12

Visual processing of pattern

ADRIAN HORRIDGE

12.1 INTRODUCTION

The more I study their visual systems, the less I believe that any invertebrate sees shapes by putting together outlines or contours, or that they distinguish between neighbouring shapes or objects against a background except by relative motion. I doubt that they see the spatial relations in the image, other than by use of relatively simple and independent preformed feature detectors that act like filters for certain specific features. The filters are groups of neurons, and like neurons everywhere, they are labelled with their positions on the body, in this case retinotopic. There are arrays of filters of different kinds for different cues in parallel behind the retina. The features or cues that are detected are only small components of the signals known to ethologists as sign stimuli. Each has a direction, sometimes also a range. Because several cues are detected in parallel, a great variety of patterns can be discriminated. Finally, I doubt whether any stationary pattern is detected other than by these cues. Even so, the vision of some invertebrates beats robots hollow.

In summary, my thesis is that the only units of vision are cues that are detected by filters. If that is so, even the champion invertebrate visual systems only detect how much of each cue is present. The cues that have been characterised turn out to be similar to the stimuli to which the neurons of the optic lobe respond. All the cues together turn out to be a small part of the original image, which is a sufficient reason to conclude that the neural signals do not re-assemble the image or detect shape. Therefore it is an anthropomorphic error to suppose a priori that any invertebrate 'sees' as we ourselves see. Anyway, nothing like that can be demonstrated experimentally.

12.2 DO INSECTS SEE IMAGES OR DO THEY SEE CUES?

You will need convincing, so let me start with two examples from recent work. Bees are used because similar training and testing experiments have not been done on other invertebrates. For definitions of terms used here and throughout this chapter, see Table 12.1.

Bees were trained in a Y-choice apparatus (Fig. 12.1) to distinguish between a yellow spot and a blue spot, each subtending 20° at the

Table 12.1. *Definitions of terms*

Patterns and targets	*Patterns* are displayed on *targets* when working with bees (Fig. 12.1). The two targets are exchanged in position every 5 min to make the bees look at them, rather than merely come to the right place. To teach the bees not to use undesirable cues, the patterns can be changed regularly, keeping the selected cue constant; for example, radial and tangential cues are revealed when the patterns are rotated during training.
Fixed pattern	A pattern that is fixed relative to the place where the animal makes its choice. The cues are best discriminated when they lie in only one direction from the point of choice, so they are in corresponding positions on the two targets in Fig. 12.1.
Shuffled pattern	A pattern that is changed every few minutes during the training to teach the bees to ignore all cues except the one that is kept constant.
Cue	An abstracted part of the pattern that is detected by a filter in the visual system. In a discrimination between two patterns, the cue usually has two possible states, one of which is rewarded.
Parameter	A scalar or vector measurement of some aspect of the pattern, e.g. the area. Some cues are parameters, but all parameters are not cues.
Template	A predetermined shape that is used by engineers as a gauge, or to be copied. In vision, it has come to mean an innate or learned hypothetical filter or other mechanism that detects a fairly complicated cue.
Sign stimulus	An older and more general term not restricted to vision, e.g. the call of a bird. It is the human idea of the essential stimulus, not necessarily the same as the cue abstracted by the neural filters.
Receptor modulation	The change in the light intensity in the receptor, and the amplitude and frequency of the electrical signal there.

Table 12.1. (*cont.*)

Image	The *image* (as detected by the visual system) is the pattern of modulations in the receptor layer.
Filter	A neuron, or group of neurons, coarsely tuned to detect a cue. It can be represented as a mathematical operator that extracts a component from a spatial pattern.
Field	The region in space and time where a signal from the image is detected by a filter or neuron.
Generalisation	Originally the recognition of a pattern in an unfamiliar place. It now means the recognition of familiar cues in unfamiliar patterns.

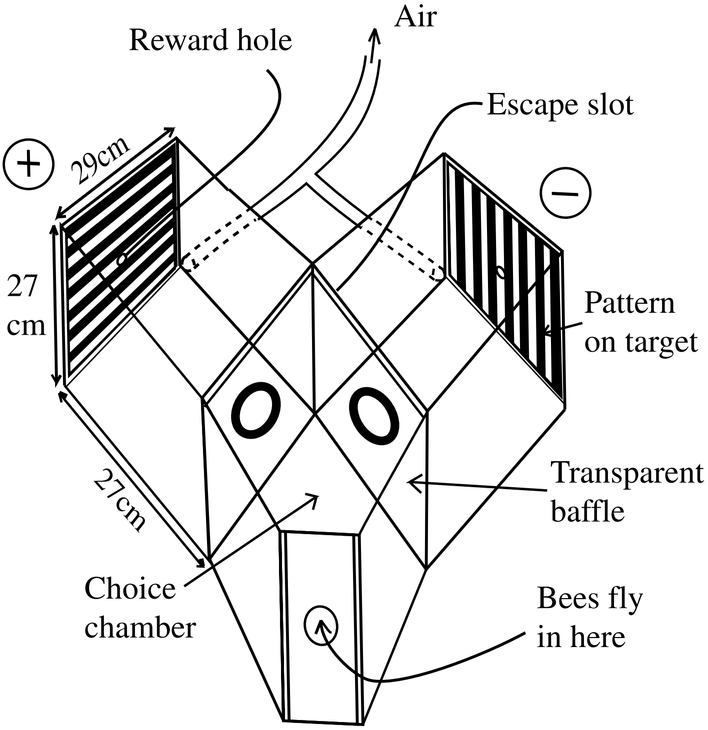

Fig. 12.1 The Y-choice apparatus, which stands on a table in indirect sunlight. The bees enter the choice chamber from which they can see both targets and they choose one of the baffle orifices, 5 cm wide. To make the bees look which side to go, the rewarded target with the reward changes sides every 5 min. +, rewarded training pattern; −, training pattern without reward.

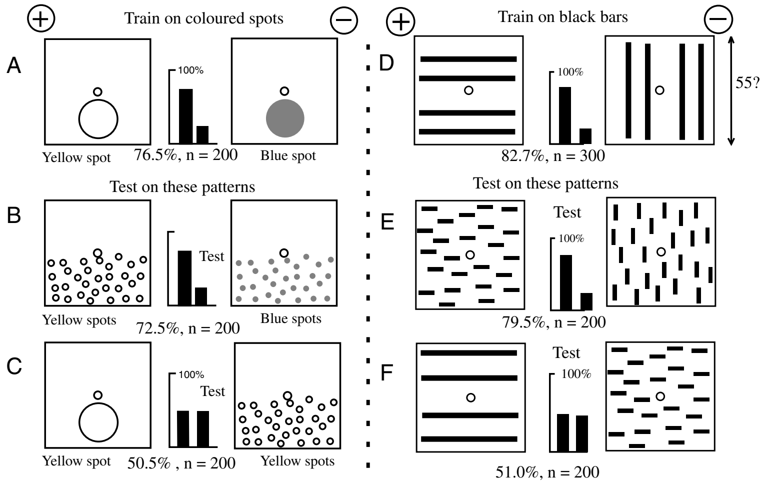

Fig. 12.2 The bees detect the cue, not the pattern. A. The bees are trained to detect a 20° yellow spot (*empty circle*) versus a similar blue spot (*filled circle*). B. The trained bees distinguish the colours in small spots. C. The trained bees cannot distinguish the yellow 20° spot from a scattering of small yellow spots of the same total area. D. New bees are trained on four horizontal thin black bars versus four similar vertical bars. E. The trained bees distinguish the orientations of short bars. F. They cannot distinguish the long bars they were trained on from short bars with the same orientation and area. The small circle at the centre of each target is the reward hole. +, rewarded training pattern; −, training pattern without reward.

choice point of the bees (Fig. 12.2A), an easy task. The trained bees were tested with a scattering of 40 small yellow spots on one target versus 40 small blue spots on the other on a white background (Fig. 12.2B). They distinguished correctly, showing that they discriminated the colour but it did not have to be in the shape of a spot, implying that the learned colour of a flower may be transferred to the same colour in many scattered flowers further away.

Now we come to the interesting result. The trained bees were also tested with the original yellow spot versus a scattering of 40 small yellow spots with a total area the same as the area of the large spot (Fig. 12.2C). They did not distinguish the targets, showing that they did not remember that they had been rewarded on the large spot. They had learned only the colour cue, not the pattern.

In the second example, bees were trained to distinguish between horizontal and vertical long black parallel bars (Fig. 12.2D). The trained bees discriminated the orientation irrespective of the pattern (Fig. 12.2E), but they failed when tested with four long horizontal bars versus many short horizontal bars of the same total length and area (Fig. 12.2F) because the only cue they had learned was present in both patterns. In these examples, the failures were an essential part of the analysis. The actual shape or pattern on which they were trained was not distinguished from quite different patterns with the same cues. They respond as if they never saw the pattern. So, how does an animal with this kind of vision avoid confusion in tasks involving recognition?

12.2.1 Explanation in terms of cues

The concept of the cue is central to the explanation of the mechanism of vision. The properties of cues or feature detectors explain subtle aspects of vision such as transfer and generalisation. In his thoughtful book and earlier works, Jacob von Uexküll (1909) distinguished the units of perception as the meaningful signals from the outside to the animal's innerworld, but for the past century, it has been hard to discover what these units really are. The actual signals in the image that are detected irrespective of the rest of the pattern have now been found experimentally to be a few kinds of cues. A cue, for example, the area or colour, may be recognised in other patterns. Each kind of cue is qualitatively different from other cues and can usually be measured as a number or a vector. In the examples above (Fig. 12.2) the cue is the colour or the angle of orientation of the bars. The cue is the only part that is detected, so the pattern cannot be re-assembled in memory to be consulted at another time.

12.2.2 Cues are detected by neural filters

With the development of electrophysiology and cybernetics after 1945, the idea of filters as the components of the visual system developed rapidly in the second half of the twentieth century. Enthusiasm ran high. Recordings in the visual systems of crayfish and frogs revealed neurons that responded to complex stimuli that were related remarkably to the repertoire of the animal. Unfortunately, it was not a universal panacea. Some remarks from a cry of despair (Vowles, 1964) are apposite at this point: 'The experiments were started partly in response to a suggestion by Bullock (1958) that one might be foolish, when studying electrophysiological responses in higher centers, to use simple stimuli like short flashes of light. He (Bullock) stressed that, as Maturana *et al*. (1960), had shown, the frog's retina already performs quite a complex analysis of the visual world, and the optic nerve carries specific information about various higher-order stimulus characteristics'. Vowles continued, 'The method of stimulating the bee was to project either a white or colored background (Kodachrome floral scene) onto a screen and to place or move against either of these backgrounds various types of stimulus objects'. Later Vowles says 'The experiments on the optic ganglia of the bee showed that *specific neurons may react to very specific stimulus characteristics*' (italics mine).

Probably not, because the neurons were not identified, the actual cues were not defined, and specificity was not demonstrated. However, this illustrates how erroneous conclusions arise because they are compatible with experimental results. Forty years later, we are no further forward in identifying the part played by any neuron for anything in the visual repertoire of the honey bee (or anything much in the other invertebrates). Some large-field neurons in other insects have been related to the directional control of flight, or avoidance of collision (Rind, 1997), but the rest made little sense. For 50 years, there has been negligible electrophysiology of pattern vision of invertebrates because no one knew the units of vision, i.e. what actual cues they detect. Recording from only a single neuron at a time blinkered our understanding, but the principal error was to confuse the sign stimulus, which is the signal that humans abstract on behalf of the bees, or even the pattern, with the actual cues that were detected.

The way forward was indicated by the single neurons in the primary visual cortex in mammals that responded to extremely simple generalised features, notably moving edges at different orientations (Hubel and Wiesel, 1962). In this case, every neuron could be represented as a spatio-temporal filter, but the meaningful signal was carried by the coincidences of many neurons in parallel. The concept was enthusiastically embraced by the artificial vision fraternity (Hinton *et al*., 1986), but robot vision needed much more.

Filter theory, like electrophysiology or neuron anatomy, is not sufficient to explain vision. First, it is essential to discover how many different kinds of parallel paths are active at any one time, exactly what cues they carry, the excitation in each, how they are interconnected, the field sizes, the destinations, the time delays and finally the central reckoning. This leads slowly to an understanding of the way the parts work together. A model with interacting boxes can be made only after the difficult behavioural analysis has been done, not before. Secondly, the visual system relies upon the visual feedback from the movements that it controls; vision is active.

12.3 PERFORMANCE VERSUS MECHANISM

For most lower animals, the visual responses are wonderfully adapted, but successful performance tells us almost nothing about the mechanism within. Researchers who observe the behaviour and postulate a mechanism are soon lost in the neural jungle without a map. However, there were clearer insights when patterns that could be discriminated were compared with related patterns that could not. For example, Sutherland (1960) and Young (1961) showed that the octopus distinguishes between some simple shapes but not other related ones. von Frisch (1914) found that bees could distinguish between several flower-like patterns, but not some geometrical shapes of the same size and colour centred at the reward hole (Fig. 12.3A). He explained this by reference to the repertoire required of a forager searching for flowers. One can find patterns that are effective in causing visual responses, and related patterns that are ineffective (Fig. 12.3), and then infer which components are the cues (see also Fig. 12.10). When the animal fails to discriminate, the cue must be absent and the rest of the pattern is ineffective. This is the proper use of negative evidence in order to map the boundaries of the fields. It is the first step towards identifying cues when there are many pathways for area, position, radial edges, modulation and so on, in parallel. Training, followed by varied and extensive testing, is required for every possible cue, a task that has taken many years in the case of the honey bee. Similar efforts with other invertebrates have not advanced in 40 years.

12.3.1 Cues used by the honeybee

Over almost a century, some of the cues used by bees have been identified consistently, largely irrespective of the methods used for training and testing; others are more recent. The method is to train a group of bees on carefully selected minimal patterns, then give them a variety of interleaved tests with possible cues between periods of continued

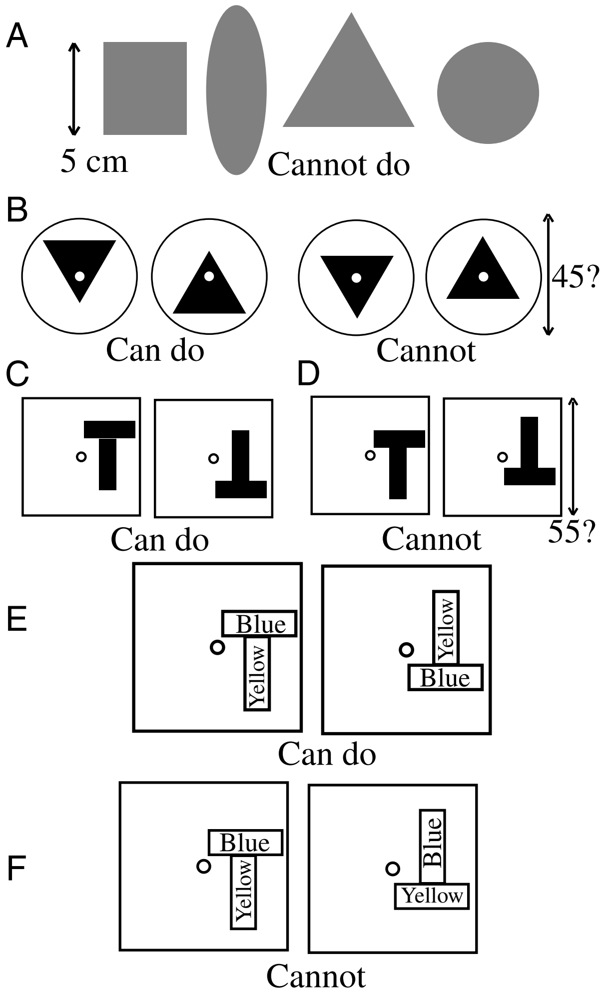

Fig. 12.3 Position is discriminated but not shape. A. In five days of training, von Frisch (1914) failed to train bees to discriminate these shapes, but recent work has shown that some pairs of large shapes can be discriminated by the cues already described (Horridge, unpublished). B. Bees discriminate the inversion of the triangle if the centres are at different positions in the vertical direction (Horridge, 1999a). C, D. They cannot discriminate the black T shapes when the centres are at the same height. E, F. They discriminate the positions of the colours in the vertical direction but not the T shape itself (Horridge, 2003d).

training. It soon becomes clear whether they can or cannot pass the tests. The rest is logical inference; statistics are unnecessary.

Modulation

The simplest cue is the frequency and amplitude of the modulation induced in the photoreceptors by relative movement. It was originally identified as the contour length or disruption of the pattern (Hertz, 1933; Zerrahn, 1933). In the Y-choice apparatus, bees discriminate differences between textures by the modulation irrespective of the pattern (Horridge, 1997). Patterns that differ in spatial frequency are discriminated by both blue and green receptors. Near the resolution limit, bees discriminate between vertical and horizontal gratings by the difference in modulation, not by the orientation. The modulation difference can be avoided by use of oblique gratings and the bees then use orientation detectors (Horridge, 2003c). When modulation is the cue, the resolution depends on the resolution of the individual receptors and is independent of the interommatidial angle.

Colour

The highest priority cue is a colour difference (von Frisch, 1914) and colour is used in the recognition of landmarks (Collett, 1992). In the Y-choice apparatus (Fig. 12.1), colour is learned rapidly. Bees trained on a 20° blue spot versus a similar yellow one (Fig. 12.2A) discriminate between smaller spots with a resolution down to a spot size of about 4° in bright light. When there is a colour difference, the bees do not learn the size or shape of the spot (Fig. 12.2C). When trained on a single coloured spot versus a blank target, bees learn the colour down to a spot size of about 5° in diameter, if the eye is stabilised in the horizontal plane, for which green contrast is essential. In an artificial situation without green contrast, a spot less than about 20° is not discriminated from a blank target (Giurfa *et al*., 1996), not because it is too small but because there is no stable frame of reference and the cue is not presented each time at the same retinotopic place. A harder task, the exchange in position of two spots of different colour in the horizontal direction, is discriminated with spots down to 6° if the frame of reference is stabilised by green contrast (Horridge, 1999b). In the wild, green contrast is everywhere, so that the relative positions of at least two colours are easily learned.

Bees trained on one colour fail when tested with the same pattern in a different colour, or with a texture added. Each time the colour or background is changed during the learning process, they start to learn over again. On the other hand, a coloured spot, unlike most cues, has some salience, in that the bees detect it after it has been moved.

The position of the centre

The position of the centre of a black area (the centroid) in the vertical direction, irrespective of the pattern, is an important cue that was not recognised until recently (Horridge, 2003d, e). The areas must be quite broad; thin black bars are not effective. We can infer that the centroids were all at the same height when von Frisch (1914) tried but failed to train bees to discriminate different shapes of the same size and colour (Fig. 12.3A). I will not name the experimenters who have been convinced that bees discriminate between shapes when the heights of the centres (as measured by the bees) were not controlled.

It has long been known that when bees have learned to discriminate between two fixed black patterns, they may fail to discriminate if a part of the pattern is moved up or down on the target in a test (Friedlaender, 1931). The memory is not lost until they begin to retrain; they just fail to detect the displaced black area. Many have interpreted this and similar results as showing that bees learn a copy of the retinotopic projection of the image on the eye. That is not a very useful idea (Horridge, 1999a; 2005a); it gets us no further forward because the copy must still be processed, and the whole pattern would have to be copied for comparisons with other patterns later. Direct tests of trained bees fail to reveal a copy; instead, they show that cues are detected in the expected places (Figs. 13.2–13.4).

The vertical position of a centre of black in the vertical direction is calculated by the bee, but we do not know exactly with what weighting factors and therefore we cannot place the centre exactly. The vertical and horizontal differences in angular extents are ignored unless that is the sole difference. The bees learn the position of the common centre of black irrespective of pattern (Fig. 12.3B–D), and appear to be unable to distinguish separate positions of individual black areas on the same side of the target (Fig. 12.4A–C). However, two coloured areas on each target can be separately located, and exchange of positions of colours in the vertical direction is easily discriminated even with no green contrast (Fig. 12.4D). Exchange in the horizontal direction (Fig. 12.4E) is discriminated if the frame of reference is

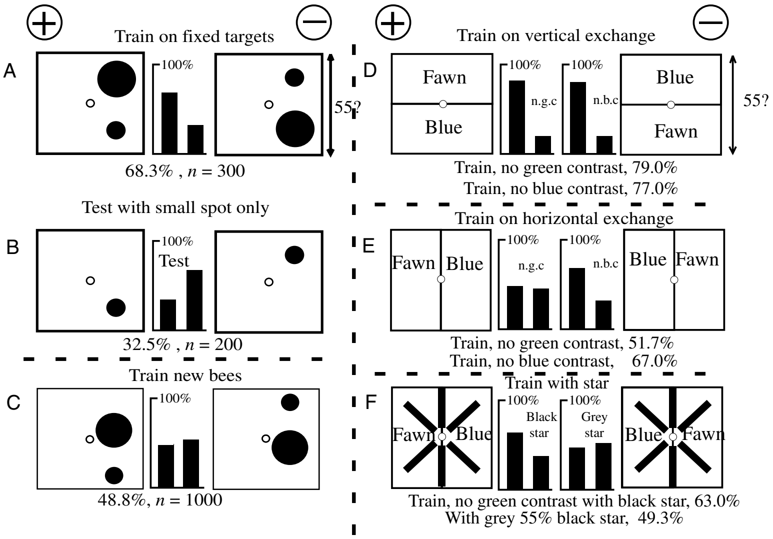

Fig. 12.4 Discrimination of the place of black or coloured areas. A. Two black spots are easily discriminated from the same two spots inverted. B. Preference is reversed when the large spot is omitted. C. Discrimination is lost when the common centroids are at the same level in the vertical direction. D. Discrimination of the exchange of position of two coloured panels in the vertical direction does not require green contrast. E. In the horizontal direction green contrast is required. F. Addition of a black star provides a reference frame in the horizontal direction. +, rewarded training pattern; −, training pattern without reward. After Horridge (1999b, 2000, 2003d).

stabilised by black bars that add green contrast (Fig. 12.4F). The bees discriminate the positions of the colours, not the shapes or the patterns (Fig. 12.3E, F).

Nearness in flight

Bees learn to approach an object at a given range versus a similar object at a different range, even when the positions and apparent sizes of the objects are shuffled during the training. The range of nearby contrasts is calculated from the continual measure of modulation of the receptors caused by the animal's own movement. The range can be learned when flying freely over a horizontal surface (Lehrer *et al.*, 1988) or in the Y-choice apparatus with the targets on vertical surfaces

(Horridge *et al*., 1992). There are indications that range is measured from the angular velocity irrespective of the direction of motion (Srinivasan *et al*., 1993). Measures of the optic flow control the speed in flight and other innate piloting behaviour such as landing and turning, all of which are colour blind and require green contrast (see Chapter 11). However, these experiments also show that the nearness in each direction is remembered in parallel with information about the cues and their directions, so that the bees return to the correct range, even when seeing the cue with the side of the eye (Lehrer, 1990).

Bees therefore have the processing power to survey the surrounding cues and remember a sequence of simple cues with their direction and range, as indeed they must to find their way.

Size

A fundamental cue for a bee is size or area, which is discriminated in colour by both blue and green receptors. Most authors have been careful to control against differences in size of their patterns, so that they could study other cues. When the criterion for a correct choice was landing on the reward, however, the angular size of the pattern increased continually as the bees approached, but they still discriminated size. Apparent (angular) size and range are used in the recognition of familiar landmarks relative to the bees' desired position (Cartwright and Collett, 1983; Collett, 1992).

When presented on a vertical surface, a black spot on one target is readily discriminated from another that differs by 50% in size in the corresponding position on the other target even when the positions are regularly shuffled. The absolute size is recognised, not the relative size (Fig. 12.5). One large black spot can be discriminated among a group of small ones. Green contrast is not necessary and size appears to be measured as area, not as vertical or horizontal extent. When bees are trained to discriminate the size of a fixed 10° black spot from a similar but larger fixed 20° spot on a vertical surface (Fig. 12.5A), they fail in a test when the spots are moved. They look for the cue only in the expected place.

Edge orientation on a vertical plane

The cue studied in greatest detail is the orientation of black bars presented on vertical white targets. When fixed bars are centred at the

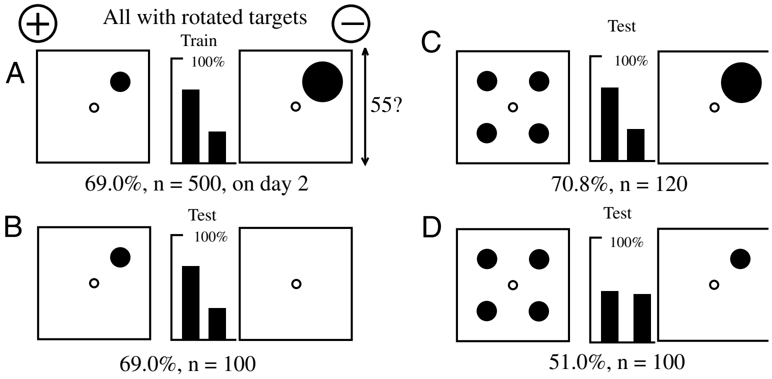

Fig. 12.5 Discrimination of spot size after training with targets rotated by 90° every 10 min. A. Train with a black spot of 10° versus a black spot of 20°. B–D. Three tests showing that the trained bees prefer the small spot in each case, not the smaller or larger amount of black. +, rewarded training pattern; −, training pattern without reward.

same place on the two targets, but differently oriented, the bees learn the orientation cues at the places where they occur during the training. As an example, the bees were trained to distinguish between a single horizontal bar on one target and a single vertical bar on the other target, with both bars across the centre (Fig. 12.6A). When the trained bees were tested with two black lines that represent just the edges of the training bars, they responded well (Fig. 12.6B). The trained bees failed when the orientation cue was destroyed by large square steps (Fig. 12.6C), showing that no cue remained although the positions of the black areas had not changed. When the training bars were moved, the trained bees failed to distinguish them (Fig. 12.6D). The bees discriminated the orientations of edges in the expected places, irrespective of whether the rest of the bar was there or not (Horridge, 2003a, e).

However, if broad black bars are fixed in different places on the two targets, the bees learn their locations. A test with the orientations reversed reduces but does not reverse the preference (Fig. 12.7A, B). If the bars are moved in the vertical direction after training, the bees fail to discriminate in a test (Fig. 12.7C–D). The orientation cue must be in the expected place in the vertical direction. If a black spot or a dappled background is added to disrupt the frame of reference, they fail. The learning of cue direction is less exact in the horizontal direction, and some preference remains when the bars are moved horizontally in tests

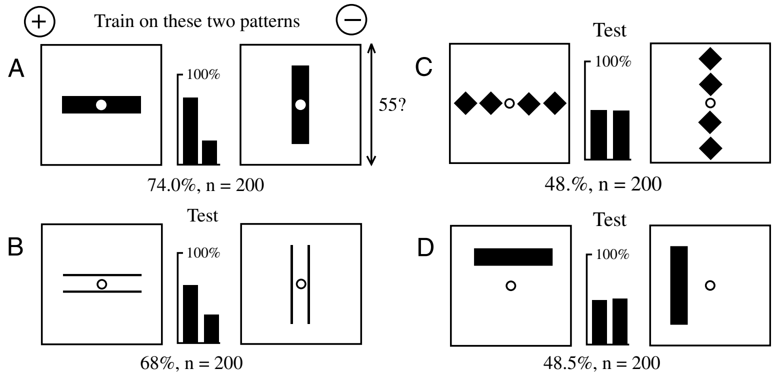

Fig. 12.6 An example where the cue is the edge orientation, and it is detected only in the expected place. A. Training on orthogonal bars centred at the same place. B. Test with only the edges represented by thin bars. C. The bees fail when the bars are broken into diamonds to spoil the orientation cue. D. They also fail when the bars are shifted to new positions. +, rewarded training pattern; −, training pattern without reward. After Horridge (2003a).

(Fig. 12.7E, F). When differently oriented black bars are centred at different places on the two targets and moved about, the bees learn nothing because the position cue is not constant and the orientation cues are not at corresponding places on the two targets (Horridge, 2003a, e).

From 1990 until recently, I thought that when the bars were moved during the training (van Hateren *et al*,. 1990), the bees learned the orientation cue irrespective of place on the target, but recently found that the trained bees recognise the orientation cue only in the places where it occurred during the training (e.g. Figs. 13.6, 13.7).

To learn the orientation cue, the differently oriented edges must be in corresponding positions on the two targets, so that the bees look in one direction from the point of choice. The learning or recognition of the orientation cue requires green contrast at the oriented edges, so it is easily eliminated by use of two suitable colours. Care must be taken with fixed patterns because the bees may also learn the positions of the colours or modulation of the blue receptors, as they do with horizontal versus vertical gratings with no green contrast (Horridge, 2003c). Differences in modulation are avoided by training with oblique gratings (Figs. 13.7E, 13.8A).

The orientation detectors have the interesting property that when there are equal lengths of edges at right angles in the field of one eye, the orientation cue is cancelled, so a square cross, staircase,

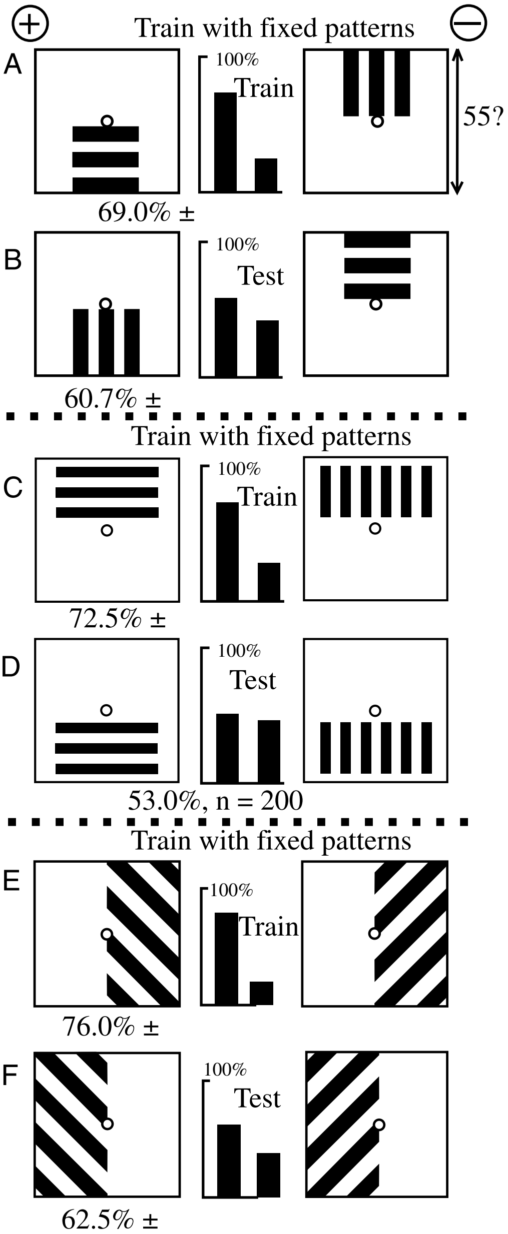

Fig. 12.7 The position of black, and the position of the orientation cue. A. Train with three horizontal bars below the reward hole versus three vertical bars above it. B. A test with the orientations reversed reduces but does not reverse the preference. The position was the stronger cue. C. Train with horizontal bars above the reward hole versus vertical bars

or square carries no orientation cue (Fig. 12.6C). The failures show that they detect no differences between these patterns. This would be the effect if the orientation detectors have angular sensitivity plots 90° wide at the 50% level of sensitivity and large spatial fields (Srinivasan *et al*., 1994). Certainly some orientation detecting neurons of the lobula have very wide fields (Yang and Maddess, 1997), but there are alternative mechanisms for the cancellation of orthogonal orientations.

It was recently confirmed that bees do not detect the orientation of rows of round spots, illusory edges, or patterns of squares if the separate items are resolved (Horridge, 2000a; 2003a). If squares in a row are so close together that they are not separately resolved, they are seen as a line that carries an orientation cue. Unlike human vision, however, if the squares are moved apart until they are separately resolved, the orientation cue disappears (Fig. 12.8B–C). There are no global orientation detectors that span a gap of more than three interommatidial angles (Horridge, 2003b).

When the bees have learned the orientation cue, they fail in tests to distinguish the rewarded training pattern from a different pattern composed of the same total length of shorter edges with the same orientation (Fig. 12.2F). They recognise the cue, not the pattern. The bees also fail to distinguish the two edges at corners; they take the summed orientation cue over a large field that is limited to each eye. The failures of the trained bees to discriminate show that they do not recognise the bars or gratings as patterns. They have filters for cues, not pattern vision.

The minimum length of the orientation detectors was measured directly by training the bees to orientation with shuffled oblique bars (Fig. 12.8A). The trained bees were then tested with targets that were filled with short parallel bars (Fig. 12.8D), which were all reduced in length until the orientation was no longer resolved. This gave a minimum length near 3° for the orientation detector (Horridge, 2003f).

We can now make a model of the circuit (Fig. 12.9). In the first stage, green receptors feed into large lamina cells that detect temporal

---

Caption for Fig. 12.7 (cont.)
(of the same total area) also above it. D. Moving the bars down spoils the preference. The orientation cue must be in the expected place in the vertical direction. E. Train with oblique bars versus orthogonal oblique bars. F. Some preference remains when the bars are moved horizontally. +, rewarded training pattern; –, training pattern without reward.

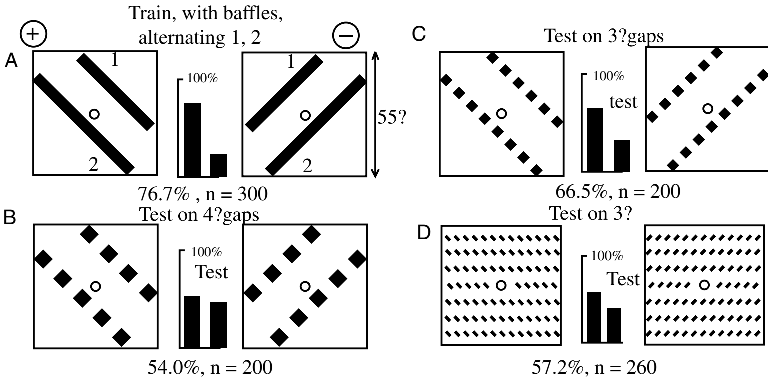

Fig. 12.8 The maximum and minimum length of the local orientation detectors. A. Train with oblique bars that are shuffled in location (*1, 2 at top*). B. Test with 4° squares separated by 4° (the smallest gaps that are resolved). C. Test with 3° squares separated by 3° (After Horridge (2003b)). D. Test near the lower limit of the effective edge length with bars 3° long (after Horridge (2003f)). +, rewarded training pattern; −, training pattern without reward.

modulation of individual ommatidia. Axons on adjacent axes connect to units that detect simultaneous modulation, and are therefore local orientation detectors. The direction that is detected depends on the spatial arrangement and sign of the inputs. Being so short, the local detectors are coarsely tuned to an axis of orientation, and noisy. The axes are aligned in at least three directions (Srinivasan *et al*., 1994). Large numbers of these local detectors with parallel axes (across the whole of one eye) feed into global orientation detectors (Fig. 12.9). Therefore, the global detectors are also coarsely tuned with the same axes of orientation, but with improved signal/noise ratio when many parallel edges are summed. Other combinations of the local orientation detectors also feed into quite different large-field detectors of circular and radial edges (see Fig. 12.11).

The width of the orientation detectors has not been measured directly, but must be at least three ommatidia wide (Fig. 12.9) because a simple flash does not excite and the orientation is unchanged when a black/white edge is reversed (Horridge, 2005b). Trained bees can detect the orientation of a gradient as low as 2% per degree (Horridge, 2000a). This extreme sensitivity to the orientation of a very small intensity gradient suggests that the orientation detectors are designed to overcome visual blurring of distant edges.

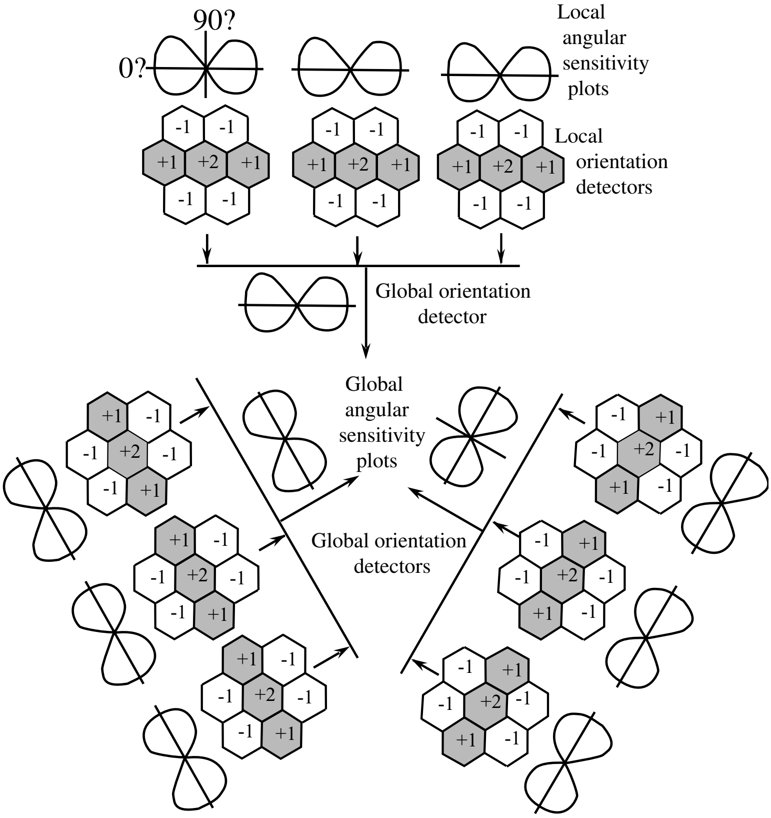

Fig. 12.9 Suggested arrangement of inputs of the orientation detectors. Seven receptors in a group feed into each local orientation detector. There are three orientations of local orientation detectors, with axes at 120° to each other. Local orientation detectors with axes in common feed into the global orientation detectors, which therefore have the same broad angular sensitivity plots as the local ones. The resolution of orientation is that of the detectors, not the interommatidial angle (Horridge, 2005b).

In 1995, when I started working with equal lengths of edges at right angles on the same target, there was a view that the detection of the orientation of edges enabled bees to discriminate between patterns. In fact, those experiments showed that neither the bars nor the orientations were detected separately. The orientation detectors actually sum the orientations and destroy pattern (Fig. 12.10D–F). Therefore they cut out the details of the background foliage with green contrast that is encountered everywhere. Detection of edge orientation

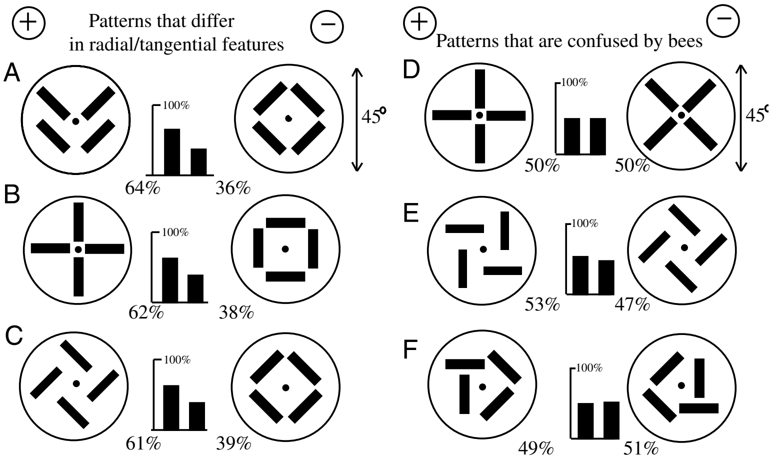

Fig. 12.10 When the orientation cues have been eliminated by use of patterns with pairs of orthogonal bars, the bees are able to discriminate if there is a difference in the radial/tangential cues. A–C. Pairs with a radial/tangential difference that is discriminated. D–F. Pairs that are confused by the bees. +, rewarded training pattern; −, training pattern without reward. After Horridge (1996).

is certainly not concerned with texture or with several different orientations simultaneously, but with an average orientation.

The radial/tangential cue

von Frisch was impressed by the bee's ability to discriminate between two flower-like patterns of similar size, in contrast to his failures with other closed shapes. Zerrahn (1933) and later Lehrer *et al.* (1995) found a preference for radial patterns by untrained bees and aversion towards circular patterns. Hertz (1933) trained bees to discriminate between different radial patterns that were laid out flat on a white table, when edge orientation was useless as a cue. Friedlaender (1931), with vertical targets, demonstrated that a radial pattern provides a reference point that can be found by the bees when a part of the pattern is moved in tests.

When the two targets are rotated at random during the training, the position and orientation cues are shuffled, but radial and tangential cues remain. The bees learn rapidly to discriminate between a radial pattern of sectors or bars and a pattern of concentric circles, both 50% black, 50% white. The trained bees respond correctly when tested with radials versus checkerboard and circles versus checkerboard.

By 1995, it was known that the orientation cue was cancelled in patterns with pairs of orthogonal bars, even in fixed patterns. With the orientation cue removed, the bees could discriminate any remaining radial and tangential cues, irrespective of pattern (Fig. 12.10A–C). These results, together with the positive and negative preferences, pointed to the existence of separate cues from radial and tangential edges, irrespective of pattern. These cues require green contrast. There is evidence for filters with three axes at 120° and non-specific filters that detect any radial arrangement, but radial patterns with six axes are preferred (Horridge, 2000c).

The bilateral symmetry cue

Bees approaching bilaterally symmetrical flowers land in line with the axis of symmetry (Jones and Buchmann, 1974). Flowers evolved with radial symmetry to assist detection by insects that had already been flying with a radially symmetrical flowfield for hundreds of millions of years. Bilaterally symmetrical flowers evolved to match an already existing ability to land in the middle of a runway. Untrained bees prefer bilateral symmetry to asymmetry in patterns that they have never seen before (Lehrer *et al.*, 1995). Bumble bees prefer to land on symmetrical real flowers rather than disfigured ones (Møller, 1995), and bees recognise bilateral symmetry in many simple patterns of black bars of equal area, size, edge length, and lacking other cues (Horridge, 1996). We cannot conclude that the bees 'see' the symmetrical patterns, or that they detect symmetry in general, only that they detect some symmetrical spatial arrangements of cues. As for symmetry detection in man or other animals, the mechanism has not been demonstrated, though theories abound.

12.3.2 A limited variety of cues

The same cues re-appear in studies by a variety of training and testing techniques. Bees fail to learn to discriminate between a variety of different patterns lacking all the above cues (see Fig. 12.10D–F), and trained bees fail in tests if these known cues are removed. Some cues have been recognised for decades, although described in different terms; further examples, such as modulation of polarisation or certain chromatic contrasts, may be discovered. Most of the cues, including the absolute size, are relatively insensitive to the range, which is probably why they have been selected in evolution.

Cues for learning the angular velocity of an object or the relative velocity of two edges appear not to be available (Lehrer and Srinivasan, 1992). Bees will learn to come to an edge that displays parallax (Srinivasan *et al*., 1990) and they will then land at right angles to an edge, but tests for learning edge orientation from parallax fail (Horridge, 2003a). The place of motion perception in discrimination of other cues is not known, largely because it is little studied.

12.3.3 Do they see more than they remember?

Cues that are detected but not remembered are useless, so we can expect that memory tests reveal all the cues. The animals respond innately or they seek the reward vigorously. The above cues apply strictly to the memory, because the results are based on tests of animals that have been trained. However, when resolution is measured, the limits of discrimination appear to be determined by the peripheral visual system. The identification of any cue is itself a survey of the boundaries of a filter, of which the measurement of resolution is a special case.

12.3.4 How do they know where to look?

Before the experiment, sugar syrup was provided at both blank white targets in the Y-choice apparatus (Fig. 12.1). The bees arrived at the choice chamber with no indication which side to go but they were familiar with the geometry of the apparatus. Early in the morning of the experiment, the training patterns were placed on the targets, only one of which was rewarded, so at first only 50% of the choices were correct. The rewarded pattern with the reward changed sides every 5 min so that the bees were obliged to look which side to go. The geometry of the apparatus provides several reference points and contrasting edges that are the same in each arm, except for the patterns. The bees' steady posture in flight allows a measure of the positions of the centres of the black areas in the vertical direction but bees yaw in flight in the horizontal direction, so that green contrasts must provide the frame of reference there (Fig. 12.4D–F).

After 2 to 4 h, depending on the task, the bees were sufficiently trained. Trained bees failed to discriminate orientation in tests in which a broad bar was moved to a new place on the target, so neither the bar nor the orientation cue had salience, otherwise the bees would have detected them when they were moved (Fig. 12.6D). This is a primitive example of 'attention blindness'. Discrimination was spoiled by the addition of a black spot or a patterned background (Horridge, 2003a, e), as if these additions modify the reference coordinates. However, when one or more thin black bars were shuffled in corresponding positions on the two targets during the training (Figs. 13.2D and 13.8A), the bees learned to expect the orientation cues within the range of places where they had occurred during the training.

When the bees failed to detect a cue from the choice chamber, they turned to the other arm, but when they detected a cue, they headed towards the reward hole. We have no evidence that the bees fixate only on the reward hole in the Y-choice apparatus. Instead, they learn reference coordinates from the geometry of the whole apparatus, and at the point of choice they learn to look for the coincidences of consistent cues in the expected directions (Horridge, 2003a,e).

12.3.5 Shape versus particular shapes

Ants, bees, and wasps can be trained to come to a particular shape versus a blank target, and they certainly discriminate between many (but not all) large shapes of similar size and colour when presented as alternatives. It is an error to conclude that they discriminate shape in general, or even that they detect the pattern. Nothing can be concluded about mechanisms when the discriminations are all successful. The cues are discovered only by designing tests in which discrimination fails. When trained on a particular shape versus a blank target, the bees learn cues such as the size, colour, modulation (length of edge), and position of the centre. When trained to discriminate from a distance between two different shapes of the same size and colour, they learn cues that are suitable for that particular pair of shapes, not for shape in general (Horridge, 2003a, e; 2005a). Very large patterns are more easily discriminated because they extend over more than one field of the detectors. It is notable that research on shape discrimination has failed to demonstrate a retinotopic memory but has omitted to test for the proposed parameters or to mention the problem of re-assembly of the different edges.

As said, the different orientations of the edges of a pattern are not detected separately but are combined to an average (Srinivasan *et al*., 1994; Horridge, 2000a). The resulting failures in the discriminations of many pairs of patterns show that there is no separate pathway in parallel that looks after shape discrimination when the known cues are lacking (Fig. 12.10D–F).

I know of no evidence for pattern or shape discrimination or re-assembly of the pattern that cannot be explained by coincidences of the known cues, usually a difference in modulation, position of the centroid in the vertical direction, position of a large oriented edge, or radial versus tangential edges. When the cues and their field sizes are better known quantitatively, it should be possible to predict the responses to any combination of patterns by working forward from the cues. Working the other way, the inference of mechanisms from successful discriminations, is an impossible exercise because there may be many cues in parallel in unknown proportions.

The final nail in the coffin of pattern perception is the type of experiment illustrated in Fig. 12.2. The bees are trained to distinguish between two patterns that differ in an obvious cue, and then they fail when presented with the rewarded pattern versus a third pattern that also contains the same cue. They cannot remember which pattern they were trained on. Experiments of this type have not been done on other simple visual systems.

12.4 THE VISUAL PATHWAYS

We can draw a map (Fig. 12.11) of the separation of pathways in the visual system with one type of filter for each kind of cue (Horridge, 2000b). This map suggests the design of further training experiments and the appropriate stimuli for electrophysiology. Such maps are badly needed for other examples of visual systems with an intermediate level of complexity, such as octopus, jumping spider, crab, fish, frog, and chicken, but the work has hardly started.

In this map, there is no provision for the detection of orientation of edges from parallax. When this point was recently tested, no evidence was found (Horridge, 2003a). There is no path for a transfer between green and blue receptors. Indeed, if there were, colour discrimination would be impossible and orientation discrimination would not be colour blind. There is no way that whole patterns can be discriminated because the cues are only small components of the image and the large fields of the detectors average out the local orientations. Apart from a few preformed templates that detect radial and bilateral symmetry, and perhaps small spots, form and shape appear to be irrelevant.

This scheme accounts for the cues discovered in the Y-choice apparatus, for maze running and also for natural landmarks. It is an essential part of the total theory that the landmarks are recognised

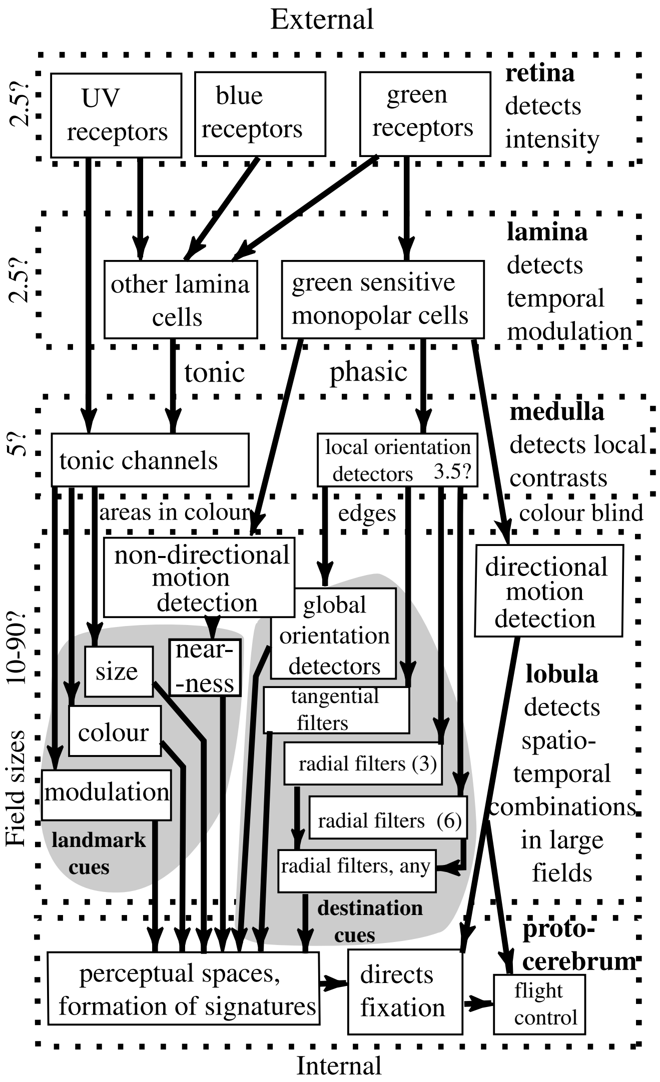

Fig. 12.11 The arrangement of separate pathways in parallel in the visual system of the bee behind every local place on the eye, as inferred from discrimination experiments. This system converts the spatial lay-out of a pattern into scalar measurements of retinotopic cues that are later averaged in large fields. Tonic channels with colour vision are on the *left*; phasic colour blind channels on the *right*. Revised from Horridge (2000b).

by the same cues as in the experiments with fixed patterns. When the expected cues in the expected retinotopic directions coincide within the animal's moment of time, then it makes a choice. No re-assembly of patterns or of the scene is required.

12.4.1 Why is experimental learning so slow?

Bees learn landmarks very rapidly but the natural environment is inconvenient for detailed experimental manipulation of the stimulus; for example, it is hard to randomise landmarks. Learning of patterns in the Y-choice apparatus (Fig. 12.1) may take 15–20 trials, for several reasons. First, the angle scanned by the bees is only about 55°. Wehner (1969) had targets that subtended 130°, and in many early experiments the bee landed on the target as the criterion, so that the visual angles were huge. In the freely flying bee, the angle is nearer to 300° and a fix on landmarks at large angles to each other is far more useful than at small angles. Second, the Y-choice apparatus is lined with white paper, originally to make the bees concentrate on the targets in front of them. I eventually discovered that white paper reduces the frame of reference and landmarks within the apparatus. Third, when learning to discriminate in the Y-choice apparatus, the bees have to look in either arm of the apparatus, they must learn to look in the right direction, and also identify the cue. The learning appears to be done by trial and error, and every error delays the learning process.

12.4.2 The global coincidences of cues

The Y-choice apparatus (Fig. 12.1) teaches us that bees learn to use fixed cues as if they are landmarks. If we extend this result around the eye, we generate an array up to 300° wide that detects cues in different directions (Fig. 12.12). The panoramic view includes a measure of the modulation in each area and the expected range and direction of each cue relative to the head. In combination, the cues define that place for the bee which can also respond appropriately to each cue individually. Local detail is lost by averaging over large fields, with one averaged cue in each field at any one time. The visual system generates a sparse map of the surroundings with a delicate compromise between too much and too little information.

The pathways in Fig. 12.11 fall into two groups, with and without colour vision, the former concerned with recognising a location from the directions of familiar landmarks. The landmarks provide

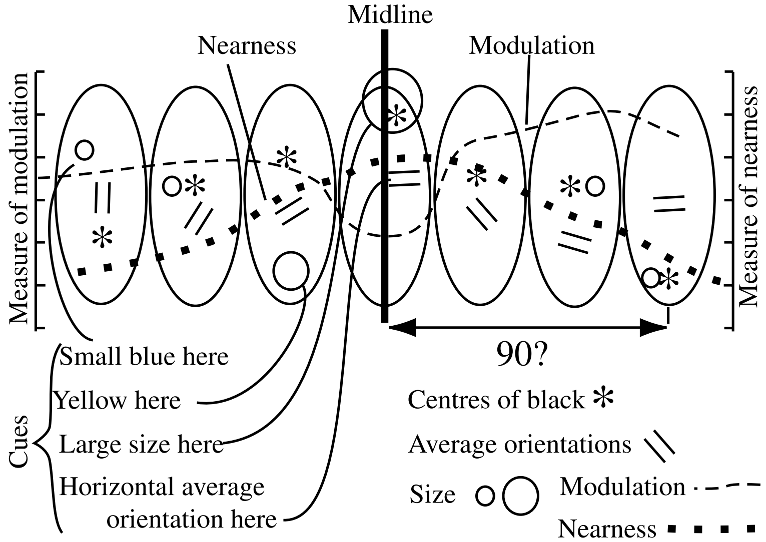

Fig. 12.12 A map of the surroundings is composed from the direction and nearness of each identified cue. The large fields of both eyes are divided into regions along the horizontal midline. Each of these regions can detect any of the cues including a smoothed measure of modulation (*dotted line*) and a measure of nearness = 1/range (*thin dashed line*). Learned cues, as shown by the *symbols*, are expected in certain directions relative to the midline. The bees look for the coincidences of the learned cues at the expected directions and ranges with the cues that they detect in the surroundings. At each correct conjunction of cues, the bees have learned which direction to go towards the next familiar place.

the frame of reference for the memory of the direction of the next cue. Pathways without colour vision are derived from the local edge detectors and cancel mixed orientations but detect flower-like patterns. Radial/tangential and bilateral symmetry cues have been omitted from Fig. 12.12 because they indicate landing destinations rather than landmarks.

With sparse panoramic vision of cues, it is possible for the bee to maximise the fit between the distribution of cues on the retina and the memory of it, as described by researchers on landmarks (Fry and Wehner, 2002). The analysis of pattern perception has now converged with the study of navigation by landmarks until they amount to the same topic. What we assumed was that pattern perception turns out to be the detection of the useful cues from landmarks and the conversion of a rich scene of contrasts to a sparse map for place recognition with a frame of reference based on the surroundings.

12.4.3 The strategy of vision with a small brain

There was a flurry of excitement 45 years ago when Maturana *et al.* (1960) described the fly detector of the frog. It was supposed that a small brain could have economical processing of a few complex but specific features by use of highly adaptive feature detectors, which would be sufficient to initiate a response. These would be the units of vision. 'Extraction of particular aspects of stimuli obviously must occur to account for observed behaviour. But to see it in single units early in the afferent pathway is impressive' (Bullock and Horridge, 1965, p. 280). However, research on mammals and the bee has taken us in an entirely different direction. The detectors of visual cues are simpler units but they come in arrays and require coincidences.

For years, we have wondered why the vast detail of the image on the insect retina is digested into relatively few deep optic lobe neurons with large fields but little sign of pattern discrimination. When we thought that bees were discriminating patterns on targets we were fooled by the total subservience of the visual processing system to detect a sparse panorama of averaged cues as if they were landmarks. It is at last apparent why the deep optic lobe neurons have large fields, with multimodal inputs from motion detectors and other modalities. They detect coincidences. Far from being interested in the maximum number of pictures that the eye could distinguish (Snyder *et al.*, 1977), the bee divides the surrounding world into a few large highly smoothed fields with reduced noise, and looks for the coincidences of a relatively small number of cues with their expected directions (Fig. 12.12). The high receptor resolution is only the initial efficient capture of what is available. However, the fields of the deep detectors are large only in the context of a bench experiment and not in a 360° scene. The optimum for remembering a familiar place to forage seems to be a fix in three or four directions, in each with two or three cues at different ranges. We are now better placed to devise experiments on place recognition and to guess what stimuli to use for electrophysiology.

Although the cues detected by the filters are very simple, they are adequate because the system relies on coincidences of parallel pathways, each of which is labelled with its position on the body as everywhere in all nervous systems. Different scenes are distinguishable because, given a few variables in a few directions, there are sufficient permutations. This strategy not only avoids the enormous task of re-assembling the spatial pattern that is split up by the retina, it is insensitive to the reduced resolution in low light levels and to differences in the numbers of ommatidia. We can see it at work whenever a butterfly, dragonfly, or wasp returns to a familiar resting place.

12.5 PATTERN PROCESSING IN OTHER ANIMALS

Almost all work has concerned the discovery of sign stimuli by observation of successful normal behaviour, with almost nothing on the actual cues detected by the visual system. There are a few exceptions.

12.5.1 *Drosophila*

This small fly is attracted to the centre of a broad black short vertical bar, but to the edges of a similar tall one (Wehner, 1972), which is one example of the common distinction between edges and areas in insect vision (see Fig. 12.11). When free to control the motion of the simulated scene around it, *Drosophila* detects a few cues, notably the size and vertical position of the centre of gravity, the vertical and horizontal extents, and the vertical distance between two black areas. Otherwise two black areas on the same target are treated as one (Ernst and Heisenberg, 1999). Discrimination is lost if the cue is moved upwards on the target, as in the bee. The problem with this animal is not the elegant collection of data, but the probability that the study of trained flies is limited by the memory, not by the visual processing.

12.5.2 Octopus

The octopus and its relatives can also be trained to a visual stimulus. The problem is the paucity of data despite those dozens of sea-water tanks in the Stazione Zoologica in Naples, where the animals were kept in school for years. Octopus recognise crabs, and they learn to discriminate between horizontal and vertical bars, size, black versus white, and squares versus round discs or triangles, but they are poor at discriminating the separation of spots (Young 1961). Sutherland (1960) used a variety of geometrical patterns, like Z, T, and N, and had a theory that two of the cues were the vertical and horizontal extents of the patterns, and that if the vertical and horizontal projections were similar, the patterns would not be discriminated. The problem was that data were slow to acquire and never sufficient, but it is clear that the visual system does not re-assemble the pattern.

12.5.3 Jumping spiders

These animals jump upon any small object that moves in a jerky way, except that they must mate and avoid battles with larger males. Much of this behaviour is fascinating in that it shows that males recognise females, prey and other males visually; in *Corythalia* even in a mirror. In this species, legs of the expected form, a yellow patch, and movement are essential for dummies to elicit the courting reaction (Crane, 1949). Salticid spiders can be trained to not jump on a black triangle or a cross, and then they distinguish other patterns from the training pattern (Drees, 1952), but again the problem is to collect sufficient data to begin understanding the system.

12.5.4 How do we analyse visual processing in other animals?

First, find the repertoire. You take your spider, octopus, mantis or whatever, and study it with infinite patience until every kind of visual stimulus to which it responds is discovered. This is sometimes a frustrating puzzle, for example, the question of colour vision in octopus and squid that have coloured displays. The problem is to separate the various effective stimuli in the repertoire and use the animal's own responses in such a way that the boundaries of the cues can be defined. We use as models the elements that have been painstakingly identified in other visual systems and used in robot vision.

The boundaries of each cue are then mapped by progressive exploration. There's nothing like a resolution test to trick us into believing that the animal can 'see'. Finding the failures of discriminations is the only way that we can say that we have found the limits of the filter; for example, some measure of the cue is reduced until the response fails. The spatial resolution is that of the cue, not the interommatidial angle (Horridge, 2005b).

Finally, there is the question of whether the cues are reassembled into the original image. The conclusion so far, for all invertebrates, and up the vertebrate scale to the amphibians, is that nothing like this is possible.

Vision based on peripheral detection of very simple cues, such as modulation and local edge directions in small fields, may be universal in all animals, and the use of coincidences may be the usual strategy. Like a human patient with blindsight, they detect only the cues and their directions. Insect vision, and perhaps all vision with a small brain, must be like our sensation when hearing, taste, smell, and touch act together; that is a set of coincidences, not a spatial panorama. The key to visual mechanisms appears to be the coincidences of simple directional cues.

12.6 CONCLUSIONS

It seems a terrible thing to say that a painstaking description of an effective visual stimulus is only one step along the road to discovering the visual cues. That beguiling magician, Anthropomorphism, has cast a spell on human scientists to see through animal eyes as if they were seeing through their own. Insects, crabs, octopus, and goldfish appear to look at things and see them, but we have evidence only that they detect something simple that is relevant to their special repertoire. To us it may be a jerky fly, a moving shadow, a red wriggle, a colour, or a circular shape like an eye; in other words, the sign stimulus. For the animal, however, the cues within the sign stimulus are far more abstract. The four patron saints of research, Observation, Experiment, Analysis, and Refutation, assist us to find what the cues really are. They certainly do not suggest that animals with small brains ‘see’ the sign stimulus or that they enjoy the wonderful panorama already endowed with meaning which we think we ‘see’.

REFERENCES

Bullock, T. H. (1958). Parameters of integrative action of the nervous system at the neuronal level. *Exptl. Cell. Res., Suppl.*, **5**, 323–27.

Bullock, T. H. & Horridge, G. A. (1965). *Structure and function in the nervous systems of invertebrates*. Freeman: San Francisco and London.

Cartwright, B. A. & Collett, T. S. (1983). Landmark learning in bees; experiments and models. *J. Comp. Physiol.*, **151**, 521–43.

Collett, T. S. (1992). Landmark learning and guidance in insects. *Phil. Trans. R. Soc. Lond. B*, **337**, 295–303.

Crane, J. (1949). Comparative biology of salticid spiders at Rancho Grande, Venezuela. IV. An analysis of display. *Zoologica*, **34**, 159–214.

Drees, O. (1952). Untersuchungen über die angeborenen Verhaltensweisen bei Springspinnen. *Z. Tierpsychol.*, **9**, 169–207.

Ernst, R. & Heisenberg, M. (1999). The memory template in Drosophila pattern vision at the flight simulator. *Vision Res.*, **39**, 3920–33.

Friedlaender, M. (1931). Zur Bedeutung des Fluglochs im optischen Feld der Biene bei senkrechter Dressuranordnung. *Z. vergl. Physiol.*, **15**, 193–260.

Fry, S. N. & Wehner, R. (2002). Honeybees store landmarks in an egocentric frame of reference. *J. Comp. Physiol. A*, **187**, 1009–16.

Giurfa, M., Vorobyev, P., Kevan, P., & Menzel, R. (1996). Detection of coloured stimuli by honeybees: the role of chromatic and achromatic contrast. *J. Comp. Physiol. A*, **178**, 699–709.

Hertz, M. (1933). Über figural Intensitäten und Qualitäten in der optischen Wahnehmung der Biene. *Biol. Zentralbl.*, **53**, 10–40.

Hinton, G. E., McClelland, J. L., & Rumelhart D. E. (1986). Distributed representations. In *Parallel Distributed Processing*. eds. D. E. Rummelhart & J. L. McClelland. Cambridge, MA: MIT Press, pp. 77–109.

Horridge, G. A. (1996). Vision of the honeybee Apis mellifera for patterns with two pairs of equal orthogonal bars. *J. Insect Physiol.*, **42**, 131–38.

Horridge, G. A. (1997). Pattern discrimination by the honeybee: disruption as a cue. *J. Comp. Physiol. A*, **181**, 267–77.

Horridge, G. A. (1999a). Two-dimensional pattern discrimination by the honeybee. *Physiol. Entomol.*, **24**, 1–17.

Horridge, G. A. (1999b). Pattern vision of the honeybee (Apis mellifera): The effect of pattern on the discrimination of location. *J. Comp. Physiol. A*, **185**, 105–13.

Horridge, G. A. (2000a). Pattern vision of the honeybee (Apis mellifera). What is an oriented edge? *J. Comp. Physiol. A*, **186**, 521–34.

Horridge, G. A. (2000b). Seven experiments on pattern vision of the honeybee, with a model. *Vision Res.*, **40**, 2589–603.

Horridge, G. A. (2000c). Visual discrimination of radial cues by the honeybee (Apis mellifera). *J. Insect Physiol.*, **46**, 629–45.

Horridge, G. A. (2003a). Discrimination of single bars by the honeybee (Apis mellifera). *Vision Res.*, **43**, 1257–71.

Horridge, G. A. (2003b). The visual system of the honeybee (Apis mellifera): the maximum length of the orientation detector. *J. Insect Physiol.*, **49**, 621–28.

Horridge, G. A. (2003c). Visual resolution of gratings by the compound eye of the bee (Apis mellifera). *J. Exp. Biol.* **206**, 2105–10.

Horridge, G. A. (2003d). Visual discrimination by the honeybee (Apis mellifera): the position of the common centre as the cue. *Physiol. Entomol.*, **28**, 1–12.

Horridge, G. A. (2003e). The effect of complexity on the discrimination of oriented bars by the honeybee (Apis mellifera). *J. Comp. Physiol. A*, **189**, 703–14.

Horridge, G. A. (2003f). Visual resolution of the orientation cue by the honeybee (Apis mellifera). *J. Insect Physiol.*, **49**, 1145–52.

Horridge, G. A. (2005a). What the honeybee sees: a review of the recognition system of Apis mellifera. *Physiol. Entomol.*, **30**, 2–13.

Horridge, G. A. (2005b). The spatial resolutions of the apposition compound eye and its neuro-sensory feature detectors: observation versus theory. *J. Insect Physiol.*, **51**, 243–66.

Horridge, G. A., Zhang, S. W., & Lehrer, M. (1992). Bees can combine range and visual angle to estimate absolute size. *Phil. Trans. R. Soc. Lond. B*, **337**, 49–57.

Hubel, D. H. & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond.)*, **160**, 106–54.

Jones, C. E. & Buchmann S. L. (1974). Ultraviolet floral patterns as functional orientation cues in hymenopterous pollination systems. *Anim. Behav.*, **22**, 481–85.

Lehrer, M. (1990). How bees use peripheral eye regions to localize a frontally positioned target. *J. Comp. Physiol. A*, **167**, 173–85.

Lehrer, M., Horridge, G. A., Zhang, S. W., & Gadagkar, R. (1995). Pattern vision in bees: preference for radial patterns. *Phil. Trans. R. Soc. Lond. B*, **347**, 123–37.

Lehrer, M. & Srinivasan, M. V. (1992). Freely flying bees discriminate between stationary and moving objects: performance and possible mechanisms. *J. Comp. Physiol. A*, **171**, 457–67.

Lehrer, M., Srinivasan, M. V., Zhang, S. W., & Horridge, G. A. (1988). Motion cues provide the bee's visual world with a third dimension. *Nature (Lond.)*, **332**, 356–57.

Maturana, R., Lettvin, J., Pitts, W., & McCulloch, W. (1960). Anatomy and physiology of vision in the frog (Rana pipiens). *J. Gen. Physiol.*, **43**, 129–75.

Møller, A. P. (1995). Bumblebee preference for symmetrical flowers. *Proc. Nat. Acad. Sci. USA*, **92**, 2288–92.

Rind, F. C. (1997). Collision avoidance: from the locust eye to a seeing machine. In *From Living Eyes to Seeing Machines*, eds. M. V. Srinivasan & S. Venkatesh. Oxford: University Press, pp. 105–25.

Snyder, A. W., Stavenga, D. G., & Laughlin, S. B. (1977). Spatial information capacity of compound eyes. *J. Comp. Physiol.*, **116**, 183–207.

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**, 331–50.

Srinivasan, M. V., Zhang, S. W., & Chandrashekara, K. (1993). Evidence for two distinct movement detecting systems in insect vision. *Naturwiss.*, **80**, 38–41.

Srinivasan, M. V., Zhang, S. W., & Witney, K. (1994). Visual discrimination of pattern orientation by honeybees. *Phil. Trans. R. Soc. Lond. B*, **343**, 199–210.

Sutherland, N. S. (1960). The visual discrimination of shape by Octopus. Squares and rectangles. *J. Comp. Physiol. Psychol.*, **53**, 95–103.

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

von Frisch, K. (1914). Der Farbensinn und Formensinn der Biene. *Zool. Jahrb., Abt. allg. Physiol.*, **35**, 1–182.

von Uexküll, J. J. (1909). *Umwelt und Innenwelt der Tiere*. Springer: Berlin.

Vowles, D. M. (1964). Models and the insect brain. In *Neural Theory and Modeling*. ed. R.F. Reiss. Stanford, CA: University Press, pp. 377–99.

Wehner, R. (1969). Die Mechanismus der optischen Winkelmessung bei der Biene (Apis mellifera). *Zool. Anz., Suppl.*, **33**, 586–92.

Wehner, R. (1972). Spontaneous pattern preferences of Drosophila melanogaster to black areas in various parts of the visual field. *J. Insect Physiol.*, **18**, 1579–91.

Yang, E.-C. & Maddess, T. (1997). Orientation-sensitive neurons in the brain of the honey bee (Apis mellifera). *J. Insect Physiol.*, **43**, 329–36.

Young, J. Z. (1961). Learning and discrimination in the octopus. *Biol. Rev*. **36**, 32–96.

Zerrahn, G. (1933). Formdressur und Formunterscheidung bei der Honigbiene. *Z. vergl. Physiol.*, **20**, 117–50.