

Acquiring long-term representations of visual classes following extensive extrastriate damage

Orna Rosenthal^{a,*}, Marlene Behrmann^b

^a UCLA, Department of Psychology, 7531 Franz Hall, UCLA, Los Angeles, CA 90095, USA

^b Carnegie Mellon University, Department of Psychology, Pittsburgh, PA, USA

Received 1 April 2005; received in revised form 20 June 2005; accepted 20 July 2005

Available online 30 August 2005

Abstract

Different areas of human visual cortex are thought to play different roles in the learning of visual information: whereas in low/intermediate cortical areas, plasticity may be manifested by enhanced selectivity to learned visual features, in higher-level areas, plasticity may result in generalization and development of tolerance to degraded versions of the learned stimuli. The most effective tolerance to degraded information is presumably achieved in the case of cooperation between the different forms of plasticity. Whether this tolerance to degraded information also applies when the visual input is degraded as a result of a lesion to lower levels of the visual system remains an open question. To address this, we studied visual classification learning in a patient with an extensive bilateral lesion affecting intermediate/low-level visual areas but sparing higher-level areas. Despite difficulty in perceiving the stimuli, the patient learned to classify them, albeit not as quickly as control participants. Moreover, the patient's learning was maintained over the long term and was accompanied by improved discrimination of individual stimuli. These findings demonstrate that degraded output from lesioned, lower areas can be exploited in the service of a new visual task and the results likely implicate a combination of bottom-up and top-down processing during visual learning.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Visual agnosia; Visual learning; Extrastriate lesions; Plasticity; Rehabilitation

1. Introduction

Visual learning has multiple manifestations: it may be reflected in improved visual perception (e.g. fine discrimination), in enhanced familiarity with the trained stimuli or even in the ability to generalize across stimuli and to create new visual categories. Findings from recent human and animal studies reveal that visual learning involves modifications of neuronal responses in low and intermediate level cortices (striate and extrastriate cortex; e.g. Crist, Li, & Gilbert, 2001; Furmanski, Schluppeck, & Engel, 2004; Ghose, Yang, & Maunsell, 2002; Gilbert, Sigman, & Crist, 2001; Rainer, Lee, & Logothetis, 2004; Schoups, Vogels, Qian, & Orban, 2001; Schwartz, Maquet, & Frith, 2002; Yang & Maunsell, 2004) as well as pronounced modifica-

tions in high-level areas (e.g. inferotemporal and prefrontal cortex; e.g. Baker, Behrmann, & Olson, 2002; Erickson, Jagadeesh, & Desimone, 2000; Freedman, Reisenhuber, Poggio, & Miller, 2001, 2003; Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999; Jagadeesh, Chelazzi, Mishkin, & Desimone, 2001; Kobatake, Wang, & Tanaka, 1998; Miyashita, 1988; Miyashita & Hayashi, 2000; Rainer & Miller, 2000; Rossion, Gauthier, Goffaux, Tarr, & Crommelinck, 2002; Sakai & Miyashita, 1991; Sheinberg & Logothetis, 2002; Sigala & Logothetis, 2002).

Interestingly, this widespread effect of learning appears not to be monolithic in that learning-related plasticity may manifest itself differently in low/intermediate and in high-level visual areas (at least in the ventral stream; e.g. Rainer et al., 2004). Plasticity in low/intermediate areas seems to induce selective amplification of feature-specific neuronal activity that corresponds to the relevant features of the trained stimuli. In contrast, plasticity in higher-level areas may result

* Corresponding author. Tel.: +1 310 267 2141; fax: +1 310 267 2141.
E-mail address: ornar@psych.ucla.edu (O. Rosenthal).

in the formation of invariant neuronal responses to the different individual exemplars of the same visual category. This latter form of plasticity, thus, appears to involve generalization across multiple stimuli rather than enhanced selectivity to specific stimulus or display features (Mumford, 1992).

One positive outcome of learning is the development of tolerance to partial information resulting from stimulus degradation: humans and animals can learn to identify highly degraded images (e.g. Rainer & Miller, 2000). Learning to recognize degraded information may benefit from cooperation among plasticity processes at different levels of processing: low and intermediate levels may enhance the selective processing of task-relevant features, amplifying the relevant information transferred to higher visual processing levels, whereas high visual processing levels may improve generalization across exemplars of a particular visual object. Such cooperation was recently shown as a result of training monkeys to recognize degraded versions of familiar stimuli. Following training, information about relevant features provided by the response of neurons in intermediate level area V4 was increased (Rainer et al., 2004), and, at the same time, neurons in higher-level visual areas, such as prefrontal cortex and the fusiform gyrus, showed a remarkable increase in the similarity of their response to highly degraded and non-degraded versions of same stimuli (Cox, Meyers, & Sinha, 2004; Rainer & Miller, 2000). These findings suggest that the development of tolerance to stimulus degradation is likely to be coordinated by simultaneous top-down and bottom-up plasticity.

An interesting question is whether this tolerance to degradation can take place not only under conditions of degraded physical stimuli, as discussed above, but also under conditions of degradation arising from the propagation of partial information from lesioned areas of early visual cortex. In both instances, with enough repetition, one might expect that some level of generalization and enhanced processing of the relevant information might ensue, leading to “completion” of the degraded information. It is also possible, however, that, in order to become effective, plasticity may depend critically on sufficiently complete and robust representation at all levels of cortex and, thus, would not be possible in the presence of impaired input from lower-level areas. The purpose of this investigation is to explore whether learning to tolerate residual information can be accomplished in the face of poor visual processing at earlier, lesioned regions of visual cortex, but intact high-level visual areas.

The potential for visual retraining in case of lesions in early visual cortex has not been extensively explored. Though there is increased interest in visual learning and recovery following extensive lesions in high-level regions of ventral (e.g. Behrmann, Marotta, Gauthier, Tarr, & McKeeff, 2005; Merigan & Saunders, 2004) and dorsal (Huxlin & Pasternak, 2004; Pasternak & Merigan, 1994) cortex both in animals and in human, surprisingly, there have been very few studies examining the extent of learning in cases with extensive lesions to low/intermediate levels of visual cortex. In the few

relevant studies, it remains unclear whether learning implicated recovery of or tolerance to the degraded processing or whether the changes were limited to non-lesioned subsystems that might not be optimal for the task at hand. For example, Merigan (2000) showed that monkeys with a focal unilateral lesion in V4 failed to learn to discriminate complex stimuli, which required grouping, despite improved discrimination of other kinds of texture stimuli. Though it seems obvious that grouping is the limiting factor for discrimination learning following a V4 lesion, the learning of the textured stimuli might have been mediated by early visual cortices without involvement of higher-level regions.

In humans, there have been a few studies of patients, who were hemianopic following occipital cortex lesions, demonstrating partial recovery of visual competence in the blind field following training to detect small visual targets (Kasten, Wüst, Behrens-Baumann, & Sabel, 1998; Zihl & Von Cramon, 1985). However the reliability and the interpretation of these findings remains controversial (Balliet, Blood, & Bach-Y-Rita, 1985; Horton, 2005; Plant, 2005). Also, in some cases, recovery could not be definitively verified by the investigators when using more meticulous perimetry, which enables the investigators to track fixations and eliminate trials during which there were eye movements or inadequate fixation location (Reinhard et al., 2005). Even if the recovery in these cases is indeed attributable to improved visual processing rather than compensatory eye movements, the extent of the recovery was rather small with a roughly 5° shift in transition zone even after very long and intensive training (hundred of thousands of trials, daily training for at least 6 months; e.g. Kasten et al., 1998; Kasten, Poggel, & Sabel, 2000). One possible reason for this limited improvement may be the very early level of the lesions, leading to the absence of aware visual perception in the contralesional field. A second possible reason for the limited training effects may be that the use of a simple ‘low-level’ detection task required minimal higher-level visual processing. Counterintuitively, our hypothesis is that, even in cases of damage to earlier parts of the cortical visual system, tasks that demand greater involvement of higher-level processing (e.g. cross-stimulus association and memorization) may actually result in a stronger interaction between low and high-level areas, and, thereby, facilitate recovery to a greater extent.

In the present study, we explore whether learning – which involves ‘high/intermediate level’ as well as ‘low-level’ processing – is possible in the face of extensive damage to intermediate/low-level visual areas in the case of JW, a profoundly agnostic patient. JW suffered an extensive bilateral occipital extrastriate lesion, including area V2 and some of striate cortex (Fig. 1), but sparing high-level visual areas such as the fusiform region, the temporal lobes and prefrontal areas. The particular approach involved determining whether JW could learn to classify simple stripe stimuli. As classification learning involves generalization – or at least pooling – across different stimuli and long-term representations, it is thought to demand higher-level visual processing. It was



Fig. 1. JW's structural brain scan. CT scan in axial plane at the level of occipital pole. The bi-lateral ischemic wrap around the occipital pole is indicated by arrows.

shown previously that extensive high-level lesions (resulting from bilateral occipital-temporal and postero-lateral frontal infarcts) impaired generalization across similar exemplars and the formation of long-term storage of class prototype representation, despite the fact that individual exemplars were well detected (Scheindler, Landis, Rentchler, Regard, & Baumgartner, 1992). In our study, the situation is the reverse: JW had spared higher-level vision and we hoped to exploit this in a task involving discrimination between similar simple exemplars (stripes) that belong to different classes. Note that, despite their simplicity, the stripe stimuli were still difficult to perceive by JW. If JW could learn to classify the stimuli even in the face of the ongoing impaired perception of the stimuli resulting from massive damage to lower-level areas, this would provide support for the idea that tolerance for lesion-related degraded information has been developed.

The research was conducted in three experimental phases: in Experiment 1 we demonstrate that JW can indeed learn the visual classes, and that learning can be maintained even in the absence of immediate feedback. In Experiment 2, which was carried out 6 months later, we demonstrate that JW can maintain the learned classes over a relatively long period. We also explore whether his initial learning was limited to the learned stimulus-response pairings or was more general, by testing his performance in a different class discrimination procedure. Finally, we determine whether JW's classification

learning involves recovery of lower-level processes (Experiment 3).

2. General methods

2.1. Participants

JW, a 45-year-old individual with visual agnosia, and four healthy male control participants with no history of neurological disorder and normal or corrected-to-normal vision participated in this study. Control participants were matched in gender, age (39–48 years old) and education to JW. All participants provided written informed consent, and the consent form was read aloud slowly to JW before he signed.

JW has an extensive lesion implicating V2 in both hemispheres but with no evidence of damage at high-level cortices of the ventral stream. A detailed description of his medical and neuropsychological history can be found in several other papers (Mapelli & Behrmann, 1997; Marotta, Behrmann, & Goodale, 1997; Vecera & Behrmann, 1997; Vecera & Gilds, 1997) and so we only outline this briefly here. His visual deficits also mirror those of other cases with similar lesions (Farah, 2004; Ricci, Vaishnavi, & Chatterjee, 1999). At age 35, JW experienced an event of anoxic encephalopathy associated with cardiac abnormality. Based on CT scans taken 2 years later, he was diagnosed as suffering from generalized atrophy and ischemic infarction in both occipital lobes, extending slightly into the right parietal lobe. An extensive ischemic wrap around the occipital pole indicates widespread V2 and V1 injury (Fig. 1). This diagnosis was provided independently by two neuroradiologists, neither of whom reported any damage at higher-level cortices. Unfortunately, we have not been able to obtain an MRI scan on JW as he has a prosthetic metal sternum, inserted during cardiac surgery.

Since the event (roughly 10 years prior this study), JW has regained relatively normal performance in non-visual tasks. His speech and movement abilities are near to normal and he shows no sign of memory impairment. At the same time, he suffers from a severe and lasting visual impairment that fits the diagnosis of visual form agnosia (or apperceptive agnosia, Farah, 2004; Vecera & Gilds, 1997). JW has corrected retinal vision but has an incomplete upper left scotoma, although this, alone, cannot explain his profound object agnosia and prosopagnosia.

Table 1 summarizes JW's performance on a range of tests of visual processing, conducted prior to and during the current study. Briefly, JW shows some impairment in low-level processing (fine orientation discrimination; see also Fig. 6B demonstrating low contrast sensitivity) as well as a severe impairment in intermediate level processing (e.g. size and shape discrimination, grouping, image segmentation; see also Fig. 6A, C and D). JW was also profoundly impaired at object and face recognition and at word recognition. However, several visual functions remain intact, including object memory and imagery (assumed to reflect high-level processing), color

Table 1
Summary of JW's performance on visual testing

Function	Performance	Remarks
Visual acuity	OLD 20/200 ^a	Acuity evaluation underestimated because test is based on digit detection
Orientation discrimination		
Fine	Highly impaired ^b	For oriented bar difference <45°
Coarse	Preserved ^b	Horizontal vs. vertical bars
Contrast sensitivity	Impaired ^c	
Color discrimination	Preserved ^a	
Color-orientation adaptation—McCullough effect	Preserved ^d	
Figure-ground segregation	Highly impaired ^b	For noisy/crowded background
Shape segmentation	Highly impaired ^b	For overlapped shapes
Gestalt grouping	Impaired ^{c,e}	Grouping by proximity, good continuation, closure
Contour integration	Highly impaired ^{c,d}	Illusory contours and segmented shapes
Symmetry judgment	Highly impaired ^e	
Object, face and letter recognition	Impaired ^b	
Aspect ratio discrimination	Highly impaired ^b	Rectangle stimuli (Efron test; stimuli > 4°)
Shape discrimination		
Fine	Highly impaired ^b	For example, circle vs. oval shape
Coarse	Preserved ^b	For example, square vs. circle
Parallel search (pop-out)		
Orientation	Impaired ^e	
Color	Preserved ^b	
Object imagery and memory	Preserved ^b	
Visual-motor coordination (grasping)		
Monocular	Impaired ^f	Cannot use 2D depth cues
Binocular	Preserved ^f	Can use stereo visual cues

^a Mapelli and Behrmann (1997).

^b Mapelli and Behrmann, Unpublished data.

^c See Fig. 6.

^d Behrmann, Unpublished data.

^e Vecera and Behrmann (1997).

^f Marotta et al. (1997).

vision, visual-motor control with binocular vision (Marotta et al., 1997), and spatial attention (Vecera & Behrmann, 1997).

2.2. Design and procedure

Stimuli and tasks were designed and run using E-prime software (Psychology Software Tools 2002; see Schneider, Eschman, & Zuccolotto, 2002) unless otherwise mentioned.

3. Experiment 1—classification learning

JW's extensive extrastriate/striate lesion led to impaired perception of both simple and complex visual stimuli (Table 1). In this main experiment, we examined JW's ability to learn to classify simple stripe stimuli into three classes, by their stripe width, despite his difficulty in discriminating size and fine shape differences. Classification learning involves the acquisition and subsequent retrieval of new long-term class representations, as well as the association over many stripe exemplars. Because of the need to develop representations over time and stimuli, this task is thought to implicate higher-level visual processing and to involve areas

that engage stimulus association and generalization as well as long-term storage and retrieval.

3.1. Methods

3.1.1. Stimuli

Each stimulus consisted of a pair of identical vertical bright stripes presented on a gray screen (1280 × 1024 pixel; 90 pixels/in.; 21 in. monitor; Fig. 2A). The viewing distance was 40 in., and was maintained using a chinrest. The stripes ran the full height of the computer monitor, and the center of each stripe was offset $\pm 4.3^\circ$ from central fixation. Serially presented stimuli varied only in stripe width (~ 0.24 – 8.6°).

Overall, the distribution of the sampled stimuli was a sum of three Gaussians (Fig. 2B). The stripes were related to one of three predefined classes, according to stripe widths, expected to be learned by the participants. The dictated classification pattern is indicated in Fig. 2B. Each class corresponded to one third of the full stripe width range (0– 8.6°). The dictated classes "A", "B" and "C" corresponded to narrow, medium and broad stripes, respectively. The boundaries between classes "A" and "B" (bound._{AB}) and between "B" and "C" (bound._{BC}) corresponded to stripe widths of 2.9 and 5.7°, respectively. For the purpose of sampling and presentation,

stripe width was binned into 36 equal bins across the range of widths. To facilitate classification learning, the stimuli in bins near each class center (“prototype”) were frequently sampled (70 stimuli/session from class center bins) whereas stimuli from bins near the boundaries were rarely sampled (2 stimuli/session from class boundary bins). During each ses-

sion, stimuli were randomly sampled from the prespecified distribution. Previous studies have shown that participants naturally tend to locate class boundaries near the range of the less frequent stimuli and to center classes at the most frequent stimuli (Rosenthal, Fusi, & Hochstein, 2001). Pilot experimentation confirmed the ease of this task.

3.1.2. Procedure

Participants performed the classification task in which each stimulus display was assigned to one of the three classes. The time course of a typical trial is presented in Fig. 2C. Instructions were read to each participant before the first classification session. Participants were informed that they would see a stimulus that belongs to one of three stimulus groups and that they should classify it accordingly, using three different response keys on the keyboard. No further instructions or demonstrations were given regarding a desired class structure. Incorrect classification choices were followed by (negative) feedback, which included a visually displayed “wrong” message and a specific sound. Correct classification was not followed by feedback.

A blank gray screen preceded each trial. Participants pressed a key to start the trial. After 1 s, a small fixation cross appeared. Following a delay of 100 ms, a stripe stimulus appeared and, after 500 ms, the screen turned blank again. Participants had 1.5 s from stimulus onset to classify the stimulus. Late responses were discarded and signaled to the participants. Each classification session included 1024 trials. All participants completed (at least) eight classification sessions. The last (8th) session did not include any feedback, but was similar to the other sessions in all other respects. For all participants, session 2 was completed within 2 days of session 1. For the rest of the experiment, the time interval between sessions was 5–7 days for JW (with the exception of 43 days between sessions 6 and 7) and 1–10 days for the control participants.

In order to familiarize participants with the experimental setting and stimuli, a single observation session of 256 trials preceded the first classification session. Conditions were identical to the classification sessions except that participants were not required to classify the stimuli and merely responded by pressing any of the three available keys. No feedback

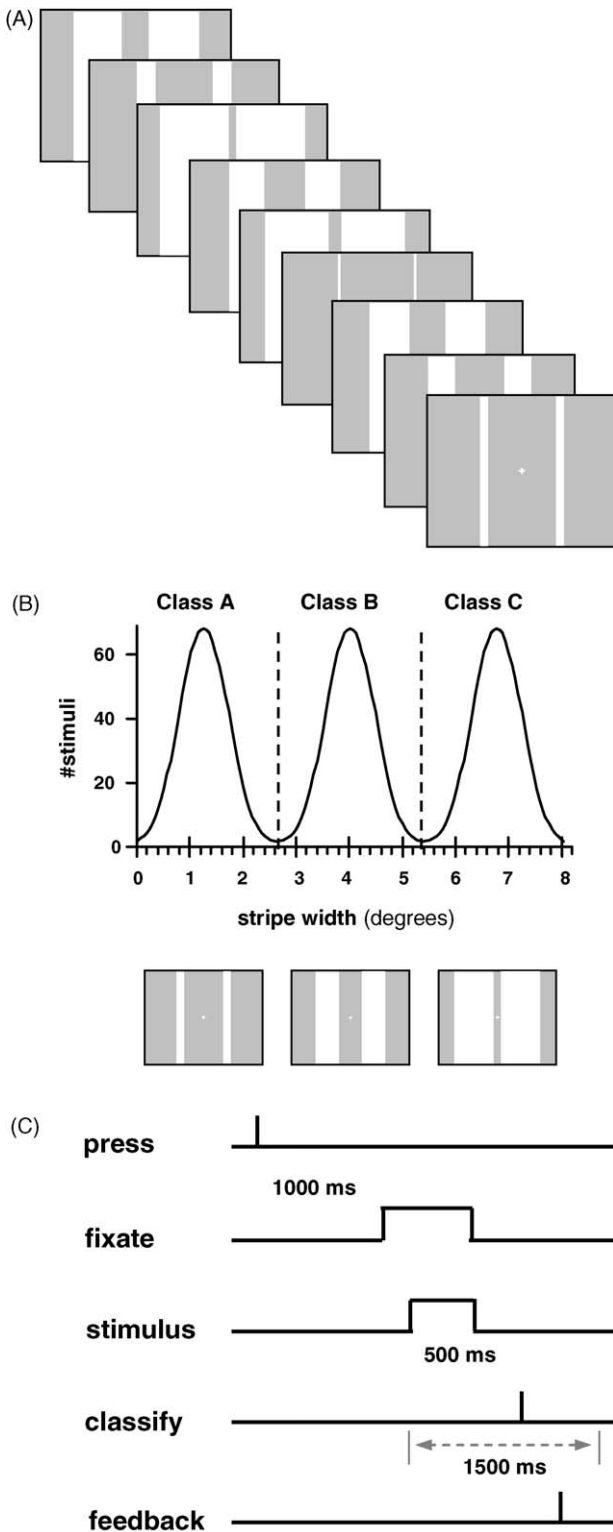


Fig. 2. Classification task. (A) Example of stripe exemplars. The whole stripe width range (~ 0.24 – 8.3°) was randomly sampled during each session. (B) Dictated class pattern. Stimulus frequency is presented as a function of an exemplar's stripe width. Vertical lines represent the dictated location of class boundaries. Miniatures of exemplars at class' centers (prototypes) are presented below the graph. Participants were required to classify the stripe exemplars into three classes (A–C) by pressing one of the three available keys. Each dictated class ranged along one-third of the total range. Each session included 1024 trials. Gray line shows stimulus frequency (denoted by the Y-axis) and, as is evident, more stimuli were presented around class prototype (centers). (C) Trial sequence. Each trial started with a key press. Participants had 1.5 s from stimulus (stripe exemplar) onset to decide about class membership of the exemplar. Incorrect answers were followed by visual and auditory feedback.

was given for this task, except in the case of very slow responses.

3.1.3. Data analysis

Classification data were summarized in terms of ‘sorting coherence’ (SC), namely, the fraction of stimulus bin samples, which were assigned by a participant to a given class. SC values ranged between ‘0’ (never choosing the class for a given stimulus bin) and ‘1’ (always choosing the class for a given stimulus bin). Any value between these reflects some level of confusion between classes. SC raw data were fitted by a generalized adaptive model, assuming that SC for each class follows a binomial distribution, and smoothed using smoothing splines with 5 degrees of freedom (using S-plus 6 software, Insightful 2003). Since, in the vast majority of cases, the confusion was mainly between two adjacent classes, class boundary location was defined as the stimulus bin that elicits SC of 0.5 (full confusion between two classes).

For each class, we also calculated the ‘class discriminability’ as follows: starting from the class boundary (SC = 0.5), we traced the smoothed SC curve up and down to the points where the slopes approached zero (i.e. a slope smaller than a threshold of 0.17 deg^{-1}), reflecting “inside” and “outside” the class, respectively. This procedure resulted in two zero-slope locations, Z_1 and Z_2 . For each of these locations, we computed “hits” (SC level “inside” the class, H) and “false alarm” (SC level “outside” class, F) rates. Class discriminability (CD) is, then, defined as

$$CD = \log \alpha(Z_1) \times \log \alpha(Z_2) \times \left(1 - \frac{|Z_1 - Z_2|}{8.58} \right),$$

$$\text{where } \alpha(Z) = \sqrt{\frac{H(Z) \times [1 - F(Z)]}{[1 - H(Z)] \times F(Z)}}$$

This measure reflects the degree of certainty in classification near a given class boundary.

Note that SC reflects the discrimination between classes assigned by the subjects. To measure the degree of match to the dictated classification pattern, we also computed the

percent of correctly classified stimuli to each dictated class, and averaged it across classes.

To evaluate JW’s learning dynamics, we fitted a power-function model to his class discriminability data across sessions, using Levenberg–Marquardt non-linear regression approach (SPSS). This analysis allowed us to estimate learning parameters (asymptote, coefficient and power) and compute the asymptotic standard error of these estimations. Because JW did not show learning between the first and second sessions, the data from the first session were not included in the estimation.

Response time (RT) data were also used in subsequent analysis and, for these analyses, RT data points that fell more than 2.5 standard deviations from the mean were excluded.

3.2. Results

Fig. 3 presents the results for JW and for each of the four control subjects separately. Classification is plotted in terms of sorting coherence (the fraction of stimulus bin samples, which were assigned to a given class; see Section 3.1) as a function of stripe width. Control subjects (Fig. 3A) show highly accurate classification performance even in the first session (overall $93 \pm 1\%$ correct/class), showing near to maximal sorting coherence for most of the stimulus range, and relatively sharp boundaries between all three classes. By contrast, as can be seen in the top graph of Fig. 3B, in the first session, JW was highly impaired in this task (overall 72% correct/class): he was largely unable to correctly classify stimuli with stripes broader than 3° (classes “B” and “C”; 64% correct/class, compared to $93 \pm 1\%$ correct/class in control participants). Interestingly, his performance with stimuli with narrower stripes was much better (88% correct) than for the other two classes, though perhaps not as good as that of the control subjects ($94 + 3\%$ correct). In fact, in the case of JW, the transition between stripes, whose widths are easily perceived to those whose widths are difficult to perceive, seems to be relatively sharp and to be around the boundary between classes “A” and “B” (bound_{AB}).

Despite the severe impairment in the first session, further training revealed a surprisingly large improvement in JW’s

Fig. 3. Classification performance. (A) Classification patterns of the four control participants during the 1st session. Each graph presents the classification pattern of an individual participant. Different gray scale curves denote different classes. Classification pattern is presented in terms of the sorting coherence as a function of stripe width (see Section 3.1). Stripe width range was binned into 36 bins, here and in the following figures. Data were smoothed (thick lines) using natural splines. The locations of the boundaries between A and B classes (bound_{AB}) and between B and C classes (bound_{BC}) are denoted by dotted vertical lines. One session suffices to elicit good classification learning in the control subjects. (B) Four out of eight classification sessions of patient JW. In contrast to the control participants, JW’s performance was highly impaired in the first session. During the next six training sessions, JW exhibited gradual improvement. (C) Summary of class discriminability between classes A and B (filled symbols) and between classes B and C (empty symbols) of all participants as a function of the training session. Shaded areas indicate area within 95% confidence level for control participants’ data (darker and lighter areas for the data of discrimination between classes A and B and classes B and C, respectively). Full and dashed thick lines indicate the estimated asymptotes for the learning data for classes “A”, “B” and “B”, “C”, respectively. Their corresponding thin lines represent 1 asymptotic standard error below each estimated asymptote (see Section 3.1). Control participants were already able to discriminate well between classes in the first session and did not show much improvement in later sessions. In contrast, JW showed slow and gradual improvement. Note that, for bound_{BC}, discriminability could not be computed during the first sessions, due to low performance. For all participants, class discriminability was not affected by the absence of feedback (8th session). (D) Summary of class boundary locations as a function of training session, for all participants. The locations of the dictated boundaries are denoted by horizontal dashed lines. Control participants completed learning of boundary locations in the first session. Similarly, JW was able to learn the location of class boundaries as soon as he started to be able to discriminate between the classes (C).

performance (Fig. 3B, lower three graphs). Slight improvement may already be evident in the second session (overall 75% correct/class). The emergence of sorting coherence >0.6 at the ranges corresponding to the dictated classes of the broader stripes (classes “B” and “C”) leads to 71% correct/class for that range. As training proceeded, classification performance gradually improved further. Fig. 3C presents the learning curves in terms of class discriminability around each class boundary (see Section 3.1). This further reveals that, whereas control participants’ performance was already at ceiling in the first classification session, JW improved slowly with the additional sessions. Also, during all the sessions, JW discriminated better between classes “A” and “B”

than between “B” and “C”, unlike the control participants who showed similar performance around both boundaries. His discriminability learning data were well fitted with a power function ($R^2 = 0.86$ and 0.92), with an asymptote of 2.42 ± 0.42 (S.E.) and 1.9 ± 0.36 for discriminability data between classes “A”, “B” and “B”, “C”, respectively (see Section 3.1). As can be seen in Fig. 3C, it took JW four sessions to approach asymptote performance level in discriminating between classes “A” and “B”, and seven sessions for approaching asymptotic discrimination between classes “B” and “C”. However, in both cases, even by the 7th session, his performance had still not reached the lower 95% confidence bounds of the control data.

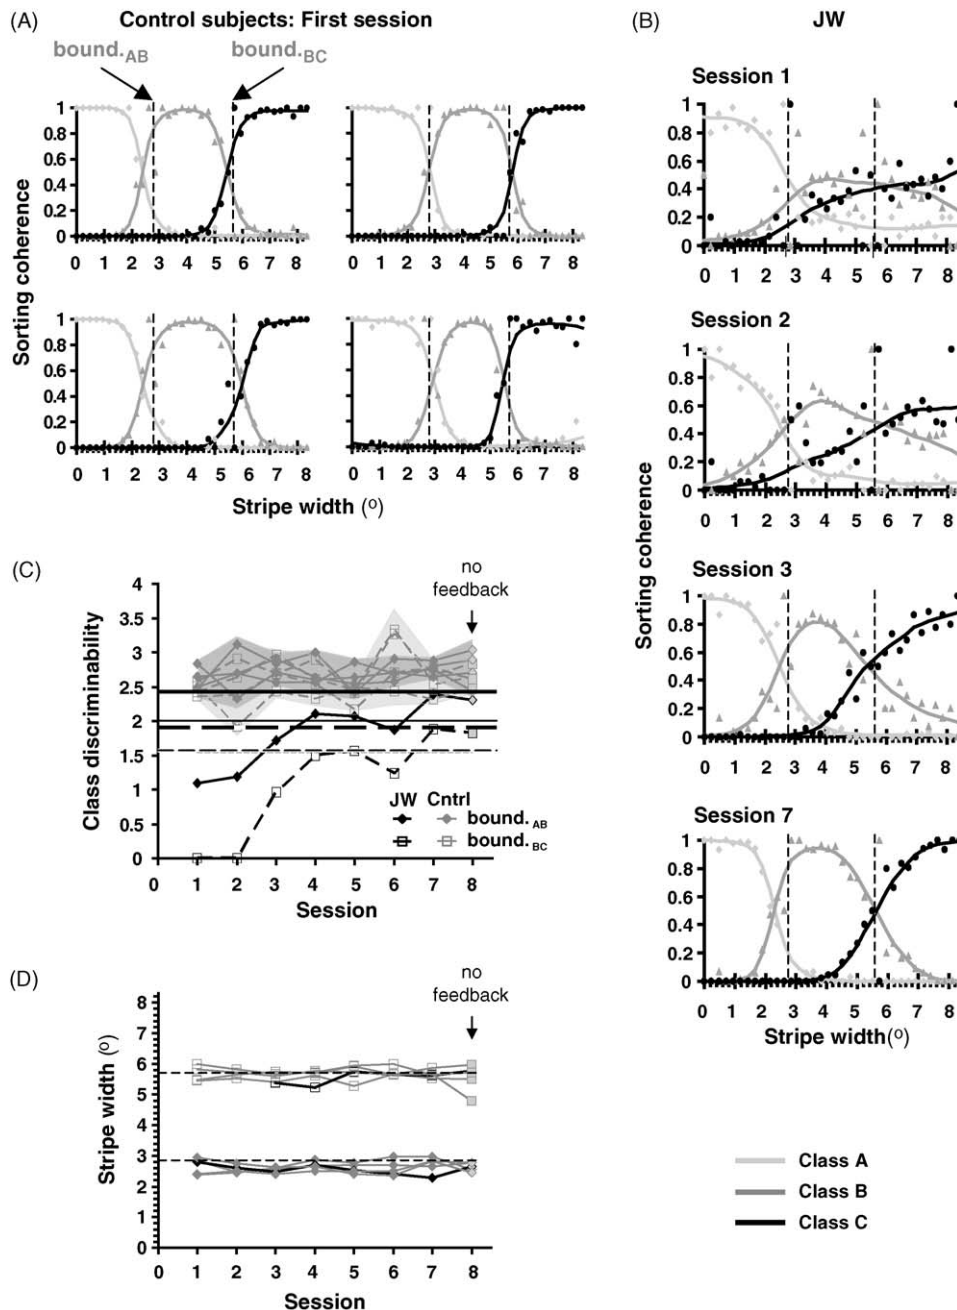


Fig. 3D demonstrates that, once JW started to discriminate between class members, his classification patterns were similar to those dictated by the feedback, as reflected by the boundary location. Note that, as evident from JW's slow learning curves, the between-class discriminability around these boundaries was still very small in the earlier sessions and increased slowly with further training (see Fig. 3C). Note that the learning curves in Fig. 3C, which are based on class discriminability, do not depend on the subject's boundary locations—in principle, one could show good discriminability between classes even if they do not reflect the prespecified boundary, and vice versa.

During the 8th session, in contrast with the previous sessions, in order to evaluate the maintenance and robustness of the learning that had occurred, the participants were required to classify the stimuli, but no feedback was provided. As can be seen in Fig. 3C and D, during this final session, JW's performance was comparable to the last session which included feedback (i.e. the 7th session), and class discriminability values did not significantly differ from the estimated asymptotes. As is apparent from the figures, there is neither any deterioration in class discriminability (Fig. 3C, 8th session) nor a shift in class boundary locations (Fig. 3D, 8th session).

These findings suggest, importantly, that JW was able, eventually, to form fairly well-separated classes corresponding to those dictated by feedback. Interestingly, he was also able to maintain these class distinctions at least over one session in the absence of feedback.

3.3. Discussion

Despite his extensive lesion to lower-level visual areas and his difficulty in size discrimination – especially of wide stimuli – JW learned to classify the stripe stimuli, and was able to maintain the class distinctions for at least one session (session 8) without feedback in order to correctly retrieve the classes. These findings indicate that JW learned to tolerate the poor residual visual input and make use of this input for the purpose of classification. JW's slow learning (several sessions) compared to the immediate learning shown by control participants, serves as an additional indication of the extent of the impaired processing of the stimuli. Of note, in the first session JW already showed much less impairment in classification with narrow stimuli, showing 88% correct classification performance in that range. That he could do relatively well with narrow but not wider stripes may possibly be suggestive of a defect in spatial summation across visual input.

It is possible, however, that JW's learning does not reflect learned tolerance for partial input. An alternative explanation is that JW simply found a way to bypass his difficulties by adopting an alternative strategy. For example, exemplars vary not only in their stripe width but also in the location of the stripe edges. JW could potentially have learned to attend to the location of one of the stripe edges and to classify the stimulus according to the edge location. However, as dis-

cussed below, JW had severe difficulty in detecting the stripe edges so this possibility seems untenable. A further possibility is that since the gray gap between broader stripe pairs is narrow and since JW perceives narrow stripes better, he might have learned to attend to the narrow gap instead of the stripes, and thus, learned to classify wide-striped stimuli by the gray gap. There are several indications against this possibility, too. First, it is reasonable to expect that improvement, which results from a change in strategy such as using gap width, will be expressed by a sudden change in the learning curve. However, as revealed from the learning curve in Fig. 3C, improvement evolved smoothly and gradually. Second, during the whole learning period, the learning curve for exemplars near bound_{BC} was lower than the learning curve for exemplars near bound_{AB}. Yet stripe widths of exemplars around bound_{AB} are similar to gap widths between the stripe pairs of corresponding exemplars around bound_{BC}. Thus, if JW learned to attend to the gaps in case of broad stripe exemplars, the two learning curves should have converged. This is not the case here. Third, classification also improved for exemplars of class "C". For most of the exemplars in this class, both stripe widths and between-stripe gaps are broader than 3°, and, based on data from the first session (Fig. 3B, top), one would expect JW to be severely impaired in perceiving stimuli in this range. Finally, JW did not report any change in his classification strategy and explicitly denied using the gap between the stimuli as an additional cue.

Although we cannot definitively rule out that JW circumvented the classification learning by adopting an alternative strategy, the results are more compatible with the idea that JW learned to tolerate poor low-level input and to utilize it for improved performance of a visual task that requires the association and formation of long-term representations. The purpose of Experiment 2 is to verify that JW's improvement does indeed reflect the formation of class representations that index class membership and that are separate from the representations of the individual class members.

4. Experiment 2—assessment of the representations of classes

The results of Experiment 1 demonstrate JW's ability to tolerate degraded input to a level that allowed him to improve substantially in a classification task. However, it is unclear exactly what kind of processing modifications allow for the development of this tolerance to noisy and partial input. One can think of the possible underlying mechanisms as if they are scaled along a continuum, between two extremes. One extreme possibility is that JW's improvement reflects the enhanced contribution of visual inputs, which bypass the lesion (e.g. via thalamo-cortical paths). In this scenario, similar to blindsight, learning is expected to be implicit and low-level, enhancing the relationship between individual stimuli and the dictated motor responses. Thus, learning might be expected to be specific to the task procedure used. At the

other extreme, learning might reflect the formation of new high-level representations of the class members. In this scenario, learning may be explicit and should not be selective for the particular stimulus-response pairing used in training. For example, the acquired class representation may be retrieved even by symbolic reference to the class, and not only by the presentation of the specific exemplars, as in the procedure used during the preceding classification training.

To evaluate whether JW did indeed acquire representations of the classes during training and did maintain these representations over time, we conducted a follow-up experiment 6 months after the completion of the training. In this experiment, JW and the control participants were asked, first, to associate a color with each of the original three classes (note that JW has good color vision; see Table 1). Thus, the colors served as symbolic labels of the classes. Having established that all participants could label the three classes by color, they were then tested with a class membership discrimination task in which they were presented with a color and a stripe stimulus and were asked to decide whether the exemplar was a member of the class denoted by the color (Fig. 4A; see Section 4.1). The purpose of this arrangement is to remove the simple stimulus-response association acquired during the initial training and to assess whether JW could determine the class membership through a symbolic referent.

4.1. Methods

4.1.1. Stimuli

In each trial, two displays were presented sequentially (Fig. 4A). The first display consisted of a centrally presented colored square against a gray background. The square was either red, yellow or green, serving as labels for the classes “A”, “B” and “C” of Experiment 1, respectively, and the second display consisted of a pair of stripe stimuli, as was used for the classification experiment (Experiment 1). Only stripe stimuli at the center of the classes (class prototypes) were used to ensure high sorting coherence for both JW and the control participants. Both the colored square stimuli and the stripe stimuli were pseudorandomly sampled.

4.1.2. Procedure

Participants were tested with a 2-alternative forced choice class membership discrimination task. This test was conducted 6 months following the final classification session for JW and 2.5–4 months following classification for the three control subjects, respectively. Before the test was taken, participants practiced associating a specific color with each of the classes for 10 min. Practice involved the classification task, where correct answers were followed by the color of the exemplar’s class. Classification data in the practice block were also used for verifying that participants maintained the classes they had learned during the classification learning phase (Experiment 1).

Each trial started with the onset of a small colored square stimulus. Participants had unlimited time to observe the color

and to associate it with its class. When this was done, the participants pressed a key and a stripe stimulus was presented for 500 ms. The participants then had to make a yes/no judgment using two response keys to indicate whether or not the stripe exemplar was a member of the class, represented by the colored square. There was no time limit for responding. Auditory and visual feedback was given at the end of each trial to encourage high accuracy. The total number of trials was 480. The experiment was run in a dark room. Participants’ distance from the monitor was kept at 40 in. using a chinrest.

4.2. Results

As shown in Fig. 4B, JW performed well in this task ($d' > 2.5$; though significantly and marginally significantly more poorly than the control participants; $p < 0.01$ and $p \sim 0.05$, for A,B and B,C class pairs, respectively). Trial-by-trial analysis revealed no change over the course of the experiment in JW’s performance, excluding the initial within-session training effect. Thus, JW was able to retrieve a class by its associated color and to ascertain the relationship between that class and the presented exemplar.

Despite relatively accurate performance (Fig. 4B), we note that this task was very difficult for JW, as demonstrated by his much slower hit as well as correct rejection (CR) responses, compared to control subjects (Fig. 4C–E). Additionally, as can be seen in Fig. 4C, JW’s RT for hits depends critically on the exemplar/class parameter (slope: ~ 200 ms/deg). Because the task requires both retrieval of classes and exemplar processing, it is unclear whether the steep slope is a result of difficulty in processing the broader exemplars, retrieving their classes or both. Analyzing the RTs for the CRs reveals that it is the exemplar processing per se, rather than the class retrieval, that is giving rise to the slowing in response times. This conclusion is based on a comparison of correct rejections of class “B” exemplars when presenting class label “A” or “C” (Fig. 4D presents the effect of class processing on RT and on a comparison of correct rejections of classes “A” and “C” exemplars when presenting class “B” label, because, in both cases, the only difference in the two conditions is in the classes which are required to be retrieved). Whereas JW’s RT depends significantly on the exemplar processes (Fig. 4E; $p < 0.01$; two-tailed t -test labels “A” and “C”) – indicating impaired exemplar processing – no significant effect is found on JW’s RT due to class processing (Fig. 4D; $p > 0.1$), indicating similar computational effort in the processing or retrieving of different classes. Taken together, these findings indicate that JW was able to acquire good higher-level class representations and retrieve them despite his persistent difficulty in lower-level processing of class’ exemplars.

4.3. Discussion

JW was able to retrieve the classes in response to a symbolic referent and without pre-presentation of an actual class

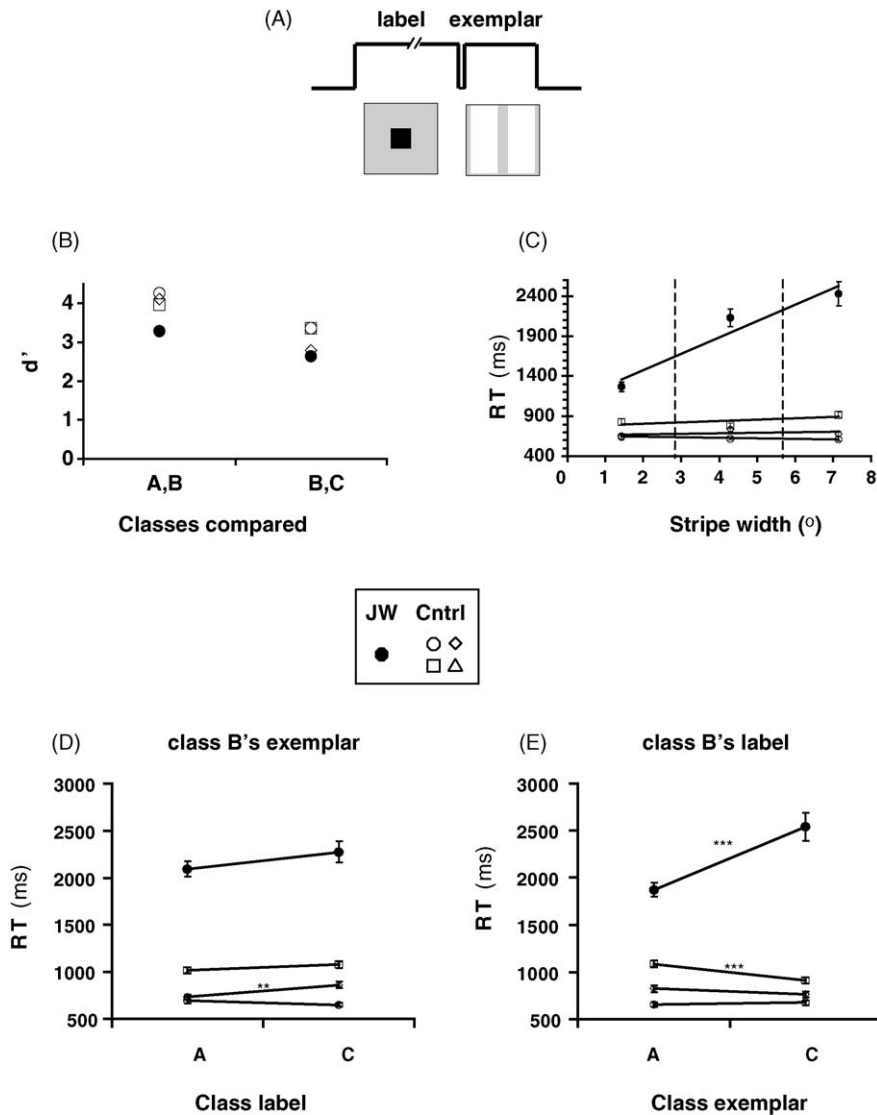


Fig. 4. Class membership discrimination test. (A) Sketch of a trial. Trials started with the appearance of a small green, yellow or red square in the center of the gray screen, symbolizing class A, B or C, respectively (denoted, here, by a dark square). After class retrieval, a stripe exemplar appeared for 500 ms. Participants decided whether the exemplar is a member of the class symbolized by the color. Feedback about participant's choice was given at the end of each trial. (B) Performance levels (d') of all participants in the membership discrimination task. d' may be interpreted here as the distance between the retrieved representation of two "adjacent" classes. Results are presented as a function of the pair of classes compared. Both control participants and JW show high performance levels ($d' > 2.5$) in both compared conditions. (C–E) Class membership discrimination response times (average), demonstrating prolonged responses in the case of JW. Error bars denote standard error from the mean. (C) RT for hit responses (i.e. when correctly answering that the stripe stimulus is a member in designated class). Note the steep slope ($\sim 200 \text{ deg}^{-1}$) as well as the increased intercept for JW. (D) RT for correct rejection (when correctly answering that stimulus is not a member of the symbolized class)—responses to presentation of class B exemplar when the presented class label corresponded to either class A or class C. (E) RT for correct rejection—responses to exemplar presentation of either class A or class C when the presented class label corresponded to class B. For JW, RT significantly differed for two different stripe stimuli (E) but not for two different retrieved classes (D).

exemplar (stripe stimulus), and maintained this ability over the long term (at least several months). In contrast with this, he continued to be severely impaired at processing the individual stripes stimuli, as evident from the RT analysis. The seemingly intact ability to retrieve classes by symbolic cues, despite the impaired perception of the individual exemplars, reveals that class representations formed by JW may be separate from the representation of individual stripes exemplars, and may indicate involvement of high-level visual processing in class representation and retrieval. It is still unclear, though,

what processes mediated the formation of class representation despite the poor perception of individual stimuli and it is this question that we address in the next experiment.

5. Experiment 3—visual perception assessment

One possible explanation mechanism for JW's ability to learn to classify poorly perceived stimuli is that there is (partial) training-related recovery of some residual processing

in the lesioned area, allowing for improved perception of some of the individual stimuli. To explore the possibility that perceptual modification accompanied JW's classification learning, we tested him and the control participants on several perceptual tasks. JW was tested only after the completion of classification training because we did not want to 'contaminate' the classification learning by exposing him to related stimuli prior to the training phase. Although we did not test directly JW's visual perception prior to classification learning, there are several indications that he was severely impaired perceptually. First, his inability to classify exemplars of classes "B" and "C" during the first classification session in Experiment 1 suggests that his width discrimination was highly impaired at the onset of classification training. Also, JW has repeatedly performed at chance on tasks requiring size (aspect ratio) discrimination, grouping and detection (Table 1) such as the Efron task in which squares/rectangles of equivalent area but differing in aspect ratio must be discriminated (Marotta et al., 1997; Vecera & Behrmann, 1997). Control participants were tested both prior to and following training to provide information about training effects on perception: note that, in normal individuals, classification learning is remarkably fast and learning is completed within several trials in the very first session. Thus, effects of a preceding task on the rate of classification learning in control subjects is likely to be negligible but it is in the context of this post-training performance that we examine JW's data.

5.1. Methods

5.1.1. Width discrimination test

To determine whether JW was able to perceive stripe width using the same stimuli as in the classification task, we used a 2AFC test for width discrimination. In this task, trials started with a key press and then, after 100 ms, a fixation cross appeared followed by two different stripe stimuli, presented sequentially. Each stimulus was presented for 500 ms with an inter-stimulus-interval (ISI) of 300 ms (see schematic description at the top of Fig. 5). Participants were asked to decide which of the two stimuli included wider stripes, by pressing one of two available keys. Time for response was unlimited. Auditory and visual feedback were given at the end of each trial. The test involved 300 trials. Stimuli were randomly sampled around class centers and class boundaries of the classification task (see above). Because of the large variability in JW's performance, we were not able to test for discrimination thresholds within a reasonable number of trials and so, based on a pilot experiment, we chose to study discriminability with a specific stimulus pair, whose stripe width differ by 0.7° . The procedure was identical to Experiment 1.

5.1.2. Edge perception test

To learn more about the characteristics of JW's impairment in stripe perception, we tested him in a stripe edge perception task. In this task, a single stripe stimulus was pre-

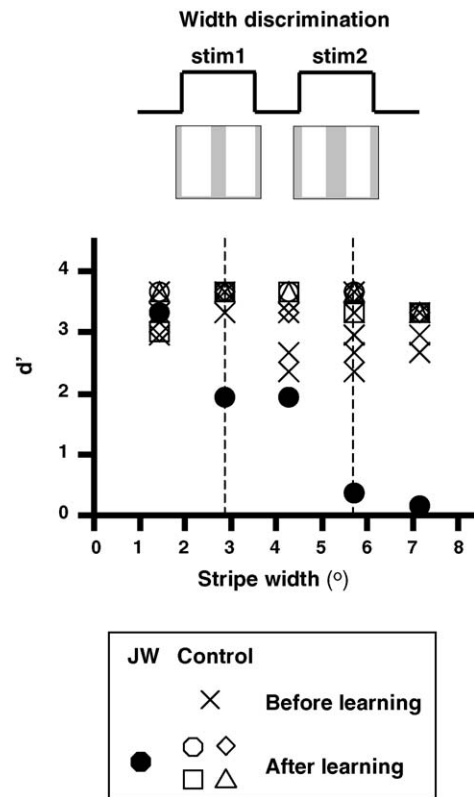


Fig. 5. Width discriminability test. (A) A sketch of a trial. Following key press, two different stripe stimuli were presented sequentially for 500 ms each, with an ISI of 500 ms. Participants decided which of the two contained the broader stripes. Time to response was unlimited. Participant's choice was followed by feedback. Width difference was always 0.7° for the data presented. (B) Discrimination performance (d') as a function of the width of the compared stripes. Note JW's inability to discriminate between wide stripes—i.e. members of class C. Also, note the improvement in discriminating medium and wide stripes by control subjects following classification training (JW was tested only following the training). Dotted vertical lines indicate dictated class boundaries during the classification training.

sented and participants decided whether the stripe edges were blurred or sharp. An example of the two possibilities is shown in Fig. 6A. The stimulus was presented until response. Blurring was done by a sigmoid change in gray level near the edge and the amount of blurring was manipulated. Half the trials were blurred and half sharp and stripe width was fixed ($\sim 4^\circ$; at the middle of the stimulus range tested).

5.1.3. Contrast sensitivity test

As will be evident below, JW exhibits a severe impairment at perceiving stripe edges. To examine whether this may be a function of poor contrast sensitivity, we evaluated this with a contrast sensitivity task. In this adjustment task, stimulus contrast is gradually adjusted to the just noticeable level, which was regarded as $1/\text{contrast sensitivity}$. Stimuli were full-screen vertical sinusoidal stripes with mean luminosity of 13.5 cd/m^2 . We studied gratings of six different spatial frequency conditions. For each condition, the adjustment procedure was repeated 10 times (experimenter manually made the adjustment) and the thresholds were averaged across

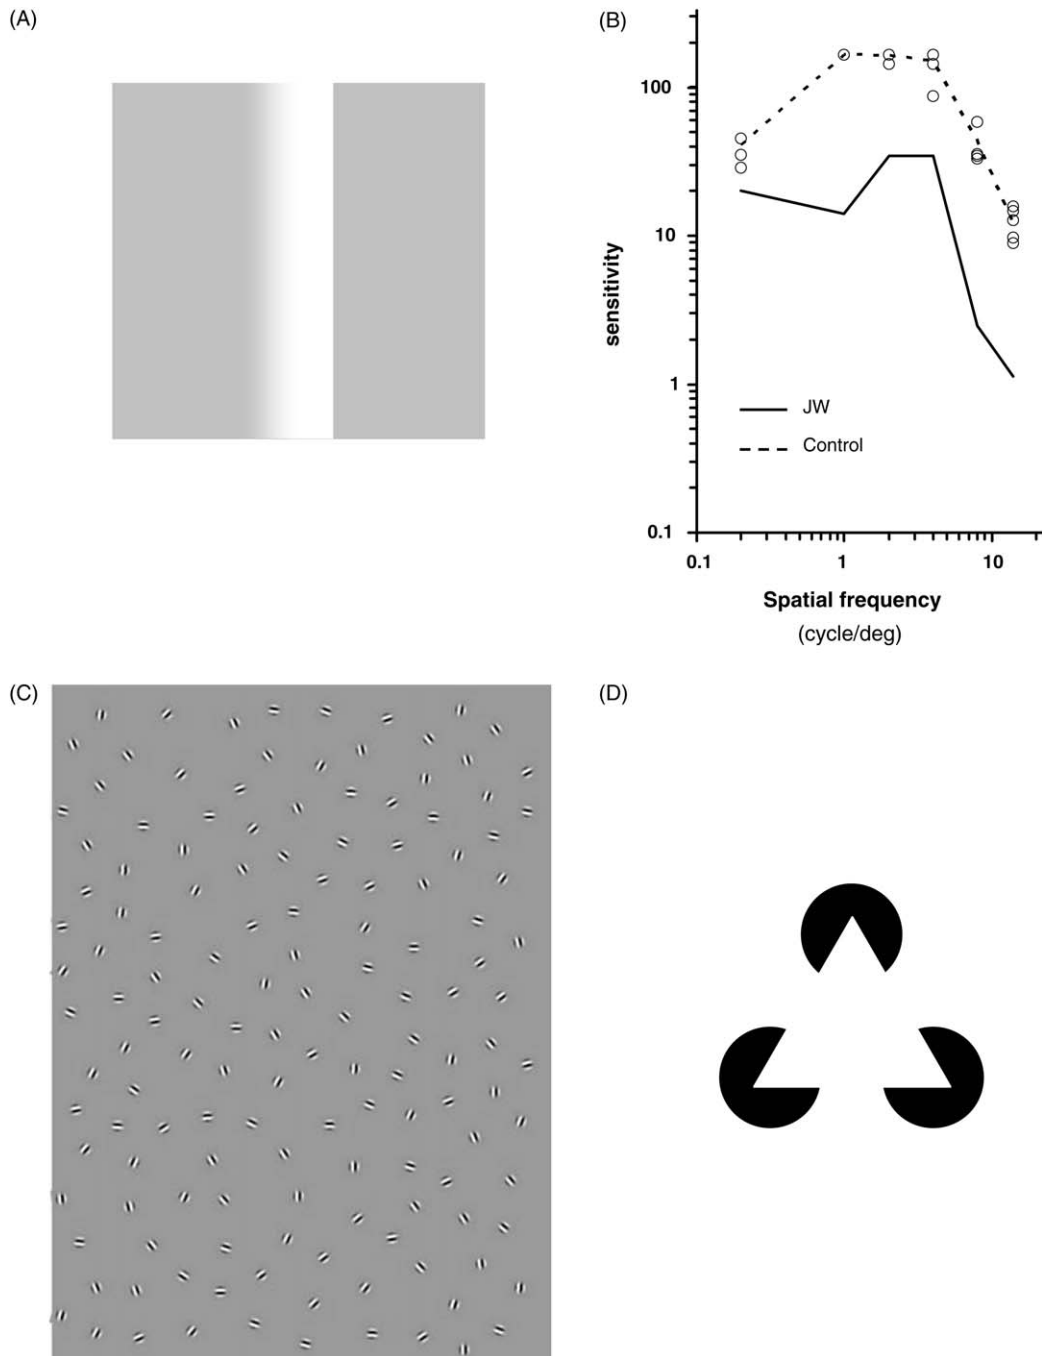


Fig. 6. Tests examining low-level processing. (A) Example of a blurred and sharp stripe edges used in the edge detection test. In the actual experiment, both stripe edges were either sharp or blurred in a given trial. JW had unlimited time to observe the stimuli. Nevertheless, JW could not discriminate between sharp and blurred edges at any level of blurring nor point towards the exact location of the sharp edge. (B) JW showed low contrast sensitivity, especially for high spatial frequencies, which are important for sharp edge processing. (C and D) Example of stimuli used in the contour integration test. JW could not perceive contours defined by aligned Gabor patches embedded within noise (C) nor could he perceive the illusory contours in Kanizsa triangles of any shape (D), even when the expected contours were explicitly described to him.

repetitions. The participant was seated 32 in. from the 17 in. monitor (45 pixels/in.) in a dark room and the distance from the monitor was fixed using a chinrest. We used Little Stimulus Maker software (<http://faculty.washington.edu/jokelly/>) to create the stimuli and contrast adjustment. With the

devices used in this experiment, the software was able to produce stimuli with contrast levels higher than 0.5%. Contrast was adjusted by pressing on one of the two keys to increase or decrease contrast (by 0.5% in each step), until the stimulus was detected repeatedly in six reversals.

5.1.4. Contour integration tests

Tests of contour integration were presented only to JW. The stimuli were printed on a sheet of letter paper and JW was given unlimited time to observe them. In the illusory contour test, JW was shown a set of Kanizsa's triangles (Fig. 6D) differing in overall dimension and inducer sizes. JW was asked, first, to report what he sees. Then he was asked to identify a triangle in any of the pictures. Additionally, several stimuli with misaligned inducers were shown and he was asked to discriminate between displays where the inducers create a shape and displays where they do not. In a second test, we tested contour detection in stimuli contained a closed path (contour) of high contrast Gabor signals embedded in a random background of Gabor signals. The stimulus was chosen from the easiest level stimulus set used by Pennefather, Chandna, Kovacs, Polat, and Norcia (1999) (Fig. 6C). JW was asked first to describe what he sees. Then he was asked to point at items that he sees. Finally, he was asked whether he can see a circle (closed contour of local elements) embedded within the elements, and, if so, to trace it with his finger.

5.2. Results

5.2.1. Width discrimination test

Fig. 5 presents the performance in the width discrimination test for JW and the control participants for stripes with width difference of 0.7° . As is evident, JW performed similarly to the control participants in discriminating exemplars of narrow stimuli (members of class "A"). This is consistent with the fact that he, like the control subjects, was able to classify members of class "A" in the first classification session. More importantly, considering his very poor discrimination between the middle-size and broad stripe classes (classes "B" and "C"), his fair post-training performance with the middle-size width seems to suggest that JW gained some ability to discriminate exemplars with middle-size width (i.e. members of class "B"). Interestingly, as can be seen, control subjects also showed improved discriminability of broader stripes following classification training. At the same time, JW failed at discriminating the broadest stripe exemplars (i.e. members of class "C"). Therefore, classification training seems to improve processing of the size of the individual exemplars, in both JW and the control participants. However, for JW, the range of enhanced perception seems to be largely limited to the narrow and medium stripe width, at best. Note, though, that perceiving the width of the stripe of exemplars in two classes should suffice for classifying the whole exemplar range into three classes.

5.2.2. Edge perception test

In contrast to the somewhat improved stripe perception at some widths, JW was unable to detect stripe edges at any of the blurring levels. Indeed, he was so frustrated with the task and so distressed at his inability to perceive sharp edges that we terminated the experiment before completion. Note that,

not surprisingly, this task was trivial for the control participants even prior to classification training.

5.2.3. Contrast sensitivity test

As shown in Fig. 6B, JW showed much lower contrast sensitivity at all tested spatial frequencies, relative to the control group. It is also the case that, for the control group, contrast sensitivity function was very similar before (not shown) and after classification training.

5.2.4. Contour integration tests

JW's impaired edge processing may also relate to poor contour integration, which may underlie perception of the continuation of stripe's edges. Despite his ability to detect the local elements in both the Kanizsa and good continuation Gabor tests, JW could not identify the triangle in Kanizsa test, was unable to integrate contours to shapes and failed to identify the closed shape made of Gabor stimuli.

5.3. Discussion

Similar to that of the control group, JW's classification training seemed to improve processing the size of the individual exemplars, though his improvement was limited to stimuli of narrow and medium stripe width. Though incomplete, JW's perceptual improvement may potentially suffice for reasonable classification performance to the narrow, medium and broad stripes (where the broad stripes are regarded as the class of the ill-perceived stripes).

Contrary to the somewhat improved size perception, classification training did not improve JW's ability to detect stripe edges. This lack of improvement is consistent with JW's low contrast sensitivity compared to the control group, and may be attributed to processes as low as V1. It is also consistent with his failure to perceive illusory contours and to derive good continuation (Fig. 6C and D), reflecting impaired processing of grouping, attributed to intermediate processing levels, including V2 (e.g. Merigan, Nealey, & Maunsell, 1993). These failures demonstrate the severity and persistence of his impairment in lower/intermediate level shape processing. In light of this profound impairment, it is particularly striking that he could learn to classify based on stripe size.

6. General discussion

This study adopts the novel approach of training a patient with profound visual agnosia on a task that requires classification of simple stripe stimuli into three different classes. In contrast with most other studies, which have retrained individuals with lesions to visual cortex, the approach adopted here was to use simple inputs and to attempt to engage higher-level visual cortex in a task that entails generalizing and retrieving class representations. The central finding of this study is that JW, a severely agnostic individual, was able to

learn to classify visual stimuli in the service of a classification task, despite the extensive lesion to lower-level visual areas and the persistent deficit in discriminating the very same stimuli. This result is particularly intriguing given that JW's lesion was bilateral and extensive, affecting extrastriate cortex, especially area V2, which serves as the major source of visual input to higher cortical areas, and some of striate cortex. In addition to learning to classify the stimuli, JW was able to acquire and maintain the class representations over the relatively long term and was able to retrieve them in response to symbolic cues at this later stage. These findings are remarkable as the classification demanded the association of different stimuli to a single representative and differentiation between different individual stimuli that are assigned to different class representatives. Interestingly, while improving at classification, at the same time, JW remained largely impaired at perceiving the simple individual stimuli. Importantly, JW's ability to learn to classify the stimuli – despite the massive loss of input to high-level visual areas and the resulting impaired perception – supports the possibility that the visual system can be modified to tolerate degraded information resulting from degraded lower-level visual processing.

What mechanism might account for JW's improved ability to use the degraded low-level input for classification? One possibility is that the multi-session training induced high-level visual processes to adapt and tolerate the degraded input from lower levels. Another possibility is that partial recovery of residual lower-level neural networks during training improved the robustness of the input sufficiently to allow high-level visual processing. Two findings, namely, the abnormal slowness of JW's classification learning and his seemingly improved stripe discrimination, are consistent with this latter possibility. A combination of both high-level acquisition and improved residual low-level processing may also be a possibility, of course. We address the likelihood of each of these possible mechanisms in turn.

6.1. Learning based exclusively on high-level input completion

Given the extent and severity of JW's lesion in lower-level areas – especially V2, which serves as a primary source of input to high-level visual areas – it seems unlikely that lower-level areas alone can provide reliable information for use by the higher-level areas. Yet, the reliability of the information may be increased by re-sampling the residual lower-level inputs, and by adapting to the degraded information by generalization across sampling iterations. Such a process of learning by re-sampling partial input seems to be compatible with high-level visual processing. In cases of severe degradation of the input, as in the case of JW, tolerance to partial information, supposedly, requires many iterations of sampling of the degraded information, making learning gradual.

The difficulty in acquiring flexible and persisting class representations after a lesion affecting more anterior regions (as well as more posterior regions) is also illustrated in a training

study with a patient with visual agnosia following bilateral medial occipito-temporal infarcts (and a right postero-lateral infarct; Scheindler et al., 1992). Although this patient, Mr. M, could learn to classify simple stimuli such as checkerboards and Gabor patches into classes, the learning period was very protracted and even when learning did take place, it was very unstable. Given that Mr. M continued to have difficulty assigning stimuli to classes even when he had achieved criteria in some cases led the authors to conclude that the patient suffers from a weakness in constructing reliable internal prototypes for pattern classification. That JW had spared high-level cortex whereas this was not the case for Mr. M may explain the difference in their training outcomes with the sparing of higher-level visual areas being critical for the acquisition of the class structure.

However, it is unclear whether high-level plasticity alone suffices for developing the tolerance to degraded input. As discussed above, even in the case where the degraded information is due to perturbation of the visual input per se, tolerance of high-level visual areas to partial information in monkeys is accompanied by improved processing at lower levels (Rainer et al., 2004). Moreover, on its own, high-level tolerance to the degraded input seems unlikely to explain the findings of JW's improvement in discrimination of the individual stripe stimuli during classification learning, as high-level areas show poor selectivity to simple individual stripes (e.g. Malach et al., 1995).

6.2. Training-related recovery in residual low-level networks

An alternative explanation for JW's ability to utilize degraded information for visual classification is that of training-related recovery of residual low-level networks in the lesioned area. Such recovery may allow sufficient input from low-level areas for classification processing by high visual processing levels. This training-related recovery also seems a likely explanation for the observed improvement in discriminating individual stripe widths. This partial recovery in the residual low-level network may be mediated by means of, for example, enhanced existing connections between the residual neurons or even growth of new axons and synapses, as a result of repeated exposure to the stimuli. The partial recovery at low levels of processing may, in turn, provide better input to higher visual processing levels, to an extent that allows classification learning.

One possible form of partial recovery at low levels in JW's case, although speculative, is an increase in receptive field size of residual extrastriate neurons, which are most likely narrowed by the extensive ischemic lesion, due to a significant reduction of input summation in the local circuits. Narrowed receptive field size could potentially underlie JW's highly impaired global vision (e.g. grouping see Table 1), as well as his initial inability to classify stripes wider than 3°. However, repeated and consistent exposure may help to modify or even create new connections between the residual

neurons, allowing for a gradual increase in their input summation and, thus an increase in their receptive field size (e.g. Schweigart & Eysel, 2002; Spillmann & Werner, 1996). Accordingly, the slow and gradual classification learning in the case of JW matches the learning dynamics in perceptual learning that is thought to be linked to low/intermediate areas (e.g. Fahle, 1993; Hochstein & Ahissar, 2002; Karni & Sagi, 1991). This explanation predicts increased islands of activity in the damaged extrastriate area following classification training. Such a prediction may possibly be verifiable using fMRI. Unfortunately, this cannot be done with JW because of the scanning contraindications mentioned above.

6.3. *Bi-directional improvement*

The fact that JW improved both in a task that requires high visual processing level (classification) and, apparently, in a task that requires low processing level (stripe discrimination) suggests that plasticity may, actually, be guided by residual recurrent loops of top-down and bottom-up interactions (e.g. Foxe & Simpson, 2002; Gilbert et al., 2001; Grossberg, 2001; Hochstein & Ahissar, 2002; Juan & Walsh, 2003; Juan, Campana, & Walsh, 2004; Lamme & Roelfsema, 2000; Lee & Mumford, 2003; Mumford, 1992; Pollen, 1999; Ullman, 1984). It is possible then that, during training, degraded input from the residual neurons in lower-level areas may initially cause some limited modification at high visual processing levels, which, even if initially insufficient for performing classification, may induce some plasticity there. In turn, via top-down projections, these small higher-level changes may affect some change in residual lower-level networks, relevant to the task processing. These lower-level changes may lead to improved processing of the individual stripes which, again, could improve processing at high levels. Such a recurrent process may require many iterations to enable processes of formation and maintenance of class representation in high-level visual areas. Note that this suggestion is one of bootstrapping, which is initiated as soon as some, even minimal, bottom-up input is propagated. Though speculative too at this stage, the appeal of such a dynamic, reciprocal process is that it does not require intact low-level processing to get the system underway. Partial input can trigger the onset and, through gradual mutual refinement, learning can be expressed at both higher- and lower-levels in tandem.

One implication of this account is that the dependence of low-level plasticity on top-down input is contingent on the amount and quality of learning, since top-down processes are assumed to “highlight” task-relevant neuronal processes. Because of this task-relevant emphasis, then, the task used during training is especially critical. Previous studies with training after visual cortex lesions have tended to use detection or discrimination tasks that do not rely on high-level visual processing. Such training has led to improvements, albeit with somewhat slow and limited learning (hundred thousands trials, hundreds of hours; e.g. Huxlin & Pasternak, 2004; Kasten et al., 2000; Widdig, Pleger, Rommel, Malin,

& Tegenthoff, 2003). In contrast, in the current study, where the task depended critically on generalization and long-term representation, fewer than 7 h (about 7000 trials) sufficed for inducing significant improvement in JW’s processing. Of course, factors other than task demands may underlie the difference in the learning rates across the different studies. Note that in the current situation, simple stimuli were adopted as we reasoned that more complex stimuli would not be processed sufficiently by JW. Whether similar acquisition of classification skills would succeed with more complex inputs remains to be determined.

There are several supportive indications for the involvement of top-down processes in the pattern of JW’s learning. As previously suggested, top-down processes may guide lower-level processes to be focused on those stimuli that are relevant for the improved performance. Indeed, it seems that JW’s improvement in width discrimination was limited to the stripe exemplars of the intermediate width class (class “B”). This may also reflect a “minimal-necessary” improvement “policy”. That is, since perception of individual members of two classes suffices for classifying the whole exemplar range into three classes (i.e., by regarding the third class as the class of the difficult stimuli), there is no real need to exert additional effort to improve processing of members in the third class. Such class-related improvement supports facilitation by top-down input from areas that process stimulus classes (the relevant context) to areas that process the individual stimuli.

The dependence of low-level plasticity on task relevance (and thus guiding of top-down effects) may also explain the lack of evidence for improvement in processing form or “shape”. Recall that JW’s detection of stripe edges and contours following the lengthy exposure to the stripes during classification training was very poor. The absence of change to this visual ability may be an artifact of the experimental design: the classification task was based on comparison between different stripe widths where the overall shape of the stimuli was fixed. Following learning, JW was able to make size judgments while exhibiting impaired shape perception. It is possible that the seemingly lack of improvement in shape perception may be a consequence of the fact that shape was not a relevant attribute for the task. It is also possible, of course, that shape processing may demand richer and more specific computations, compared to size, and that these more complex computations are beyond the processing limitations of the residual network. To adjudicate between these accounts would require developing a training regimen for JW with shape as its focus and then assessing any observable changes in shape perception.

In summary, we found that learning of a task that is thought to require higher visual processing, involving the acquisition and maintenance of class representations, can occur even in the face of severe and extensive damage in low/intermediate level visual areas. This result is interesting as these earlier areas serve as an important source of input to the high-level visual areas and a standard feedforward account of visual

processing would not be tenable. In parallel, we find evidence for training-related improved perception of individual stimuli, which may reflect plasticity and some recovery at low and/or intermediate stages of visual processing, as well. These findings suggest that even in a case of an extensive lesion, learning involves modification in both high and residual low-level networks, and is mediated by top-down and bottom-up interactions. These findings are encouraging and suggest potential directions for rehabilitation, namely the recruitment of intact higher visual processing levels, which can adapt to the partial or degraded visual information even in cases with profound visual agnosia.

Acknowledgements

We thank JW for his participation in this study, Drs. Jason Barton and Keith Thulborn for reviewing JW's CT scan, Drs. Carl Olson, Ione Fine, Tai-sing Lee and Galia Avidan for their insightful comments. This work was supported by NIH grant 54246.

References

- Baker, C. I., Behrmann, M., & Olson, C. (2002). Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. *Nature Neuroscience*, 5(11), 1210–1216.
- Balliet, R., Blood, K. M. T., & Bach-y-Rita, P. (1985). Visual field rehabilitation in the cortically blind? *Journal of Neurology, Neurosurgery and Psychiatry*, 48, 1113–1124.
- Behrmann, B., Marotta, J., Gauthier, I., Tarr, M. J., & McKeef, T. J. (2005). Behavioral change and its neural correlates in visual agnosia after expertise training. *Journal of Cognitive Neuroscience*, 17(4), 554–568.
- Cox, D., Meyers, E., & Sinha, P. (2004). Contextually evoked object-specific responses in human visual cortex. *Science*, 304, 115–117.
- Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, 4, 519–525.
- Foxe, J. J., & Simpson, G. V. (2002). Flow of activation from V1 or frontal cortex in humans. A framework for defining “early” visual processing. *Experimental Brain Research*, 142, 139–150.
- Erickson, C. A., Jagadeesh, B., & Desimone, R. (2000). Clustering of perirhinal neurons with similar properties following visual experience in adult monkeys. *Nature Neuroscience*, 3(11), 1143–1148.
- Fahle, M. (1993). Visual learning in the hyperacuity range in adults. *German Journal of Ophthalmology*, 2(2), 83–86.
- Farah, M. J. (2004). *Visual agnosia* (2nd ed.). Cambridge, MA: MIT Press.
- Freedman, D. L., Reisenhuber, M., Poggio, T., & Miller, E. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science*, 291, 312–316.
- Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2003). A comparison of primate prefrontal and inferotemporal cortices during visual categorization. *The Journal of Neuroscience*, 23(12), 5235–5246.
- Furmanski, C. S., Schluppeck, D., & Engel, S. A. (2004). Learning strengthens the response of primary visual cortex to simple patterns. *Current Biology*, 14(7), 573–578.
- Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform ‘face area’ increases with expertise in recognizing novel objects. *Nature Neuroscience*, 2(6), 568–573.
- Ghose, G. M., Yang, T., & Maunsell, J. H. (2002). Physiological correlates of perceptual learning in monkey V1 and V2. *Journal of Neurophysiology*, 87(4), 1867–1888.
- Gilbert, C. D., Sigman, M., & Crist, R. E. (2001). The neural basis of perceptual learning. *Neuron*, 31, 681–697.
- Grossberg, S. (2001). Linking the laminar circuits of visual cortex to visual perception: Development, grouping, and attention. *Neuroscience and Biobehavioral Reviews*, 25(6), 513–526.
- Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36(5), 791–804.
- Horton, J. C. (2005). Disappointing results from Nova Vision's visual restoration therapy. *British Journal of Ophthalmology*, 89, 1–2.
- Huxlin, K. R., & Pasternak, T. (2004). Training-induced recovery of visual motion perception after extrastriate cortical damage in the adult cat. *Cerebral Cortex*, 14(1), 81–90.
- Jagadeesh, B., Chelazzi, L., Mishkin, M., & Desimone, R. (2001). Learning increases stimulus salience in anterior inferior temporal cortex of the macaque. *Journal of Neurophysiology*, 86, 290–303.
- Juan, C. H., Campana, G., & Walsh, V. (2004). Cortical interactions in vision and awareness: Hierarchies in reverse. *Progress in Brain Research*, 144, 117–130.
- Juan, C. H., & Walsh, V. (2003). Feedback to V1: A reverse hierarchy in vision. *Experimental Brain Research*, 150(2), 259–263.
- Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences of the United States of America*, 88(11), 4966–4970.
- Kasten, E., Poggel, D. A., & Sabel, B. A. (2000). Computer-based training of stimulus detection improves color and simple pattern recognition in the defective field of hemianopic subjects. *Journal of Cognitive Neuroscience*, 12(6), 1001–1012.
- Kasten, E., Wüst, S., Behrens-Baumann, W., & Sabel, B. A. (1998). Computer-based training for the treatment of partial blindness. *Nature Medicine*, 4(9), 1083–1087.
- Kobatake, E., Wang, G., & Tanaka, K. (1998). Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *Journal of Neurophysiology*, 80, 324–330.
- Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neuroscience*, 23, 571–579.
- Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America A: Optics, Image Science and Vision*, 20(7), 1434–1448.
- Malach, R., Reppas, J. B., Benson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A., et al. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 92(18), 8135–8139.
- Mapelli, D., & Behrmann, M. (1997). The role of color in object recognition: Evidence from visual form agnosia. *Neurocase*, 3, 237–247.
- Marotta, J. J., Behrmann, M., & Goodale, M. A. (1997). The removal of binocular cues disrupts the calibration of grasping in patients with visual form agnosia. *Experimental Brain Research*, 116, 113–121.
- Merigan, W. H. (2000). Cortical area V4 is critical for certain texture discriminations, but this effect is not dependent on attention. *Visual Neuroscience*, 17, 949–958.
- Merigan, W. H., Nealey, T. A., & Maunsell, J. H. (1993). Visual effects of lesions of cortical area V2 in macaques. *The Journal of Neuroscience*, 13(7), 3180–3191.
- Merigan, W. H., & Saunders, R. C. (2004). Unilateral deficits in visual perception and learning after unilateral inferotemporal cortex lesions in macaques. *Cerebral Cortex*, 14, 863–871.
- Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, 335, 817–820.

- Miyashita, Y., & Hayashi, T. (2000). Neural representation of visual objects: Encoding and top-down activation. *Current Opinions in Neurobiology*, 10(2), 187–194.
- Mumford, D. (1992). On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biological Cybernetics*, 66(3), 241–251.
- Pasternak, T., & Merigan, W. H. (1994). Motion perception following lesions of the superior temporal sulcus in the monkey. *Cerebral Cortex*, 4, 247–259.
- Pennefather, P. M., Chandna, A., Kovacs, I., Polat, U., & Norcia, A. M. (1999). Contour detection threshold: Repeatability and learning with 'contour cards'. *Spatial Vision*, 12(3), 257–266.
- Plant, G. T. (2005). A work out for hemianopia. *British Journal of Ophthalmology*, 89, 2.
- Pollen, D. A. (1999). On the neural correlates of visual perception. *Cerebral Cortex*, 9, 4–19.
- Rainer, G., Lee, H., & Logothetis, N. K. (2004). The effect of learning on the function of monkey extrastriate visual cortex. *PLoS Biology*, 2(2), 0275–0283.
- Rainer, G., & Miller, E. K. (2000). Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron*, 27, 179–189.
- Reinhard, J., Schreiber, A., Schiefer, U., Kasten, E., Sabel, B. A., Kenkel, S., et al. (2005). Does visual restitution training change absolute homonymous visual field defects? A fundus controlled study. *British Journal of Ophthalmology*, 89, 30–35.
- Ricci, R., Vaishnavi, S., & Chatterjee, A. (1999). A deficit of intermediate vision: Experimental observations and theoretical implications. *Neurocase*, 5, 1–12.
- Rosenthal, O., Fusi, S., & Hochstein, S. (2001). Forming classes by stimulus frequency: Behavior and theory. *Proceedings of the National Academy of Sciences of the United States of the America*, 98(7), 4265–4270.
- Rossion, B., Gauthier, I., Goffaux, V., Tarr, M. J., & Crommelinck, M. (2002). Expertise training with novel objects leads to left-lateralized face-like electrophysiological responses. *Psychological Science*, 13(3), 250–257.
- Sakai, K., & Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. *Nature*, 354, 152–155.
- Scheindler, W., Landis, T., Rentschler, I., Regard, M., & Baumgartner, M. (1992). A pattern-recognition approach to visual agnosia. *Clinical Vision Sciences*, 7(3), 175–193.
- Schneider, W., Eschman, A., & Zuccolotto, A. (2002). *E-Prime user's guide*. Pittsburgh: Psychology Software Tools Inc.
- Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, 412, 549–553.
- Schwartz, S., Maquet, P., & Frith, C. (2002). Neural correlates of perceptual learning: A functional MRI study of visual texture discrimination. *Proceedings of the National Academy of Sciences of the United States of the America*, 99(26), 17137–17142.
- Schweigart, G., & Eysel, U. T. (2002). Activity-dependent receptive field changes in the surround of adult cat visual cortex lesions. *European Journal of Neuroscience*, 15(10), 1585–1596.
- Sheinberg, D. L., & Logothetis, N. K. (2002). Perceptual learning and the development of complex visual representations in temporal cortical neurons. In M. Fahle & T. Poggio (Eds.), *Perceptual learning*. MIT Press.
- Sigala, N., & Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*, 415, 318–320.
- Spillmann, L., & Werner, J. S. (1996). Long-range interactions in visual perception. *Trends in Neuroscience*, 19, 428–434.
- Ullman, S. (1984). Visual routines. *Cognition*, 18(1–3), 97–159.
- Vecera, S. P., & Behrmann, M. (1997). Spatial attention does not require preattentive grouping. *Neuropsychology*, 11(1), 30–43.
- Vecera, S. P., & Gilds, K. S. (1997). What is it like to be a patient with apperceptive agnosia? *Consciousness and Cognition*, 6, 237–266.
- Widdig, W., Pleger, B., Rommel, O., Malin, J. P., & Tegenthoff, M. (2003). Repetitive visual stimulation: A neuropsychological approach to the treatment of cortical blindness. *NeuroRehabilitation*, 18(3), 227–237.
- Yang, T., & Maunsell, J. H. R. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *The Journal of Neuroscience*, 24(7), 1617–1626.
- Zihl, J., & Von Cramon, D. (1985). Visual field recovery from scotoma in patients with postgeniculate damage. *Brain*, 108, 335–365.