

Abrupt learning and retinal size specificity in illusory-contour perception

Nava Rubin^{*†}, Ken Nakayama^{*} and Robert Shapley[‡]

Background: In behavioral studies of learning, a distinction is commonly made between gradual and abrupt improvements in performance. The learning of perceptual and motor skills is often characterized by gradual, incremental improvement, and is found not to generalize over stimulus manipulations such as changes in the size or location of the retinal image. In contrast, marked improvement in performance can occur suddenly—a phenomenon which has been termed ‘insight’. Consequently, the brain mechanisms subserving the two types of learning are commonly thought of as distinct. Here, we examine learning of a perceptual task in which improvement appears to exhibit characteristics of both gradual and abrupt learning.

Results: We describe experiments on illusory-contour perception in which the observers underwent an abrupt, dramatic improvement in performance, resembling an incident of insight. At the same time, however, the phenomenon showed a degree of stimulus-specificity that was previously thought to characterize incremental, gradual learning. The improvement was triggered only by specific visual stimuli, whereas other, quite similar, stimuli were found to be ineffective for training; the learning did not generalize to a new retinal image size, and re-training was necessary for different-sized images.

Conclusions: The juxtaposition of abrupt and stimulus-specific learning that we observed suggests that the distinction between the two forms of learning needs to be revised. Rather than postulating two distinct mechanisms, incremental and insightful learning need to be addressed within a single framework. In particular, the findings suggest that learning may involve interactions between multiple levels of representations of the stimulus.

Background

Learning is a fundamental property of living organisms. It allows us to adapt, so that when we are faced with a situation a second time we can cope with it better. Although behavioral and physiological studies of learning have taught us much about the underlying mechanisms, many central issues still await resolution. One of these issues regards a distinction that is commonly made in behavioral studies, between gradual and abrupt learning. Gradual, incremental improvement is often observed in the acquisition of perceptual or motor skills [1–9]. With other tasks, however, dramatic improvement can occur abruptly. This is commonly observed in ‘problem-solving’ tasks and has been given the name ‘insight’ [10]. Besides the different timescales of the two types of learning, another characteristic differentiates them. Many forms of perceptual learning are found to be stimulus-specific: improvement does not transfer across stimulus attributes such as the size or location of the retinal image [1–9]. In contrast, improvements due to insight appear not to be susceptible to such ‘superficial’ changes of circumstances. Consequently, the two types of learning are commonly thought to involve

different underlying brain mechanisms. It has been suggested that the incremental, stimulus-specific improvement observed in perceptual skills involves synaptic modifications in early cortical areas (such as visual areas V1 and V2) [5–6,9]. In contrast, the sudden improvement in performance observed in insight phenomena is usually taken to indicate a cognitive event which occurs more centrally.

Hebb [11] argued that this dichotomy may be artificial, and that the two forms of learning may share common mechanisms. He observed that insightful behavior may be present even in cases in which an initial phase of poor performance is not observed, and the animal finds the solution immediately. He went on to ask “is insight or hypothesis — or, in the broadest terms, intelligence — something distinct from the mechanism of association?” Addressing this question experimentally has proved to be a difficult task. One of the reasons for this difficulty is that insight and incremental learning are commonly studied in different behavioral domains — typically, using problem-solving for the former and perceptual or motor performance for the

Addresses: ^{*}Vision Sciences Lab, Harvard University, 33 Kirkland Street, Cambridge, Massachusetts 02138, USA. [†]Center for Neural Science, New York University, 4 Washington Place, New York, New York 10003, USA.

Present address: [‡]Center for Neural Science, New York University, 4 Washington Place, New York, New York 10003, USA.

Correspondence: Nava Rubin
E-mail: nava@cns.nyu.edu

Received: 28 February 1997

Revised: 18 April 1997

Accepted: 8 May 1997

Published: 6 June 1997

Current Biology 1997, 7:461–467

<http://biomednet.com/eleceref/0960982200700461>

© Current Biology Ltd ISSN 0960-9822

Figure 1




Abrupt transitions in perceptual organization. On first exposure, most observers perceive the figure as a random collection of black spots; after a suggestion that a Dalmatian dog is embedded in the picture, however, most observers recognize it. In this new interpretation of the image, the organization of the scene changes: the spots comprising the dog are seen as grouped together, and segregated as a figure from the background. An illusory outline of the dog's body is often reported. The transition from one perceptual state to the other is abrupt, resembling an incident of insight. (Photograph by Ronald C. James.) The insert shows a Kanizsa-square [15] supported by inducers which are small relative to its size; in this situation, the illusory shape is perceptually weak. In our experiment, observers underwent an abrupt transition into perceiving salient illusory shapes in such displays.

latter. Another difficulty is gaining experimental control over precisely when the unique event of insight occurs, so that data from many such events can be collected and analyzed. Perceptual learning offers an appealing framework in which to address these issues. First, abrupt transitions in performance can occur in perception—consider the well-known example of the camouflaged Dalmatian dog (Figure 1). Second, we show here that, with careful experimental titration, it is possible to gauge the stimuli so that the perceptual transition will occur in a predictable time.

Results

While running an experiment in a previous study about illusory-contour perception, anecdotal observations suggested that subjects sometimes showed a remarkably fast improvement, which occurred over the course of a few trials. At the same time, it was observed that, when the subject was exposed to the same stimulus but with a new retinal image size, performance fell back down, followed by a reoccurrence of fast improvement with the introduction of 'easy' trials. We were intrigued by this observation because it suggested that the theoretical claims that stimulus-specificity indicates an early site of plasticity would have to be applied here, too. This seemed at odds

Figure 2



A shape-discrimination task based on the perception of illusory contours. The inducers of a Kanizsa square [15] were rotated by a variable degree, resulting in the perception of curved illusory surfaces of (a) fat or (b) thin shapes. The subjects were required to choose between the two alternatives. The range of curvatures used was varied from one experimental block to the other, thus allowing us to control for the level of difficulty and the onset of the abrupt learning.

with the extreme rapidity of the improvement observed. Furthermore, a number of observers spontaneously reported that their improved performance followed a perceptual transition in how the stimulus was organized (see below). We therefore set out to test whether it is possible to use the illusory shape discrimination task to obtain insight-like learning— an initial phase of no-improvement followed by an abrupt jump in performance— while maintaining the stimulus-specific nature of the learning. In this paper, we report results from two experiments that demonstrate that such a juxtaposition can indeed be induced.

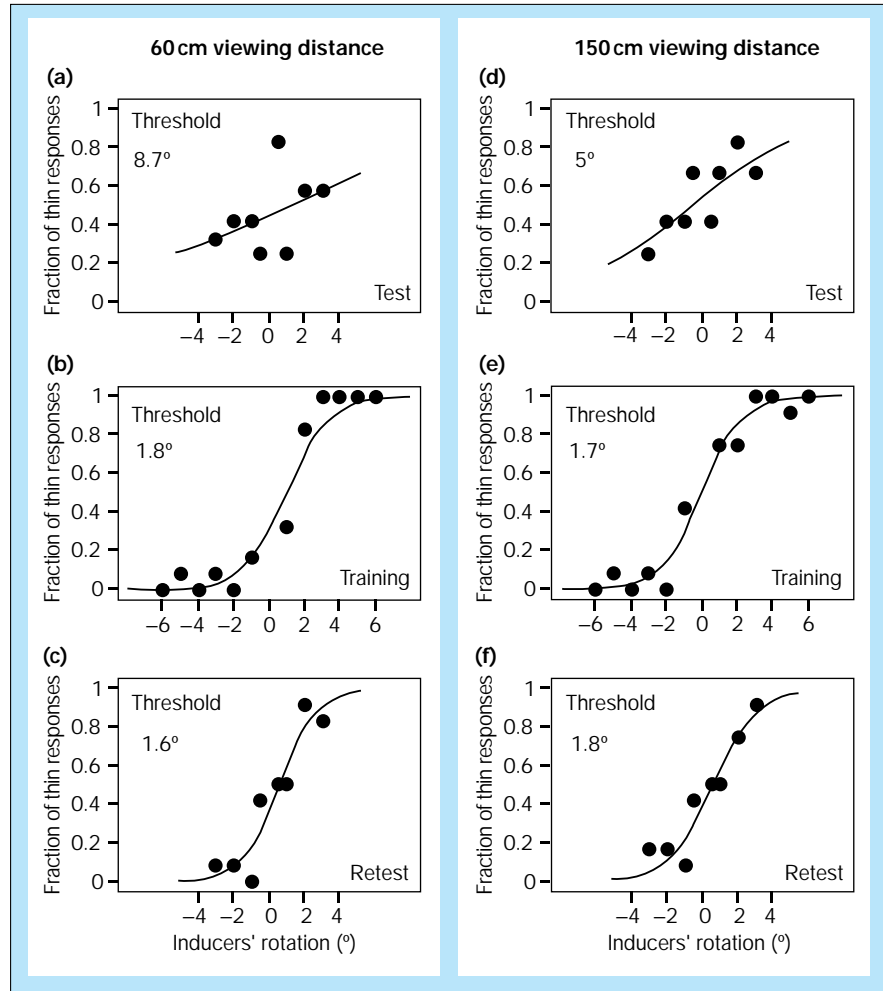
Experiment 1: abrupt learning and its specificity to the retinal image size

In the first experiment, the task used was similar to one we have used before [12–14]; observers discriminated between two classes of illusory shapes [15,16] that were globally defined by four inducers located at the corners of a square. Rotating the inducers about their centers made the illusory shapes appear 'thin' or 'fat' (Figure 2). Performance was measured by the angle of rotation of the inducers needed to yield reliable discrimination. Previous work showed that perception of the curved illusory contours significantly facilitates performance, compared with discriminating the orientation of the inducers [12–14].

In order to establish that they understood the task, subjects were given a 'practice' session, in which the illusory shapes were highly visible as a result of the large size of the inducers. The support ratio, defined as the ratio between the luminance-supported part of the illusory edge and the total edge length, was 0.4 in this practice block. Subjects were required to give at least 17 correct responses out of 20 trials in their first or second practice block in order to participate in the experiment. Once they passed this criterion, they were given the experimental blocks, in which the illusory-contour stimuli were less salient—because of the smaller size of the inducers (support ratio 0.25)—making the task more demanding.

Figure 3

Performance of an individual subject (AH) in six blocks of the illusory shape-discrimination task. For each block, the fraction of times that the subject judged the stimulus to be thin was plotted as a function of the inducers' rotation angle. A sigmoid curve was fit to the data (see Materials and methods), and the threshold (shown in each panel) was estimated from the fitted curve. (a–c) Performance of AH in three consecutive blocks of the task at a viewing distance of 60 cm. (a) In the 'test' block, when the range of rotation of the inducing elements was $0.5\text{--}3^\circ$, performance was poor. (b) In the 'training' block, when large-curvature stimuli ($4\text{--}6^\circ$, as well as longer-duration stimuli; see text) were added to the set, AH improved markedly on the 2° and 3° stimuli, which were identical to those used in the test block. (c) The 'retest' block was a repeat of the stimulus set used in (a); after exposure to the training block, the subject was able to discriminate these stimuli reliably. (d–f) The subject was moved to a viewing distance of 150 cm; performance in the test block (d) revealed that the learning exhibited in (b,c) was specific to the retinal image size used. (e) Exposure to high-curvature stimuli again triggered marked improvement, and (f) the subject's final performance in the retest block was similar to that in the 60 cm viewing condition.

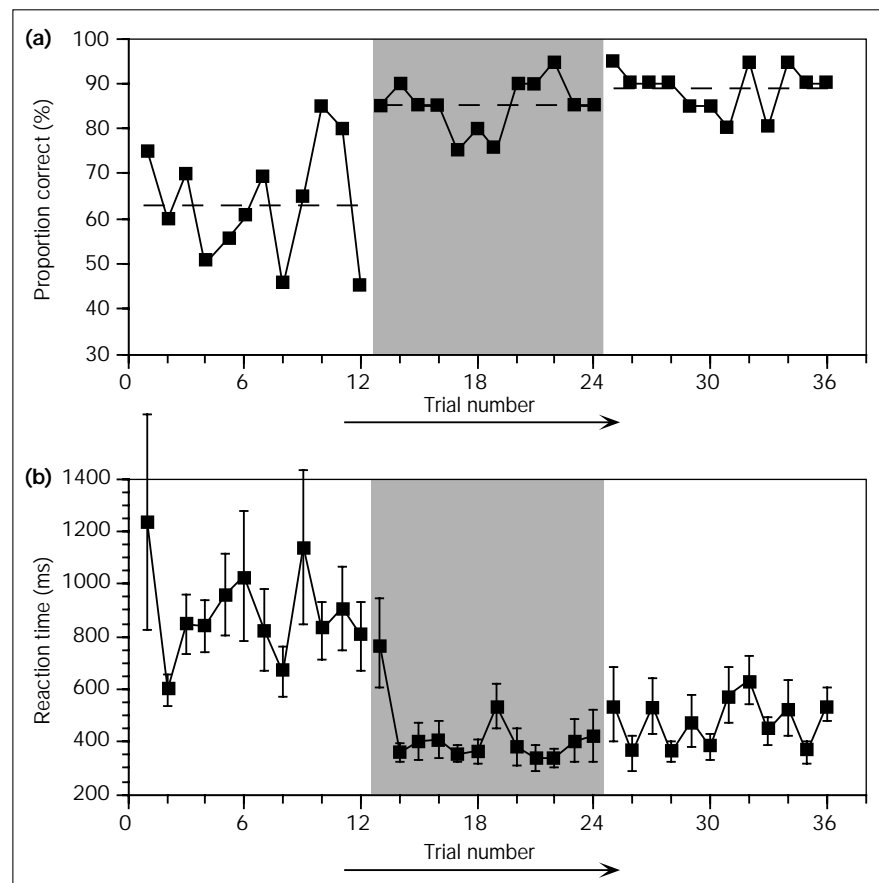


Repeated presentation of stimuli with low curvature ($< 3^\circ$) led to little improvement in performance; however, adding illusory-contour stimuli that had greater curvature produced a sudden and lasting improvement. Figure 3a–c shows the performance of an individual observer (AH) in three consecutive blocks. The first block (Figure 3a) was a 'test' block consisting of stimuli in which the inducers' rotation angles were small ($0.5\text{--}3^\circ$); in this block, performance was poor (the threshold was 8.7°). In the second, 'training' block (Figure 3b), the small-angle stimuli were embedded in a set of new stimuli, with larger inducer rotation angles ($4\text{--}6^\circ$). In addition, large-curvature stimuli ($4\text{--}6^\circ$) of longer exposure duration (194 ms) were intermixed. The observer's performance improved dramatically (the threshold was 1.8°). Note that AH correctly discriminated in this block between thin and fat stimuli with curvature values of 2° and 3° in 92% and 96% of the trials, respectively, compared with only 58% and 63% for identical stimuli in the previous block. The sudden improvement was not due to a lack of cognitive understanding of

the task in the first block — AH got 20/20 trials correct in the practice block. The third, 'retest' block (Figure 3c) consisted of a stimulus set identical to that of the first block. The good performance seen for the training block was maintained; the learning was also retained in a repeat of the retest block several days later (see below). From our previous studies [12–14] we can deduce that the level of performance after the learning indicates that subjects were basing their judgments on perceived illusory contours, whereas the level of performance before training, in the test block, is characteristic of a strategy based on the local inducers' orientation.

How abrupt is the improvement in performance? To answer this question, we performed a trial-by-trial analysis of the performance of a group of 10 subjects for repetitions of the same stimulus. We selected the pair of stimuli which had inducers' rotation angles of $+2^\circ$ and -2° , and examined how the subjects performance for this pair of stimuli evolved with time. Datapoints in Figure 4a are the

Figure 4



A trial-by-trial analysis of performance as a function of time. Performance on each successive pair of $+2^\circ$ and -2° stimuli was separately tabulated, averaged over 10 subjects, and plotted as a function of time. The first 12 trials were in the test block, the next 12 trials (shaded area) show the performance when identical $\pm 2^\circ$ stimuli were embedded in the training block, and the final 12 trials were in the retest block. (a) The proportion of correct answers abruptly jumped, from 63% to 85%, right at the onset of the training block ($P < 0.0001$), after no indication of improvement during the first 12 trials of the test block (slope 10^{-3} , $r < 10^{-4}$). (b) Reaction times showed a sharp drop of 460 ms on average after presentation of the training stimuli.

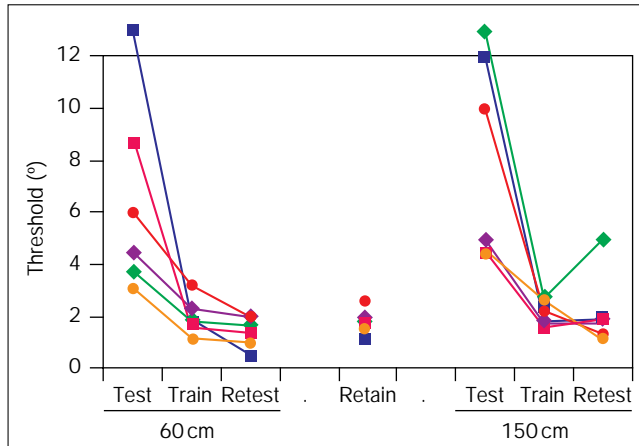
averages of the proportion of correct responses the group of subjects gave to the $+2^\circ/-2^\circ$ stimulus pair as a function of time (trial number). The group performance showed no indication of improvement during the test block, and then abruptly jumped right at the transition between the test and training blocks. This means that exposure to very few large-curvature stimuli (2.7 trials on average) was sufficient to trigger the sharp improvement in performance. This is a different form of improvement from that usually reported in perceptual learning studies—a gradual (even if sometimes fast [8,9]) increase in performance [1–9]. We shall return to the implications of this distinction in the Discussion.

Figure 4b shows the trial-by-trial analysis of the reaction times to the $\pm 2^\circ$ stimuli. A sharp drop of roughly 460 ms occurred at the beginning of the training block. Note that the subjects were not told that their reaction times were being recorded; the only emphasis in the instructions was on the correctness of responses. Thus, the sharp drop in the mean and variability of the reaction times reflects a facilitation in performing the task—which is consistent with the subjective experience of the observers. They

often reported that in the first block they did not see the global illusory shape and were basing their judgments on the local inducers. In the second block they suddenly started seeing the global shapes (sometimes noting the well-known brightness effect associated with it [15–16]).

Sudden improvements in performance are commonly thought to involve a cognitive insight into the nature of the task. It would therefore be surprising if the abrupt learning did not generalize to another retinal image size. However, this is precisely what we found. In the second phase of experiment 1 (Figure 3d–f), subject AH was moved to a viewing distance of 150 cm—the global stimuli now subtended a visual angle of 5.7° (the visual angle of the inducers was 1.4°). The stimulus set of the first block (Figure 3d) only contained illusory contours of small curvatures, like the test and retest sets in Figure 3a and Figure 3c, respectively. Performance was quite poor, as it had been for the test block at the 60 cm viewing distance. Thus, the abrupt improvement observed before (in Figure 3b,c) did not generalize to the new (smaller) retinal image size. Subsequent exposure to high-curvature stimuli in this new retinal image size (Figure 3e) again

Figure 5



Results from six naive observers summarized in terms of threshold performance as a function of block type; all showed a marked improvement between the performance in the identical stimulus sets of the test and retest blocks, which was triggered by the training block. Data in the 'retain' block show performances in a repeat of the retest block 1–7 days after the first session, and demonstrate that the learning was long-term. Finally, the learning was found not to generalize to a new retinal image size, but occurred anew after retraining at this new size. The data for the first three blocks in the 60 cm viewing condition were collected in one session. The data from the retain block in the 60 cm viewing condition and the three subsequent blocks in the 150 cm viewing condition were collected in a second session on a different day (1–7 days after the first session), to avoid fatigue of the naive subjects and to establish that the learning was long-term. Results similar to those presented here were obtained with two experienced observers who performed all blocks on the same day.

triggered rapid learning, leading to similar performance in the retest blocks for the two retinal sizes (compare Figure 3c and Figure 3f).

Figure 5 summarizes the results for six naive observers who were given the same sequence of blocks as AH: all show sharp improvement in the transition from the test to the training block, and a lack of generalization of the learned performance to the new retinal image size. The figure also shows that the learning was long-term: between one and seven days after their first session, the subjects came for another session and were given a block that consisted of a stimulus set identical to that used in the test and retest block. Their performance in shown in the 'retain' block, and is very near that of the training and retest blocks

When is abrupt learning observed?

A natural question to ask at this point is why abrupt learning of the form exhibited in Figure 4 was not found in previous studies of perceptual learning. The reason is likely to lie not in one factor, but in a combination of factors. One of the main sources of difficulty had already

been brought up by Hebb [11] who observed that, in order to induce insight, one needs "tasks ... of just the right degree of difficulty ... [they] must neither be so easy so that the animal solves the problem at once, thus not allowing to analyze the solution; nor so hard that the animal fails to solve it except by rote learning in a long series of trials."

In order to find this 'right degree of difficulty', before running the experiments described in this paper, we ran a pilot study on a large group of subjects ($n = 34$), which was aimed at characterizing the distribution of performance in our subject population. We varied the degree of difficulty of the stimuli (among subjects) by changing various stimulus parameters (primarily exposure duration and curvature). Each subject received a block of difficult trials followed by one or more blocks of easier stimuli; subjects were tested in 1–4 sessions. Under those different conditions, the onset of learning varied from abrupt to more gradual, or did not occur at all. Thus, with the very same task and training protocol, abrupt learning occurs only under certain parametric conditions. Subsequently, the stimuli for the test and training blocks were selected to match the goal formulated by Hebb ('not too easy' and 'not too difficult', respectively) for our subject population. (The subjects who participated in the two experiments were different from the ones used in the pilot study.) Fifteen of the subjects who participated in the pilot experiment, all of whom showed robust learning (either gradual or abrupt), were tested for transfer across retinal size. None of them showed transfer. The fact that stimulus-specific learning can occur across a continuum of learning rates supports the idea that abrupt and incremental learning may not have distinct underlying mechanisms.

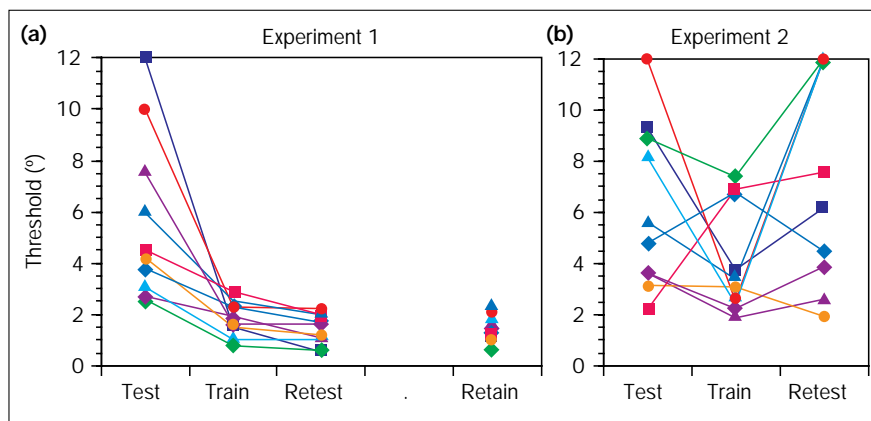
Experiment 2: specificity of the learning to the training procedure

The purpose of the second experiment was to test the specificity of the learning to the particular training stimuli. We asked whether the abrupt transition to perceiving illusory contours and the improvement associated with it require the large-curvature stimuli to be of the same support ratio (see Materials and methods) as the test stimuli, or whether any kind of 'easier' stimuli could serve for training.

A group of 10 new subjects was given three blocks of the thin/fat task. In the second (training) block, the local inducers of the high-curvature stimuli were increased in diameter, so that the support ratio was 0.4. Other than that, experiment 2 was identical to the first phase of experiment 1. Note that the size of the illusory shapes remained unaltered, and that the test (low-curvature) stimuli always had a support ratio of 0.25. Figure 6a,b presents the results for the two groups who participated in experiments 1 and 2, respectively. (The thresholds for the

Figure 6

Specificity of the training stimuli. (a) The results from experiment 1, in which 10 naive observers were given a training block with large-angle stimuli of the same inducer-size as the small-angle (test) stimuli. All subjects showed a dramatic improvement in their performance, which was maintained after the training stimuli were again removed (retest), and even several days later (retain). (b) The results from experiment 2, in which 10 subjects were given a training block with large-angle stimuli that were of larger inducer size than the test stimuli. In this case, subjects showed large individual differences. Many did not improve during the training block at all, and those who did improve during the training block did not retain the good performance once the large-angle stimuli were taken away (retest block). Thus, the abrupt learning found



in experiment 1 was triggered only by specific visual stimuli, and not by any 'easy' stimuli.

training blocks were estimated on the basis of the data from the 1–3° stimuli only, in both experiments.) Whereas all the subjects in experiment 1 improved in the training block and retained the learning in the retest block, the subjects of experiment 2 showed large individual differences in their performance. Moreover, even those subjects who improved during the training block deteriorated back to their initial (test) level of performance in the third, retest block. We conclude that the improvement in performance in experiment 1, and the accompanying transition in the perceptual organization of the small-curvature (test) stimuli into illusory surfaces, can only be triggered by large-curvature illusory contour stimuli with inducers of a similar size.

Discussion

Previously, stimulus-specificity was taken to imply that the learning occurs in early, retinotopically organized cortical areas [5–6,9]. In previous studies, the learning had two important characteristics which made the idea of an early site of plasticity seem plausible. First, the tasks were of a local nature — they involved interactions between image points 1° apart or less [1–9]. Second, the improvement was incremental: often it took place over hundreds or even thousands of trials [1–5,7], and even in cases in which fast learning phases were observed [8,9], performance showed a steep, but gradual, improvement over a few tens of trials.

In the study presented here, neither of these two characteristics is present. First, because of the large retinal distances between the inducers (more than 10° visual angle in the 60 cm viewing condition), the relevant information was stored on widely separated neurons in early visual cortical areas. Second, the learning occurred not only fast but abruptly: the subjects underwent a transition between a

'don't know' and a 'know' state, which did not show the intermediate-performance stages observed in other studies. Taken together, the abruptness and the global nature of the task make it unlikely that a model in which learning depends exclusively on quick synaptic modifications of local connectivity in early cortical areas, such as has been suggested previously for other tasks [9], could account for our results.

Yet the size specificity of the learning in our study suggests an involvement of early visual cortical areas in which retinotopic organization still exists. Thus, it is difficult to conceive of the learning as occurring at a single site, whether it be retinotopic or some higher area. Rather, our findings suggest a more global rerouting of information and knowledge structure [17]. Our preliminary results from other experiments also seem to indicate that there is a 'high-level', general component to the learning, as well as stimulus-specific components. As such, the results suggest that the dichotomy made between insight and gradual, incremental learning may need to be revised. Rather than postulating two distinct mechanisms for the two forms of learning, our findings may be better understood within a single framework, suggesting a major role for interactions between multiple levels of representation of the stimuli [18–21].

This view has already been put forward by Hebb [11], who suggested that "insight ... continually affects the learning of the adult animal," and that "it is not wholly separate from rote learning." Hebb proposed a unitary mechanism, based on the associations of co-occurring internal states, within which to understand all learning phenomena. However, he emphasized that the sequence of internal states is not merely determined by external events, but is rather an active process in which the animal is attempting

continually to discover structure and meaning in the incoming information. This is a very different view from the incremental and 'unsupervised' form of learning with which Hebb is usually associated today [22,23]. Our findings, as well as other recently reported evidence [24], vindicate Hebb's original ideas and call for a more integrative approach to studying the organization of learning.

Materials and methods

Subjects performed a shape-discrimination task [12–14] in which the global shapes were illusory, in that they were defined by local inducers separated by large gaps [15,16]. The four inducers were located at the corners of a square. Rotating the inducers about their centers made the illusory shapes appear 'thin' or 'fat' (Figure 2). Performance was measured by the angle of rotation of the inducers needed to yield reliable discrimination. Each experimental block contained stimuli from a fixed range of inducers' rotations. Each stimulus was repeated 12 times within the block, in pseudo-randomized order. The fraction of times that the subject judged the stimulus to be thin was plotted as a function of the inducers' rotation angle. A sigmoid curve ($[1 + \tanh(\beta(x-\alpha))]/2$) was fit to the data, with the slope (β) and bias (α) as free parameters. The threshold, defined as the inducers' rotation angle needed to reach 82% correct discrimination, was estimated from the fitted curve.

The side of the global (illusory) shapes was 15 cm, leading to a visual angle of 14.3° and 5.7° in the 60 cm and 150 cm viewing distances used, respectively. The support ratio, defined as the ratio between the luminance-supported part of the illusory edge and the total edge length, was either 0.25 for all stimuli in experiment 1 and for the low-curvature (1–3°) stimuli in experiment 2, or 0.4 for the high-curvature (4–6°) stimuli in experiment 2. The stimuli were presented for 97 ms, followed by a blank screen (56 ms) and a mask (250 ms). In the training blocks of experiments 1 and 2, one half of the high-curvature stimuli (a third of the total number of stimuli) were of longer exposure duration (139 ms and 56 ms, respectively, for the stimulus and blank screen). The stimuli were generated by a Silicon Graphics Indigo II computer and presented on a CRT screen with resolution 1280 × 1024 pixels. The refresh rate was 72 Hz.

Subjects were Harvard undergraduate or graduate students aged 18–30 years (the mean age was 21.7). All subjects were naive about the purpose of the experiment. Before collection of the experimental data, each subject was given a practice session in which the illusory shapes were highly visible due to the large inducers' size—the support ratio was increased to 0.4. The practice block consisted of 4 presentations of long duration (700 ms) stimuli, followed by 20 presentations of short duration (97 ms plus 56 ms blank screen) masked stimuli. Subjects were required to give at least 17/20 correct responses in their first or second practice block in order to participate in the experiment (2 subjects out of 22 were rejected from the experiment because of failure on this criterion). Subjects were given feedback in the form of a computer 'beep' after correct responses throughout the practice and all experimental blocks.

Acknowledgements

We thank Merav Ahissar, Shaul Hochstein, Nancy Kanwisher, Michael Oren, Dario Ringach and Daphna Weinshall for helpful discussions. Special thanks to Anne Grossetete for experimental assistance and many insightful suggestions. Supported by the McDonnell-Pew program in Cognitive Neuroscience and the Sloan Foundation (N.R.), the McKnight Foundation and AFOSR (K.N.) and the US National Institutes of Health (R.S.).

References

- Ramachandran VS, Braddick O: Orientation-specific learning in stereopsis. *Perception* 1973, 2:371–376.
- Ramachandran VS: Learning-like phenomena in stereopsis. *Nature* 1976, 262:382–384.
- Florentini A, Berardi N: Perceptual learning specific for orientation and spatial frequency. *Nature* 1980, 287:43–44.
- Ball K, Sekuler R: A specific and enduring improvement in visual motion discrimination. *Science* 1982, 218:687–698.
- Karni A, Sagi D: Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proc Natl Acad Sci USA* 1991, 88:4966–4970.
- Sagi D, Tanne D: Perceptual learning: learning to see. *Curr Opin Neurobiol* 1994, 4:195–199.
- Ahissar M, Hochstein S: Learning pop-out detection: specificities to stimulus characteristics. *Vision Res* 1996, 36:3487–3500.
- Karni A, Sagi D: The time course of learning a visual skill. *Nature* 1993, 365:250–252.
- Poggio T, Fahle M, Edelman S: Fast perceptual learning in visual hyperacuity. *Science* 1992, 256:1018–1021.
- Köhler W: *The Mentality of Apes*. London: Routledge and Kegan Paul Ltd.; 1925.
- Hebb DO: *Organization of Behavior*. New York: John Wiley and Sons; 1949.
- Ringach D, Shapley R: Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Res* 1996, 36:3037–3050.
- Rubin N, Nakayama K, Shapley R: Rapid propagation speed of signals triggering illusory contours. *Invest Ophthalmol Vis Sci* 1995, 36 (suppl):1037.
- Rubin N, Nakayama K, Shapley R: Enhanced perception of illusory contours in the lower versus upper visual hemifields. *Science* 1996, 271:651–653.
- Kanizsa G: *Organization in Vision*. New York: Praeger; 1979.
- Petry S, Meyer G: *The Perception of Illusory Contours*. New York: Springer-Verlag; 1987.
- Mollon JD, Danilova MD: Three remarks on perceptual learning. *Spatial Vision* 1996, 10:51–58.
- Ullman S: Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. *Cereb Cortex* 1995, 5:1–11.
- Dayan P, Hinton GE, Neal RM, Zemel RS: The Helmholtz machine. *Neural Computation* 1995, 7:889–904.
- Grossberg S: Competitive learning: from interactive activation to adaptive resonance. *Cogn Sci* 1987, 11:23–63.
- Edelman GM: *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books; 1987.
- Rumelhart DE, McClelland JL: *Parallel Distributed Processing*. Cambridge, Massachusetts: MIT Press; 1986.
- Brown TH, Kairiss EW, Keenan CL: Hebbian synapses: biophysical mechanisms and algorithms. *Annu Rev Neurosci* 1990, 13:475–511.
- Ahissar M, Hochstein S: Eureka: one shot viewing enables perceptual learning. *Invest Ophthalmol Vis Sci* 1996, 37 (suppl):3183.

Because *Current Biology* operates a 'Continuous Publication System' for Research Papers, this paper has been published on the internet before being printed. The paper can be accessed from <http://biomednet.com/cbiology/cub> – for further information, see the explanation on the contents page.