

# 09

## FEATURE DETECTORS AND CUES<sup>1</sup>

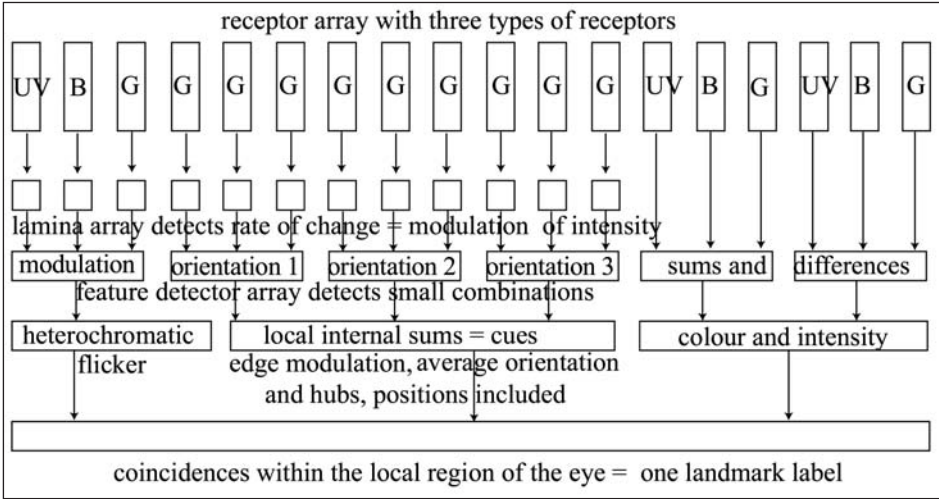
---

This chapter traces the effort from 1990 to 2008 to identify and characterise the parallel pathways of feature detectors and cues at the heart of the mechanism of visual processing in the bee. Bees have a few different kinds of feature detectors in large arrays that respond to parts of parameters in the pattern. The features in the parameters are edges or areas of black or colour; that is all. The analysis has been done with patterns subtending 30–45° at the eye, so the responses are limited to a small part of each eye. The responses of each kind of array of feature detectors are summed into a cue that lies within the bee and the cue can be learned together with its position on the eye. The bee tends to look at the reward hole at the centre of the target and the summation is done separately in each eye, so each eye picks up one set of cues from its own side of the pattern.

The common cues are area, modulation (total edge length), position of the centre of an area, radial edges, average edge orientation, tangential edges and the absence of a cue. The coincidence of the different cues in each local region of the eye can be remembered as the label on a landmark, whether or not an isolated landmark lies in that direction. Visual recognition of the place of the reward is nothing more than the coincidences between a few landmark labels at large angles to each other and the corresponding positions of simple parameters in the panorama.

In the most recent period, since 1990, the cues have been identified and characterised in the bee's visual system. Each cue has its own story. Perception of the configurational layout in patterns of this size (30–45°) was ruled out on logical grounds because the responses of edge detectors were summed on each side of the target so that the orientations of separate edges were lost, and it was ruled out experimentally because there were many quite different patterns for which bees showed equal preference, despite being trained to go to one of them.

Figure 9.1 A summary of the visual processing system for some of the cues in a local region of the eye. Most of the input is from green-sensitive receptors. The lamina detects the rate of change of intensity for the feature detectors for contrasts. The three types of orientation detectors are summed together. Cues relating to bilateral symmetry and the position of the centre are not illustrated here.



Box 9.1 Glossary of terms

The parameters are outside the eye as part of the pattern or panorama. The image is the distribution of excitation on the retina. The feature detectors behind the eye respond to the parameters in the image. A cue is the sum or count of the responses of one kind of feature detector in a local region of the eye, and is therefore a quantity inside the bee. The cue is derived from a part of the image in the local region, but the process of summation destroys the local layout. If rewarded, the bees learn the cues in their retinotopic positions. A landmark is recognised as the coincidence of several different cues in a local region of the eye.

The feature detectors are the units of perception of modulation, edge orientation, black, white or colour. They are small, about 3 ommatidia across on the retina, and all respond independently in parallel. The responses of the feature detectors are summed to form cues and the bee remembers the totals and their positions, not the individual detector responses.

The field of a filter or neuron is the region in space and time within which a signal is detected.

A fixed pattern, as opposed to a shuffled one, has the pattern fixed as seen from the choice point of the bee.

A generalised parameter is one that is recognised in a context other than in the training pattern. Originally, it was merely in a different position on the target, but later it was in a different pattern.

The image is the pattern of excitation in the array of receptors in the retina.

The label is the coincidence of cues in a local region of the eye, by which the bee recognises a landmark and its position.

The modulation of a receptor is the change in the light intensity in the receptor and the consequent electrical signal. The motion of the eye over contrasts generates the modulation of the receptors. The modulation of a pattern is roughly equal to the total length of edges in it.

Orientation of an edge is usually the angle to the vertical in a vertical plane.

Bees can be trained to remember the retinotopic position of some cues within the local region of the eye.

A parameter is a scalar or vector measurement of some aspect of the pattern outside the eye—that is, the area, total length of edge or averaged edge orientation.

The patterns are displayed on the targets during training and tests.








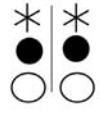




Place for bees is a geocentric term, like the place on a map; position and direction are usually retinotopic terms for the direction relative to the axes of the head. Location or position refers also to the position of a parameter on the target, a shift in position of a pattern or a shuffle of the locations of boxes, targets or bars during training and tests.

Point of choice is the place where the bee detects a cue and makes a choice by moving away from or towards the reward or the next target.

A sign stimulus is an older and more general term that is not restricted to vision—for example, it also applies to the call of a bird. It is the human idea of the essential stimulus outside the animal, not the parameter that is eventually identified, and certainly not the cue formed by the feature detector responses within the animal.

A template is a hypothetical mechanism that detects a fairly complicated pattern that has been identified by the human observer; it can be innate or learned. In vision, a spatial copy is usually implied. Templates are useful in pre-programmed robot vision.


Figure 9.2 Representations of the cues that are similar to common parameters, in order of preference during the learning process. The cues in fact consist of excitation in groups of neurons, but this illustration might assist the reader to understand the text. The bees do not see the parameters; they detect edges and areas with feature detectors and the cues are the various separate sums of the responses.

feature	weak	strong
modulation		
area		
radial		
bilateral symmetry		
orientation		
tangential		

## A new apparatus for measurement of resolution

To measure the resolution of bee vision required a new apparatus in which flying bees chose between two targets at a known distance (Srinivasan and Lehrer 1988). In earlier experiments, the bees made their choice with the external panorama around them and, except in Wehner's experiments (Figures 1.1c and 4.4), they learned the cues at an unknown range while they prepared to land on the reward hole. The Y-choice apparatus (Figure 9.5) can be used for a great variety of experiments in carefully controlled conditions. The bees detect the patterns on the targets at a fixed range and fixed angular size, so calculations of resolution are possible. By chance, in the apparatus constructed in Canberra, the angular sub-tense of the target from the point of choice was 35–50°, and later this turned out to correspond with one local region at the front of each eye. The bees therefore learned two landmark labels that were usually identical if the pattern was the same on each side of the reward hole.

Figure 9.3 The Y-choice apparatus. The bees enter through the hole 5cm in diameter at the front and pass through one of the transparent baffles. The targets and their patterns with the reward change sides every five minutes, to prevent the bees from learning which side to go. The air pipe extracts odours. As in all the figures, (+) and (–) indicate the rewarded and unrewarded patterns.



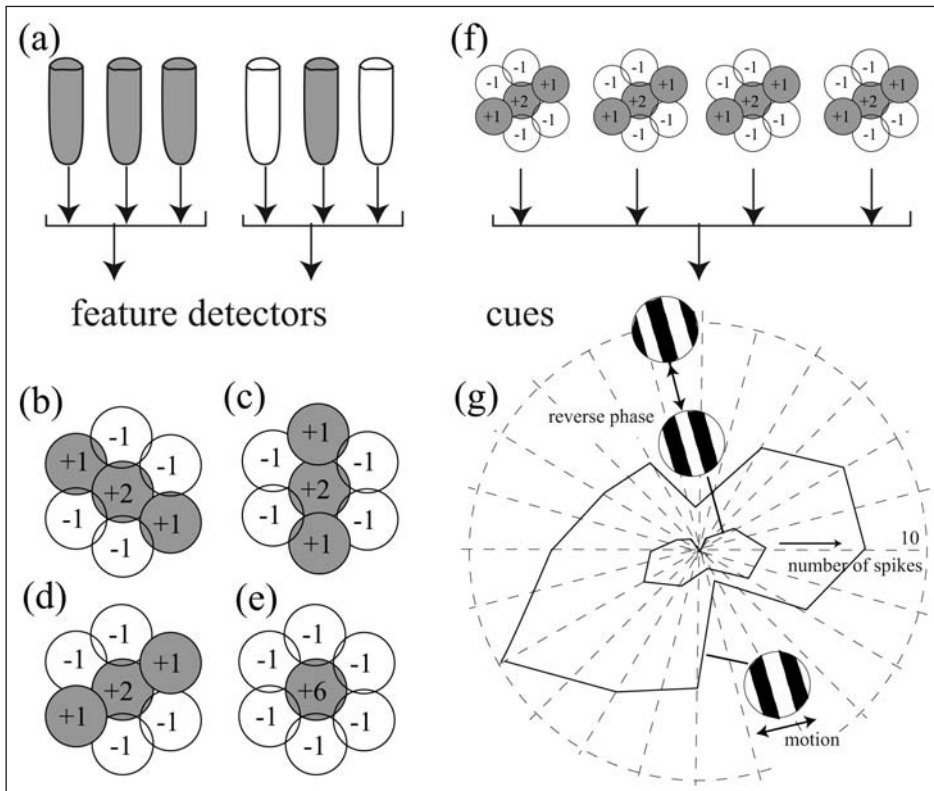
At first it was not realised that this new apparatus restricted the visual angle of the bees like blinkers, so that they were restricted to one or two local regions of the eye. The bees could no longer learn several landmarks and, unlike bees trained on large targets, they became tolerant of test patterns that differed from the training pattern. This change in their behaviour made possible the identification and characterisation of the cues.

The apparatus was placed under a shelter so the bees could not use the sun compass when inside and was lined with clean white paper, which was originally intended to help the bees find the patterns, but left plenty of vertical and horizontal edges that the bees could use to stabilise themselves in flight and orient themselves with reference to the patterns. The bees must look for the expected positions of the parameters by use of a frame of reference within the box—usually the internal edges and the reward hole at the centre of the target. I did not realise the significance of these changes in the task until the position of the hub in radial patterns was recognised as a parameter—about 1998.

After 1995, a transparent baffle was placed across the entrance to each arm to slow the bees and define the range at which they made their choice. No adverse effects of the baffle were noticed, but it gave the bees more time to look and favoured the use of the front of the eye. The bees came more easily into the apparatus if the entrance had the bees' attractant odour. After 1995, a stream of air was drawn out at each side so the bees made their choice in clean air without interference from new odours. During training, the two sides were interchanged every five minutes (10 minutes before 1995) so the bees could not learn which side to go and must look at the targets, which could both be seen from the choice chamber. The bees learn the geometry of the inside of the apparatus while in flight. Usually two hours of training is sufficient and it is important to consider why training takes so long.

To investigate what the bees had learned, they were given a variety of different tests with unfamiliar patterns on the targets. Several different tests were intercalated so that they saw a given test only once or twice a day. When they arrive, they look at one test pattern and then at the other if the first is not recognised. By watching them in the choice chamber, one can see whether they decide quickly or whether they spend a long time looking. In each experiment, the aim is to see whether the bees can do the task or not after a reasonable period of training. The test patterns changed sides after five minutes, which allowed only one visit on each side, then training resumed for 20 minutes before a different test was displayed. It is preferable to use many tests in a sequence so that the bees cannot learn any one of them. Usually a small group of individually marked bees was trained on a Monday morning and each experiment lasted all week.

Figure 9.4 Orientation detection at two steps in visual processing. a) A circuit that detects the coincidences between receptor responses. b–d) Feature detectors for edge orientation. They are symmetrical about one axis, with a field size of  $3^\circ$ , and are therefore 3 ommatidia long. e) A modulation detector of the same size (compare with Figures 4.1b and 4.1c). f) The circuit that detects a cue by the coincidences of feature detectors. g) Responses of a detector neuron in the lobula of the bee to moving edges at different angles. The stimulus was either motion at right angles to the edge or alternating phase of bars, as shown in the insets. The field size was about  $20^\circ$  across.



## The cues

### The modulation cue

Long ago, Hertz showed that bees distinguished between many patterns by something related to the total length of edge (Figures 1.2 and 1.3), a finding later confirmed by all who examined this point. The bees must have made a quantitative measure of something, but no-one asked over what area of the eye the measurements were made. It seemed to be assumed that the bee detected an object or shape, then measured its edge length. Until about 2003, when the feature detectors were measured, it was not clear what the bees really detected.

Resolution tests going back to Hecht and Wolf (1929), and also repeated several times, gave an absolute limit of the resolution of the modulation caused by a regular grating near  $2^\circ$ , irrespective of the orientation. The most likely modulation detector is therefore a single ommatidium with a symmetrical inhibitory surround (Figure 9.4e), which is smaller than that proposed by Jander (Figures 4.1c and 4.1d). The modulation cue would then be the total number of simultaneously excited modulation detectors in a local region of the eye, maybe in the field of a neuron. The difference in modulation between two patterns is a preferred and common cue that is measured quantitatively. In a test, the trained bees make the same measure irrespective of a change in the pattern (Hertz 1933; Horridge 1997a).

There is experimental evidence that modulation is detected in colour, but it is not clear whether there are heterochromatic detectors or separate blue and green detectors—or both. Bees discriminate between a horizontal and a similar vertical grating in colour with no green contrast; therefore, modulation is discriminated via the blue as well as via the green receptors. With a grating at  $45^\circ$  versus one at  $135^\circ$ , with no green contrast, bees fail to discriminate even if the period is large, because there is no modulation difference and the orientation detectors have inputs only via the green receptors.

## Discrimination of oriented edges

Before 1988, bees were trained either with very large patterns, with a single pattern versus a blank, or with landing on the patterns as the criterion of success. Therefore, the significance of edge orientation was not discovered because it was not a preferred cue when modulation or an area of black was available. In 1990, with the Y-choice apparatus (Figure 9.5) without the baffles, bees learned to discriminate between two gratings—one with horizontal bars versus a similar one with vertical bars when the stripes were continually shuffled in position and width during the training and tests to eliminate the position of black as a cue (van Hateren et al. 1990). Memory of eidetic images was therefore not tested, let alone refuted. The authors assumed that the orientation of edges was learned, although they already knew that in this situation the bees learned the difference in the modulation caused by scanning the eye across horizontal versus vertical edges (Srinivasan and Lehrer 1988). The part played by the orientation of the central axes of the bars, as suggested long before by Wehner (1967), was not considered. For 10 years, no-one worried about the missing controls because the randomisation technique suggested many new experiments. Later, the difference in modulation was eliminated by the use of oblique gratings at  $45^\circ$  versus  $135^\circ$  to the vertical, so orientation detection was substantiated.



At this point, let us look at the unregenerate beliefs of the time (Srinivasan 1994):

The 'vocabulary' of the visual system continues to be a mystery, and we are still largely in the dark as to what 'tokens' the visual system uses to represent, analyse and recognize patterns...While there is little doubt that bees use some kind of neural 'snapshot' to remember and recognize patterns and landmarks, it is hard to imagine that this is all there is to pattern recognition.

The concept of a snapshot came from the work of Collett, with several artificial landmarks at different positions relative to a reward out in the open. It was given a new meaning by the demonstration that orientation was recognised in brief 2ms flashes of illumination, so the relative motion of the eye and pattern was unnecessary (Srinivasan et al. 1993b).

Srinivasan then pointed to a number of publications from 1987–93 suggesting that the visual processing of the bee resembled that of the mammalian cortex. Unfortunately, none of the examples he quoted now substantiate this claim. Illusory contours have not been confirmed, despite efforts. Simultaneous recognition of two or more different orientations is explained by the large size of the targets that spread over adjacent local regions of the eye. Different cues are processed in different pathways that learn separately.

Srinivasan concluded that

recognition is probably mediated by at least two kinds of processes. One kind...involves the participation of long-range mechanisms which evaluate general properties of the object, such as colour, size and orientation. These mechanisms...rapidly exclude objects which do not possess the right attributes. Another kind of process involves mechanisms which operate at short ranges and require fixation. These mechanisms are more precise and work on the basis of a memorized template. (Srinivasan 1994)

Srinivasan had a concept of mechanisms in parallel, but both of his processes were guesses and both required a global perception of the whole image.

## Edge orientations are not the basis of pattern vision

In the work on discrimination of the rotation of a square cross in the 1960s, the patterns were huge, subtending  $130^\circ$  at the eye (Figure 4.2a). The response to rotation of a cross was not a sine-squared function of the angle. The smallest detectable rotation of the cross was only about  $4^\circ$  and the edges could be cut into steps with no effect, so that edge orientation was not involved. Instead, the bees learned the positions of separate areas of black in neighbouring regions of the eye (Wehner 1967, 1969).

In 1994, however, Srinivasan et al. discovered that bees were not able to learn to discriminate between a smaller square cross and the same cross rotated by  $45^\circ$  (Figures 4.2c and 4.2d). They suggested that the detectors of the orientation of edges had very large fields and an angular orientation sensitivity curve that was  $90^\circ$  wide at the 50 per cent level of sensitivity. The response of the detector to the rotation of an edge or thin bar would then be a sine-squared function that rises from zero to a saturation of 100 per cent as the angle changes from  $0^\circ$  to  $90^\circ$ . The response to the rotation of the orthogonal edge or bar would be a cos-squared function. With two equal bars in the form of a square cross, the total response would be constant so that rotation of the cross has no effect. Later, many other examples were found where equal lengths of edges at right angles on each side of the target mutually cancelled the orientation cue.


Irrespective of the theory, the data showed that different orientations in the same local region could not be detected simultaneously and therefore patterns could not be recognised by the combinations of orientations of their edges. The mutual cancellation of orientations in fact destroys pattern, including texture, but not the modulation detected by the orientation detectors. Srinivasan's mechanism was not, however, so certain. The large fields were not demonstrated and alternatives were possible. Later, we found that the fields of the orientation detectors were restricted to the eye on their own side and neurons with corresponding properties were found in the deep optic lobe of the bee (Figure 9.4g). We all guessed—wrongly as it turned out—that the responses of the edge detectors were strung together to make continuous lengths, as observed by humans. When the orientation detectors were measured (see below), they turned out to be small ( $3^\circ$  long), independent and not strung together to span gaps. The large fields were therefore summations of many small parallel orientation detectors.

It will be noticed that neither large nor small patterns support the idea that patterns with several edges at angles to each other are discriminated by the orientations of the edges, although this is the almost universally popular belief.

## Orientation of fixed bars

In 1998, after a long delay, Wehner's 1966–72 method of training on one pattern versus a blank target was repeated. It turned out to be the beginning of a new theme. The bees were trained in the Y-choice apparatus (Figure 9.5) with a single oblique black bar versus a plain white target. The bar was offset on the target to allow for a subsequent shift to a new place. Two results were startling.

Figure 9.5 Training on a single fixed vertical bar versus a blank white target; the bees learn only that there is something black in the expected place. a) The training targets. b) Reduced response with the bar moved down. c) Reduced response with the bar rotated by 90°. d) No discrimination with the training bar versus a similar bar with the edge cut in steps, so they did not look for edge orientation. e) No discrimination with the training bar versus a similar bar rotated by 90°. f) Good discrimination of the expected position of the bar.



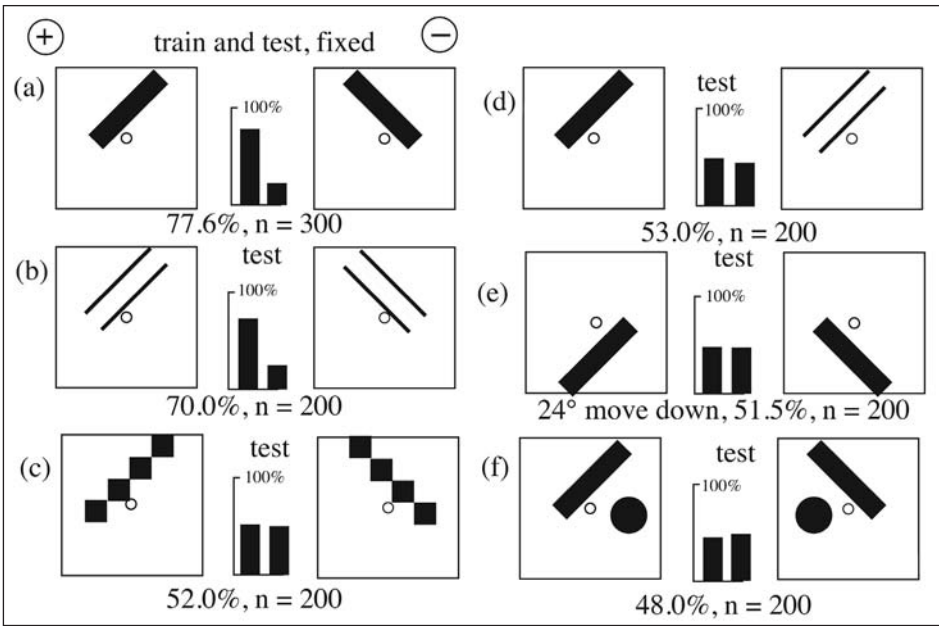
Source: Horridge (2003a).

First, when the training bar was simply moved to a new place in a test, the trained bees no longer recognised it. Second, when the trained bees were tested with the training bar versus the same bar rotated through 90°, they showed equal preference (Figure 9.5). In the training, they had learned the position of something black of a certain area and edge length, but the edge orientation and shape were not preferred parameters. The strongest cue within the bees was the position of the centre of black on the target, especially in the vertical direction. Although they could not remember the bar they had been trained on, they knew its position. This was powerful stuff that suggested new experiments that eventually showed that each cue was a measured quantity of a certain quality with its position on the eye. Of course, these are the properties of a neuron.

On the other hand, when they were trained to discriminate between a fixed black bar on one target and a similar bar at right angles centred on the same place on the other target, they responded only to orientation and its expected position (Figure 9.6). In tests, they were less able to discriminate the more the bars were displaced. When trained on a broad black bar in one position on one

target versus another bar in a different position on the other, the bees learned the difference between the vertical positions of the centres and ignored all other differences. Orientation was a cue only if there was no other available.

Figure 9.6 Learning the orientation cue when nothing else is available. a) Training with an oblique bar versus an orthogonal bar, centred on a place that is neither radial nor tangential. b) They discriminate two thin oblique bars versus two thin orthogonal bars. c) They fail in tests with bars with stepped edges to remove the edge orientations. d) They fail to distinguish the training bar from two thin oblique bars with the same orientation. e) They fail in tests with both bars moved down. f) They fail when a black spot is added to the training targets.



Source: Horridge (2003a).

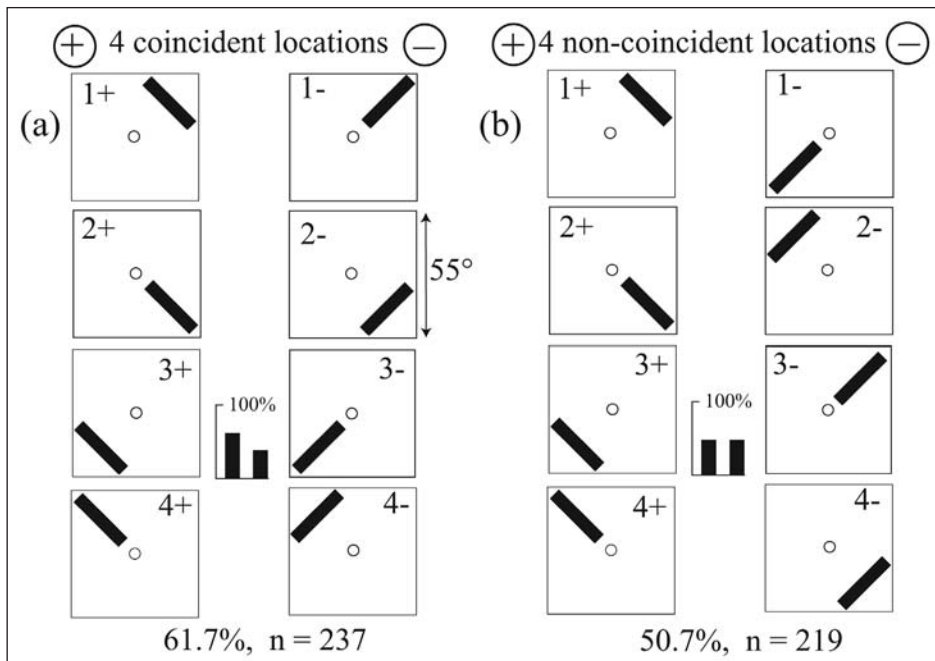
The bees used the scene from the choice point in the apparatus to fix a frame of reference for the expected direction of the cue. In a test, they lost a cue that was not in its expected place. They therefore did not fixate on the black bars or the cues, which had no salience for them. They were not interested in the pattern, only in the cues in their expected positions as a way to identify the place of the reward. Therefore, when cues were shuffled in position during training, they were learned in the range of places where they occurred during the training.

Making a fixed edge fuzzy, even extremely fuzzy, has no effect on the orientation cue. A gradient from black to white in  $20^\circ$  is still detected as an edge with an orientation.

When trained on a single black bar versus another, both shuffled in position, the bees learn the orientation cue only if the bars on the two targets are in

corresponding positions at all times (Figure 9.7). Each bar is then detected in the same expected positions on the eye. The local region of the eye in which the memory is formed and recovered is about  $15\text{--}20^\circ$  across.

Figure 9.7 The bars on the two targets must be in corresponding places for the orientation to be recognised. Every five minutes, the bar on both targets was moved to a new position, 1, 2, 3, 4, so that the bees learn not to use the location of the bar or the radial or tangential cues. a) With the bars in corresponding positions on the two sides, they learn. b) With the bars in non-corresponding positions, they fail to learn.



Source: Horridge (1998a).

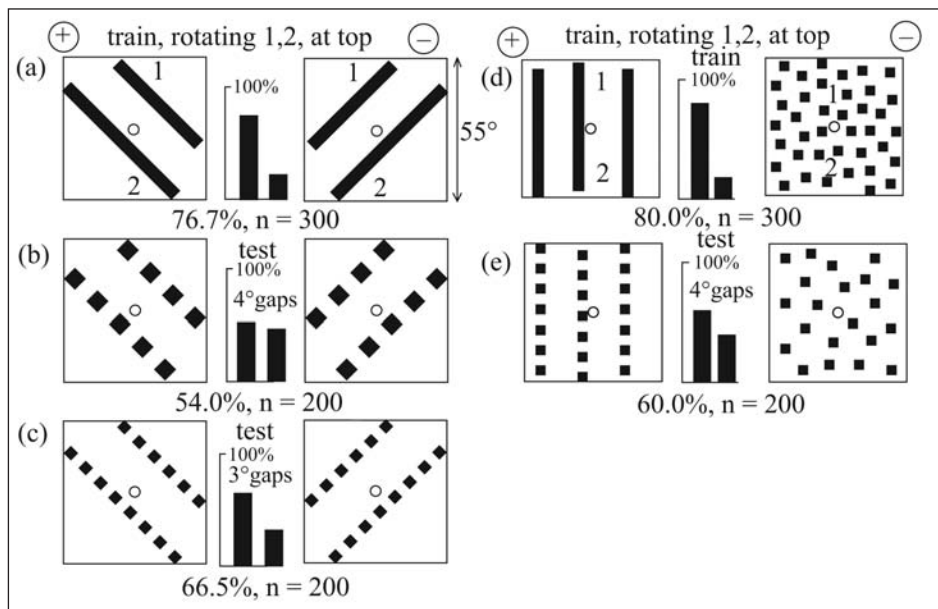
## Orientation of gratings

To a bee, a grating is a place where cues are found, not a collection of bars. When the task is to discriminate between two black and white gratings at right angles, the bees learn rapidly to a high standard above 80 per cent correct, and they are about 65 per cent correct when tested with single bars placed anywhere on the targets. The bees learned the average position of black, the difference in modulation caused by scanning and the difference in edge orientation. There was no evidence that they remembered or even detected the positions of the individual bars. When the criterion of success was landing on the reward hole, the main cue was black or white immediately below the reward hole (Baumgärtner 1928; Giger and Srinivasan 1995). On the other hand, when bees were trained on single bars and tested on gratings, the result depended on what

cues they had learned. They failed when they had learned only the position of the training bar, but were successful if they had learned an orientation difference.

When the gratings have a period of less than about  $10^\circ$ , it is quite unnecessary to shuffle the positions of the bars. Randomising the width of the bars during the training, however, confuses the modulation cues generated by the horizontal movements of the eye. This is an important consideration because modulation is the preferred cue and the bees learn it quantitatively.

Figure 9.8 Measurement of the maximum length of the feature detectors for orientation. a) Training on orthogonal bars that are shuffled in position. b and c) Testing with rows of squares with gaps of controlled width between the squares; the orientation is discriminated when the feature detectors can bridge the gaps. d) Training with shuffled black bars versus squares of the same total area. e) Testing with rows of squares with gaps of controlled width; in each case, the limit was near  $3.5^\circ$ .



Source: Horridge (2003b).

Since 1967, it had been uncertain how fixed gratings were discriminated (Chapter 4). Bees trained with a horizontal or a vertical black-and-white fixed grating versus a grey target of 50 per cent black discriminated down to periods of about  $2^\circ$  irrespective of the direction of the edges, from which it was inferred that the bees learned the modulation difference by scanning, not the spacing of the bars (Srinivasan and Lehrer 1988). Although at the time the colours were used to remove motion signals, they also found that bees discriminated the horizontal/vertical orientation in the absence of contrast to the green

receptors, because the cue was the difference in modulation. Later, when the bees were trained with randomly shuffled bars or gratings, it was found that the orientation cue required green contrast (Giger and Srinivasan 1996). Bees cannot discriminate equal but orthogonal oblique gratings (at  $45^\circ$  and  $135^\circ$ ) with no green contrast because there is no modulation difference and no orientation cue. Black oblique orthogonal gratings at  $45^\circ$  and  $135^\circ$  present no difference in modulation induced by active vision and the resolution of them is now that of the orientation cue—about  $3^\circ$ . So, the preferred parameter changed as the gratings were rotated. All these data led me to devise ways to measure the size of the edge orientation detectors.

### The size of the feature detectors for edge orientation

When small squares in a straight row are so close that they are not resolved separately, they are detected as a bar with an orientation, but when the squares are resolved, the orientation is cancelled by the equal lengths of edges at right angles. Therefore, when bees are trained to a given orientation, they can be tested with rows of squares at different separations to find the limit of resolution of the orientation cue. The result is the maximum size of the orientation feature detectors that can span the gap between the squares—about  $3^\circ$ . Cutting long straight edges into square steps that are resolved also destroys the orientation cue with a similar result (Figures 4.2b and 9.6c). A staircase that is resolved has no net orientation as it has in human vision.

In a different method, bees trained to discriminate between vertical bars and a pattern of squares (Figure 9.8d) were tested with rows of squares of controlled separation (Figure 9.8e). When the squares were resolved separately, the orientation of the row was not detected, showing again that the maximum length of the feature detector that spanned the gap between squares was about  $3^\circ$ . It was surprisingly small, and the edge detectors acted independently; they would not join up to span gaps.

Bees can be trained to discriminate between two equal arrays of oblique orthogonal bars with no modulation difference (Figure 9.9a) and then tested on arrays of shorter bars of similar total length. The minimum length for orientation detection at the threshold is about  $3^\circ$  (Figure 9.9b). In another method, bees were trained to discriminate orientation with shuffled orthogonal long oblique bars (Figure 9.9c) and then tested with the long bars versus a pattern of short bars of the same total length and parallel to them. The bees have learned only the orientation cue and when they detect it on both targets, they cannot discriminate (Figure 9.9d). The threshold is not reached until the short bars are the same length as the feature detectors—about  $3^\circ$ . Bees trained on Figure 9.9d failed to discriminate between orthogonal oblique bars (Figure 9.9e), showing

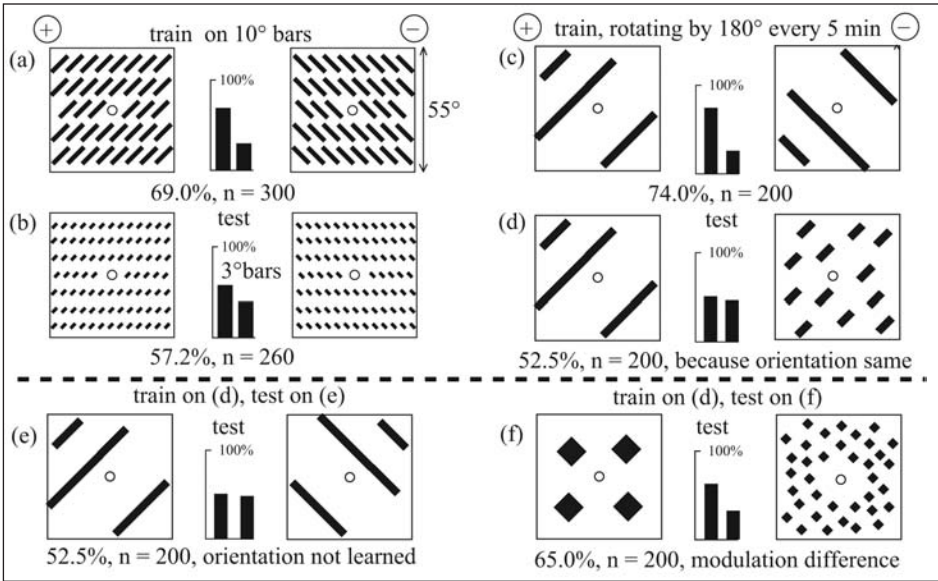


that they had not learned the pattern or the orientation that was the same on both targets. They did, however, learn the modulation difference in Figure 9.9d, as shown by a test (Figure 9.9f).

The discrimination of the orientation cue was little affected when black was exchanged for white (Figures 4.6c and 4.6g), showing that the detectors of edge orientation were bilaterally symmetrical. From these results, and assuming that the detectors of edge orientation depend on simultaneous modulation of a few adjacent receptors, we can infer that the feature detectors for orientation are three ommatidia long (Figures 9.4b–e). This result implies that there are only three types with axes at 120° to each other (Figures 9.4b–e).

Most significantly, the feature detectors were about an order of magnitude smaller than the cues and each cue was the sum of the feature detector responses in a local region of the eye, with their average position. They remind me of the small-field and medium-field neurons of the insect optic medulla.

Figure 9.9 Measurement of the minimum length of the feature detectors for orientation. a) Train on orthogonal bars. b) Test on shorter bars. c) Train on large shuffled orthogonal bars. d) Test on the same large bars versus smaller bars with the same orientation; discrimination fails when the orientation of the small bars is detected. e) The bees trained on the patterns in (d) do not recognise the orientation of orthogonal bars in a test because it was not a cue in the training. f) Bees trained on the patterns in (d) discriminated the modulation difference.



Source: Horridge (2003f).



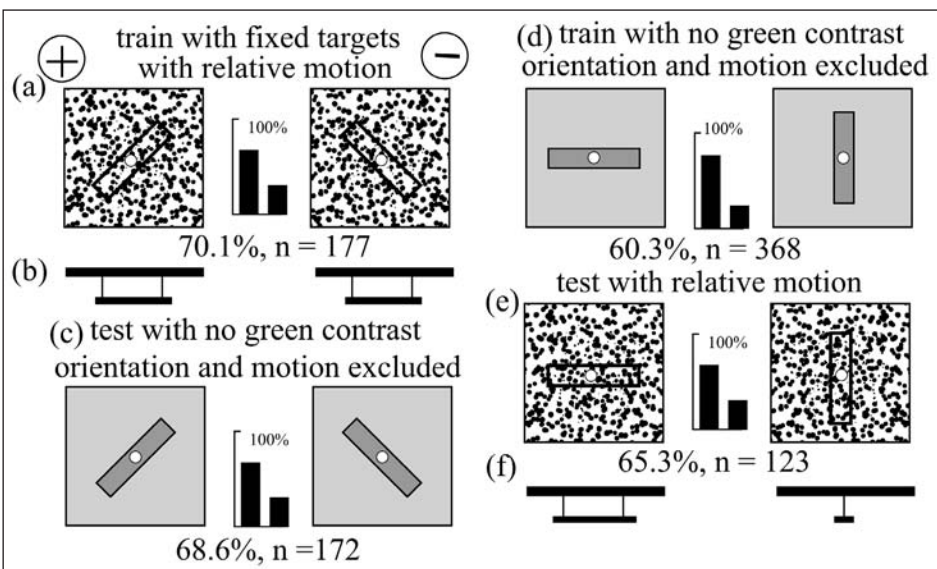
## Distinction from motion detectors

Bees discriminate between two orthogonal moving gratings that move at speeds up to 50°/s as seen from the point of choice in a Y-choice apparatus and they can do this task when illuminated by a slow stroboscope with flashes 2ms in duration (Giger 1996). The response is independent of the direction of motion of the moving grating. The edge orientation detectors are therefore quite different from the directional motion detectors described for many insects. The relation, if any, between the edge orientation detectors and the non-directional system that measures the angular velocity of the flow field has yet to be investigated.

## Misunderstandings with orientation detectors

After the experiments with gratings with randomised bar positions in 1990, it was assumed that bees learned the orientation cue. This conclusion was valid, however, only when no other cue was available, because orientation was the least preferred cue. The problems illustrated in Figures 4.2–7 were long forgotten, probably because in the second half of the twentieth century one simply did not refer to previous authors if they had published contrary conclusions.

Figure 9.10 Conclusions that now appear unjustified. a) Bees were trained to discriminate between two black-textured bars raised 6cm above black-textured backgrounds, as shown in side view in (b). c) The trained bees then discriminated the bars with no green contrast. d) Bees were trained with no green contrast, but the bars differed in modulation. e) The bees trained in (d) detected a difference in the textured bars raised over textured backgrounds.



Source: After Zhang et al. (1995).

In Canberra after 1990, we looked at the detection of bar orientation at various distances from the target to compare global and local discrimination (Zhang et al. 1992), the size of the regions in which orientation could be discriminated (Zhang and Horridge 1992) and the detection of orientation in illusory edges (Horridge et al. 1992). The last cannot be repeated and inferences of global vision based on all three of these studies were premature (see Chapter 12).

The spill-over of ideas from motion perception also misled us into thinking that edge orientation could be discriminated by the relative motion of a patterned edge against a patterned background. After a preliminary training on other patterns, Zhang et al. (1995) trained bees to discriminate between a fixed oblique textured black bar at  $45^\circ$  raised 6cm above a textured background versus a similar oblique bar at  $135^\circ$  (Figure 9.10a). The trained bees could immediately discriminate between two orthogonal oblique bars in plain colour with no green contrast to eliminate cues from motion detection (Figure 9.10b). Without further tests, it was proposed that the cue was orientation, but that could not be so. At the time, it was known that a texture of square pixels, if resolved, would destroy the orientation cue. Even worse, it was later shown that lack of green contrast also destroyed the orientation cue. Later, this and similar experiments could not be repeated despite considerable efforts, but the bees were sensitive to shadows under the raised bar (Horridge 2003a), so they had probably learned the difference in position of shadows.

Zhang et al. (1995) also trained without green contrast to eliminate motion signals, but at the time they were unaware that the orientation cue was also excluded, while the modulation difference remained. They intuitively inferred that the discrimination of horizontal versus vertical fixed bars that were textured or without green contrast (Figures 9.10d and 9.10e) was due to the orientation difference, which was unlikely, but the results were easily explained by modulation differences.

As well as the mistaken use of no green contrast to eliminate motion cues and because the real cues were not identified, most of these 1992–95 experiments required re-examination. The patterns were fixed and the bees could have learned the cues of position, modulation, area and edge length, and maybe orientation. The textures used for camouflage, with pixels 4mm square, were probably not resolved at the 27cm range. Unfortunately also, the scores were too high because the test patterns were presented for 10 minutes on each side, during which some bees made two visits and could learn which side to go. This mattered only for marginal scores. Also, it was not realised that the orientation on the left side of the target was discriminated separately from that on the right side, so some of the test patterns were inappropriate. All these errors of the day were uncritically accepted at the time and they still confuse the literature because they are quoted as support for various ideas about cognition in bees (see Chapter 12).

When colour is added to the tasks of the bees, we are still not clear whether the colours of areas are detected only by tonic blue, green and UV detectors of photon flux or whether they are detected by phasic modulation detectors of ultraviolet, blue, green or chromatic contrast at the edges—or by both. Different classes of neurons adapted to either type of input occur in the optic lobe (Chapter 6). The resolution of small areas of colour is related to intensity times the area, but some authors relate the detection of colour to the contrast at the edges. The experiments require accurate calibrations and a variety of tests in colour.

## Symmetry cues

### Possible adaptation of flowers to bees

Flowering plants evolved long after the insect visual system, so the evolution of flower colours and shapes was presumably influenced by insect vision. The colours of flowers and the colour vision of many pollinating insects are adapted to each other, but plant communities are rarely stable for long enough for an equilibrium to be reached. Free (1970) found that bees preferred symmetrical radial patterns, then bilateral symmetry and then irregular patterns, and also that bees landed at the edge of a plain target, but on a spot at the centre of a circle. Bumblebees prefer to land on flowers that are symmetrical. Hertz, Anderson and Free all showed that bees more easily learned the radially symmetrical patterns that they spontaneously preferred.

In his earliest experiments, von Frisch found that flower-like radial or concentric patterns of the same size were easily distinguished, but that triangles, squares, discs and ellipses were not (Figures 1.2 and 1.4), and a chequered pattern of squares was not distinguished from one of triangles. Later, Hertz found three classes of patterns—stars, circles and irregular blobs (Figure 1.4)—that were discriminated from each other irrespective of the length of edge, location or orientation. For 100 years, the outstanding problem was how the lopsided visual abilities of bees were adapted to their foraging needs, and mechanisms took a back seat.

### Preferences for symmetrical patterns


All the early workers made use of symmetrical shapes with radial edges for training bees. On a horizontal surface, they could be approached from any direction. They showed salience—that is, the bees found them easily on a flat, featureless white table and would learn them readily because they could fixate on them. What the bees in fact learned was a different matter. It is now clear that when the training patterns are all equally symmetrical, the bees will not learn to discriminate the symmetry, because the patterns all show it.

Hertz found that bees avoided circles (Figure 1.5). It was also noticed that when flying bees landed on bilaterally symmetrical flowers, they lined up with the direction of the axis (Jones and Buchmann 1974). Much later, it was shown that they measured the flowers for degrees of symmetry. For example, Møller (1995) found that bumblebees preferred to forage on more rather than less symmetrical flowers, and the former yielded more pollen and nectar.

In 1994, Miriam Lehrer decided to revise an old demonstration by Gertrud Zerrahn (1933) to show that bees had a preference for symmetry. We built an apparatus with 12 compartments (Figure 9.11a) and trained marked bees to come to neutral patterns. The bees entered the central arena, from which they could see into all 12 compartments. At intervals, the apparatus was rotated, so the bees could not learn the locations and had to look at the patterns. When the bees were familiar with the place, four new and different patterns (each reproduced three times) were spread around the 12 compartments with no reward, and the free choices of the bees were observed. The criterion of a choice was when a bee crossed the threshold of a compartment. With a plain black disc as the initial attractant pattern, we were able to show that bees preferred radial patterns to other patterns.

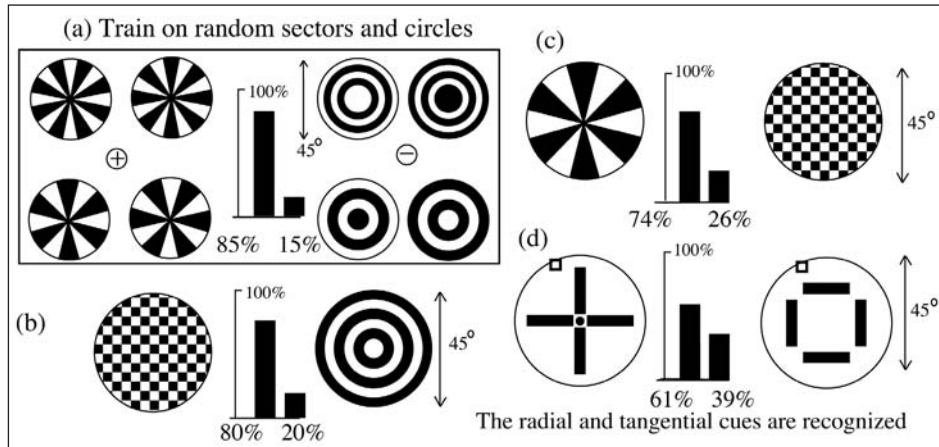
When we randomised the modulation by using one of six regularly changed 50:50 black-and-white checkerboards as the attractant pattern (Figure 9.11b), we were able to show that the bees preferred patterns of lower spatial frequency and radial flower-like patterns to random or regular textures, circular patterns and even over the checkerboards to which they were attracted in the first place. They preferred radial patterns of any sort and a vertical axis of bilateral symmetry, but avoided concentric circles. No preferences were found with other patterns, although many were tried. These results launched us into a search for mechanisms of discrimination of radial and tangential edges and symmetry.

Figure 9.11 Bees' preferences for unfamiliar patterns. a) The bees were trained to come to the apparatus with 12 partitions, which displayed various shuffled checkerboard patterns, (b) some of which were rewarded. They were then presented with several hundred choices between four patterns (three of each). Their choices were counted and reduced to percentages. c) Radial or tangential was preferred to random. d) Radial preferred to tangential. e) Bilaterally symmetrical preferred to asymmetrical.



Source: Lehrer et al. (1995).

Figure 9.12 Radial and tangential cues are recognised in unfamiliar patterns.  
 a) In the square, training on rotated and shuffled radial patterns versus tangential ones. b–d) Tests of the trained bees with circles, sectors or bars.



Source: Horridge and Zhang (1995).

## Discrimination of sectors and circles

Following this demonstration of preferences for different parameters, bees were trained with a pattern of black and white sectors or spokes, versus one with concentric circles or a spiral, with no average orientation in either target. To control against differences in areas of black, length of edge or location of black areas, the sectors in one target and the circles in the other were randomised by substituting a different target every 10 minutes, so that nothing remained constant except the kind of pattern and the position of its centre (Figure 9.12a). Bees so trained then discriminated pairs of unfamiliar patterns with radial versus tangential contours, such as a cross and a hollow square and also parts of circles or patterns of spokes. Rather than learning to recognise a circle, the bees preferred to learn to avoid the unrewarded target even if it was blank.

These results, and those of Hertz with patterns presented on a flat surface, led naturally to the proposal that bees had global filters for radial features and other global filters for concentric circles or tangents, and that these filters detected any part of their own pattern that coincided with their field of view. How easy it was to imagine global filters, but how wrong!

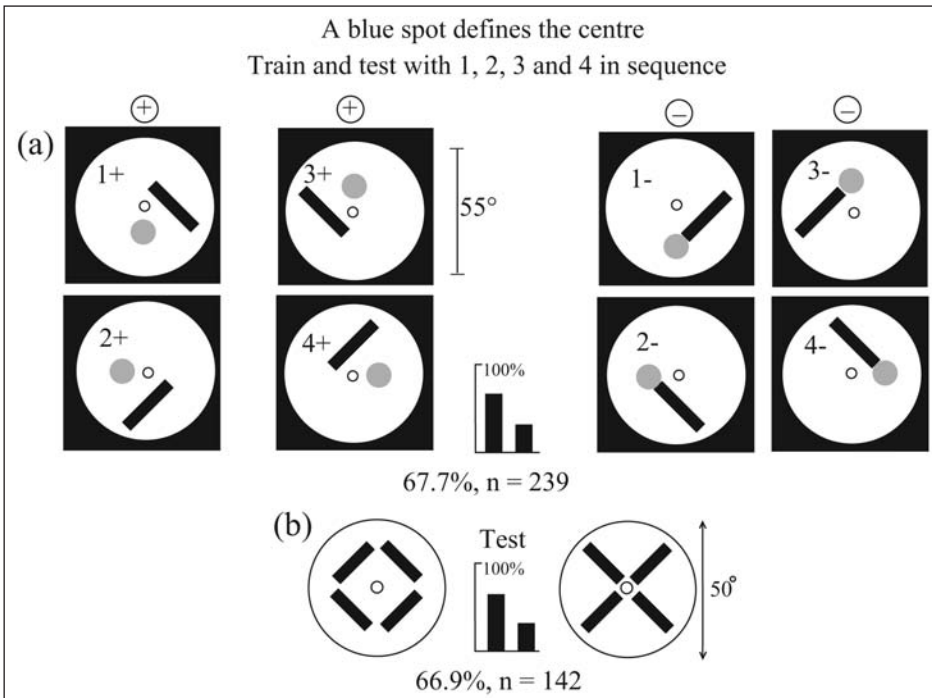
## Directing recognition with a coloured spot

The location of an added coloured spot can influence the bees to treat a bar as a radius or as a tangent, depending on the bar location and orientation relative to the spot. To demonstrate this, the positive target had a blue spot at the side of the bar; the negative target had a similar blue spot at the end of the bar, with



the spots in the same position on each target (Figure 9.13a). Both targets were rotated by  $90^\circ$  in the same direction every 10 minutes between positions 1, 2, 3 and 4, so that modulation and the locations and orientations of the bar and spot were useless as cues. The trained bees are then able to discriminate a pattern of tangents versus a pattern of radials (Figure 9.13b).


Figure 9.13 a) A bar is recognised as radial or tangential with reference to a blue spot although the whole pattern is shuffled in location during the training. b) The trained bees discriminate tangential from radial in quite different patterns.



Source: Horridge (1997c).

Whether the single black bar was accepted as a radius or as a tangent depended on the bees accepting the spot as the centre. The bar could be rotated about the centre and was still discriminated as a radius or a tangent. In other experiments, strong symmetry, a strong outline of the target, occurrence of several radial or tangential features or a coloured spot and the geometry of the situation all influenced whether the bee detected a single bar as a radius, tangent or orientation.

Figure 9.14 Examples of discriminations between patterns that displayed two pairs of orthogonal bars, in which the orientation cues usually cancelled out. a–d) Patterns that differed in radial/tangential cues. e) A difference in average orientation between the two sides. f–j) Patterns that look different to humans but display no differences in cues for the bees. Possible cues are indicated as follows: H = horizontal; O = oblique + and O-- = orthogonal oblique orientations; R = radial; T = tangential; V = vertical. To the bees, the patterns in (e) differ, but (j) when rotated they are similar.



Sources: (e) from Stach et al. (2004); others from Horridge (1996a).

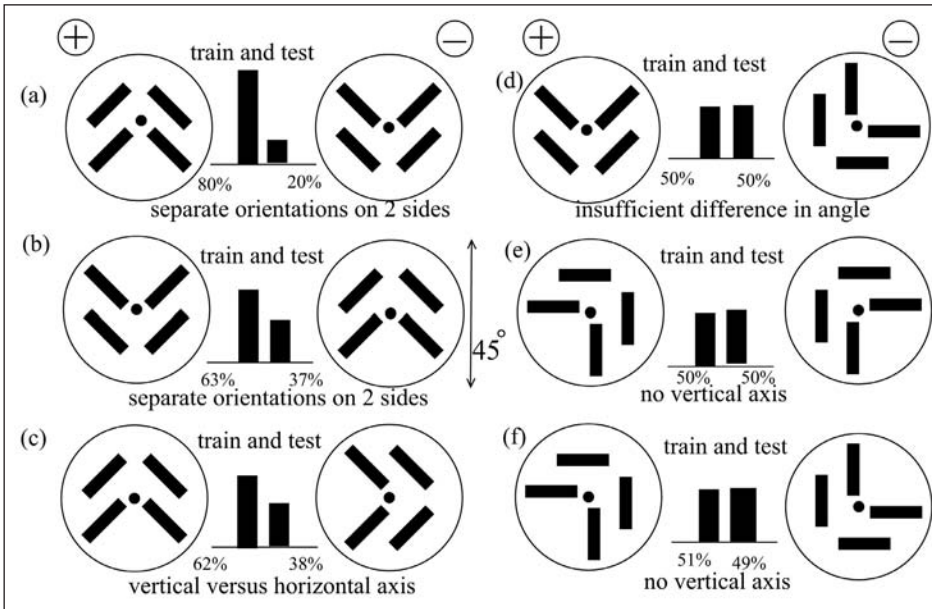
## Strategies for listing the cues

Successive efforts progressively defined the limited number of cues in the repertoire of the bees. First, in 1995, the radial, circular or spiral patterns were rotated at intervals during training to remove cues derived from orientations, leaving the radial and tangential cues intact. In 1998, these were found to be colourblind. In 1999, radial and tangential patterns with radial symmetry based



on three or six spokes were easily discriminated, but those based on four, five or seven spokes or sectors were not. This result is explained by the existence of three orientations of edge detectors.

Figure 9.15 Separate training experiments with various pairs of chevron patterns. a) and b) They detect the orientation cue on each side of the targets. c) The rewarded target is still discriminated when the other is rotated by 90°. d) A difference of 45° in the axis is not discriminated. e) and f) Discrimination requires one of the axes of symmetry to be vertical.



Source: Horridge (1996a).

In the second strategy, the arms of a square cross were rearranged to make many fixed patterns of two pairs of orthogonal bars, all the same size, area of black, length of edge and average position of black on the target (Figure 9.14). These patterns of four bars could not be discriminated from the same pattern rotated by 45° unless one of them had a vertical axis of bilateral symmetry (Figure 9.15). In that case, the bees could discriminate a rotation of the axis of bilateral symmetry by 90°, even if the test patterns were different from the training patterns.

The patterns with two pairs of orthogonal bars could be roughly divided into groups. The first group differed greatly in their content of radial, tangential or bilateral symmetry cues (Figures 9.14a–d) and the bees were easily trained to distinguish them from each other. Those in the second group were quite different from one another but were not distinguished (Figures 9.14f–j). It was inferred that they displayed similar cues. In tests, the trained bees accepted any of these patterns displaying the expected cues but no unexpected cue.

The bees measured the cues quantitatively. The real pattern was irrelevant. These patterns and experiments illustrate the value of Mill's rules of logic in the search for causation (Chapter 2).

Many pairs of patterns that differed from each other were not discriminated. For example, when the bees had learned the orientation cue, they failed in tests to distinguish the rewarded training pattern from other patterns with the same total length of edge and average orientation (Figure 9.9d). The bees looked for the orientation cue and found it equally in both patterns, irrespective of differences in layout. It was curious that the authors who still supported the eidetic image in the 1990s were blind to similar examples where patterns were different but not distinguished.

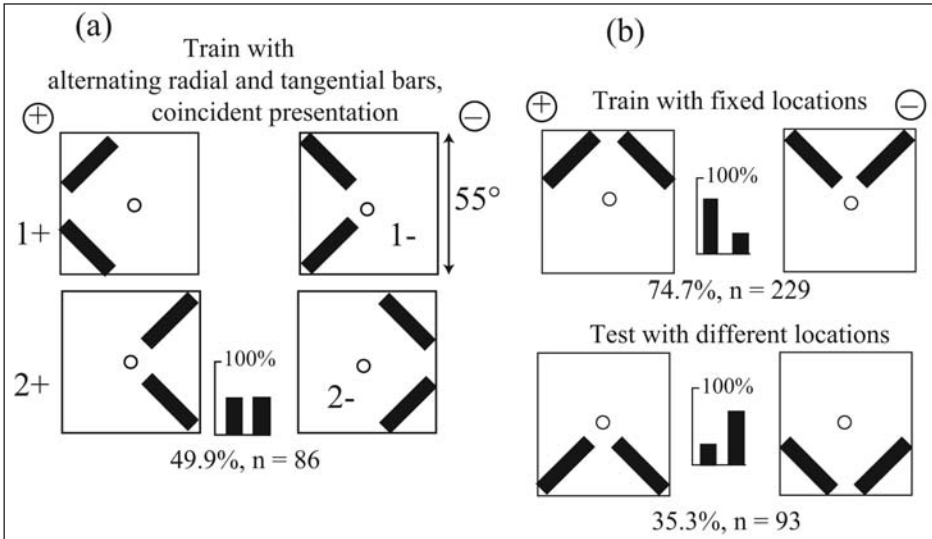
The third strategy turned to the recognition of position. A fixed pattern composed of two different colours was discriminated from the same with the colours reversed in position (von Frisch 1914; Gould 1986). Bees easily detected a vertical shift of the centre of an isolated area of black or colour relative to the reward hole as a cue of high priority. They also discriminated the exchange of two different colours in the left/right direction if there was green contrast or a radial pattern to stabilise the eye in the horizontal plane.

Finally, a chevron pattern with its axis of symmetry vertical was discriminated from itself rotated by  $180^\circ$  or by  $90^\circ$  (Figures 9.15 and 9.14i). The chevron pattern is curious in having no net orientation and its two radial bars cancel the effect of its two tangential bars. I concluded: 'The result with the chevron suggests that bees have a filter beyond those for circles or radial patterns, or for average orientation, and that it is related to bilateral symmetry, which is already known to have a broad biological significance for bumblebees' (Horridge 1996a). Of course, a single global filter was a bad idea, but at the time I could not model a bee filter that would detect the axis of bilateral symmetry irrespective of pattern. The widespread occurrence of symmetry in animals and plants, and the fast response to it, implies that there are many innate visual mechanisms for detecting symmetry.

### More experiments with two bars at right angles

Bees detected the orientation cue separately on the two sides of the target (Figure 9.14e) and the radial or tangential edges on either side, but failed to respond to the global pattern (Figures 9.14j and 9.16). First, bees were trained with two bars on each target, alternating between radial and tangential, in corresponding positions on the two targets. The bees could not learn to discriminate because the orientation, radial and tangential cues were cancelled. The bees could not detect the consistent global pattern of an arrowhead pointing to the left on one target and to the right on the other (Figure 9.16a).

Figure 9.16 Examples of training with two orthogonal bars on each target.  
 a) Shuffling between radial and tangential bars in corresponding positions on the targets; there was no remaining cue at all. b) Training with fixed patterns; the preferred cue is the rad/tan difference. The shift of the bars causes a reversal of choice because radial and tangential bars are interchanged.



Source: Horridge (1997b).

Next, fixed bars in corresponding positions formed an arrowhead pointing upwards on the rewarded target and downward on the other (Figure 9.16b). The bar orientation was the least preferred cue. When the patterns were moved down, the radial and tangential cues persisted but the orientation cues were lost. The bees reversed their choices because the rad/tan cues were reversed although the arrowheads were unchanged. From this, it was clear that bees did not detect a global pattern of even two bars.

## Analysis of cues in radial patterns

Bees easily spot the difference between radial and tangential edges, almost as though the visual system operates in radial coordinates. The real pattern was of no importance and, in tests, the trained bees accepted other patterns displaying the rad/tan cues on either side of the target (Horridge 1996a).

Similarly, bees also discriminate between a fixed ring of spots or sectors, with up to six spots in the ring and the same ring of spots rotated by half the angle between the spots (Horridge 2000c:Fig. 6). The cue was the position of a key black area, and unlike the situation with radial edges, the number of axes of symmetry was less important for discrimination than the size of the individual areas.

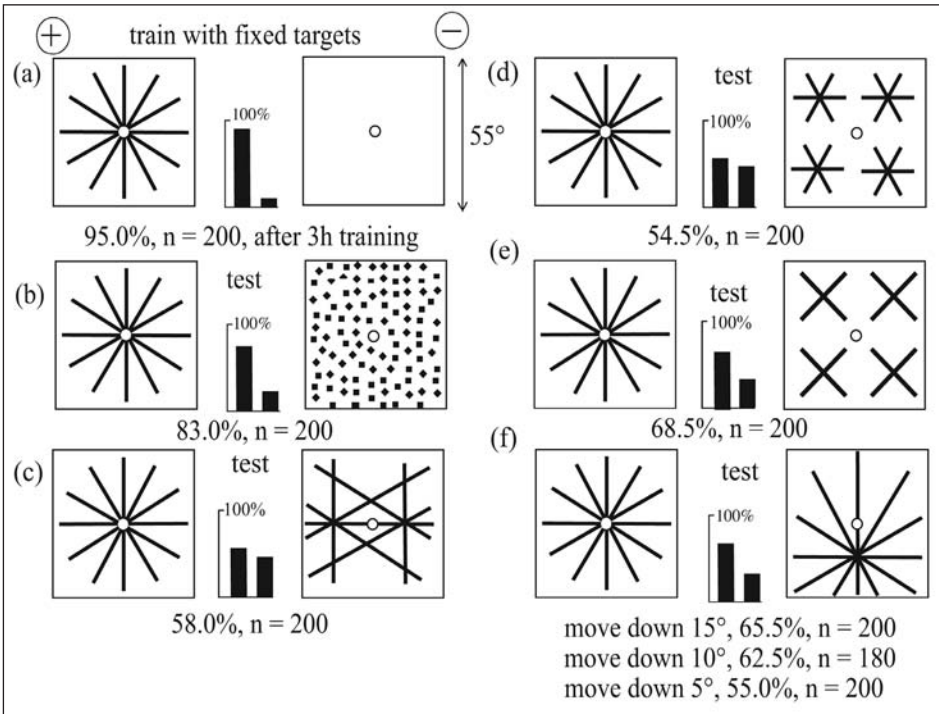
When targets are rotated at random during the learning process so that positions and edge orientations are shuffled, the cues must be presented as radial or

tangential edges, not as spots or areas of black. Radially symmetrical patterns of spokes have salience for bees, but they lose it when green contrast is removed, which shows the reliance on edges. In conclusion, there were two ways for a bee to detect the rotation of a radial pattern: by the orientations of edges of spokes and by the positions of areas with spots or sectors.

### Feature detectors for radial and tangential edges

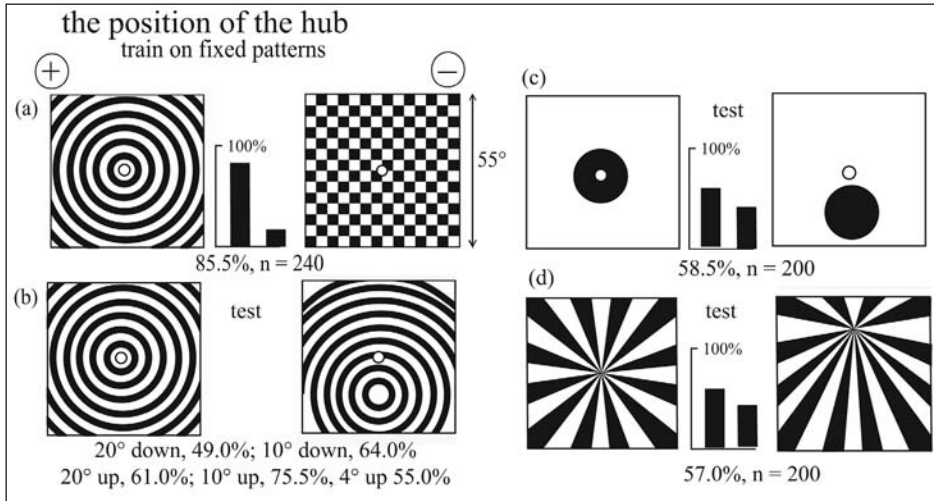
When analysed by the methods used for the orientation cue—by cutting the edges into short lengths or into square steps that were separately resolved—the feature detectors for radial and tangential edges were the same as those proposed for the orientation cue, 3° long, and therefore spanning three ommatidia in a row (Figure 9.5).

Figure 9.17 With radial spokes, bees detect the cues of ‘black’, ‘radial’ and ‘position of the hub’. a) Training pattern. b) With black on both sides, the score is reduced, so black contributes. c) With the bars rearranged, the bees detect little difference. d) They detect radial on both targets, but not much difference. e) Square crosses, or angles at 90°, make a difference. f) The trained bees notice a difference in the position of the hub down to 5°.



Source: Horridge (2006a).

Figure 9.18 Bees learn the position of the hub with concentric circles. a) Training pattern. b) The trained bees notice a difference in the position of the hub down to 5°. c) and d) With quite different patterns, they prefer the hub at the centre.



Source: Horridge (2006a).

## A new cue: the position of the centre of symmetry

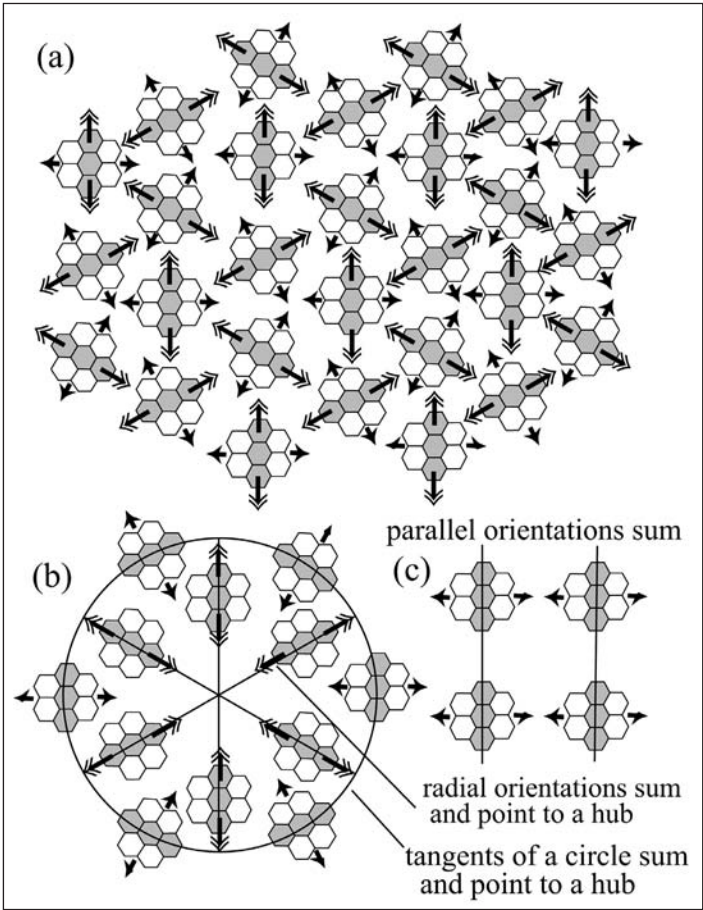
Bees could also learn the position of the centre of a ring or of concentric circles when trained versus a blank or a neutral pattern (Figure 9.18). They even detected the position of the centre of concentric curved lines, which implied that they detected the convex or concave curvature of the edges. There was no evidence for the idea that the bees detected the layout of whole rings or long curved edges, and much evidence against it.

When trained with a symmetrical pattern versus a blank target and then tested with two patterns at different heights (Figure 9.17a), the bees discriminated the expected position of the hub by as little as 5° (Figure 9.17f), in some cases with unfamiliar test patterns (Figure 9.18). A pattern of spokes or rings also stabilised the vision of the bees in the horizontal plane so that the position of a plain black area could then be learned (Figures 1.6b and 1.6c).

Bees discriminated half of a pattern of radial spokes or concentric circles from the other half, cut either vertically or horizontally, and irrespective of scale. This was the observation showing that radial and circular patterns were not detected by pre-formed combinations of orientation detectors or global filters like templates, because with a single output such filters could not distinguish the separate halves of the pattern. Instead, the bees detected edges as radial or circular by the coincidence of numerous local edge detectors converging to a hub from anywhere in the array, irrespective of the real pattern (Figure 9.19).

The binding that defined the cue as radial or tangential was therefore not hard-wired but depended on the coincidences of responses of similar edge detectors anywhere in the local region of the eye. This is a diffuse mechanism with no template and no memory of the layout of the pattern. It also explains the mutual cancellation of radial and tangential edges or orthogonal orientations. Because it depends on coincidences, such a system gives the impression of having taken a snapshot as it detected a hub. There was no global template that detected the positions, angles or numbers of spokes, a circle of a given size or a right angle as a whole. Instead, there was a distributed administration that would identify any incomplete or partially obscured symmetrical pattern and find the position of its hub. The mechanism is similar to that summing orientation.


Figure 9.19 The distributed mechanism with no fixed template for locating and identifying a hub as radial or circular. a) The three orientations of the axes in the array of feature detectors for edge orientation were inferred from the retinal arrangement of the ommatidia on the retina. b) The coincidences of the radial vectors and the tangential vectors at the position of the hub. c) Parallel orientations are summed.





In conclusion, the bees identified radial and circular patterns by the regional coincidences and convergence of local vectors extending from edge orientation detectors, and the position of the hub was also a cue. We saw that every possible use was made of the various ways that the coincidences of the positions and vectors of the edge detectors could be counted in a local region, or their absence noted, but there was no mechanism to reassemble the pattern.

Figure 9.20 Detection of a vertical axis of bilateral symmetry irrespective of pattern. a–c) The bees were trained on seven bilaterally symmetrical patterns simultaneously, taken successively in pairs for 10 minutes on each side in the choice maze. Only three of the patterns are shown here. The pattern with the vertical axis in each pair was not rewarded, so the training was against the preference. Training scores from each pattern were collected separately. d–f) The trained bees were tested on the same seven pairs of patterns rotated through 180°. The same three are shown. The tests were done in random order between periods of continued training, which improved the performance.



Source: After Horridge (1996c).

## Bilateral symmetry

Bumblebees prefer to settle on flowers that have a more perfect bilateral symmetry than their neighbours. They measure the perfection of the symmetry. Untrained honeybees spontaneously prefer a vertical axis of bilateral symmetry in arbitrary unfamiliar targets irrespective of pattern (Figure 9.11e). The special properties of bilateral symmetry are illustrated by the chevron pattern (Figure 9.15), which, as a whole, displays no average orientation, radial or tangential cue. Bees readily discriminate it from itself rotated by 90° when one of the patterns has a vertical axis of bilateral symmetry, but the resolution of angle is poor—the same as for orientation differences.

The detection of the axis of bilateral symmetry irrespective of pattern was demonstrated with patterns of two pairs of orthogonal bars in the following way. The bees were trained with seven different bilaterally symmetrical patterns of four bars taken in succession, to train the bees to ignore the real pattern (only three patterns are shown in Figure 9.20). The positive (rewarded) target was the pattern placed on its side and the negative targets all had a vertical axis of bilateral symmetry (Figures 9.20a–c). The training was therefore against the innate preference. The bees learn to discriminate the symmetry irrespective of pattern in a few hours of training. On successive days, with continued training, they were tested on the same seven patterns rotated through  $180^\circ$  (only three patterns are shown in Figures 9.20d–f). Although all the patterns in the tests were unfamiliar, the trained bees still picked out the asymmetrical one of each pair, and with an improved performance, because they had more training. These trained bees also discriminated the axis of bilateral symmetry in completely different unfamiliar patterns with different numbers of bars. It does not matter whether the various patterns can be discriminated from each other by other cues, because only the orientation of the axis of symmetry was learned in this experiment, all other cues being inconsistent or the same on both targets.

The mechanism is not such a puzzle as it is in humans, because there is no evidence that the bees really see the patterns. My expectation is that the bees do not recognise the abstract property of bilateral symmetry about an axis in general any more than they recognise pattern or shape in general. As with shape, they find a way to use their feature detectors for the task in hand, irrespective of the pattern. Some time ago, Jones and Buchmann (1974) found that bees landed on zygomorphic flowers in line with the axis of symmetry, even when tipped away from the vertical. Therefore, as the bee scans from side to side, she detects the same sequence from her feature detectors with either direction of the scan, whereas an asymmetrical flower sends back a signal that is different in the two directions. This would be sufficient to identify and measure the bilateral symmetry with a vertical axis.

As another mechanism with distributed administration, it is known that some cues are detected separately on the two sides, probably by the two eyes. If the bee compares the colour, area, height of the centre, radial or tangential, and perhaps other cues, by the two eyes, then a measure of bilateral symmetry can be made in a large number of patterns. In fact, any filter that has two spatially separated pass bands can detect some bilateral symmetry about an axis drawn between them and it is possible that symmetry is detected by several sets of coincidences in the overlapping forward parts of the two eyes.



## Other cues

### The measurement of size and area

Larger objects project to more facets on the eye and size is measured as the solid angle subtended. The total photon flux within the area is a separate measure and bees can be trained to discriminate a large grey spot that is 50 per cent black from a smaller spot displaying the same amount of black. Bees spontaneously prefer a large black spot to a small one, but the preference can be reversed by training. When there are two or more spots or bars in the same local region of the eye, the bees lump them together and cannot distinguish their separate sizes, but when they are more than  $15^\circ$  apart, bees distinguish them separately, like landmarks.

The feature detectors for areas are probably related directly to the photoreceptors of the retina. The measure of size might be as simple as the summation of responses in a local region of the eye, not necessarily the same size as for edges.

The angle at the eye is combined with the range to give the bee a measure of the absolute size of a black or coloured area, as shown by randomising the angle subtended and the range while keeping the absolute size constant (Figure 7.5).

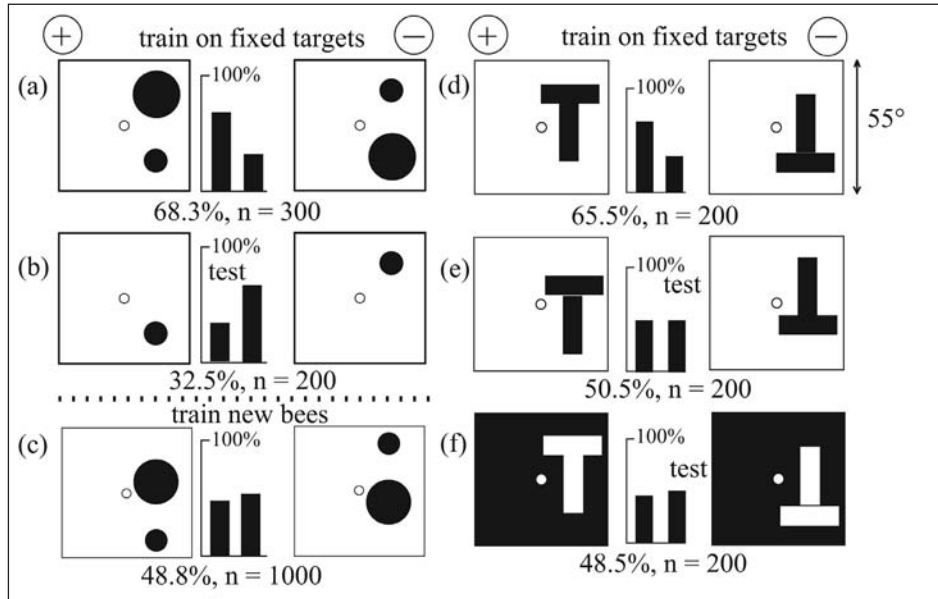
### The reward hole

As the bees were familiarised with the apparatus, they flew into the reward hole many times in the days before the training began. They detected it from the baffle and looked towards the point where they had previously landed on its lower lip. Bees easily remember whether the place just below the reward hole is dark or light. Many years ago, Baumgärtner (1928) found that bees discriminated the relative positions of small coloured rectangles only when displayed close to this point of landing (Figure 1.3). Friedlaender (1931) found that when bees had learned to discriminate an area of black near the reward hole, they lost it when it was moved up (Figures 1.6e and 1.6f).

### A filter for the height of the centre

Let us consider what we know about how we locate things in space. Simple tasks that humans take for granted, such as grasping, stepping and tracking, are dependent on an ability to locate objects. Several studies have found that humans locate either the centroid or the midpoint between opposite edges and that the least change in position that can be detected is proportional to the linear separation of the objects. This is Weber's law of separation (Whitaker et al. 2002). The experimental results for human vision can be explained by two-dimensional spatial filters with fixed coordinates on the eye, which detect intensity and operate at several different field sizes to locate position.

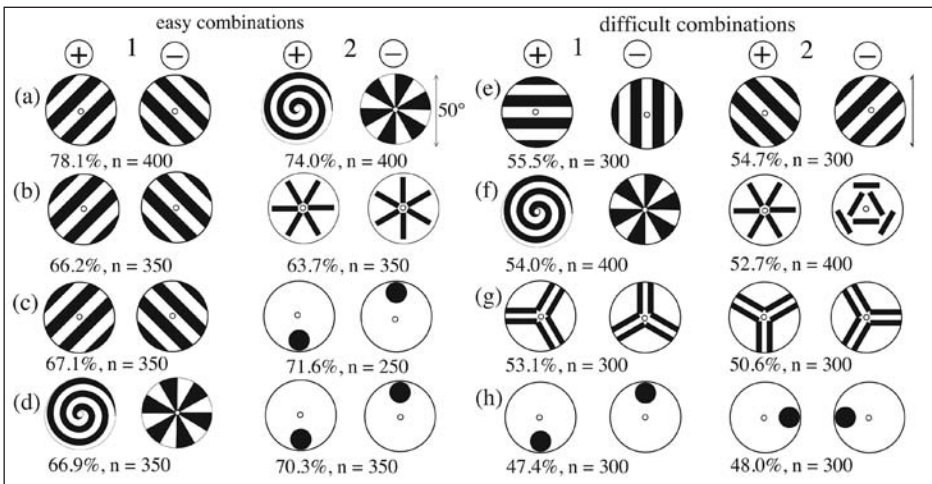
Figure 9.21 In these examples, the edge orientation cancelled to zero and the most preferred cue was the location of the centre of black within the local region.  
 a) Training with two separate spots on each target. b) Test with the small spots only; the bees preferred the black at the top. c) There was no preference in a test with the common centres of gravity at the same level. d) Training with T patterns. e) Failure with the centres of gravity moved to the same level, showing that the pattern was not the cue. f) Failure when black was exchanged for white, because the bees had learned the positions of the centres of black.



Source: After Horridge (2003b).

The position of the centroid is a feature of any pattern and is used in machine vision as a uniquely defined point. In a lucky discovery, it was found that bees could not discriminate the rotation of an equilateral triangle when the centroid remained at the same position in the vertical direction (Horridge 1997a). The difference in the vertical direction in the positions of the centres of two otherwise equal patterns is a sufficient cue for discrimination—a very small part of the pattern indeed. When there are two separate spots or bars on the target within the local region of the eye, they are not detected separately, but the position of their common centre can be learned (Figure 9.20). Spots or bars further apart on the target (in adjacent local regions of the eye) are discriminated separately. Therefore there is only one filter for position of an area in the vertical direction in each local region of the eye.

Figure 9.22 a–d) Patterns that bees easily discriminate when trained alternately on pairs 1 and 2. a) Orientation at  $-45^\circ$  versus  $+45^\circ$  and spiral versus sectors. b) Orientation at  $-45^\circ$  versus  $+45^\circ$ , and the rotation of six spokes. c) Orientation at  $-45^\circ$  versus  $+45^\circ$  and two spots, one at the top versus one at the bottom. d) Spiral versus sectors (randomly rotated) and two spots, one at the top versus one at the bottom. e–h) Patterns that bees do not discriminate when trained alternately on pairs 1 and 2. e) Horizontal versus vertical gratings, and  $+45^\circ$  versus  $-45^\circ$ ; these patterns were fixed in position during the training. f) A spiral of period  $8^\circ$  versus 12 sectors, and six radial versus six tangential bars; these patterns were shuffled by rotation during the training. g) Radial patterns based on symmetry of three and two other positions of the same patterns; these patterns were fixed in position during the training. h) A spot (sub-tense  $16^\circ$ ) at the top of the pattern versus the same spot at the bottom, and the same spot at right versus at left.



## Bees learn one cue of each type in each local region

Bees were trained, first on one pair of patterns for 10 minutes, then on a second pair for 10 minutes, and then back to the first, which was repeated for two hours (Figure 9.22). The pairs of patterns were selected to test the hypothesis that, in each local region of the eye, there was only one channel of each kind and each processed its own cue. If so, two different pairs of patterns that display different states of the same cue would interfere with each other during the learning process. The cues tested were: average orientation of patterns of sectors, edges, radial and tangential edges based on a symmetry of three or six, the position of a black spot and the exchange of black and white.

In Figures 9.22e–h, the patterns all look different to human vision, but the cues in each pair excite the same set of filters in the bee. In Figure 9.22f, the positive cue is tangential for the spiral/sector patterns but is radial for the patterns of six bars, so the bees learn not to use the tangential and radial cues and are left with nothing else, although the pairs of patterns are quite different. The positions of

the spots in Figure 9.22h are all different, but the bees are unable to learn that they are positive in two positions and negative in the other two, when they are seen in the same context. The bees do not find a consistent cue in either the positive or the negative targets.

On the other hand, bees easily learn to discriminate when one alternated pair of patterns is of one type and the other pair is of a quite different type (Figures 9.22a–d). The inference is that there is one processing channel for each type of cue in each local region of the eye (Figure 10.1).

The bees failed when they were faced with two simultaneous tasks involving the same type of cue in different states, although all sixteen of the pairs were readily learned individually. They do not learn one of the pairs and ignore the other, which would improve their chance of a reward. Instead, they start to learn afresh each time the patterns are changed, as if each cue channel cannot learn two tasks at the same place. Of course, in a different context, at a different place, the bees might be using the same cues for a different choice because other landmarks are different.

## Detection with and without memory of it

In the experiments in which the bees discriminated gratings, they detected the modulation and orientation differences but not the grating pattern. Similarly, they detected spokes or parts of circles as radial or tangential and located the centres but did not recognise the patterns. This helps to clarify the difference between discrimination of patterns and memory of them.

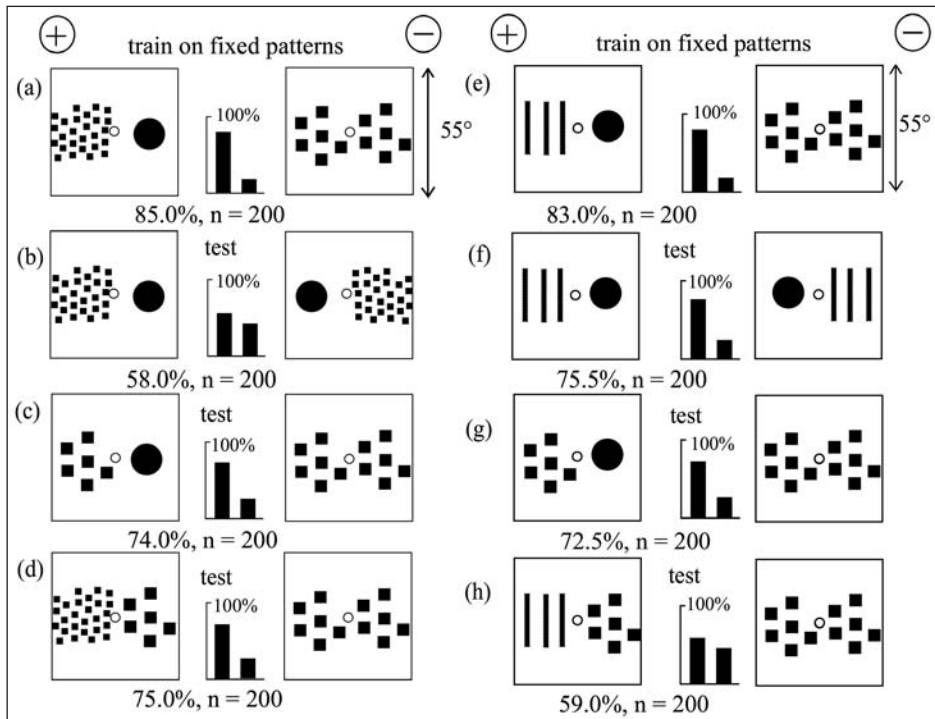
The essential first step is the simultaneous detection of all edges in the local region, but there is no memory at this stage. In each region, the number of excited edge detectors gives the total modulation and their orientations are integrated to give the average orientation (Figure 9.5) and to identify spokes and circles and locate their centres (Figure 9.19b). Other cues, such as the position of the centre of black, colour and the area or size, are abstracted by other pathways in parallel and remembered according to a scale of preferences. All these cues are remembered separately if rewarded, but processing of the image stops short of reassembly.

## The preferences for the cues

In Hertz's earliest efforts, preferences were observed when untrained bees selected one pattern from a variety and they learned most readily the patterns that they spontaneously preferred. Preferences were also revealed in differences in the rate of learning and the maximum score achieved. In most discrimination tests in the past decades, the bees had no control over the choice of the images and it was often not clear whether they learned the rewarded pattern, the difference between the rewarded and the unrewarded patterns, to avoid the

unrewarded target, or to avoid unexpected cues that were not in the training patterns at all. More to the point, it was not clear how many cues were learned in parallel and in what order. This situation was due to the lack of sufficient tests of trained bees. The preferences were not explored because most of the cues had not been described. Because the preferences were ignored, it was not possible to understand how patterns displaying several cues were discriminated.

Figure 9.23 A way to demonstrate preferences for cues during the learning process; bees were trained to discriminate between a rewarded target displaying two various cues versus a neutral pattern. a) Train with modulation and a large spot. b) Test against the mirror image, with a poor score, showing that these cues are both preferred. c) Test showing that the bees used the spot as a cue. d) Test showing that they also used the modulation. e) Train with parallel bars and a large spot. f) Test against the mirror image, with a high score showing that the position of something was remembered. g) The bees used the spot as a cue. h) The bars were poorly remembered.



Source: Horridge (2007).

To list the preferences, bees were trained to discriminate between a rewarded target with one pattern on its left side and a different one on the right, versus a neutral pattern (Figure 9.23). This arrangement gave the bees a choice of what to learn on a single target. Tests showed that in some cases they learned two or three cues simultaneously; in other cases, the bees learned one or they preferred to avoid the unrewarded target.

By working with different combinations of patterns, it was possible to put the cues into an order of preference. When two or more were displayed at the same time, the bees learned one first and more strongly than the others. The order of preference during learning was: 1) total area, 2) position of the centre of area, 3) total modulation, 4) radial edges, 5) average local orientation, 6) positions of hubs, and 7) tangential edges. Single black spots and strongly modulated patterns were easily learned. Large black spots were preferred over small ones. Radial spokes and parallel edges were weak cues. Symmetry in a pattern of bars was preferred as a cue over the edge orientations that generated the symmetry. When a weak and a strong cue were presented together, the weak one was scarcely noticed. Various patterns such as a spot, a square cross, a group of small squares and many complex patterns, provided as cues only the area of black, modulation and position of the centre of black. The bees could learn not to avoid circular patterns. When two colours were presented side by side on the rewarded target or on separate targets, the bees had difficulty learning both at the same time. They learned blue in preference to fawn or yellow, even if they had to learn to avoid the blue.

When presented with a pattern on each target, they ignored the cues that were displayed on both targets. When no preferred cue was associated with the reward, they learned to avoid the unrewarded target even if stuffed with cues or blank. In general, they learned to avoid the negative target when the most preferred cues were displayed there.

## Salience versus retinotopic cues

In previous work, when a broad black bar or spot was moved more than  $10^\circ$  after training, the bees did not recognise it in its new place, showing that the bees had learned the place. In the choice above, with three parallel bars or a large black spot (Figure 9.23f), the training pattern was distinguished from its mirror image. This was explained by the change in the attraction of the spot in an unexpected position. In other cases, even though the training was successful, the preference was equal on the two sides, so mirror images were confused (for example, Figure 9.23b).

The sensitivity to displacement was related to the field size of the cue. Cues of orientation and the position of a black area were more retinotopic and therefore detected in smaller fields and were not salient. Modulation was detected over larger displacements and therefore in larger fields, so that it was more salient. Small fields implied some failures to detect; large fields implied some failures to localise, but improved the salience. Each cue had its own compromise field size. The most salient cue was a small black spot, and was detected in a large field.

## Resolution of the feature detectors

In general, the sizes of the fields of the different cues have not yet been measured. The bees estimated the cues quantitatively and learned absolute size, relative size or angular size of a spot, depending on how they were trained. The minimum detected difference in modulation between two textured patterns was about 30 per cent (Horridge 1997a). When trained to a particular orientation of an edge, and then offered a choice between two others, they preferred the orientation that was nearest to the rewarded one, with minimum detectable difference about 30°.

The widths of the angular sensitivity curves for the orientation of an edge (Srinivasan et al. 1994) or of the axis of bilateral symmetry (Horridge 1996c) were about 90° at the 50 per cent level, because the edge orientation detectors were short. The minimum detected difference in positions of an area of black or the centre of a radial hub in the vertical direction was about 10° (Horridge 2006b). Long training improved the precision of discrimination. It is probable that the field sizes, minimum difference and resolution limit depend also on the pattern. Bees are particularly effective at discriminating the transposition of two coloured panels in the vertical direction on the target, even with no green contrast, and easily discriminate differences of 6° (Gould 1986; Horridge 1999a, 1999b, 2000b). Probably more than two can be learned, but tests of the trained bees with the individual colours in their separate positions have not gone beyond two. Bees make simultaneous use of landmarks in different directions but the minimum angle between them for each of the cues is not known.

## Avoidance of a parameter not in the training

It was accidentally discovered that in a test, discrimination is lost when a black spot is added at the same place on both targets. Later, it was noticed that the addition of any parameters that were not displayed in the training caused the bees to act as though they were in the wrong place when they detected a cue that should not have been there. As a result, a small black spot is characterised mainly by parameters that are absent. The decision process makes full use of the available options provided by the repertoire of feature detectors. On arrival at a new place, the bees behave as though they have a list of cues marked as familiar or not, so increasing the variety of labels and useful landmarks. There is less effect when a parameter is duplicated or when an expected parameter is omitted from a test, because they had learned several. All these conclusions were logical inferences from a variety of tests.

## Why do they learn more than one cue?

The very high scores obtained when training on a single pattern versus a white target are due partly to the fact that the bees easily detect black (and yet



they still make some errors). This soft option was not, however, validated by the subsequent tests, which showed in every case that the bees learned three preferred cues in parallel—notably, the area, the modulation and the position of the centre. High training scores are misleading at the start of the training when the bees first learn to go to anything black. The high scores show that the bees have an easy choice, not that they see the patterns.

Learning several cues in order of preference, with their positions, has two advantages. First, in a natural situation, the coincidence of several cues makes it less likely that they mistake the place. Second, the more cues they learn, the more likely they are to find the reward although some part of the scene is changed.

### It was cues all the way

By the 1960s, feature detectors were the best explanation for image processing in vertebrates, as vindicated by work on computer vision. By 1994, the idea had reached bee vision. A different type of experiment, however, was required to show that there was no additional memory of the pattern as distinct from the cues. Instead of shuffling the patterns to eliminate unwanted cues, the patterns now had to be fixed on the targets to give the bees an opportunity to form an eidetic image (or retinotopic memory). The trained bees were then given a large variety of tests to see what they had learned.

A rewarded black square was easily discriminated from an unrewarded oblique bar of the same area (Figure 2.1). This was exactly the pair of patterns used by Wehner (1981:Fig. 86) to illustrate the detection of areas of overlap and non-overlap in the theory of the eidetic image. Tests showed that the real cue was the position of the orientation of the edges of the bar on the unrewarded target and there was no indication that the bees noticed the rewarded black square at all.

Even more significant results emerged after training on a single black bar versus a plain white target (Figure 9.5). The trained bees showed equal preference when the alternative choice was a square, a rotated bar or a line of small squares of the same area, centred on the same place on the target. They could not recognise the training bar when it had been moved to a new place on the target. The cue was anything black of the expected size at the expected place. The bees were sensitive to an additional parameter or change in the magnitude of the cue, but not to shape or pattern.

In numerous further examples, there was no evidence that the bees had remembered or even detected the patterns, only the cues. The cues had won the day by default, but they were limited when alone. Each feature detector is interested only in a field of  $3^\circ$ , but the cues, like the orientation-detector neuron



(Figure 9.4g), have been summed over a field up to  $20^\circ$  wide in a local eye region. The local eye regions collaborate to recognise a familiar place, as described in the next chapter.

This analysis has been exposed with all the tedious detail that has to be explored to reach the simple model (Figure 9.1). It was made possible by the use of patterns that subtended less than  $40\text{--}50^\circ$  as seen from the point of choice of the bees, so that the bees could not make use of the configurational layout of widely separated parameters. It is significant that the type of system is not like wax, which moulds to any shape to make a memory, but is a varied collection of innate boxes that collects running totals of a few types of units of data from their local region. The mechanism illustrates how a picture or panorama can be recognised by a simple mechanism although the information in it is greatly reduced.

## Endnotes

1. One way to understand this chapter is to read the illustrations, starting with the training patterns at the top, and then consider yourself in the position of the trained bee in the tests, looking in the expected place for the cues learned in the training. It then becomes apparent why the bee succeeded or failed in the test.