Pattern Recognition in Honeybees: Analysis of Orientation

S. W. Zhang and M. V. Srinivasan

Phil. Trans. R. Soc. Lond. B 1994 **346**, 399-406
doi: 10.1098/rstb.1994.0157

Pattern recognition in honeybees: analysis of orientation

S. W. ZHANG AND M. V. SRINIVASAN

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

SUMMARY

The discrimination of pattern orientation in freely flying honeybees was examined by testing their ability to discriminate orientation in a variety of patterns composed of bars, edges or textures. The results indicate that orientation discrimination improves as (*a*) the length of the oriented element is increased, (*b*) the number of similarly oriented elements is increased, and (*c*) when a number of similarly oriented elements are arranged in collinear fashion. The orientation of bars is discriminated better when the bars are spatially homogeneous in intensity than when they are randomly textured. Furthermore, orientation discrimination is better with textured bars on a homogeneous background than vice versa. These findings suggest that orientation is analysed globally and that the strength of the orientation signal increases with the length and collinearity of the pattern's constituent elements. Furthermore, in analysing orientation, boundaries of objects seem to be less important than internal detail.

1. INTRODUCTION

The study of pattern recognition in honeybees has a long history. It is well known that bees can be trained to distinguish between two different patterns by associating one of them with a reward of sugar water (von Frisch 1915). Early experiments suggested that, for the bee, an important cue in distinguishing patterns is what Hertz (1933) called the degree of 'brokenness' or 'disruption'. A disrupted pattern gives a stronger flickering percept than a solid pattern as the bee flies past it, and is preferred over the latter. This led Hertz to suggest that, contrary to humans, form perception in bees is based on temporal cues. More recent work has suggested that discrimination is based on some type of visual spatial memory, in which the pattern to be learned is stored in the form of a 'template' or eidetic image, and is matched against the pattern that the bee is currently viewing (Wehner 1981; Collett 1983; Gould 1988; Cartwright & Collett 1983; Zhang & Horridge 1992). To acquire a template of the pattern that is to be learned, and to compare it with other patterns, the bee has to 'fixate' the shape by hovering in front of it in a specific position (Wehner & Flatt 1977). Recently van Hateren *et al.* (1990) found that bees can extract information on orientation from patterns and use this as a parameter to distinguish between patterns that they have never encountered previously. This finding, together with those of Wehner (1971) and Srinivasan *et al.* (1993), suggests the existence of a visual mechanism that can abstract general properties of patterns, such as orientation. Such mechanisms were previously believed to be confined to vertebrate vision (Blakemore & Campbell 1969; Henry *et al.* 1974).

How does the visual system of the bee extract the orientation of a pattern? What are the features or tokens that are analysed? Edges seem to play an important role in the detection of objects (Lehrer & Srinivasan 1993; Lehrer *et al.* 1990). Are edges important in signalling orientation as well? Here we describe experiments to unravel the nature of the cues that bees extract from patterns to analyse orientation.

2. MATERIALS AND METHODS

(*a*) *Apparatus*

Worker honeybees (*Apis mellifera*) were marked and trained to enter a Y-shaped, dual-tunnel apparatus (figure 1), similar to that described in Srinivasan & Lehrer (1988). Bees entered the apparatus through an aperture in the window of the laboratory, and could simultaneously view two stimuli, each presented in the vertical plane on the end wall of a tunnel. One of the stimuli (termed positive) offered a reward of sugar water, dispensed by a feeder located in a box behind the pattern, accessible to the bee via a small tube. The other stimulus (termed negative) offered no reward. In the experiments described here, the positive stimulus was a pattern composed of horizontally oriented elements (such as edges or bars) and the negative stimulus a pattern composed of vertically oriented elements. The significance of the design of the apparatus is that the bees must make a choice at the entrance to either tunnel, which is at a considerable distance (27 cm) from the pattern being viewed. At this distance, the bees do not fixate the patterns and are therefore unlikely to discriminate between them on the basis of a 'photographically' memorized image (van Hateren *et al.* 1990).

Phil. Trans. R. Soc. Lond. B (1994) **346**, 399–406
Printed in Great Britain

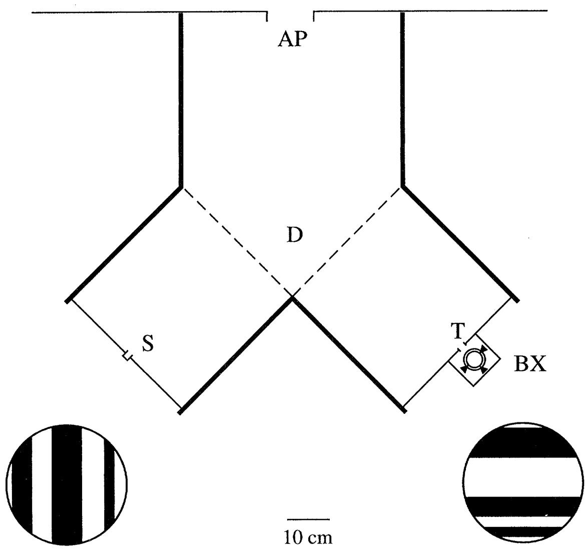

Figure 1. Experimental apparatus. Overhead view of training and testing apparatus, shown here in plan. Bees enter the apparatus through an aperture AP in the laboratory window and proceed to the decision chamber D where they choose one of two tunnels, each bearing a stimulus presented in the vertical plane. Behind the positive stimulus is a reward box BX containing a feeder that is accessible to the bees through the tube T. The stimuli are placed symmetrically at the ends of the tunnels (redrawn from Srinivasan & Lehrer (1988)). S, negative stimulus.

(*b*) *Training and testing procedure*

Each experiment was begun by training a fresh, naive group of four to seven bees to enter the apparatus and collect the reward. On a warm day, all of these bees would visit the apparatus roughly twice every 10 min (the hive was about 50 m from the laboratory in which the experiments were conducted). The positions of the positive and negative patterns were interchanged every 10 min. The reward box was also moved, so that it stayed with the positive pattern. This interchanging was performed continually throughout the experiment, to ensure that the bees did not associate the reward with a particular tunnel.

In the testing procedure we presented the bees with two patterns, the rewarded one called positive and the other called negative. The reward was offered behind the positive pattern. Tests were of three kinds. (i) In tests to assess whether learning had occurred (learning tests), the positive and negative patterns were identical to those used in the training. (ii) In control tests the two patterns were identical, to check for the possibility that the bees were using olfactory cues from the reward box. Such tests, and similar tests conducted in an earlier study (van Hateren *et al.* 1990), assured us that olfactory cues did not play a significant role in the bees' choice behaviour in our experimental set-up. (iii) In critical tests the positive, the negative or both patterns were different from those used in the training. The duration of each testing session was restricted to 10 min (allowing two rewards per bee on average). Furthermore such tests were interspersed with long training intervals and tests with different pairs of patterns.

The tests were conducted for short periods of about 10 min, during which each bee was rewarded only four times on average: twice in one tunnel and twice in the other. The brevity of the tests, together with the fact that tests were interspersed by lengthy periods of training, ensured that the bees did not learn to discriminate the test patterns on the basis of being rewarded at one of them. (Several controls against this possibility, showing that the behaviour is not influenced by the brief tests, are described in van Hateren *et al.* (1990).) In the tests, a bee's choice was scored as correct if it entered the tunnel leading to the rewarded pattern, or as incorrect if it entered the other tunnel. Only the first choice of each bee on each visit was taken into account, to eliminate the possibility that the second choice might be influenced by the outcome of the first. This was particularly important if the first choice happened to be incorrect, so that the bee inspected one pattern at close range and then went to the other.

(*c*) *Data analysis*

The bees' responses were analysed in terms of the choice frequency, α, in favour of the positive pattern. Thus, $\alpha = 50\%$ implies that the bees do not discriminate between the two patterns, while $\alpha = 100\%$ indicates perfect discrimination. A χ^2 test was used to determine whether a measured α was significantly different from random choice behaviour ($\alpha = 50\%$). Further details are given in van Hateren *et al.* (1990).

(*d*) *Stimuli*

Stimuli were prepared on cardboard disks, 24 cm in diameter. Patterns were presented on a white background of photocopying paper or a grey background composed of 50% grey Letratone. A variety of patterns were used: random gratings, bars, patterns of oriented bars, and textured bars. The random gratings and bar patterns were produced by cutting stripes out of thin matt cardboard (or 50% grey Letratone where required) and gluing them onto the background. The dimensions and spacings of the stripes were such that individual stripes should have been clearly resolvable by the bee's visual system, based on its known visual acuity (see Srinivasan & Lehrer 1988). The textured bars were designed on a computer and printed on a laser printer. All patterns were carefully screened to avoid flaws. The details of each pattern are specified in the description of the corresponding experiment in the results section.

3. RESULTS

In experiments 1–4, bees were first trained to discriminate patterns composed of horizontal stripes from patterns composed of vertical stripes. The training patterns were random, one-dimensional, black and white gratings. Each grating consisted of 12 bars, each 2 cm wide, with each bar having an equal probability of being black or white. Two examples are shown in figure 1 (see van Hateren *et al.* (1990) for details). During training the gratings were chosen randomly, in pairs, from a pool of ten such randomly constructed gratings. With each pair of gratings, each bee was rewarded four times on average: twice with the positive pattern in the left tunnel and twice with the positive pattern in the right tunnel. The training was then continued with another randomly chosen pair of gratings. This procedure ensured that the bees learnt to discriminate the patterns on the basis of orientation, and not on the basis of a 'photographically' memorized image of the positive pattern. It is known from earlier studies that bees trained in this way can abstract orientation from the training patterns and use this as a parameter to distinguish orientation in other patterns that they have never previously encountered (van Hateren *et al.* 1990; Srinivasan *et al.* 1994). The bees learned to discriminate the orientation of such random gratings quite well (figure 2*a*). Bees trained to discriminate orientation in this way were then tested on a variety of patterns, as described below.

(*a*) *Experiment 1. Influence of element length on orientation discrimination*

Bees trained as described above were tested with three pairs of patterns, composed of single bars of various lengths (figure 2*b*–*d*). The bees were unable to discriminate the orientation of 2 cm × 4 cm bars (figure 2*b*), marginally able to discriminate the orientation of 2 cm × 8 cm bars (figure 2*c*) and clearly able to discriminate the orientation of 2 cm × 16 cm bars (figure 2*d*). Thus, for a pattern composed of a single bar, the bee's ability to discriminate orientation increases with the length of the bar. This is in agreement with the earlier findings of Wehner (1971).

(*b*) *Experiment 2. Cumulative effects in orientation discrimination*

Does the orientation signal in a pattern strengthen as the number of oriented elements in it is increased? To examine this question, bees were trained as above and subjected to two types of critical tests. In one type of critical test, the patterns were composed of identical, similarly oriented randomly positioned bars, each 2 cm × 4 cm (figure 3*a*) or 2 cm × 8 cm (figure 3*b*). From the results, it is clear that orientation discrimination improves as the number of bars is increased (compare the results of figure 3*a* and figure 3*b* with the results of figure 2*b* and figure 2*c* respectively). Furthermore, and in agreement with the results of figure 2, orientation discrimination improves as the length of the bars is increased (compare the results of figure 3*b* and figure 3*a*).

In another type of critical test, the patterns were composed of elements arranged in a collinear fashion. Each element was a 2 cm × 2 cm square (figure 3*c*) or a 2 cm × 4 cm bar (figure 3*d*). The gaps between the elements were 2 cm in either case. The results of figure 3*c* show that bees are able to extract orientation in a series of collinearly arranged elements even when each element on its own carries no orientation information. The patterns of figure 3*a* and figure 3*d* are each composed of eight identical parallel bars. However, bees can discriminate orientation better when the bars are arranged in collinear fashion (figure 3*d*) than when they are positioned randomly (figure 3*a*).

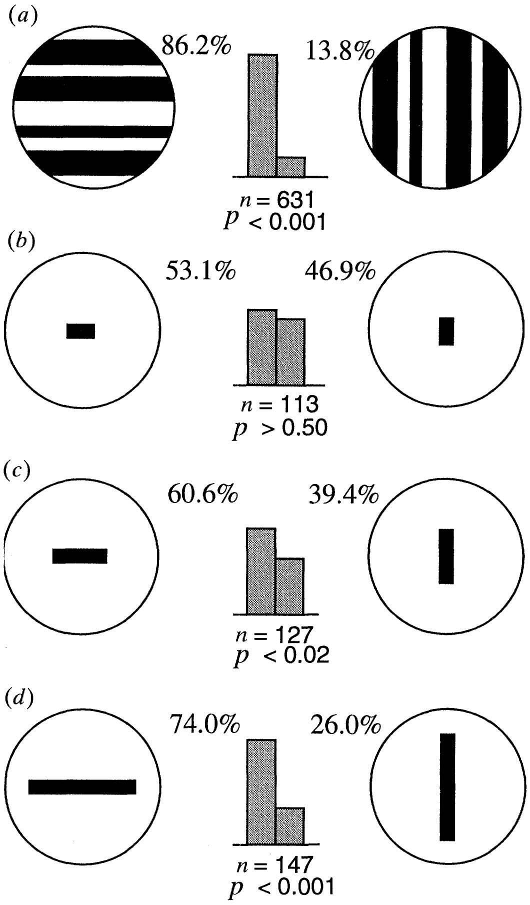

Figure 2. Orientation discrimination for single bars. Honeybees, trained to distinguish a horizontally oriented random grating from a vertically oriented random grating, learnt to discriminate orientation very well (*a*). They were then tested in their capacity to discriminate the orientation of single bars of various lengths. The bees were unable to discriminate the orientation of 2 cm × 4 cm bars (*b*), marginally able to discriminate the orientation of 2 cm × 8 cm bars (*c*) and clearly able to discriminate the orientation of 2 cm × 16 cm bars (*d*). This and subsequent figures show the choice frequencies in favour of each pattern, the number of choices (*n*), and the *p* value in a χ^2 test for significant departure from random choice.

(*c*) *Experiment 3. Effects of contrast and texture on orientation discrimination*

In another series of experiments, we examined the influence of the contrast and texture of the bars on orientation discrimination. Bees, trained to discriminate orientation on random gratings as above, were tested on pairs of patterns composed of three bars, each 2 cm × 10 cm (figure 4). The background was white. When the bars were black, the bees

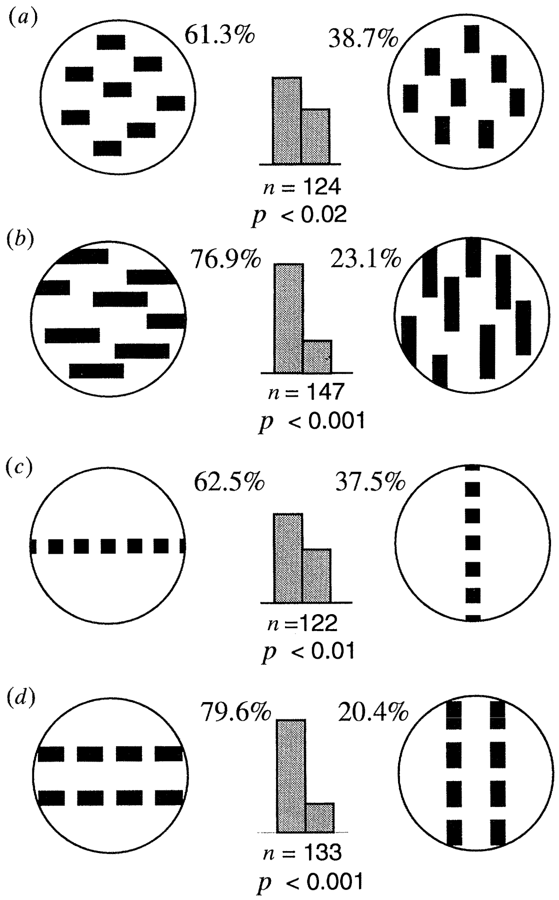

Figure 3. Cumulative effects of orientation discrimination. Bees, trained as in figure 2*a*, were subjected to two types of critical tests. In the first type of critical test, the patterns were composed of identical similarly oriented but randomly positioned bars, each 2 cm × 4 cm (*a*) or 2 cm × 8 cm (*b*). The ability to discriminate orientation improves as the number or the length of the bars is increased. In the second type of critical test, the patterns were composed of similarly oriented elements, arranged in collinear fashion. Each element was a 2 cm × 2 cm square (*c*) or a 2 cm × 4 cm bar (*d*). The gaps between the elements were each 2 cm. The results show that, even with patterns composed of non-oriented elements, bees can extract orientation information if the elements are arranged collinearly (*c*). Furthermore, with patterns composed of oriented elements, orientation discrimination is much better when the elements are positioned collinearly (*d*) rather than randomly (*a*).

discriminated orientation very well (figure 4*a*). Discrimination was somewhat poorer but nevertheless good when the stripes were 50% grey (figure 4*b*). Bees continued to discriminate orientation well when each stripe was composed of a fine grating of period 0.59 cm (figure 4*c*). Performance with these patterns was not significantly different from that with the grey bars (*t* test, $p > 0.15$). However, orientation discrimination became much poorer when the bars were composed of random Julesz textures, of pixel size 0.5 cm × 0.5 cm (figure 4*d*). Evidently textured bars do not provide a strong orientation signal. We shall discuss the implications of these results in a later section.

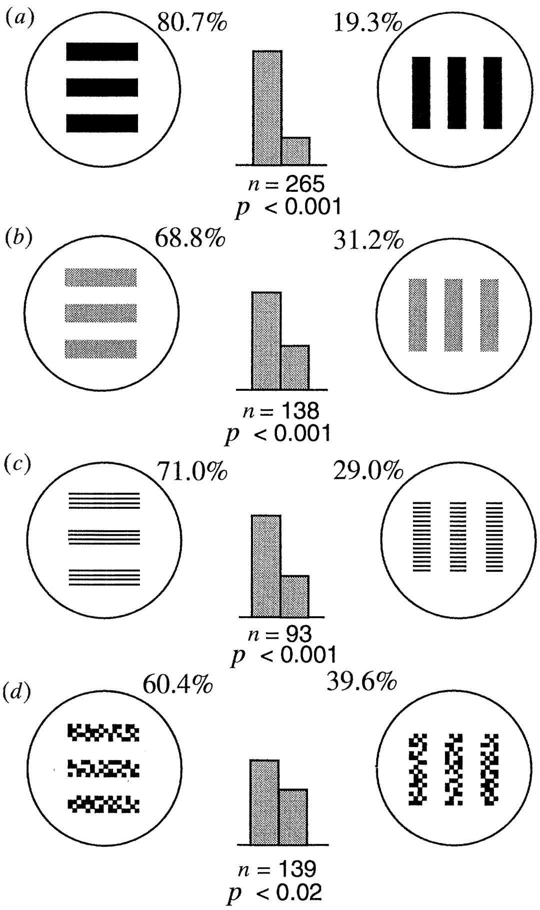

Figure 4. Influence of internal texture. Bees, trained to discriminate orientation on random gratings as above, were tested on pairs of patterns composed of three bars, each 2 cm × 10 cm on a white background. The bees could discriminate orientation very well when the bars were black (*a*). Discrimination was not as good when the bars were grey (*b*) or composed of a fine grating of period 0.59 cm (*c*). Discrimination was significantly worse when the bars were composed of random Julesz textures, of pixel size 0.5 cm × 0.5 cm (*d*).

(*d*) *Experiment 4. Effect of background on orientation discrimination*

Here we examined the effects on orientation discrimination of interchanging the composition of the figure and the background. Bees, trained to discriminate the orientation of random gratings as described above, were tested on three pairs of patterns: textured bars on a white background (figure 5*a*), textured bars on a grey background (figure 5*c*) and textured bars on a black background (figure 5*e*), and on the counterparts of these patterns in which the textures of the bar and the background had been interchanged (figure 5*b*,*d*,*f*, respectively). In all cases, it is evident that orientation discrimination is significantly poorer when the background is textured (compare figure 5*b* with figure 5*a*, figure 5*d* with figure 5*c*, and figure 5*f* with figure 5*e*).

(*e*) *Experiment 5. Orientation information from a single edge*

Can bees acquire orientation information from a single edge? To examine this question, we trained bees

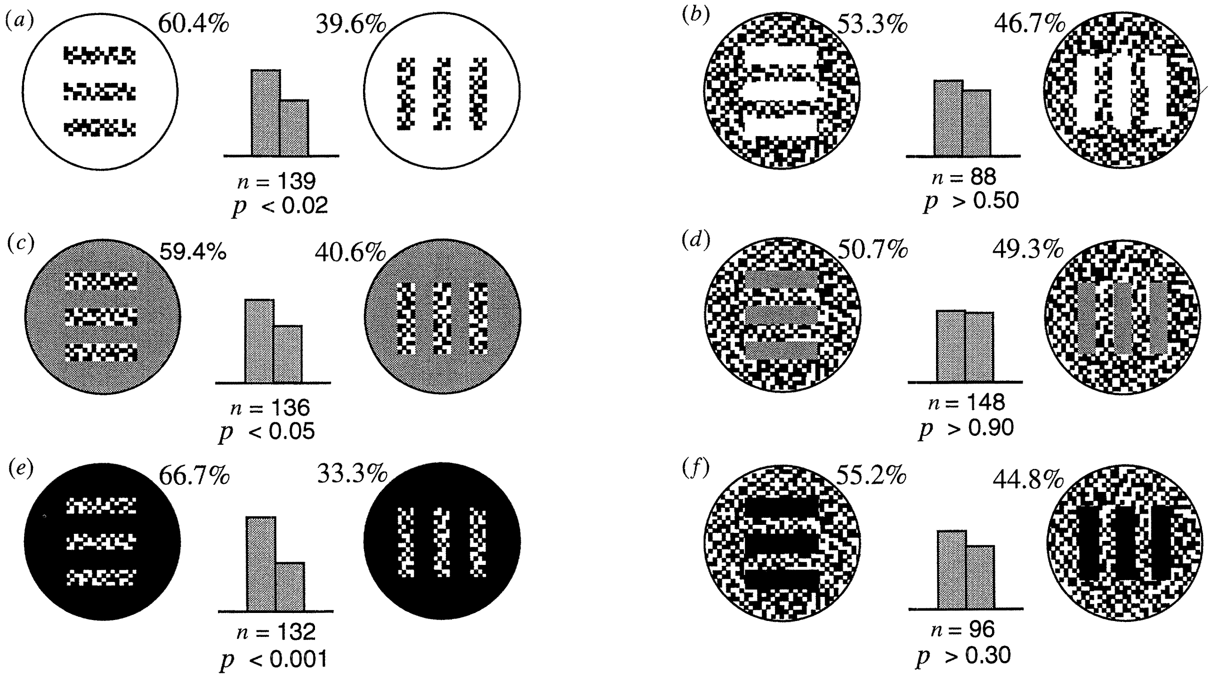

Figure 5. Effect of background on orientation discrimination. Bees, trained to discriminate orientation on random gratings as above, were tested on pairs of patterns: textured bars on a white background (*a*), textured bars on a grey background (*c*) and textured bars on a black background (*e*), and on the counterparts of these patterns in which the textures of the bars and the background had been interchanged (*b*, *d* and *f* respectively). Orientation discrimination is significantly poorer when the background, rather than the bars, is randomly textured.

on patterns composed of semicircular black and white regions, which presented a single contrast edge (figure 6*a*). Bees were trained to distinguish between a horizontally oriented edge (positive) and a vertically oriented edge (negative). To ensure that the bees acquired only edge information, and not an eidetic image of the pattern, the polarities of the edges were reversed frequently. During training, this polarity reversal was combined with the usual interchanging of the rewarded tunnel, by cycling through the sequence of training stimuli *a*1–*a*4 depicted in figure 6*a*. Thus, the training procedure forced the bees to learn the orientation of an edge (i.e. to prefer a horizontal edge to a vertical one) regardless of its polarity. In learning tests, bees trained in this way showed a clear and statistically significant preference for the positive pattern (figure 6*a*). The choice frequency represents an average over all of the four pattern combinations. The results indicate that the honeybees were able to use effectively the one property that is common to all of the patterns, namely orientation of a single edge, rather than a specific eidetic image.

To check whether the bees had really learnt orientation information in this task, we tested the trained bees on patterns that they had never encountered in the training. The trained bees were able to discriminate a horizontal bar from a vertical bar quite well (figure 6*b*), even better than for the training patterns. The trained bees could also discriminate orientation very well in patterns composed of collinear bars (figure 6*c*). Thus bees are clearly able to extract orientation information from a single edge and apply it to distinguish between patterns that they have never previously encountered.

4. DISCUSSION

(a) *Cumulative phenomena in orientation discrimination*

Examination of the performance of bees at discriminating orientation reveals a number of cumulative effects. Firstly, the ability to discriminate orientation improves as the length of the oriented element in the pattern is increased. This is true for patterns consisting of a single bar (figure 2*b*–*d*) as well as patterns composed of multiple bars (figure 3*a*, *b*). We have done a control experiment to show that the reason for this is not simply that longer elements resemble more closely the stripes of the random grating on which the bees are trained. Bees trained to discriminate between short orthogonally oriented bars (each 2 cm × 8 cm) display $\alpha = 63\% \pm 0.03$ for the bar with the correct orientation. However, when the same bees are tested on longer bars (each 2 cm × 16 cm), they display $\alpha = 70\% \pm 0.03$ for the correctly oriented bar, which is significantly higher than that for the short bars (*t* test, $p < 0.05$; data not illustrated). At the neural level, longer bars could either stimulate orientation-signalling units more effectively or recruit additional units sensitive to the same orientation. Secondly, orientation-discrimination improves as the number of similarly oriented elements in a pattern is increased.

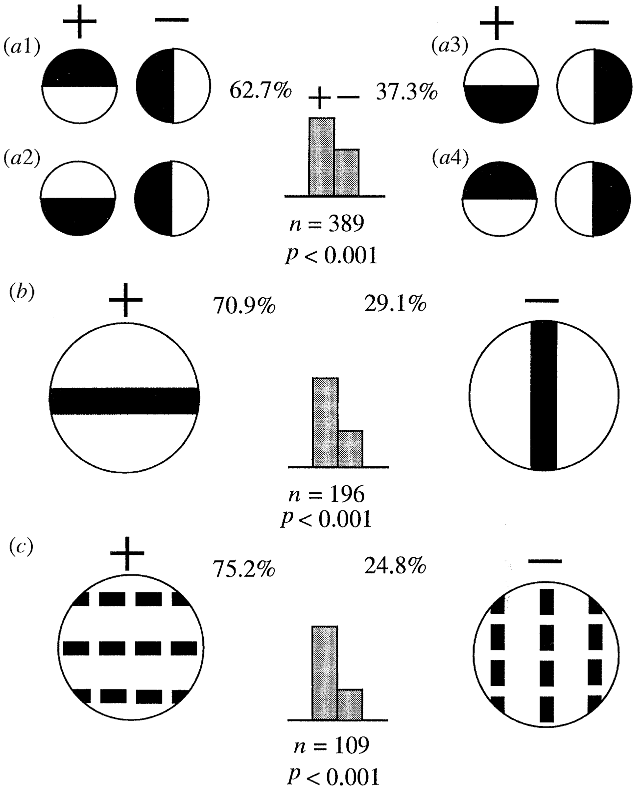

Figure 6. Orientation information from a single edge. Bees were trained to distinguish between a horizontally oriented edge (positive) and a vertically oriented edge (negative). The polarities of the edges were frequently reversed to ensure that the bees acquired only edge information, and not an eidetic image of the training patterns (*a*1–*a*4). The learning test indicated a clear and significant preference for the positive pattern (*a*). The trained bees were able to discriminate a horizontal bar from a vertical bar (each 3 cm wide) quite well (*b*) and to discriminate orientation very well in patterns composed of collinear bars, each 2 cm × 4 cm, separated by 2 cm gaps (*c*).

This suggests that the number of units signalling a given orientation increases as the number of bars is increased. Thirdly, it appears that a given number of similarly oriented bars elicits a stronger orientation signal when the bars are positioned collinearly than when they are positioned randomly (compare figure 3*a* with figure 3*d*). The 2 cm gaps between the collinear elements can be clearly resolved by the bee's visual system (see Srinivasan & Lehrer 1988). Therefore the reason for the stronger orientation signal is not simply that the collinearly arranged elements are indistinguishable from a long bar, but that collinear bars generate a stronger orientation signal. This is to be expected in orientation-tuned units such as those that exist in the mammalian cortex (Hubel & Wiesel 1962), and similar units have recently been discovered in the insect visual system (O'Carroll 1993).

(*b*) *Effect of contrast and texture on orientation discrimination*

Orientation discrimination deteriorates somewhat when the contrast of orientation-signalling elements in a pattern is decreased: bees are better able to discriminate the orientation of black bars on a white background (figure 4*a*) than that of grey bars on a white background (figure 4*b*). Discrimination remains unaltered when the grey bars in figure 4*b* are replaced by a textured bar composed of a fine grating (figure 4*c*) of the same mean luminance. Evidently the bees are unable to resolve the texture in the fine gratings at the distance from which they make their discriminations (one period of the grating would subtend an angle of 1.2° at the tunnel entrance). But orientation discrimination is significantly poorer when the bars are composed of a chunky, Julesz random-textured pattern (figure 4*d*), whose internal structure can be at least partially resolved by the bee (for estimates of honey bee visual acuity see Srinivasan & Lehrer (1988)). There could be two reasons for the deterioration of orientation discrimination with the Julesz-textured bars. Firstly, such patterns produce broken edges around the bars because 50% of the pixels, on average, are of the same intensity as the background. Thus, if the orientation of a bar is signalled by the orientation of its boundaries, one would expect orientation discrimination to be poorer with Julesz-textured bars than with grey bars (as in figure 4*b*). Secondly, if orientation is signalled not by the boundaries themselves but by the appearance of the internal texture as seen by the bee, then it is easy to see that Julesz-textured bars will produce a weaker orientation signal than the grey bars. The Julesz texture is a spatial 'white-noise' pattern with a broad spatial frequency spectrum, including power at low frequencies. If we assume that the bee's visual system sees a spatially low-pass filtered version of the pattern, it would see a number of 'blobs' within the textured area. This low-pass filtering could arise from the optics as well as the subsequent neural processing. Some of the blobs would be circular, signalling no particular orientation, and others would be elongated, signalling an orientation along their long axis. A blob that is shorter than the width of the bar can be oriented in any direction, but longer blobs will necessarily tend to be parallel to the long axis of the bar. Since the length of each textured bar is five times its width, the orientation signalled by a textured bar will tend to be biased towards its long axis. However, the net orientation signal produced by the variously oriented blobs would be considerably weaker than that produced by a grey bar, which would appear as a single large elongated blob oriented parallel to the bar's long axis.

Is orientation signalled by boundaries, or oriented blobs, or both types of features? A possible answer to this question can be obtained by considering the experiments of figure 5. They compare orientation discrimination of textured bars on homogeneous backgrounds (white, grey or black, figure 5*a*,*c*,*e*) with that of homogeneous bars (white, grey or black, figure 5*b*,*d*,*f*) on a textured background. Considering first the textured bars (figure 5*a*,*c*,*e*), we see that orientation discrimination is weakest when the background is grey, i.e. of the same mean intensity as the textured bars (figure 5*c*). This is despite the fact that the boundaries of the bars are unbroken and are therefore most clearly visible in this situation, at least to the human eye. This finding can be explained if we assume that orientation is signalled by oriented blobs within the texture, with no particular importance attached to boundaries. The contrast of the blobs against the background would be stronger when the background is black (figure 5*e*) or white (figure 5*a*) than when it is grey (figure 5*c*). This is consistent with the experimental finding that orientation discrimination is poorest when the background is grey (figure 5*c*). Turning now to the textured backgrounds (figure 5*b*,*d*,*f*), we find that orientation discrimination is poorer with homogeneous bars on a textured background than with textured bars on a homogeneous background. This is true regardless of whether the homogeneous part of the figure is grey, white or black (compare, in figure 5, (*d*) with (*c*), (*b*) with (*a*) and (*f*) with (*e*). This finding is again readily explained if we assume that orientation is signalled by blobs. With the texture in the background, the orientation signalled by the blobs representing the bars would tend be swamped by the randomly oriented blobs in the large, textured background. Thus, any mechanism that computes dominant orientation in a global sense would produce weaker orientation signals with homogeneous bars on a textured background than vice versa, as the experimental results demonstrate. Recent evidence suggests that the bee's visual system computes orientation in patches of diameter *ca.* 20° (Zhang & Horridge 1992). A quantitative computer simulation involving examination of low-pass filtered versions of the stimuli of figure 5 leads to the same conclusions as the qualitative argument presented above (S. W. Zhang & M. V. Srinivasan, unpublished calculations). A similar asymmetry has also been observed in human perception, and models incorporating spatial low-pass filtering and global processing have been proposed to account for this phenomenon (Gurnsey & Browse 1989; Rubenstein & Sagi 1990).

In summary, the results of figure 4 and figure 5 are consistent with the following notions.

(*a*) Under the conditions of these experiments, bees discriminate between the patterns in terms of their dominant orientation, obtained by analysing a fairly large patch of the visual field. This is in agreement with the findings of Srinivasan *et al.* (1994).

(*b*) It appears that orientation is analysed primarily in terms of the dominant orientation of the blobs that are generated by low-pass filtering the image.

(*c*) *Orientation from a single edge*

Given the above discussion, it is natural to ask whether bees can abstract orientation in stimuli that present only boundary information. The experiments of figure 6 were designed to address this question. The results demonstrate that bees can abstract orientation from a single edge, even under conditions in which the training procedure is designed to eliminate all information on the spatial distribution of intensities within the training patterns as well as on the polarity of the edge. However, discrimination of orientation under these conditions is fairly poor: the choice frequency in favour of the rewarded orientation is *ca.* 63% (figure 6*a*). It is unlikely that the bees are learning this task by memorizing the two positive (and two negative) patterns separately in a photographic way, because bees that have been trained on these patterns can use this information to distinguish orientation in patterns that they have never encountered previously. It is intriguing that bees trained on single edges in this way are *better* at discriminating the orientation of bars than they are at discriminating the orientation of the stimuli on which they have been trained. There could be two reasons for this. Firstly, a bar may be better than a single edge simply because the former is composed of two edges, rather than one. Alternatively, and as suggested above, a bar may be more effective because it is intrinsically more effective as an orientation 'token'. Further experiments are needed to address this question.

We thank Miriam Lehrer and Tom Collett for many useful discussions and comments on the manuscript.

REFERENCES

Blakemore, C. & Campbell, F.W. 1969 On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *J. Physiol., Lond.* **203**, 237–260.

Cartwright, B.A. & Collett, T.S. 1983 Landmark learning in bees. *J. comp. Physiol.* A **151**, 521–543.

Collett, T.S. & Cartwright, B.A. 1983 Eidetic images in insects: their role in navigation. *Trends Neurosci.* **6**, 101–105.

Frisch, K. von, 1915 Der Farbensinn und Formensinn der Bienen. *Zool. Jb.* **35**, 1–188.

Gould, J.L. & Gould, C.G. 1988 *The honeybee*, Scientific American Library, a division of HPHLP, pp. 222–225. New York: Scientific American Library.

Gurnsey, R. & Browse, R.A. 1989 Asymmetries in visual texture discrimination. *Spatial Vision* **4**, 31–44.

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

Henry, G.H., Bishop, P.O. & Dreher, B. 1974 Orientation, axis and direction as stimulus parameters for striate cells. *Vision Res.* **14**, 767–777.

Hertz, M. 1933 Ueber figurale Intensitaeten und Qualitaeten in der optischen Wahrnehmung der Biene. *Biol. Zbl.* **53**, 10–40.

Hubel, D.H. & Weisel, T.N. 1962 Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol., Lond.* **160**, 106–154.

Lehrer, M. & Srinivasan, M.V. 1993 Object detection by honeybees: why do they land on edges? *J. comp. Physiol.* A **173**, 23–32.

Lehrer, M., Srinivasan, M.V. & Zhang, S.W. 1990 Visual edge detection in the honeybee and its chromatic properties. *Proc. R. Soc. Lond.* B **238**, 321–330.

O'Carroll, D. 1993 Feature-detecting neurones in dragonflies. *Nature, Lond.* **362**, 541–543.

Rubenstein, B.S. & Sagi, D. 1990 Spatial variability as a limiting factor in texture-discrimination tasks: implications for performance asymmetries. *J. opt. Soc. Am.* A **7**, 1632–1643.

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., Zhang, S.W. & Rolfe, B. 1993 Pattern vision in insects: 'cortical' processing? *Nature, Lond.* **362**, 539–540.

Srinivasan, M.V., Zhang, S.W. & Witney, K. 1994 Visual discrimination of pattern orientation by honeybees: performance and implications for 'cortical' processing. *Phil. Trans. R. Soc. Lond.* B **343**, 199–210.

Wehner, R. 1971 The generalization of directional visual stimuli in the honeybee, *Apis mellifera*. *J. Insect 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.

Wehner, R. & Flatt, I. 1977 Visual fixation in freely flying bees. *Z. Naturf.* **32c**, 469–471.

Zhang, S.W. & Horridge, G.A. 1992 Pattern recognition in bees: size of regions in spatial layout. *Phil. Trans. R. Soc. Lond.* B **337**, 65–71.

Received 21 March 1994; accepted 18 July 1994