Pattern Recognition in Bees: Size of Regions in Spatial Layout

S-W. Zhang and G. A. Horridge

Phil. Trans. R. Soc. Lond. B 1992 **337**, 65-71
doi: 10.1098/rstb.1992.0084

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Pattern recognition in bees: size of regions in spatial layout

S-W. ZHANG AND G. A. HORRIDGE

Centre for Visual Sciences, Research School of Biological Sciences, Australian National University, G.P.O. Box 475, Canberra A.C.T. 2601, Australia

SUMMARY

These experiments investigate the relation between the local features that are known to be discriminated by bees and the size of local regions in the spatial layout that is assumed to be the essence of pattern discrimination. Previous work with the Y-choice apparatus demonstrates that bees discriminate very well from a distance the angle of inclination of black–white edges of bars and striped patterns, but that the performance is progressively reduced by increasing the variety of contrasting edges in the pattern or its background. We devised a series of tests which show that bees can discriminate different black–white striped patterns on the basis of the spatial layout of different angles of stripes. The patterns were round discs divided into two, four or eight sections that were filled with stripes at different angles. With discs divided into two or four sections, the bees could discriminate between left-right mirror images and more effectively between top-bottom mirror images, but performed best when they discriminated between a pattern and itself rotated by 180°. Patterns with orthogonal stripes appear to be more difficult than the others, and discs divided into eight segments are too difficult unless some gross features are outstanding. The results are consistent with the view that the average angles of inclination are discriminated in local regions, and that the bee sees the gross spatial layout especially in the lower half of the pattern. The area required for resolution of spatial layout with this stimulus is about four stripe periods, but discrimination between vertical and horizontal can be done with a single edge.

1. INTRODUCTION

The idea that different regions of the visual field are topographically represented in the visual memory of the bee is well established in the literature (Wehner 1972). That paper demonstrated that the bee can learn the angle of a black–white edge, that disruption of the pattern with additional edges makes discrimination more difficult, and that the lower part of the forward-looking visual field was the most significant for discrimination of angle. In those experiments, however, the bee was obliged to enter a tube and fixate the pattern with the forward-looking part of the eye. The training and test situation were designed to facilitate the formation of an 'eidetic image' by the bee at the arrival at the correct target, and differed entirely from the following experiments, which tests the ability of the bee to discriminate the target from a distance while in flight.

Having examined the closely related problems of spatial acuity (Srinivasan & Lehrer 1988), range measurement (Srinivasan *et al.* 1989), detection of parallax (Srinivasan *et al.* 1990), discrimination of edge inclination (van Hateren *et al.* 1990) and the interaction between local and global effects (Zhang *et al.* 1992) in the feeding behaviour of the flying honey bee, we now turn to the more difficult question of how many receptors must view the pattern to discriminate one region of a pattern from another. Learning a pattern and then discriminating it visually from similar patterns is a task that requires a memory of spatial layout, but the number, N, of possible combinations detected by an array of sampling detectors rapidly escalates as the number of detectors, n, increases; $N = S^n$, where S is the number of possible states of each detector. Because the number of combinations becomes so great, the relatively low level visual system of the bee must have ways of reducing the data to that which is relevant, and yet retain the flexibility to learn a wide variety of possible patterns.

The threshold spatial resolution of the bee's eye is quite a different matter from the separation of spatial layout. Tests of discrimination between vertical and horizontal black and white stripes fall to 60% success at a distance of 30 cm with a period of 2 cm and results are not significant with stripe periods subtending less than 3.5° (Srinivasan & Lehrer 1988). With this spatial resolution, the bee could theoretically discriminate an enormous number of possible spatial layouts. The question is, what is the compromise achieved by the bee's visual system in throwing away the combinatorial detail and retaining sufficient of it for adequate performance. We imagine that the memory trace for visual form in bees is fairly primitive in the sense that it must rely on rather coarse spatial layout, and we wonder how big these areas are.

A fundamental feature of insect vision is that the eye requires the relative motion or the flicker genera-

Phil. Trans. R. Soc. Lond. B (1992) **337,** 65–71
Printed in Great Britain

ted by its own motion to function at all as it scans or flies along (Horridge 1986, 1987). Many types of insects including bees deliberately scan as they fly, especially when searching visually for a previously learned target. It is obvious that bees do not 'see' as we do, because they make gross errors in discriminating black–white patterns on a horizontal surface that are obviously different to us. To distinguish between target patterns on a horizontal surface the bees seem to have little more available than the temporal frequency of the flicker that the patterns generate at the eye (Hertz 1935) but the bees also see the general layout of the flicker in the visual scene or they would not be able to negotiate through it in flight. It is not clear to what extent they scramble the data as they scan in such a way that the resolvable detail is lost locally, but broadly retained region by region. In fact, in a previous study of the discrimination of the angle of inclination of a black-and-white striped pattern of random period on a *vertical surface* (van Hateren *et al.* 1990) two significant facts emerged. First, although a bee can easily distinguish a striped pattern from the same pattern tilted by 45°, she finds it difficult to discriminate between two quite different striped patterns at the same inclination as long as they cause about the same amount of flicker. This shows that some information about spatial layout is not processed by the flying bee. Secondly, a bee that is trained to prefer a target on the basis of angle of inclination will select the correct angle in a new set of patterns irrespective of the distribution of black and white or number of edges.

These results lent support to the idea that bees see the range and the local average inclination of edges but not the layout of the pattern in every detail. Many earlier results fit in with the idea of numerous broadly-tuned colour-blind templates for moving edges with overlapping fields in parallel on every visual axis (Horridge 1991). A possible explanation of the discrimination of angle of inclination in terms of templates was that the bee scans the striped patterns while in flight, and the inclinations of edges are given by the ratio of the numbers of responding templates that are sensitive to vertical and horizontal relative motion. Taking ratios of numbers of different template responses in each local area allows discriminations of average edge orientation to be independent of the number of edges, contrast frequency, spatial frequency, contrast and illumination, but destroys perception of spatial layout within that local area. Greater differences between numbers of template responses make local areas more easily separated. Acting independently within local areas, mechanisms using template responses could yield a discrimination of patterns which produce the same total combinations of directional temporal flicker but differ in average spatial layout, region by region of the pattern. On this theory, we might design experiments to measure the size of these regions.

We start by testing the bees' ability to distinguish between discs carrying differently oriented gratings, when each disc presents a single orientation. The task is then made progressively harder by dividing each disc up into two semicircles, or four quadrants and finally eight segments, and presenting differently oriented gratings in the various regions. By making mirror images or rotations of a disc with four differently striped quadrants, we are able to show that the bee can use the relative spatial locations of the different quadrants as visual cues. Similar experiments with sector wheels of differing numbers of sectors were reported by Wehner (1981, figure 59) with patterns of differing complexity by Gould (1985), and with quadrants of four different colours by Gould (1988) but in all of this earlier work the choice did not have to be made from a distance while in flight and the experimental situation was designed to allow the bee to fixate the target and form an 'eidetic image'.

2. METHODS

We followed the procedure previously described in detail for discrimination of the angles of black and white stripes in round targets which could be rotated (van Hateren *et al*, 1990), using the Y choice apparatus first illustrated in Srinivasan & Lehrer (1988). The advantage of this apparatus is that the bee must make a choice visually while still flying at a distance from both targets. In many experiments reported in the literature (see above) the bees were able to approach the targets and examine them closely before making a choice, making it difficult to identify the cues that they used. The two targets both have a hole at the centre, into which the bee can crawl to obtain a reward or not. The bee is trained to select the rewarding (positive) target irrespective of the side of the Y by randomizing the side every few visits. The bee is trained positively towards one target and also possibly to avoid the other. One must guard against these separate tendencies in subsequent tests. The bees were individually marked and each visit to the equipment was recorded.

The test procedure in all experiments was divided into two parts. First, as a control, the positive pattern was presented in both sides of the Y chamber to check that the bees had no preference for one side or the other, or for possible olfactory cues from the positive pattern. The second part of the test, to demonstrate that discrimination occurs, was attempted only if the side preference was negligible. In the test for discrimination, the patterns presented could be the original positive and negative patterns, or they could be others which test the preference of trained bees for patterns to which they had not been trained, but in that case the result could depend on the negative as well as the positive training pattern.

The test patterns were prepared on discs of 24 cm diameter, placed at a viewing distance of 27 cm from the point of decision at the meeting of the two arms of the Y, so the whole disc subtended an angle of about 45° at the eye of a bee entering one of the arms, which was taken to be the point of decision. Each marked bee was given a positive or a negative score when it first crossed the threshold of one of the arms. Bees were tested over short periods of about 20 min during which they were usually rewarded four times, twice in one tunnel and twice in the other. Rewarding was continued during testing. Short periods of testing alternated with longer periods of continued training.

Each pattern is a circular disc filled with a black and white striped pattern known to be resolvable by the bee from the point of decision in the Y. In one experiment, the entire disc was filled by a single striped pattern (figure 1); in others, the disc was divided into two semicircles (figure 2) four quadrants (figures 3–5), or eight sectors (figure 6), carrying striped patterns of different orientations. The striped patterns had a period of 4 cm, subtending 0.11 cycles per degree at 27 cm from the bees' eye at the point where they must decide. The resolution of the worker bee for black and white stripes is sufficient to give 65% correct choice at a corner frequency of 0.25 cycles per degree, and falls to nearly zero at 0.34 cycles per degree, measured with the same equipment as here. Also, the resolution is the same for horizontal as for vertical stripes for freely flying bees. The stripes were carefully prepared with equal spacing to avoid differences in average intensity. There is no doubt that the bees can resolve the stripes and discriminate the angles presented here (Srinivasan & Lehrer 1988).

Each pair of patterns is illustrated with the results of each test.

Data processing

The arrival of each marked bee and the frequency, α, with which it chose the positive pattern was noted, and each test was repeated at least three times. An α value of 0.5 or 50% implies that the bees do not discriminate between the two patterns, whereas $\alpha = 1$ would indicate a perfect performance that is never achieved. A weighted mean (depending on the number of bees) was calculated for each test. Two statistical tests for significance have been used, one based on an estimated standard deviation from the mean, and the other a chi-squared test, as detailed in van Hateren *et al.* (1990).

3. RESULTS

The bees rapidly become familiar with the entry to the Y-tunnel test apparatus to collect a reward of sugar solution. Usually the training and test were done with a group of four or five individually marked bees which then became famillier with the procedure. Because the patterns were exchanged between the two sides after every two or three visits, the bees were obliged to turn to look at each target to make a choice. Some bees look carefully several times, but others are more hasty and they may cross the threshold of the wrong arm of the Y then turn back when they realize their error. These bees get a negative mark because the criterion is the first crossing of the threshold of the Y-arm.

(*a*) ***Experiment 1*** (figure 1)

In this experiment, bees are trained to distinguish between a horizontal grating (positive) and a vertical one over the whole circular field. It is clear that the bees learn this task very well: the choice frequency in favour of the positive pattern is $\alpha = 80.4\%$; ($n = 158$; $p < 0.001$). This can be regarded as the baseline score with no distracting region of the pattern. The same spatial frequency is used in all the other experiments.

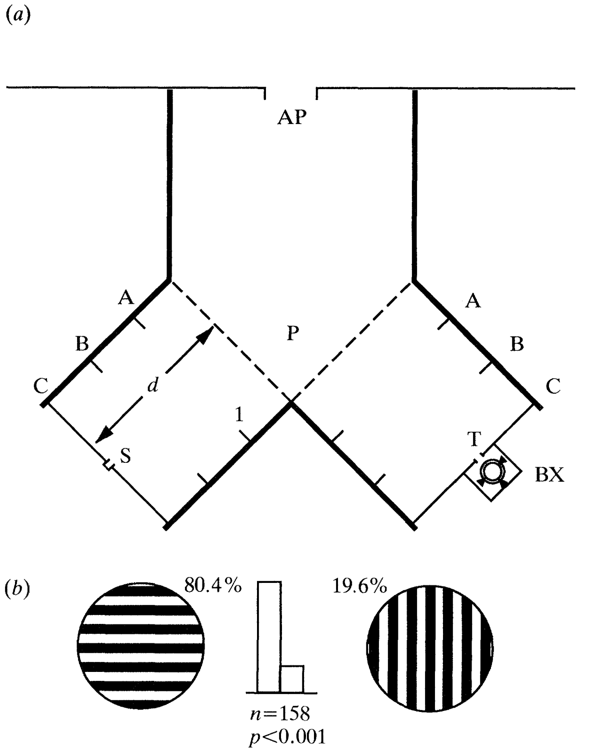

Figure 1. (*a*) The Y-shaped tunnel for training and testing. Bees enter at AP and make their choice at about P. Targets can be changed quickly at positions A, B, or C. The sugar solution reward is in the box BX. (*b*) Test of bees' ability to learn to distinguish between a disc carrying a horizontally oriented grating and a disc carrying a vertically oriented one. The bees are able to do this task very well.

(*b*) ***Experiment 2*** (figure 2)

In this experiment, the disc is divided into two halves. The positive pattern has a vertical grating in the left half and a horizontal one in the right. The negative pattern is the left–right mirror image, which in this case is the same as the positive pattern rotated by 180° (figure 2*b*). When tested, the bees discriminate very well between the two patterns ($\alpha = 64.6\%$; $n = 243$; significance level $p < 0.001$).

The bees are next presented with the positive pattern, as before, to discriminate from the negative pattern rotated clockwise by 90°, which is the same as the positive pattern rotated anticlockwise by 90° (figure 2*c*). The bees cannot make the discrimination ($\alpha = 56.8\%$, not significantly different from 50%, $p > 0.2$). We have just demonstrated (figure 2*b*) that the bees can discriminate the positive pattern from its left–right mirror image but now we have found that they fail to recognize the same positive pattern. Possibly they had been trained not to go to the negative pattern, which would explain why they fail to perform when the negative pattern is no longer available. Alternatively we might conclude that the

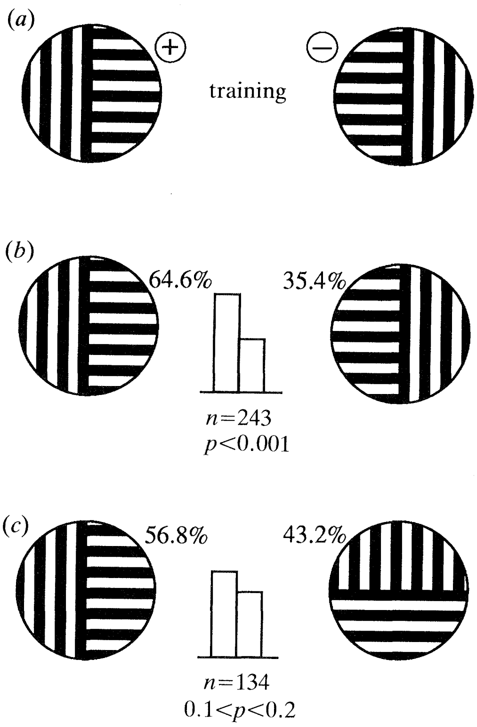

Figure 2. The disc, divided into two halves, can be discriminated from its left–right mirror image, as in (*b*) but the bee fails to discriminate significantly when trained on (*a*) and tested on (*c*), suggesting that the bees were trained not to go to the negative pattern.

bees are confused because two quarters of the patterns in figure 2*c* are the same on both, decreasing the areas on which the choice must rely. Clearly the result in figure 2*b* is inconclusive and more tests must be done.

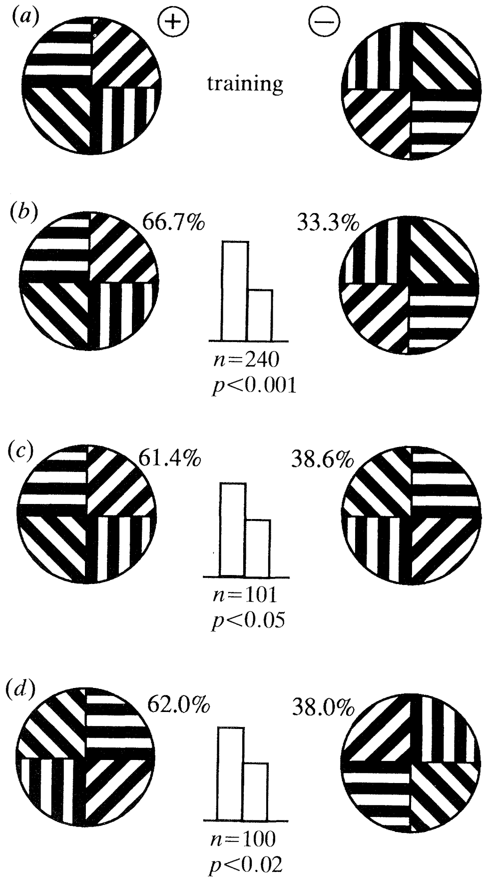

Figure 3. The quadrant pattern with a different inclination in each quarter. (*a*) The bee is trained to discriminate between the pattern and the pattern rotated by 180°. (*b*) Test of the positive pattern against the negative one. (*c*) Test of the positive pattern against the left–right mirror image of the positive one, which is also the negative pattern rotated 90° clockwise. (*d*) Test of the left–right mirror image of the positive pattern against the left–right mirror image of the negative one (which is also the up-down mirror image of the positive pattern).

(*c*) ***Experiment 3*** (figure 3)

This, and the next two experiments, were done with discs divided into four equal quadrants. In this experiment, bees were trained to discriminate between the positive pattern and the positive pattern rotated by 180° (figure 3*a*). These two targets are not left–right or up–down mirror images of each other but one can be derived from the other by making these two operations in succession.

Four tests were done: first, a control against side preference, giving a 51.2% correct performance which is not significantly different from 50% ($p > 0.8$). Secondly, the bees were tested (figure 3*b*) with the training patterns to see whether they had learned them ($\alpha = 66.7\%$; $n = 240$; clearly significant, $p < 0.001$). Thirdly, they were presented with the positive pattern to discriminate from the mirror image of the positive pattern, which is also the negative pattern rotated by 90° clockwise or the positive pattern rotated by 90° anticlockwise (figure 3*c*). The bees clearly prefer the original positive pattern ($\alpha = 61.4\%$; $n = 101$; significant, $p < 0.05$).

Finally the bees, trained on the original positive and negative patterns, are asked to discriminate between the left–right mirror image of the positive patterns (which is also the positive pattern rotated 90° anticlockwise and the negative pattern rotated 90° clockwise) and the top–bottom inversion of the positive pattern (which is also the positive pattern rotated 90° clockwise and the negative pattern rotated 90° anticlockwise) (figure 3*d*). The bees clearly prefer the former ($\alpha = 62\%$; $n = 100$; significant $0.01 < p < 0.02$). When we consider the basis of this preference, we are forced to conclude that the bees treat the left–right mirror image of the positive pattern as positive but they avoid its top-bottom mirror image (and the left–right mirror image of the negative pattern) as clearly negative. This ability to accept an unfamiliar best approximation was called 'faculatative ambiguity' by Gould (1988) in tests with differently coloured quadrants. Rather than invent a name, we propose the mechanistic explanation that the bees put more weight on the lower half of the pattern (Wehner 1972) and simply prefer the pattern with some vertical stripes in the lower half.

(*d*) ***Experiment 4*** (figure 4)

In this experiment the bees are trained to discrimi-

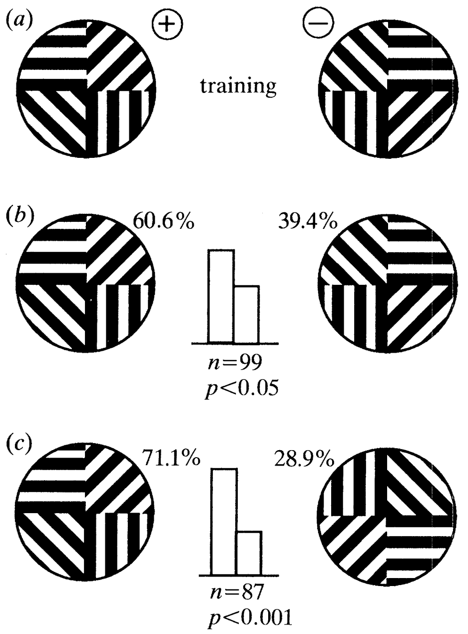

Figure 4. The quadrant pattern with a different inclination in each quarter. (*a*) Training situation; left–right mirror images. (*b*) Test of learning. (*c*) Test of positive pattern against itself rotated by 180° (which is the negative pattern rotated anticlockwise by 90°).

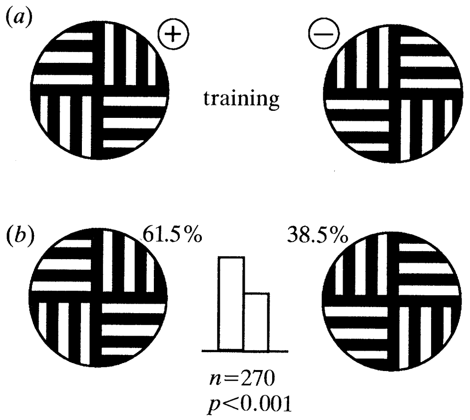

Figure 5. A quadrant pattern with stripes in two directions. (*a*) The training situation. (*b*) The positive pattern tested against its own left–right mirror image (which is also its up–down mirror image).

nate between the positive pattern and its left-right mirror image. In the control they show no side preference ($\alpha = 52.2\%$; $p > 0.6$). In the test of learning (figure 4*b*), they clearly discriminate correctly ($\alpha = 60.6\%$; $n = 99$; $p < 0.05$). Next, the positive pattern was presented with the top-bottom mirror image of the negative pattern, which is the same as the positive pattern rotated by 180°. The bees perform even better with this pair (figure 4*c*) than with the pair of patterns on which they were trained ($\alpha = 71.1\%$; $n = 87$; $p < 0.001$). Taken together with experiment 4, this result suggests that they clearly discriminate the transfer of areas between the top and the bottom of the pattern, but they can less easily discriminate left–right mirror images. Gould (1988) found a similar result with four differently coloured quadrants, and as mentioned, the reason may be the greater reliance upon the lower region of the eye.

(*e*) ***Experiment 5*** (figure 5)

The pattern is changed in this experiment so that all stripes are either horizontal or vertical. The main test, after the control for side preference, was whether the bee can discriminate between the pattern and its left–right mirror image (which is the same as its top–bottom mirror image). From the result of experiment 4, this should be an easy task but we did seven separate experimental runs with an average performance of $\alpha = 61.5\%$ ($n = 270$; $p < 0.001$) with results that suggest that the task is difficult. Possibly the bees are confused because the central cross is a gross feature which is the same on both sides. At any rate this result shows that the bees do not concentrate upon a particular quadrant when making a choice (cf. experiment 2). This experiment is also a demonstration that the bee is not making use of global differences in the distribution of low spatial frequencies which could be generated by scanning in one plane.

(*f*) ***Experiment 6*** (figure 6)

We now turn to patterns of the same scale divided into eight equal quadrants. After a control against side preference, the pattern in figure 6(*a*) was tested against itself rotated by 45°, as in the training procedure. The bees can manage this task but not very well, so four repeated runs were made (63.1% average of four runs; total $n = 125$; $0.05 < p < 0.1$).

The patterns of figure 6*a* should be more difficult to discriminate than those of figures 2 and 5 but the bees perform similarly in all of them, irrespective of the number of segments. These patterns all have stripes arranged at right angles to each other; it seems to be an experimental observation that the bees' visual system has difficulty with such patterns. The patterns in figure 6*a*, however, have a marked gross pattern which runs through the division into eight segments. Previous experience suggests that the bees look out for any prominent features when confronted by complex patterns. To avoid this problem, we selected two types of pattern with eight segments and less obvious gross cues. The pattern in figure 6*b* was presented against itself rotated by 45°. The bees fail on this test. The spiral pattern in figure 6*c* is presented against its own left–right mirror image (which is the same as its own top–bottom mirror image). Rotation of one can never make these two patterns the same. The bees fail also on this test, from which we can conclude that we now have arrived at spatial layouts with areas that are too small to be separated by the bee visual system.

4. DISCUSSION

Unlike the patterns with four different colours in four quadrants used by Gould (1988), our patterns were composed of regions of black and white stripes differing in inclination, the aim being to approach the

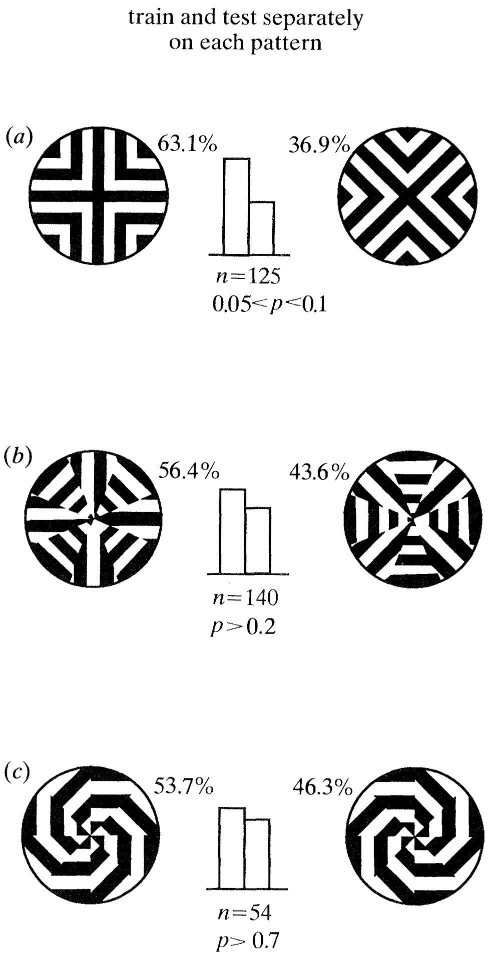

Figure 6. Tests with eight segments in the pattern, with the same spatial frequency and total size as before. The bees succeed with (*a*) but possibly use the gross cues of the central four large radial stripes. They fail with (*b*) and (*c*), suggesting that the segments are now below the critical size for determining the spatial layout of the stripe inclination.

question of form vision irrespective of colour, contrast, brightness or differences in average contrast frequency as the whole target is seen by the bee. Our intention was to use earlier results on discrimination of edge orientation to test the general idea that some spatial layout of form is preserved on a gross scale. In agreement with Gould (1985), we found that increasing the complexity of the pattern within a constant local area increases the difficulty of discrimination, but Gould worked with flower patterns that the bee could approach before making a choice. Flying bees, unable to fixate on a target when learning and testing, are apparently unable to form an 'eidetic image' and with radial gratings their performance drops off rapidly between 0 and 10 cm (Srinivasan & Lehrer 1988, figure 15). Our results suggest that the region required for discrimination of spatial layout by flying bees subtends a rather large solid angle of about 16° × 16°. At the same time the angle of inclination of long stripes is discriminated by the bee when they subtend 4° or more per period. The difference between resolution of spatial layout and of stripes gives us a measure of the extent that local features are lumped together in the effort to see spatial layout with a visual system of limited processing power.

The patterns used in these experiments were designed as well as we could so that the total numbers of edges, the flicker induced by motion, or the average brightness would not give the bees a clue to the correct solution. These experiments show that, when the test is based on the inclination of a striped pattern, bees can discriminate between left–right mirror images but they find it more difficult to discriminate between left–right mirror images than between rotations and top–bottom inversions of the quadrants. Bees are clearly able to resolve and discriminate between the quadrants of stripes at different angles and they are able to assign some information about these angles to different locations in the pattern. We have already shown that when presented with a pattern they have not seen before they make the best approximation they can (Zhang *et al.* 1992). Gould (1988) has reported similar results with asymmetrical coloured patterns, but the bees did not have to make a choice from a distance. If we separate our data into discriminations of patterns rotated by 45°, 90° and 180°, we see that the bees find the latter a much easier task.

As shown in figure 3*d*, the bees can treat the left–right mirror image as a replacement for the positive pattern but avoid the top–bottom mirror image as clearly different. Explained in terms of the required performance in a natural world, this behaviour of the visual system is entirely in keeping with the way that visual processing is designed to fit in with locomotion in the horizontal plane of an animal that stays the right way up and uses horizontal scanning by the eye. We ourselves see objects from one side then walk round and see the approximate mirror image from the other side. We see at once that Thurber's famous cast-iron lawn dog looking to the left resembles the same dog looking to the right (Thurber 1950, figure on p. 86). As noted by Gould (1988) in this context, children confuse 'p' and 'q' or 'b' and 'd' but not 'p' and 'b' or 'd' and 'q' when learning to read and humans have increasing difficulty with form discriminations as the target is rotated (Arnoult 1954) up to about 90° and then improve again for angles up to 180° (Bischoff *et al.* 1985). Other animals respond differently to these transformations depending on their visual repertoire. As the bee looks from side to side in the horizontal plane, creating a pattern of flicker at the eye, it doesn't seem to affect its performance whether it looks to the left or to the right at any one moment. This implies that the template responses (whatever the templates may be) are integrated over short periods, and then define major regions. The simplest explanation in terms of mechanisms is that the bee makes the best use of what clues it can get in the lower half of the target, and it cannot perform if the lower half is replaced (Wehner 1972).

To assist the design of future discrimination tests, let us estimate the minimum number of facets or visual axes needed to collaborate together in the discrimination of striped patterns where the task is one of recognizing the spatial layout with a stripe period subtending 9° at the eye. The quadrants subtend approximately 22° at the eye, so that about 120–150 facets are looking at a quadrant, and 60–75 facets at one eighth of the target. The only comparable number in the recent literature is an assertion by Nakayama (1990) that nodes or icons in human visual memory contain only 100 pixels and they capture the essential properties of image fragments. If there are about 40 different unit templates on each visual axis in insects, as suggested by the number of small neurons in each column of the optic medulla and only 10% of these are excited by the pattern, as suggested by implementation of the template model with natural scenes (Horridge 1991), then, very approximately an assembly of about 400 responding templates are needed to discriminate a minimum region of the bees visual world from another minimum region in this task. On this model, the greater the differences in the ratios of template responses, the more easily are the regions separated. This gives some idea of how an array of numerous templates, each individually ineffective, can collaborate together to make specific ensembles that fit the pattern sufficiently well for the bees discrimination test.

The natural selection of particular visual processing mechanisms is inevitably a compromise between the limited capacity of the insect's small processing mechanisms and the enormous demands of the combinatorial structure of the visual world. We can therefore expect that the bee can just manage to synthesize the minimal cues for the central memories of patterns that are essential for its daily tasks, and presumably no more.

These experiments are part of a series in which several assistants and co-workers, including Dr M. Srinivasan, Dr M. Lehrer and Dr H. van Hateren, have participated and contributed ideas. We also thank our numerous colleagues in Visual Sciences for their pertinent comments. The illustrations were prepared by Jenny Vaughan and the numerous drafts of the manuscript were typed by Liz Boldiston.

REFERENCES

Arnoult, M.D. 1954 Shape discrimination as a function of the angular orientation of the stimuli. *J. exp. Psychol.* **47**, 323–328.

Bischof, W.F., Foster, D.H. & Kahn, J.L. 1985 Selective internal operations in the recognition of locally and globally point-inverted patterns. *Spatial Vis.* **1**, 179–196.

Gould, J.L. 1985 How bees remember flower shapes. *Science, Wash.* **227**, 1492–1494.

Gould, J.L. 1988 A mirror-image ambiguity in honeybee pattern matching. *Anim. Behav.* **36**, 487–492.

Hateren, J.H. van., Srinivasan, M.V. & Wait, P.B. 1990 Pattern recognition in bees: orientation discrimination. *J. comp. Physiol.* A **167**, 649–654.

Hertz, M. 1935 Die Untersuchungen über den Formensinn der Honigbiene. *Naturwiss.* **23**, 618–624.

Horridge, G.A. 1986 A theory of insect vision: velocity parallax. *Proc. R. Soc. Lond.* B **229**, 13–27.

Horridge, G.A. 1987 The evolution of visual processing and the construction of seeing systems. *Proc. R. Soc. Lond.* B **230**, 279–292.

Horridge, G.A. 1991 Ratios of template responses as the basis of semi-vision. *Phil. Trans. R. Soc. Lond.* B **331**, 189–197.

Nakayama, K. 1990 The iconic bottleneck and the tenuous link between early visual processing and perception. In *Vision: coding and efficiency* (ed. C. Blakemore), pp. 411–422. Cambridge University Press.

Srinivasan, M.V. & Lehrer, M. 1988 Spatial acuity of honeybee vision and its spectral properties. *J. comp. Physiol.* A **162**, 159–172.

Srinivasan, M.V., Lehrer, M. & Horridge, G.A. 1990 Visual figure-ground discrimination in the honeybee: the role of motion parallax at boundaries. *Proc. R. Soc. Lond.* B **238**, 331–350.

Srinivasan, M.V., Lehrer, M., Zhang, S.W. & Horridge, G.A. 1989 How honeybees measure their distance from objects of unknown size. *J. comp. Physiol.* A **165**, 605–613.

Thurber, J. 1950 *The owl in the attic, and other perplexities*. London: Hamish Hamilton.

Wehner, R. 1972 Dorsoventral asymmetry in the visual field of the bee, *Apis mellifica*. *J. comp. Physiol.* **77**, 256–277.

Wehner, R., 1981 Spatial vision in arthropods. In *Handbook of sensory physiology*, vol. VII/6C (*Vision in invertebrates*) (ed. H. Autrum), pp. 287–616. Heidelberg: Springer Verlag.

Zhang, S.W., Srinivasan, M.V. & Horridge, G.A. 1992 Pattern recognition in honeybees: local and global analysis. *Proc. R. Soc. Lond.* B **248**, 55–61.

Received 30 December 1991; accepted 11 February 1992