evidence. This activity may be probabilistically interpreted and related to a classification frequency.

An evidence-based classifier first has to be trained with patterns with known class membership. Given a reservoir of unary and binary attributes an attempt is made, for each attribute combination, to train the neural network via the back-propagation algorithm, in order to identify those attribute combinations that successfully allow to separate the classes. More specifically, for the simulations we first constrained the system parameters according to a range that had proved optimal in previous work using the same type of stimulus material (see Jüttner et al., 1997): The segmentation stage used a region-analysis technique that was based on partitioning the image according to connected grey-level regions yielding 3–5 parts per image. The rule-generation stage employed K-means clustering procedure producing a set of 10–14 rules. Furthermore, the classifier was supplied with a reservoir of four unary attributes (position, luminance, aspect ratio, and size) and three binary attributes (distance, relative size, contrast). We then tested for which attribute combinations the training of the system converged, i.e., the system successfully learned to distinguish between the classes.

When used as a framework for cognitive modelling, EBS describes category learning as a successive testing of working hypotheses. Each working hypothesis corresponds to the selection of a subset of attributes that define a reference system for describing parts and their relations. Once chosen, the elaboration of such a working hypothesis will include the formation of rules and the tuning of evidence weights. Eventually, the elaboration process either results in a successful categorization, or the current working hypothesis is rejected and replaced by a different one.

Each subset of attended attributes may be regarded as a state within a search space of possible working hypotheses defined by the set of all possible combinations of unary and binary attributes. Learning speed is determined by the time required to find a solution within that search space, i.e., a set of attributes, which allows successful distinction between classes. Under the assumptions that the search is exhaustive and that the time needed to evaluate each working hypothesis is constant, learning speed should be directly proportional to N_{FS} , the number of EBS solutions within the search space, or conversely, learning duration should be proportional to $1/N_{FS}$. In that sense any variation of context that affects learning difficulty (i.e., the number of EBS solutions) should be reflected in behaviour (i.e., the observed learning duration).

The model-predicted learning durations for the five sets of patterns are summarized in the dark bars of Figure 9a. A comparison with the behavioural data shows that not only the ranking order of the empirical learning durations is preserved in the simulated values, but that also the ratios of learning durations are well approximated by the latter. Thus, the model provides a unified account

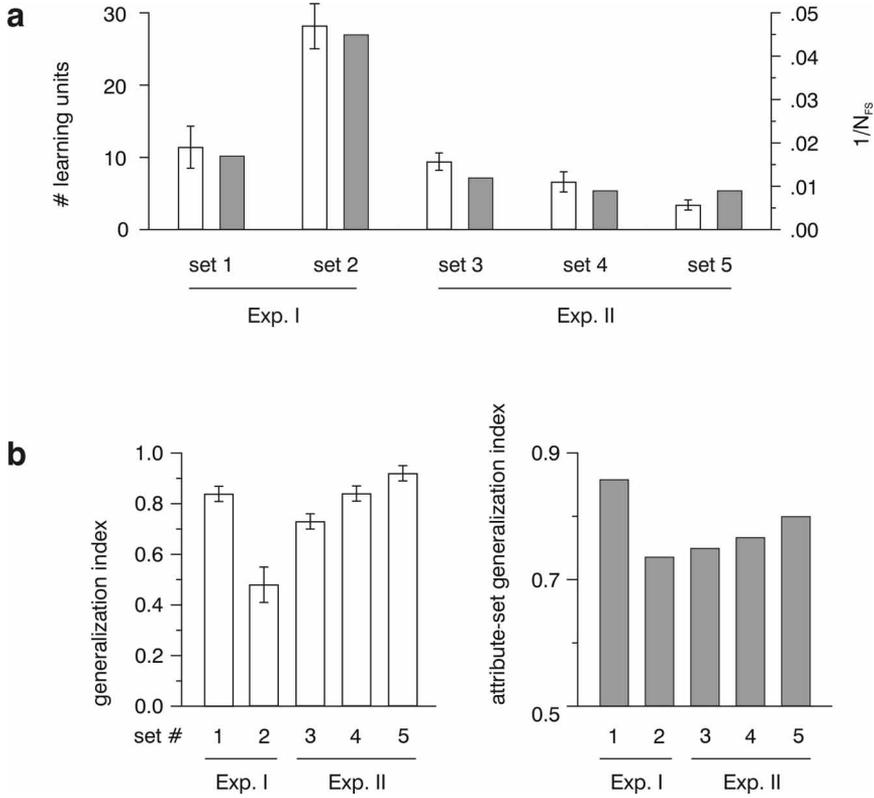


Figure 9. (a) EBS-simulated learning durations ($1/NFS$, where NFS is the number of EBS solutions in the learning space, dark bars) and observed group means (bright bars, cf. Figures 3 and 6) for learning set 1–5 in Experiments 1 and 2. (b) Empirical generalization performance (bright bars, cf. Figure 3) and EBS attribute-generalization indices (dark bars) for set 1–5 in Experiments 1 and 2.

for context effects induced by very different experimental manipulations—the alteration of local class configuration (Experiment 1) and a global variation concerning the degree of stimulus accentuation (Experiment 2).

The analysis in terms of EBS also allows predictions concerning the difficulty with which acquired class knowledge may be generalized. The set of evaluated attributes providing the “signature” of each EBS solution constrains generalization to contrast inversion, as some attributes (luminance and contrast) depend on contrast polarity, whereas others (such as size or aspect ratio) do not. As an index of the potential for generalization we computed the relative weight of contrast-independent attributes, as given by their relative frequency within the EBS solutions. As illustrated in Figure 9b the ranking of this attribute generalization index mirrors the ranking in the observed generalization

performance. Such a result suggests that the contextual changes performed in Experiment 1 and 2 not only have an impact on learning difficulty as reflected in learning duration. Rather they also affect the set of evaluated attributes thus pointing at a context-induced change in the internal representation of pattern categories.

DISCUSSION

In this study we have shown that changes in context, as given by the stimulus set to be learned, may distinctly affect the acquisition of knowledge on pattern categories, i.e., the way in which such categories are internally represented. We demonstrated such context effects by two complementary manipulations. In Experiment 1 we changed the learning context by an alteration of the local class configuration, in terms of the number of classes included in the learning set, from three classes (set 1) to four classes (set 2). We found a distinct advantage of the 3-class configuration relative to the 4-class configuration in terms of learning speed and generalization. This result seems remarkable given that the 3-class configuration involved a larger number of signals in the learning set (15 vs. 12) and that the signal clusters defining the classes were less well separated in the defining feature space in terms of the ratio between interclass distance and intraclass variance (cf. the relative spreads of the signal classes for sets 1 and 2 in Figure 1a). We have shown in previous work that *increasing* the class variance adds to the complexity of the categorization task and leads to an *increase* in learning duration (Jüttner & Rentschler, 2000). Nevertheless, these potential disadvantages of the 3-class configuration relative to the 4-class configuration were clearly outweighed by the contextual facilitation observed for the former relative to the latter, thus demonstrating the prevalence of context effects in pattern category learning.

That the greater learning difficulty found for the 4-class configuration (set 2) in Experiment 1 was not caused by particular structural properties at the level of the individual patterns was demonstrated in Experiment 2, where we changed the degree of stimulus accentuation by applying various types of thresholding operations on the grey-level spectrum of the patterns. The thresholding allowed us to selectively highlight certain components within the original patterns (for example the dark parts in set 3) without distorting their overall spatial configuration. Because a given transformation was applied to all patterns in the same way, irrespective of their individual class assignment, it represents a global contextual manipulation with regard to the original pattern set. This global change of context significantly reduced learning time and distinctly improved generalization performance, thus corroborating and complementing the findings of Experiment 1. For both experiments, the context dependency of internal category representations was demonstrated on two levels: On a behavioural level by measuring the degree of generalization to contrast reversal; and on a

computational level by means of simulations, the results of which indicate that different contexts entail category representations with different attribute signatures.

“Context” in the present investigation was given by the set of stimuli to be learned. This interpretation implies a notion of context that is complementary to—and in some respect more general than—that prevailing in previous work. As outlined in the introduction, context traditionally has been regarded in vision research as a constraint within the space–time continuum in which the occurrence of an object is embedded. This notion has been made most explicit by Chun and Jiang (1998, 1999) in a series of experiments on visual search, demonstrating that contextual cues provided by the distractor items can lead to a significant reduction in search time. Contextual cueing effects have been shown to be mediated either by spatial (configurational) information, distractor shape, or temporal contingency (distractor trajectories) in relation to the search target. According to Chun and Jiang such contextual information is learned implicitly via the statistical covariation between distractor and target items as a consequence of the repeated pairing of those stimuli. Whereas contextual cueing paradigms focus on the *learning of context given some object*, the current investigation addresses the complementary issue, i.e., the *learning of objects given their context*. In addition, a more general perspective is adopted by referring to objects at the level of their categorical identity. As categories in turn are defined within a continuum of appropriate attribute dimensions the notion of context is extended from that of a configuration within the perceptual space–time domain to that of a configuration defined by the set of learning stimuli in their generating feature space.

Similar notions of context have been proposed in the classification literature in various forms. According to the so-called context model (Medin & Schaffer, 1978) and its predecessors (for a review see Estes, 1994) the probability of a given stimulus to be assigned to a particular category is determined by the relative similarity between the stimulus and the individual exemplars of that category and its similarity to the exemplars of all alternative categories. Hence, the set of stimuli defining the categories acts as context for categorization. Various extensions of this type of model exist, for example to relate identification and categorization (Nosofsky, 1986) or to capture the time course of perceptual categorization (Lamberts, 1998). Category learning has been explored particularly within connectionist context models. According to the ALCOVE model (Kruschke, 1992), categorical knowledge is acquired by storing individual exemplars in memory, and by forming associations between these exemplars and the categories to be learned. Formally, ALCOVE represents each exemplar as a point in a multidimensional psychological space, and a mechanism of selective attention is included that allows a weighting of the dimensions of that space according to their diagnostic value. However, due to the very nature of the concept “psychological space”, the relation between

physical features and their internal representation remains largely unspecified (Shepard, 1987).

In contrast, EBS provides an *explicit* linkage between physical and internal representation. Image segmentation, attribute extraction, and rule generation are entirely defined with respect to the physical image domain. Furthermore, structural information is preserved in the representational format by describing patterns in terms of components and their unary (part-specific) and binary (part-relational) attributes. Previous work has demonstrated that the EBS approach provides a unified account for both category learning and generalization of Compound Gabor patterns (Jüttner et al., 1997). One notable feature of category learning with this type of stimulus is the occurrence and reoccurrence of distinct patterns of confusion errors, which in the conceptual space representation manifest as fixed configurations, or cognitive stereotypes (Jüttner & Rentschler, 1996; Unzicker, Jüttner, & Rentschler, 1999). Within the EBS framework such hypotheses correspond to the selection of particular sets of attributes defining the dimensions of the unary and binary feature spaces. A hypothesis may then be elaborated, by forming rules and by adjusting the evidence weights assigned to each rule and each rule combination. If successful the process of hypothesis elaboration will result in a perfect separation of the stimulus classes, otherwise the current working hypothesis has to be discarded and replaced by a different one. In the former case, i.e., upon completion of the learning task, the set of evaluated physical attributes will constrain the potential for generalization, thus becoming the signature of the underlying conceptual representation.

In the present study we explored the effect of context on category representation within a “minimum-classification paradigm” (Jüttner & Rentschler, 2000), which deliberately was restricted to three or four classes defined within a continuum of two independent feature dimensions (i.e., amplitude and phase in the generating Fourier space). Although more elaborate forms of category learning might involve additional mechanisms for the generation and evaluation of novel features (Schyns, Goldstone, & Thibaut, 1998), our results have implications for complex recognition tasks as well. The specific choice of using generalization to contrast inversion as an indicator for context dependency allows us to relate our results to the well-known phenomenon of contrast (in)variance in visual perception: Characters and simple geometric patterns are readily recognized even when seen as negative (i.e., contrast inverted) images, whereas it is much more difficult to recognize a contrast-inverted face (Figure 10). The latter difficulty is not related to the presence of a grey-level spectrum as it persists in lithographic, two-tone face representations (Phillips, 1972).

Previous explanations of this phenomenon have regarded the spectral composition of images as crucial (Hayes, Morrone, & Burr, 1986). Accordingly, it has been proposed that the recognition of characters and geometric patterns critically depends on high spatial frequency information, whereas low spatial frequencies are important for the identification of faces. As lower spatial fre-

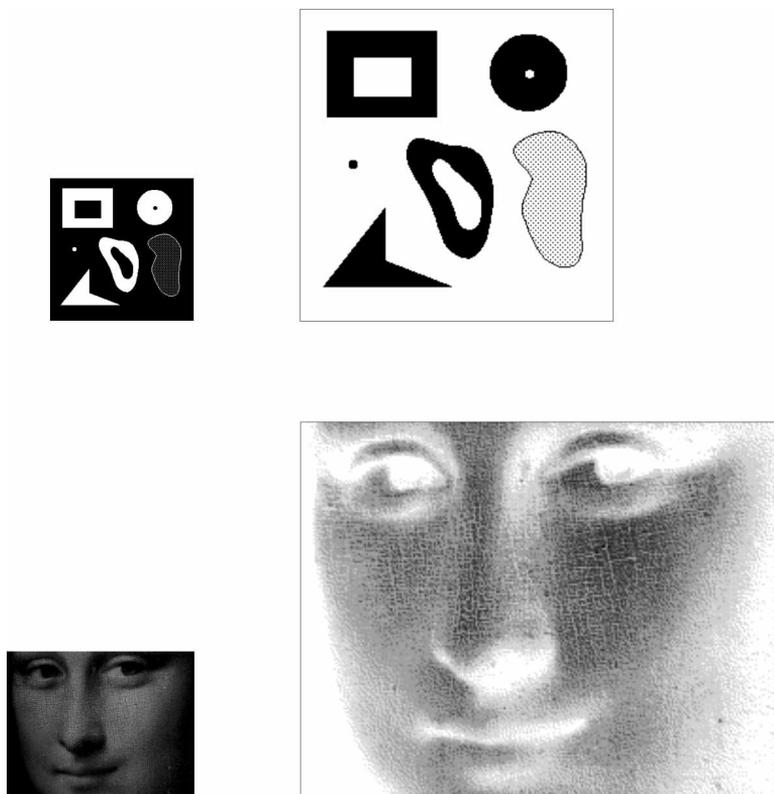


Figure 10. Examples of contrast (in)variance in visual perception. Characters and simple geometric forms are easily recognized when seen as contrast inverted images whereas the recognition of contrast-inverted portraits is exceedingly difficult.

quencies convey information about shading, contrast reversal might impair perception by disrupting the formation of a three-dimensional representation of the surface geometry of the face (Johnston, Hill, & Carman, 1992; Kemp, McManus, & Pigott, 1990).

However, the extent as to which face recognition relies on encoding 3-D surface properties has been questioned in more recent studies. Liu, Collin, and Chaudhuri (2000) found that stereo information failed to reduce shading effects and concluded that it is the 2-D shading pattern *per se* rather than (3-D) shape-from-shading information that contributes to face recognition. The impact of contrast reversal on recognition also depends on learning. Liu and Chaudhuri (1997) compared conditions in which faces were learned and tested as either positive or negative images. They observed a decline in performance in case of a contrast incongruity between learning and testing (i.e., when learning and

testing employed different contrast polarities) but there was no difference between the condition where subjects learned the positives and were tested with the negatives and the complementary condition (learning the negatives and testing the positives).

The effect of contrast reversal is similar to a change in lighting direction from above to below. Because face stimuli are typically encountered (learned) under conditions where illumination is from above shading does provide diagnostic information for the category “face” and might therefore be retained in the internal representation of that category. The necessity to distinguish between stimulus representation and category representation is illustrated by the results of Experiment I, which demonstrate that pattern sets with identical power spectrum may yield very different degrees in contrast-invariance (cf. Figure 3a). The results of the EBS simulations suggest that it is the relative proportion of contrast-invariant attributes that determines how well class concepts relying on these attributes may be generalized to contrast inversion. As a consequence, invariance to contrast is not a dichotomous property linked to certain physical features of the individual stimulus. Rather it depends on the way in which the category to which the stimulus belongs is internally represented.

Some support for the idea that invariance to contrast is an emerging property of category learning comes from neurophysiological research including single-cell recordings in monkeys and functional magnetic resonance (fMRI) imaging in humans. Ito, Jujita, Tamura, and Tanaka (1994) studied the effect of contrast reversal on the response of cells in the anterior inferotemporal (IT) cortex of the macaque monkey. Cells in the anterior IT, which represents the final stage of the ventral visual pathway, have been shown to respond to complex stimulus features in complex images, with some cells responding optimally to moderately complex shapes, others to combinations of shape and colour, or of shape and texture (for a review, see Tanaka, 2000). Ito et al. found both cells that were sensitive to contrast polarity, i.e., where contrast reversal led to a significant response reduction, and cells that were largely tolerant towards such a manipulation. This suggests that the extent of contrast (in)variance in perception may depend on the appropriate integration of activities of different groups of cells at a relatively high level in the processing hierarchy. Further converging evidence has been provided by a recent fMRI study involving human subjects. Avidan, Harel, Hendler, Ben-Bashat, Zohary, and Malach (2002) measured contrast response within a sequence of visual areas of the ventral processing stream, from area V1 to the lateral occipital complex (LOC) in the occipital-temporal cortex. Along this path they found a gradual increasing degree of contrast invariance both for face and object images, suggesting a hierarchical and gradual nature for the transition from early retinotopic to high-order areas in the build-up of abstract object representations.

In conclusion, contrast invariance might become manifest in a gradual, context-dependent way during the ontogenesis of internal representations for

object categories. According to this view contrast information will only be evaluated if necessary for a given (categorization) task, to an extent that is specified by the context provided by the set of stimuli to be learned.

REFERENCES

- Albright, T. D. (1995). "My most true mind thus makes mine eye untrue". *Trends in Neuroscience*, *18*, 331–333.
- Avidan, G., Harel, M., Hendler, T., Ben-Bashat, D., Zohary, E., & Malach, R. (2002). Contrast sensitivity in human visual areas and its relationship to object recognition. *Journal of Neurophysiology*, *87*, 3102–3116.
- Biederman, I. (1972). Perceiving real-world scenes. *Science*, *177*, 77–80.
- Biederman, I. (1981). On the semantics of a glance at a scene. In M. Kubovy & R. J. Pomerantz (Eds.), *Perceptual organization* (pp. 213–253). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, *14*, 143–177.
- Boyce, S. J., & Pollatsek, A. (1992). Identification of objects in scenes—the role of scene background in object naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 531–543.
- Bruner, J. (1957). On perceptual readiness. *Psychological Review*, *64*, 123–152.
- Bruner, J., & Postman, L. (1949). On the perception of incongruity: A paradigm. *Journal of Personality*, *18*, 206–223.
- Caelli, T., & Dreier, A. (1994). Variations on the evidence-based object recognition theme. *Pattern Recognition*, *27*, 185–204.
- Caelli, T., Rentschler, I., & Scheidler, W. (1987). Visual pattern recognition in humans: I. Evidence for adaptive filtering. *Biological Cybernetics*, *57*, 233–240.
- Carroll, J. D., & Chang, J. J. (1970). Analysis of individual differences in multidimensional scaling via an n-way generalization of "Eckart-Young" decomposition. *Psychometrika*, *35*, 238–319.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*, 28–71.
- Chun, M. M., & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, *10*, 360–365.
- Clowes, M. B. (1971). On seeing things. *Artificial Intelligence*, *2*, 79–116.
- De Graef, P., Christiaens D., & d'Ydewalle, G. (1990). Perceptual effects of scene context on object identification. *Psychological Research* *52*, 317–329.
- Estes, W. K. (1994). *Classification and cognition*. New York: Oxford University Press.
- Findlay, J. M., & Walker, R. (1999). A model of saccade generation based on parallel processing and competitive inhibition. *Behavioral and Brain Sciences*, *22*, 661–721.
- Freuder, E. C. (1986). Knowledge-mediated perception. In H. C. Nusbaum & E. C. Schwab (Eds.), *Pattern recognition by humans and machines: Visual perception* (pp. 219–236). Orlando, FL: Academic Press.
- Friedman, A. (1979). Framing pictures: The role of knowledge in automatized encoding and memory for gist. *Journal of Experimental Psychology: General*, *108*, 316–355.
- Grimson, W. (1990). The combinatorics of object recognition in cluttered environments using constrained search. *Artificial Intelligence*, *44*, 121–165.
- Hayes, A., Morrone, M. C., & Burr, D. C. (1986). Recognition of positive and negative bandpass-filtered images. *Perception*, *15*, 595–602.
- Henderson, J. M. (1992). Visual attention and eye movement control during reading and picture viewing. In R. Rayner (Ed.), *Eye movements and visual cognition: Scene perception and reading* (pp. 260–283). New York: Springer.

- Henderson, J. M., & Hollingworth, A. (1999). High-level scene perception. *Annual Review of Psychology*, *50*, 243–271.
- Henderson, J. M., Weeks, P. A., & Hollingworth, A. (1999). The effects of semantic consistency on eye movements during complex viewing. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 210–228.
- Hollingworth, A., & Henderson, J. M. (1998). Does consistent scene context facilitate object perception? *Journal of Experimental Psychology: General*, *127*, 398–415.
- Ito, M., Jujita, I., Tamura, H., & Tanaka, K. (1994). Processing of contrast polarity of visual images in inferotemporal cortex of the macaque monkey. *Cerebral Cortex*, *5*, 499–508.
- Jain, A. K., & Hoffman, D. (1988). Evidence-based recognition of objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *10*, 783–802.
- Johnston, A., Hill, H., & Carman, N. (1992). Recognizing faces: Effects of lighting direction, inversion, and brightness reversal. *Perception*, *21*, 365–375.
- Jüttner, M., Caelli, T., & Rentschler, I. (1997). Evidence-based pattern classification: A structural approach to human perceptual learning and generalization. *Journal of Mathematical Psychology*, *41*, 244–259.
- Jüttner, M., & Rentschler, I. (1996). Reduced perceptual dimensionality in extrafoveal vision. *Vision Research*, *36*, 1007–1021.
- Jüttner, M., & Rentschler, I. (2000). Scale invariant superiority of foveal vision in perceptual categorization. *European Journal of Neuroscience*, *12*, 353–359.
- Kemp, R., McManus, C., & Pigott, T. (1990). Sensitivity to the displacement of facial features in negative and inverted images. *Perception*, *19*, 531–543.
- Krieger, G., Rentschler, I., Hauske, G., Schill, K., & Zetsche, C. (2000). Object and scene analysis by saccadic eye movements: An investigation with higher order statistics. *Spatial Vision*, *13*, 201–214.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.
- Kruskal, J. B. (1964). Nonmetric multidimensional scaling. *Psychometrika*, *29*, 1–27.
- Lamberts, K. (1998). The time course of categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 695–711.
- Liu, C. H., & Chaudhuri, A. (1997). Face recognition with multi-tone and two-tone photographic negatives. *Perception*, *26*, 1289–1296.
- Liu, C. H., Collin, C. A., & Chaudhuri, A. (2000). Does face recognition rely on encoding of 3-D surface? Examining the role of shape-from-shading and shape-from-stereo. *Perception*, *29*, 729–743.
- Loftus, G. R., & Mackworth, N. H. (1978). Cognitive determinants of fixation location during picture viewing. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 565–572.
- Medin, D. L., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.
- Miller, G. A., Bruner, J. S., & Postman, L. (1954). Familiarity of letter sequences and tachistoscopic identification. *Journal of General Psychology*, *50*, 129–139.
- Morris, R. K. (1994). Lexical and message level sentence context effects of fixation times in reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 92–103.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.
- Palmer, S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory and Cognition*, *3*, 519–526.
- Phillips, R. J. (1972). Why are faces hard to recognize in photographic negative? *Perception and Psychophysics*, *12*, 425–426.
- Polyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*, 341–423.

- Rentschler, I., & Caelli, T. (1990). Visual representations in the brain: Inferences from psychophysical research. In H. Haken (Ed.), *Synergetics of cognition* (pp. 233–248). Berlin: Springer.
- Rentschler, I., Jüttner, M., & Caelli, T. (1994). Probabilistic analysis of human supervised learning and classification. *Vision Research*, *34*, 669–687.
- Riseman, E. M., & Hanson, A. R. (1987). A methodology for the development of general knowledge-based vision systems. In M. A. Arbib & A. R. Hanson (Eds.), *Vision, brain, and cooperative computation* (pp. 285–328). Cambridge, MA: MIT Press.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Schyns, P. G., Goldstone, L., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1–53.
- Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function: I and II. *Psychometrika*, *27*, 125–140.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.
- Soloman, R. L., & Postman, L. (1951). Frequency of usage as a determinant of recognition thresholds for words. *Journal of Experimental Psychology*, *43*, 195–201.
- Tanaka, K. (2000) Mechanisms of visual object recognition studied in monkeys. *Spatial Vision*, *13*, 147–163.
- Ullman, S. (1996). *High-level vision: Object recognition and visual cognition*. Cambridge, MA: MIT Press.
- Unzicker, A., Jüttner, M., & Rentschler, I. (1999). Modelling the dynamics of visual classification behaviour. *Mathematical Social Sciences*, *38*, 295–313.
- Von Helmholtz, H. (1896). *Handbuch der physiologischen Optik*. Leipzig, Germany: Voss.
- Watson, A. B., Barlow, H. B., & Robson, J. G. (1983). What does the eye see best? *Nature*, *302*, 419–422.
- Westheimer, G. (1998). Lines and Gabor functions compared as spatial visual stimuli. *Vision Research*, *38*, 487–491.

Manuscript received April 2003
Manuscript accepted February 2004

