

10

RECOGNITION OF THE GOAL¹

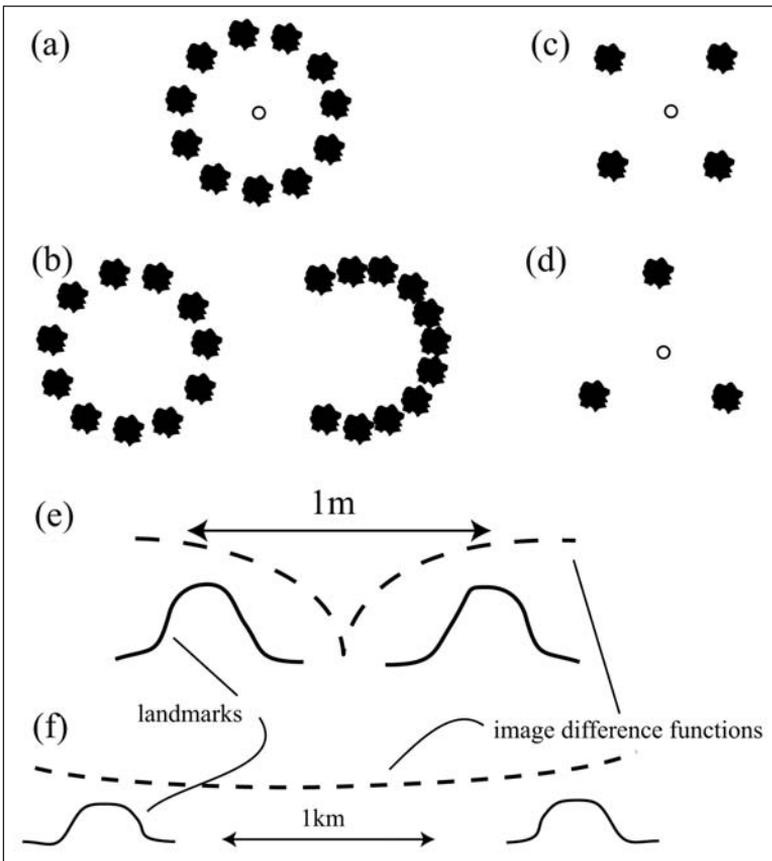
We can now return to the topic that caused Forel, Lubbock and Plateau so much trouble in the late nineteenth century: how bees locate and then recognise their destination. The roadblock to progress at that time was that the bees had already arrived at their destination, so they were as confused by the changes in the flowers and rewards as the researchers were about what the bees expected to find. Research since then has shown that bees navigate in the right direction for the correct distance using a variety of flexible mechanisms. They care little about exact appearances along the way, unless they have to search, but they care a great deal about recognition of the exact place of the reward or the hive. They persist in searching because they have not learned an alternative strategy. To study recognition, we need to know exactly what the bees have learned.

Wasps that dug out nests among the sand dunes of the Netherlands provided early indications of the mechanism. During the 1940s, Baerends, van Beusekom, Tinbergen and their Dutch colleagues placed artificial markers around nests and removed other obvious landmarks. When the wasps had learned the layout, the configuration of the markers was changed. Arriving near the nest with food for the young, the wasps made the best match they could between the altered configuration and what they remembered of the previous one. They preferred to rely on distant rather than nearer objects of the same apparent size, and at first used the whole configuration to guide them to the nest hole. With increasing experience, the wasps relied more on a few selected landmarks. Their responses showed that they approached progressively towards the place as they detected the expected landmarks at the expected angles (Figures 10.1a and 10.1b).

Bees approaching the goal along their usual track detect first the most preferred cues on the nearest landmark. This reminds them of the direction to the next landmark, and so on. They orient themselves by reference to their own body coordinates and move in the direction that increases the angles between familiar landmarks. This strategy improves the fit between the scene surrounding them and their memory of it. The whole panoramic context must be appropriate for

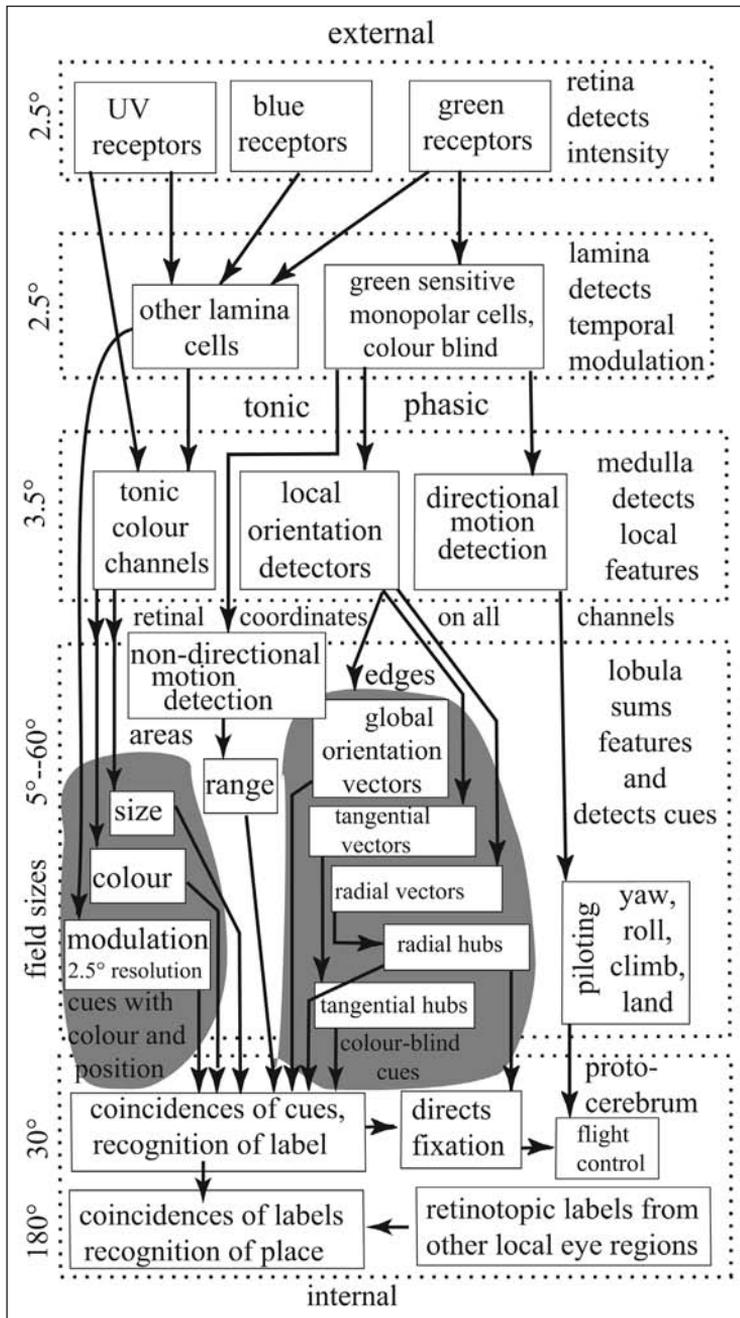
the recognition, as described over the years by many researchers on landmarks (Raubold 1928; Thorpe 1956:258; Anderson 1977b; Collett 1992; Collett et al. 2002; Fry and Wehner 2002), but they do not remember a copy of the whole scene. The phrase ‘whole panoramic context’ means that a number of expected landmarks must be recognised and no unexpected ones, otherwise something is wrong. For the bee, however, there is no ‘whole scene’, only labels on landmarks, which are all recognised independently of one another.² Although odours are significant, bees can rely solely on visual landmarks (Dyer and Gould 1981b; Geiger et al. 1995).

Figure 10.1 Recognition of the goal. a) A ring of pinecones was placed around the nest entrance. b) When the wasps were familiar with this, they were tested on their return with a choice between the original and a modified ring. c) and d) They would also use smaller numbers of cones. e) With two landmarks close together relative to the size of the wasp, the image difference function has a sharp minimum that indicates the position of the nest. f) With landmarks that are far apart, the image difference function is broad and shallow, allowing the insect to circle and wander without getting lost.



Source: (a–d) partly after van Beusekom (1948).

Figure 10.2 The arrangement of the channels in parallel behind each local region of the eye, as inferred from a wide variety of data. This local system detects one cue of each type in parallel, together forming one landmark label. These local regions are arranged around the head, as illustrated in Figure 10.7.



Source: Revised from Horridge (2000a).

The most important aspect of this homing strategy is the scale of the playing field. For the bee or wasp using distant landmarks to head towards home, the directions of cues and their heights above the horizontal change slowly, as though the insect is sliding down a gentle energy slope on which its position at any moment is not very critical (Figure 10.1f). The energy slope is that of the three-dimensional image difference function, which can be calculated from the total change of range, position and height of all the landmarks as the insect makes an incremental movement among them. When bees are using landmarks nearby, the image difference function changes direction quickly, as though the bee is sliding down a steeper energy slope with a sharp indicator of position at the lowest point (Figure 10.1e). The most effective strategy is to be able to switch between distant and close landmarks. This ability to switch between landmarks, and the use of several cues, caused some confusion for early investigators. Exactly the same principles apply to mobile machines with computer vision that recognise a place with a panoramic camera.

From the work described in the previous chapter, we have a list of cues that bees recognise, and if there are signs that further cues exist, we have methods of discovering them. The time has come to put the whole mechanism together.

Parallel pathways in each local region of the eye

Our knowledge of the eye and optic lobe (Chapters 5 and 6), together with the research that lists the cues, can be summarised by a formal diagram with a separate channel for each type of cue in a local region of the eye (Figure 10.2). In the periphery, green receptors connect with the large lamina cells that detect temporal modulation in individual ommatidia. These cells in several separate channels connect with the feature detectors, which detect the direction of local motion from the *sequential* modulation caused by a *moving* edge, and also with local orientation detectors, which detect the *simultaneous* modulation caused by a suitably *oriented* edge. Other lamina neurons detect the modulation in individual receptors (Figure 9.4e) with maximum resolution, and with blue and green-sensitive pathways.

These feature detectors span a group of seven retinula cells in bright light (Figures 9.4 and 9.19). They respond as independent units, so there is no improved detection of the modulation or the orientation angle with increased edge length. Large numbers of local orientation detectors with parallel axes feed into large-field orientation detectors, which therefore have the same axis of orientation as the edge detectors (Figure 9.4f). There is improved detection of large or parallel edges because the summation increases the signal but smooths the noise. Within the local region of the eye, the summation of different orientations destroys the discrimination of shape and measures the average orientations in local areas of patterns and textures. The vectors of the local

edge orientation detectors also feed into other detectors with large fields for the positions of the hubs of circular and radial arrangements (Figure 9.19b). Radial and circular patterns are identified separately but not visualised or reassembled in their layout. All three receptor colour types feed into tonic channels that separately detect colour, size and pattern disruption (on the left in Figure 10.2).

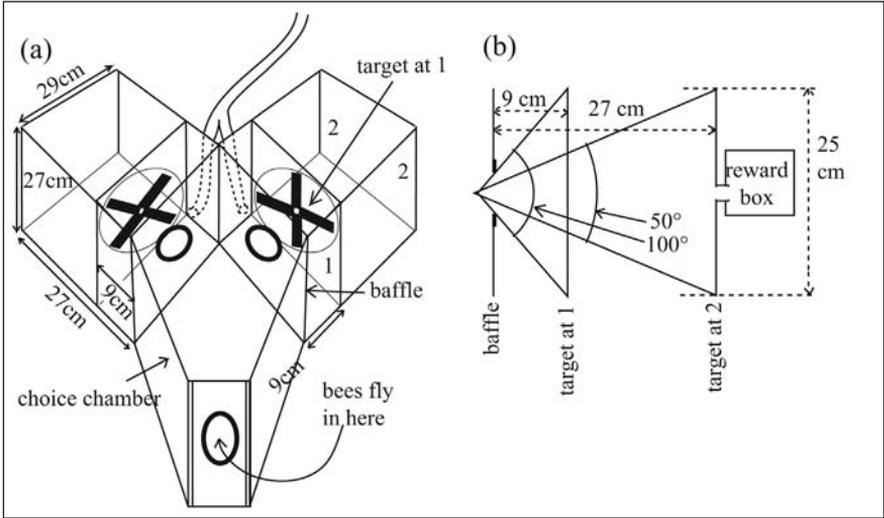
These local regions of the eye subtend about 20° , depending on the cue, and each local region sends a localised label, consisting of one cue of each type, towards the memory. There is no provision for detecting two separate sets of coincidences of cues within the local region. Whether the locally coincident cues are stored in memory depends on the context, the reward or the time of day.

The arrangement of channels has further consequences. There is no path for a transfer between green and blue pathways, otherwise orientation discrimination would not be restricted to green receptors and colour discrimination would be impossible. There is no provision for discrimination of orientation of edges from parallax and when this point was recently tested, no evidence was found (Horridge 2003a). The summation in the local eye region rejects non-coincident excitation and smoothes out the local features. The bee cannot detect two orientations, radials, tangentials, areas, positions or colours *at the same place*. All processing is done by the coincidences of responses in each array of feature detectors of each kind, all of which function independently of each other irrespective of the layout of the pattern. There is nothing special about this universally occurring mechanism of sensory processing. At the level of the local region, discrimination is like tasting a pudding that has a coincidence of flavours, or detecting an odour containing a number of different molecules, irrespective of their spatial pattern in the mixture.

Large and small patterns are differently discriminated

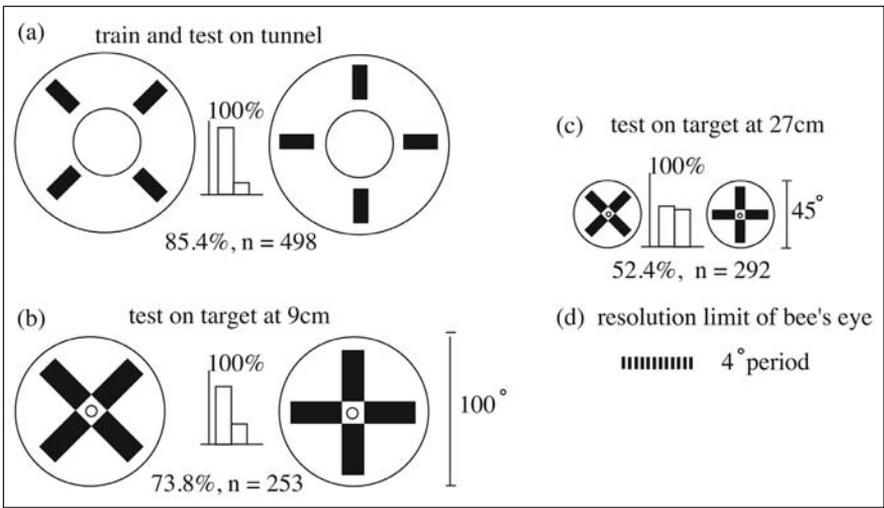
The different processing of large and small patterns has been a troublesome issue, but it provides the key to understanding how bees use the whole panorama. When Wehner (1967) trained bees to discriminate the rotation of a square cross subtending 130° at the bee's eye, he correctly inferred that they located the areas of black on the targets (Figures 4.2a and 4.5). With the same patterns subtending about 45° at the bee's eye, however, Srinivasan et al. (1994) found that the bees could not discriminate a difference in the orientation of the cross at all and they inferred that only the edge orientation could act as a cue while orthogonal orientations were cancelled by summation (Figure 4.2d). It would be some years before it was clear that both observations were correct.

Figure 10.3 The Y-choice apparatus modified by the addition of a transparent baffle in each arm. The targets can be placed at 27cm or 9cm from the baffles to control the angle subtended by the patterns. The decisions of the bees are scored when they pass the baffles.



Source: From Horridge (1996c).

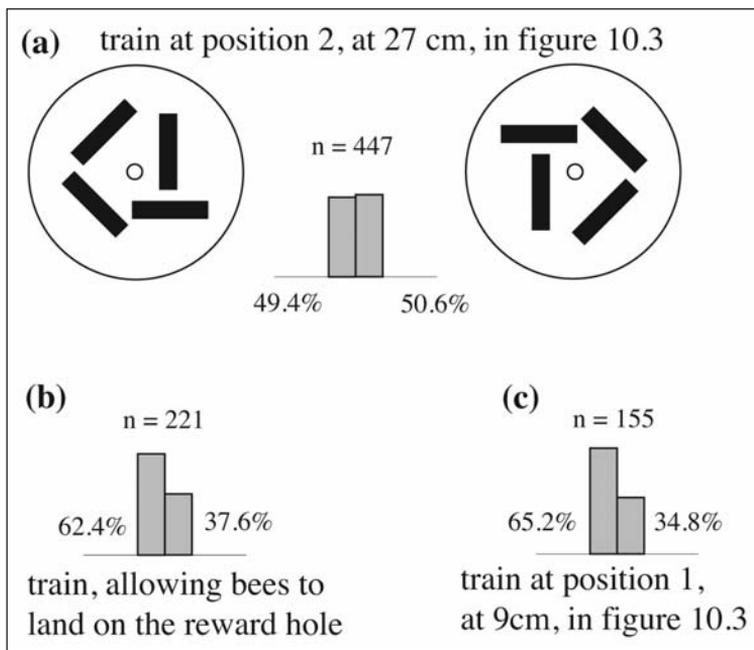
Figure 10.4 In the Y-choice apparatus, the bees pass through one of two training tunnels, each of which has four horizontal bars but one is rotated 45° relative to the other. a) The bees learn this situation very well. b) They then transfer their discrimination to targets of two crosses at 45° to each other, with the baffle at 9cm, as in Figure 10.3. c) With the baffle at 27cm, the bees cannot discriminate these targets, but this is not due to a lack of resolution of the eye, as shown by (d), a grating of period 4°.



Source: From Horridge (1996b).

The discrepancy between the two results was due to the difference in angular scale, as shown by many similar experiments and by training on patterns of one size and testing on other sizes (Figures 10.4 and 10.5). The size of the local area for the summation of the orientation cue has been measured as 15–25° across by training bees on the orientation cue and then testing them with two orthogonal bars at various distances apart. Also, within the local area, two black spots within 12° of each other were not separated (Figure 9.19c), but further apart they were separate (Horridge 2003b).

Figure 10.5 a) This pattern is not discriminated at 27cm (subtending 45°) from the same pattern rotated by 180°, because there is no difference in cues and there is no eidetic image in a local region of the eye. b) When the criterion of success is landing on the reward hole, or (c) at a range of 9cm, this pattern is easily discriminated.

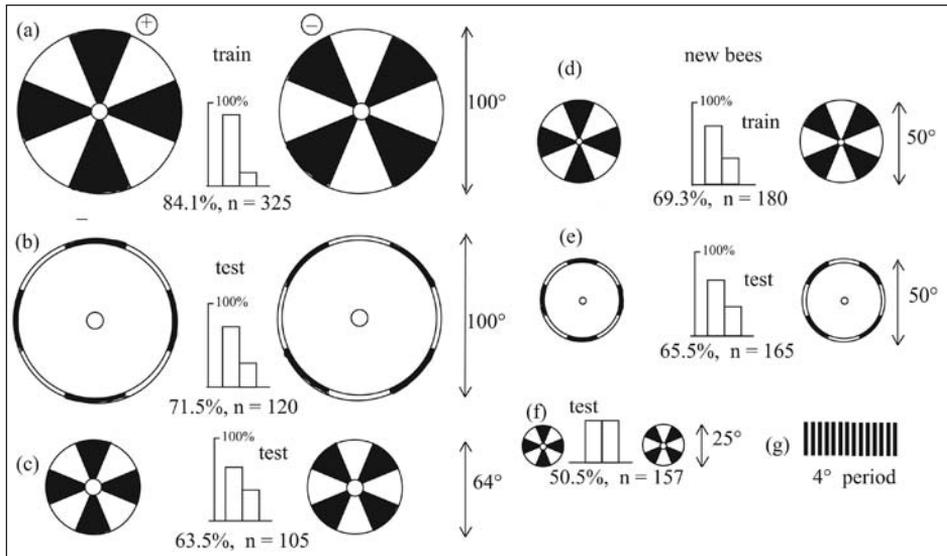


Source: After Horridge and Zhang (1995).

The sizes and separations of local regions on the eye can be measured by comparing discriminations of the same pairs of patterns at different scales (Figure 10.3b). A pattern of four bars that subtended 45° at the point of choice was not discriminated from itself rotated by 180° (Figure 10.5a) because the orientations cancelled and the only cues were the area, the position of the centre of black and modulation, and these were the same on each target. It was easily discriminated, however, when it subtended 100° or when the criterion was landing on the central reward hole (Figures 10.5b and 10.5c). In very large checkerboard patterns, the positions of squares greater than about 10° were discriminated separately (Figure 4.2f). A

pattern of plaids subtending 100° with bars 20° long and separated by 10° was easily discriminated from the same rotated by 45° (Horridge 1996b). As another example, a pattern of four gratings at 90° to each other on a target subtending 50° was discriminated from the same pattern rotated by 45° (Zhang and Horridge 1992). Thin black bars at an angle to each other were not discriminated separately and only one averaged orientation could be detected on each side of the target (Horridge 1996a, 2000b). Data such as these showed that the size of the local region for summation of orientation was smaller for gratings than for single bars. These measurements suggested that there were 10 to 15 local regions in the horizontal direction around the eye—more than enough to identify a place. The map of the local areas of the eye is not necessarily the same for each type of cue.

Figure 10.6 Discrimination of the rotation of a pattern of four black and four white equal sectors, with the pattern subtending 100° at the point of choice. a) Training produces a high score. b) The excellent performance persists when the bees see only the peripheral rims 4° wide. c) With new bees and the sectors at 27cm , the performance is still good. d) Rims only 3° wide again provide a sufficient cue. e) With only the central part of the pattern at 27cm , the bees choose at random although the patterns are well above the resolution limit of the eye, as shown by the grating (f) of period 4° at bottom right.



Source: After Horridge (1996c).

An example of how the cue can be a small part of the pattern is illustrated by the discrimination between two very large sector patterns (Figure 10.6). The bees learn the position of black only in the periphery because that is where the black areas fall on different eye regions,³ as detected from the point of choice.

In experiments with large angles between cues, bees readily detected the position of the correct reward hole by use of a cue at the side of the eye (Figure 8.2); they

learned two separate cues of orientation or colour with the two eyes (Giger 1996) or on the two sides of the target (Horridge 1997b). In many experiments over the years, they learned to distinguish two or three artificial landmarks at large angles to each other by colour, orientation or height of the centre.

The same few cues in the landmark labels

In recent experiments, bees were trained in a situation that resembled the natural task of a bee arriving at a foraging site. A black pattern on a white background was displayed in one arm of the Y-choice apparatus at a range of 27cm versus a plain white target in the other arm. The bees were obliged to use this one useful landmark. In the Y-choice maze, the pattern was nothing more than a landmark about the size of a local eye region, so the bees detected only one cue of each type. Various patterns displayed in the training experiments included the previously identified cues. They were: an oblique bar, three parallel oblique bars, an oblique grating, a square cross, six radial spokes, a large or a small spot, a spotty modulation or a ring (Figure 10.2).

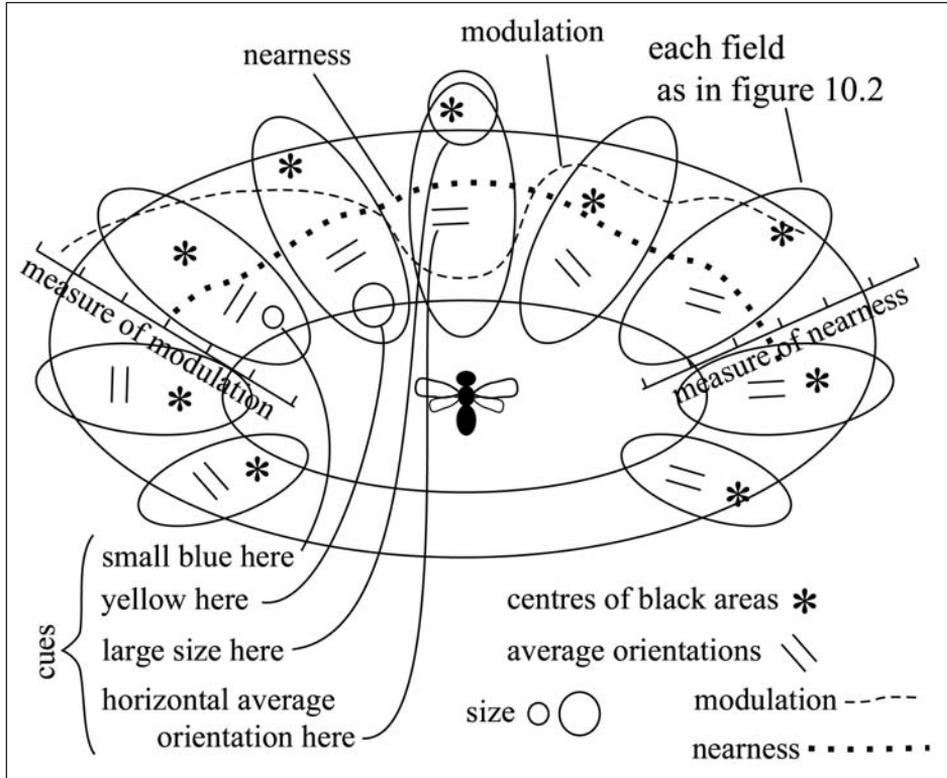
The trained bees were given a large number of interleaved tests to discover the order of preference for cues in the learning situation. They preferred to learn first the black area at the expected place, and second, modulation caused by edges at the expected place. These cues were quantified and always available. Next in preference, the orientation cue was learned from a grating that covered the target, but was ignored in a single bar. Next, the bees remembered the existence of a radial pattern and the positions of the centres of black and of radial symmetry. They preferred a blank paper to a circle. All the feature detectors behave as though they are always switched on and in tests the bees recognise and avoid unfamiliar cues that are not displayed in the training.

The cues that bees use in identification of landmarks in the local eye region turn out to be the same as those used for discriminating between fixed patterns on experimental targets. What we thought was pattern perception turned out to be the recognition of the cues displayed on a landmark. The bees in the Y-choice maze were learning the label for the correct place, not a pattern.

There were several advantages of learning several cues. First, in the natural situation, the bees are less likely to 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. Third, the redundancy improves the ability to switch from one cue to another.

It was an accident of the design, in 1987, that the Y-choice maze was approximately the size of the local region of the eye. On the other hand, in the natural panorama, the bees learn the separate labels of landmarks over a much wider range of larger angles.

Figure 10.7 The visual fields of both eyes are divided into an array of local regions around the head. Each of these regions detects any of the cues, including a smoothed measure of modulation (dotted line) and a measure of nearness = 1/ range (dashed line). Cues, as shown by the symbols, are expected in retinotopic directions relative to the midline. The bees recognise a place by the conjunction of expected cues in the expected directions.



Source: After Horridge (2005b).

The global coincidences of cues

To understand the whole mechanism in the bee, we now assemble an array of local regions side by side around the head to form a whole-eye detector of place. If we repeat the local region (Figure 10.2) about 15 times around the head, we generate an array up to 300° wide that detects up to 10 cues in each local region (Figure 10.7). In all directions around the head, the bee measures the modulation and the nearness of contrasts. For a bee, a place was recognised by a sparse but unique arrangement of landmarks and the angles between them, kept separate in the local regions around the eye, and a small fraction of this array was sufficient to define a place. All experiments revealed that each task was a separate learning experience, there was no evidence that they learned anything

more than the cues and in tests they accepted other places that displayed the same cues. By equating the cues from landmark and pattern recognition experiments, we have arrived at a synthesis.

Generalisation after training with patterns

Although few cues were involved at each landmark, the memorised coincidence of the directions of well-separated landmarks ensured that the bees would not accept the wrong place. Conversely, if the reward was moved, they would not accept that the place was wrong. Quite a different behaviour followed training with patterns. We can now explain this distinction.

The patterns were regularly moved on the target or on the flat table to make the bees look at them. This trained the bees to ignore everything outside the targets including local landmarks. Because the responses of the feature detectors were summed to form one cue of each type in each local region of the eye, and the pattern subtended about the same size as the local region of the eye, they learned only one cue in each channel. Because they were limited to one cue of each type in one local region, they would then accept quite different patterns that displayed the same cues. They generalised—that is, made errors—because they had been trained to ignore cues outside a single local eye region.

In nature, the memorised label was the only way that a bee recognised a landmark. The configurational layout of the whole wide panorama around the eye could be detected because it was divided into regions (Figure 10.7). The labels on different landmarks could be similar or not. Because there were several cues commonly available, and many different labels could be distinguished, recognition was much more precise than with a single pattern. Vision for a bee was a succession of landmark labels in different directions—some familiar, some not.

The bee moved about like a blind man navigating by a succession of familiar touch, odour and sound cues. The memory held information about only the coincidences of cues, with poor resolution of positions within the local regions. The rest of the visual input did not pass the cue detectors. There was no reassembly of pattern. In fact, the bees were not interested in patterns. Bees have no pattern perception.

Why patterns were difficult to learn

Bees made several visits before they associated a black and white pattern with a food source. The task was to select one pattern from among several displayed on the front of reward boxes (von Frisch 1914), on a flat table (Hertz 1933) or in a Y-choice maze (Figure 1.1). In all of the experiments with patterns, the rewarded pattern was moved around together with the reward to make the bees look at it, rather than where to go. In recent experiments, the pattern was moved every five minutes. When that is done, the bees must learn to ignore everything outside

the target instead of following their natural inclination to pick up several local landmarks in different eye regions. They alternate between learning to go to cues displayed on the rewarded target and to avoid the cues displayed on the unrewarded target. When most of the cues are the same on the two targets, it takes the small brain of the bee some time to grasp the difference. On the other hand, if they find a stationary food source, they make an orientation manoeuvre and immediately learn its location in relation to several convenient landmarks at wide angles, then return for more in a few minutes.

In retrospect, for the whole of the twentieth century, there was a conceptual block to understanding the relation between patterns and landmarks but no lack of experimental data. Bee vision is anti-intuitive, so it is hard to imagine that the mechanism is so simple—and even more difficult to design the right experiments. The bees did not remember the patterns or the landmarks as objects; they remembered the directions of the labels that marked the right place. In each label, the bees learned first a coincidence of modulation, area and position, then the less preferred cues, and they recognised and avoided added cues that were not in the training, but nothing more. The artificial Y-maze apparatus offered only one attractive landmark and one to avoid.

Because the bees were quick to learn to recognise a place but slow to learn a difference in the experimental training, and because it was generally believed that the bees in fact saw the patterns, bee pattern perception became a subject in its own right. For the whole of the twentieth century, however, it was anthropomorphic delusion to accept that bees perceived and discriminated patterns.

The behaviour helps explain the neuron properties

Since the early days of insect visual electrophysiology, many researchers have wondered why the image on the retina is funnelled into relatively few neurons with large fields that make no sense in terms of vision. They were unaware of the total subservience of the bees' visual processing system to a panorama of sparse retinotopic cues averaged over large fields. The fields of about 20° are large only in the context of a bench experiment, not in a compound eye with a panoramic view up to 360°. Large fields throw away detail and all chance of pattern perception within a local eye region, in favour of a few smoothed data points derived from coincidences in an array of extremely simple retinotopic feature detectors.

Endnotes

1. Bees detect something about the configurational or spatial layout of a pattern or shape when several local eye regions overlap it—for example, when the bee examines the target closely or the criterion of success is landing on the reward. If the angle subtended by the target is unknown at the point of choice, it is impossible to analyse the mechanism of discrimination. The solution to this impasse was the accidental use of the Y-choice maze, which limited the field of view to a manageable size.
2. For convergence of ideas, see Lehrer and Campan (2006).
3. Ibid.