

12

GENERALISATION AND COGNITIVE ABILITIES IN BEE VISION¹

“if that truth involves the putziness of other people or events, so be it, but if it involves the narrator’s own schmuckiness, limitations, prejudices, foibles, screw-ups at the event, etc, then these get told about too”
(David Foster Wallace, 2008).

For a century, there have been claims of something in bee vision more subtle than the coincidences of feature detectors and cues. Anthropomorphism—that is, the tendency to put human capabilities into the brains of the bees—was not openly supported, but cognition trickled down from work on higher animals.

Under the general heading of cognition in vision, the oldest belief was that the bees really saw and remembered the spatial layout of patterns. Also, it was thought that bees generalised patterns that looked similar to them. More recently, it was proposed that bees recognised patterns as a whole, that they detected patterned shapes over a patterned background and detected abstract features such as symmetry, topology and other pattern qualities irrespective of the real pattern. Indeed, they do distinguish certain global features such as size, total length of edge or modulation, average edge orientation and the presence of circles or spokes, but with only a limited repertoire of cues.

Some strange conclusions can be found in high places. For example, Giurfa et al. (2001) state that bees ‘interpolate visual information’, ‘categorize visual information’ and ‘learn contextual information’. They ‘form sameness and difference concepts’, ‘transfer to the same or a different sensory modality’, perform ‘delayed matching’ or ‘non-matching to sample tasks’, ‘learn specific objects and their physical parameters’ and ‘master abstract inter-relationships’ such as ‘sameness and difference’. These claims of cognitive abilities were based on the performance of bees that were not tested in a way that would easily have eliminated those conclusions.

Of course, anyone is free to persist with the idea that bees recognise things, rather than places. Of course, bee behaviour can be described in the terminology of the cognitive sciences, with no reference to the analytical work since that of Hertz, which showed that when presented with unfamiliar patterns, trained bees chose according to the cues that they learned in the training. Bees, however, do not recognise the patterns they were trained on when these are tested against other patterns that display the same cues (Chapter 11). This final chapter further shows that even our own experiments that supported the so-called cognition of bees rested on very shaky ground.

Generalisation

In the experiments of Mathilde Hertz (summarised, 1933), bees were trained to come for a reward of sugar solution at a flat white table where a group of patterns of similar size were shuffled in position at intervals. One of the patterns was consistently rewarded and the others were not. The bees learned to go to the rewarded one if it differed from the others in length of edge or certain other features (Figures 1.2d, 1.2e, 1.4). When tested with unfamiliar patterns, the trained bees accepted some but not others. For example, when trained on circles, the bees treated them as equal to a pattern of spots (Figure 1.5). This does not look like recognition of similarity. Hertz inferred that, although the patterns were indeed different, the bees recognised certain cues or parameters, such as a measure of the modulation or total length of edge, the area of black and the presence of symmetry. The acceptance of unfamiliar patterns was called generalisation and was attributed to two factors: the low-level recognition of parameters held in common and the existence of higher-level categories, such as a similarity detected by the bees.

Also, bees could learn to generalise when some features were shared in common between a number of training patterns, and the trained bees then recognised the same features in other patterns. When wasps (*Vespa germanica*) were trained simultaneously with different kinds of equilateral triangles, they distinguished unfamiliar triangles from squares or other shapes (Verlaine 1927). This example of generalisation during the training was said to be a remarkable performance that suggested a higher cognitive function, but there was no consideration of simple cues as the explanation.

There was not universal acceptance of generalisation. In his useful (but usually ignored) review of the topic, Carthy (1958) was equivocal. He accepted that patterns were preferred or discriminated by differences in edge length irrespective of pattern, but also gave examples of patterns of similar edge length that were discriminated and others that were not. Carthy assumed that the bees had a limited repertoire and poor recognition and he made the telling remark that 'the bees might be reacting to only parts of the pattern and not to the

whole'. Almost certainly, he had in mind the pioneering work of Lashley (1938), who showed that rats learned only a token part of the pattern that indicated the reward. Later bee researchers also ignored Lashley.

Categorisation

The term 'generalisation' comes from the way humans generalise many different shapes such as different chairs or the letter 'a' in different fonts, and cannot prevent themselves from unconsciously categorising everything that they consciously see. The members of a category can be substituted for each other without loss of understanding. Human language and vision depend on a long process of learning the useful categories detected by all the senses.

In other animals, intermediate between bees and primates, some patterns can be substituted for each other. For example, a rival male can successfully substitute almost any patch of red for the red breast of the robin or the stickleback and still initiate an attack. Because there are numerous levels of complexity and different kinds of visual systems, generalisation is hard to pin down. Bees fly about, visit flowers and navigate with landmarks, so it has been assumed that they also categorise things. This was summarised succinctly as 'patterns have to be grouped into invariance classes' (Wehner 1975). In the light of recent experiments, perhaps this should have read 'patterns are naturally grouped into invariance classes by the cues abstracted from them'.

From the earliest training experiments to the present time, there were therefore two extreme explanations of generalisation—almost opposites in their mechanism. In one, the general properties were related to categories that classified things or qualities and within which there was generalisation. On the other hand, substitutes are accepted because there are insufficient cues to distinguish them from the genuine article. Neither of these explanations was validated by tests on trained bees.

Spatial memory

In the early twentieth century, there was a variety of theories that memories were represented spatially in the mammalian brain—some even by analogy with magnetic fields. For example, following the ideas of Pavlov, 'neuron paths are established between parts of the brain'. 'We use Semon's term "engrams" to denote these physiological paths and Head's term "neural schema" as a permissible synonym' (Campion and Elliot Smith 1934). The engrams could be in or out of consciousness. The neural schema were hypothetical reassemblies of patterns in the brain.

The Gestalt theory, popular in the first half of the twentieth century and still influential today, was based on the idea that the visual image was laid out as a spatial field that would be preferred or remembered when its neural organisation

matched the previously established image in memory. Another principle of Gestalt theory was that parameters such as symmetry, roundness, star shape, coarse or fine texture and size were detected as generalised features because the visual system was adapted to detect and remember them. For most of the twentieth century, this was the dominant conceptual scaffolding and many experiments with bees were designed with this theory in mind. With improved techniques, however, modern neuroscience has not found any reassembled schema or images, even in humans, although there are neurons that look like cue detectors.

Bees certainly generalise

When Hertz shuffled the positions of patterns on a flat white table to make the bees look for them, the bees did not use information about their flight directions in relation to the orientations of the patterns. The parameters that were described were those that could be used despite the training strategy—notably, the colour, edge length, circle versus spoke, area or size, irrespective of the pattern. When trained to a pattern of a particular total edge length or modulation, versus a variety of other patterns, the bees looked for the training cue in entirely unfamiliar targets and were not interested in the real patterns (Chapter 1). In the vertical plane also, bees trained to one pattern readily accepted some unfamiliar patterns (Baumgärtner 1928; Friedlaender 1931; Wiechert 1938). From detailed experiments (Chapter 4), it was inferred that the bees simply totalled the areas of overlap of black and a measure of the edge length in a global comparison of the training and the test patterns (Cruse 1972; Anderson 1977a). This was very low-level stuff. There remained, however, a belief in something more than quantifying the parameters. As a separate mechanism, Hertz thought that radial and circular symmetry were detected as a whole, irrespective of the detail, and inferred high-level cognitive mechanisms. In contrast, from similar data, I infer a distributed low-level mechanism (Figure 9.19).

The observed generalisations of bees fell into categories that Wittgenstein would call ‘natural families’—in this case based on clearly definable simple parameters that did not overlap or merge into each other, so providing some indication of their validity. Other possible parameters, such as angles between edges or counting the corners, spots or bars, did not yield data of the same kind.

In a more suspect example, Mazokhin-Porshnyakov (1969) trained bees to discriminate between a large, hollow triangular pattern (rewarded) versus a number of ring-shaped patterns of different sizes, all presented on a horizontal surface (with the orientation randomised). The large triangles were composed of many smaller triangles and the large rings of many smaller rings, so that the bees might distinguish triangles and rings from a distance as well as from close up. The trained bees were then able to discriminate between triangles and rings of unfamiliar sizes or orientation or with different background

and form of outline. Because generalisation implied cognitive behaviour, and because memories of images were believed to be laid out spatially in the brain, Mazokhin-Porshnyakov inferred that the bees had learned the generalised concept of ‘triangularity’.

This example illustrates the flaw in all work designed to test a theory. The data were compatible with the theory, but the theory was not corroborated by further tests that could have disproved it. The bees obviously learned something from the training—possibly a small part of the pattern—but it was concluded that they learned ‘triangularity’. This faulty logic persists to this day. Later, it was shown that the bees indeed learned a few cues, but not a triangle (Anderson 1972).

Similar data; different conclusions

Following similar work with ants, Jander et al. (1970) trained wasps to discriminate an oblique black bar (Figure 4.6a) and showed that they detected the orientation when black and white were interchanged (Figure 4.6c). The trained wasps, however, confused the training bar with the white bar on a black background (Figure 4.6d). This result was interpreted in terms of rows of symmetrical detectors of modulation (Figures 4.1b and 4.1c). Generalisation was not mentioned.

At the same institute, Wehner (1971) trained bees to come to a huge oblique black bar (subtending 130° long) on a white background versus a plain white target (Figure 4.6e). Unlike the wasps trained by Jander (Figure 4.6d), the trained bees easily distinguished between the black bar on white versus a white bar on black (Figure 4.6h). Wasps and bees had learned sufficiently to respond to the edge orientation, but Wehner’s bees had learned the position of black as well. Wehner (1971) inferred that ‘the information about the direction of a visual stimulus is laid down in the central nervous system as an invariant information irrespective of the actual contrast condition’. This was in fact the experimental result expressed in different words, not an explanation. Local feature detectors were not mentioned.

Wehner then proposed that the bees must be able to distinguish between the patterns that they were observed to generalise, to exclude the possibility that they simply could not detect the differences. When several patterns are generalised, however, it does not imply that they are separately distinguishable. Indeed, they could be identical. Categories are based on usage and vary with the agent. For example, sheep distinguish between each other but humans do not distinguish between sheep. Bees distinguish between larvae that need feeding and those that do not, but probably not between individual larvae. Although

illogical, the proviso that the trained bees must be able to distinguish between the patterns that they are supposed to generalise has persisted in the literature to the present time (Benard et al. 2006).

Later, the patterns in the brain became rather volatile. For example, 'generalized information can be transferred later on to other stimulus configurations, which never occurred during the training' (Wehner 1975). This apparently destroys the idea that the memory is a shape in the brain. The inclination of a bar was discriminated 'even if the contrast was completely reversed'. 'Therefore a two-dimensional matching...has to be followed by a sampling mechanism according to invariance classes...Preprocessing of the pictorial input has to be studied first if one wants to solve the classification problem.' All this mental gymnastics, based on few results, assumed the image in the brain before recognition. In my view, however, the engram was unsupported by experiment and it was a 'devoted attempt to force nature into the conceptual boxes supplied by professional education' (Kuhn 1970:5).

Later, categories and spatial images dropped out: 'even an "experienced" bee does not seem to build up in its mind abstract search images consisting of pure geometrical forms that are invariant against other visual parameters such as hue of colour, size, contrast, or fine pattern detail' (Wehner 1981). What, then, is the way forward? One way is by more of the same. From 1995 on, several researchers found examples of transfer to unfamiliar patterns by trained bees and concluded that the patterns were generalised. Ignoring numerous examples of unlike patterns that were interchangeable and published testable explanations in terms of cues, and making no critical tests of their own, they said that the bees had cognitive abilities (Giurfa et al. 2003; Stach et al. 2004; Benard et al. 2006).

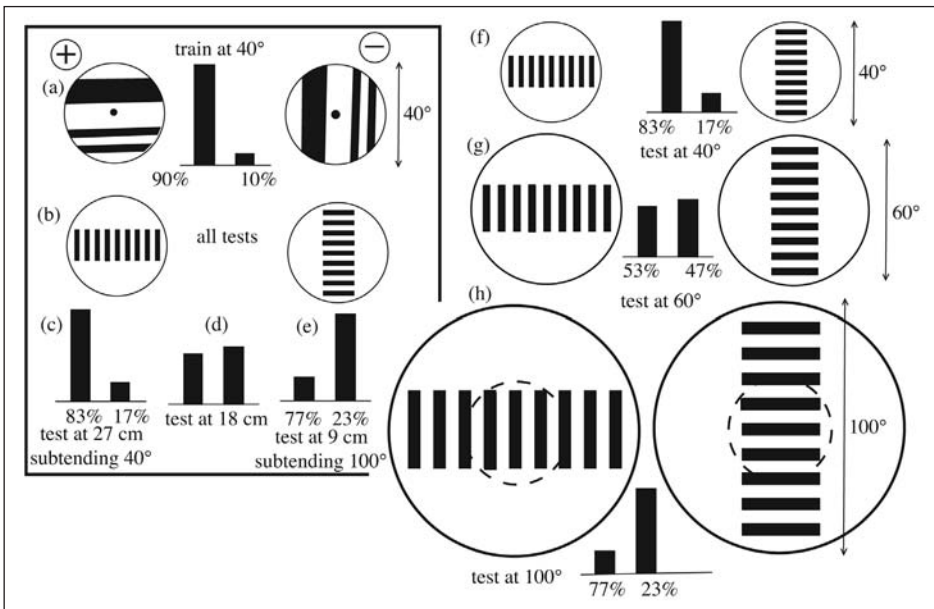
Generalisation within the training regime

Bees in flight have a very good appreciation of the sizes and ranges of contrasting objects around them. When the rewarded parameter was kept constant during the training while the other parameters were randomised, the bees could be trained to choose a black disc at a certain range irrespective of the angular size of the disc (Lehrer et al. 1988). They could also remember a disc of a certain absolute size irrespective of the apparent angular size (Horridge et al. 1992). The bees learned to generalise from the randomisation during the training. The angular size, the absolute size and the range all turned out to be parameters that could be learned.

The same strategy was used with a pattern of vertical parallel bars on one target versus a similar but horizontal pattern on the other (van Hateren et al. 1990). The positions and widths of the bars were randomised during the training, so that the bees 'made their decision on the basis of orientation only'. For a time, these results suggested that the orientation was detected irrespective of position

and that ‘specific features of the pattern, such as bars and edges, are extracted and their orientation analysed as in the mammalian cortex’ (Srinivasan et al. 1993b). As shown later, however, with a vertical versus a horizontal bar, the bees ignored the orientation and preferred to learn the modulation difference. The parameters were recognised in tests only in the places on the target where they occurred during the training (Horridge 2003a, 2007).

Figure 12.1 An error of interpretation, shown within the square. a) In the modified Y-choice apparatus (Figure 10.3), bees were trained on horizontal versus vertical random gratings, so they learned the orientation cue. b) The trained bees were tested on the composite bars at various distances. c–e) At the 9cm range, the trained bees preferred the small horizontal bars, but at 27cm, they preferred the large composite bar. Memory of local and global orientation was inferred. f–h) The illustrations are now drawn at the relative sizes detected by the bees. The bees preferred bars similar in size to those in the training patterns and the horizontal edge orientation within the small area where they had learned the modulation or orientation cue in the training, as shown by the dashed circles in (h).



Source: After Zhang et al. (1992).

Other inferences of cognition

Global versus local perception: a dog's breakfast again

In our paper (Zhang et al. 1992) that claimed to be the first attempt to examine ‘whether bees analyse patterns in terms of their local properties, global properties, or both’, our introduction was based on our reading of human

psychophysics. In the experiments, bees were trained to prefer horizontal edges by using gratings of random period (Figure 12.1a) and then tested on two large bars, each composed of many small bars at right angles to the axis of the main bar. Seen globally, there was therefore one large bar on each test target, but locally there were many smaller bars at right angles to them (Figure 12.1b).

When the trained bees made their choice at a range of 9cm, they chose the horizontal orientation of the small bars, but from 27cm, they chose the horizontal orientation of the large bar. At 18cm, the effects of the global and local orientations were supposedly cancelled (Figures 12.1c–e). In other experiments, bees trained on bars composed of smaller bars could use either the global orientation or the local orientation in tests where only one was available.

The result was not queried at the time although there were severe faults in the experiments. In fact, before baffles were introduced in 1995, the bees could have detected the global orientation from a distance and then the opposite orientation of the small bars at a later point in the flight path. Although at a range of 27cm the small bars were separated by spaces of 4° , from a greater distance, they were not separately resolved. Conversely, the bees probably detected little of the global pattern from a range of 10cm because they had been trained to expect the orientation cue within a target subtending 45° . Moreover, it has since been shown that the perceived orientation is a sum over each local region of the eye. The illustration has now been revised to clarify the situation faced by the bees (Figures 12.1f–h), but there are other problems.


In these experiments, the bees were allowed two visits on each side of the apparatus in each test, so they could have improved their score at the second visit. This is relevant only to the marginal successes. Also, vertical edges generated more modulation than horizontal edges because bees in flight scan in the horizontal plane. Luckily, our conclusions were cautious: ‘Although our experiments demonstrate the existence of local and global analysis, they do not shed light on the underlying processes’ (Zhang et al. 1992). How could they, without numerous tests of greater variety?

We in fact suggested modulation as a cue: ‘the coarse and fine gratings are detected and analysed in terms of the different temporal signatures that they produce’ (Zhang et al. 1992). Indeed, it was later found that bees preferred to learn the modulation cue rather than an orientation cue (Horridge 2007), and untrained bees and wasps preferred patterns rich in modulation to those rich in orientation (Jander et al. 1970; Lehrer et al. 1995). Other work showed that the detectors of edge orientation were only local and that they did not span gaps to detect global orientation (Horridge 2003c). ‘Global perception’ was simply a cover for ignorance, but for years we knew no better.

If we had known more at that time about detection of cues in fixed patterns, we would have tested for modulation and locations of black and orientation

cues. Finally, some of the data were suspect because there was a limited variety of 10-minute tests. In hindsight, our suggestion of global detectors was no explanation at all. It was a form of words that was consistent with human impressions of what the bees detected. This is exactly what science is supposed to eliminate.

Figure 12.2 The scores in training experiments with pairs of patterns, each with four different orientations in the four quadrants. a) With a difference in average edge orientations on the two sides of the targets. b) As before, but with the right target rotated. c) An example with no average orientation. d) With a difference in average edge orientations on the two sides of the targets. e) With radial versus tangential cues and also orientation differences on the two sides. f) With radial versus tangential cues. g) Mirror images of (f), with the same cues. h) Patterns with no detectable difference in cues.



Sources: (a–c) from Zhang and Horridge (1992); (e) from Giurfa et al. (1999); (f, g) from Stach et al. (2004); (h) from Horridge and Zhang (1995:Fig. 6a).

Separate regions of the target

In our next experiment, we planned ‘to see how many parts of a pattern could be discriminated separately, and whether discrimination was lost on rotation or inversion of the parts’ (Zhang et al. 1992). A target was divided into four quadrants with a differently oriented grating of period 8° in each quadrant.

These patterns confused subsequent researchers but not the bees. In the Y-choice maze, bees discriminated the rewarded training pattern from a similar pattern with the quadrants rearranged (Figure 12.2a). Increasing the number of sectors to eight or 16 showed that the smallest effective sectors subtended about 22° at the eye, which was more than 100 facets or a similar number of unit orientation detectors. This calculation gave ‘some idea of how an array of numerous templates, each individually ineffective, can collaborate together to make specific ensembles that fit the pattern sufficiently well’ (Zhang and Horridge 1992).

In fact, this was all rubbish because the design of our experiment and the data were faulty. First, we were unaware at the time that one side of our training targets had more horizontal edge and the other side more vertical edge, and that the bees processed the average orientation separately on each side of the target. This cue was there for all to see (Figures 12.2a and 12.2b). In another pair of similar patterns (Figure 12.2c), the orientation cues were more likely to cancel out but something was apparently discriminated. Also, we did not test what the bees really detected or even whether they remembered an ensemble at all. Third, from 1990 to 1996, the bees were allowed two visits (10 minutes) on each side of the apparatus in the tests, which was sufficient for them to add a few points to the borderline scores.

With similar naivety, and similar patterns with orientation cues in four quadrants, Giurfa et al. (1999) allowed the bees to approach close to the targets, which therefore subtended very large angles at the final choice point, so the configurational layout of areas of black could be discriminated. They concluded that when trained with a pattern of four quadrants versus a blank target, the bees learned mainly the lower half of the rewarded pattern, but when trained with one pattern versus another (Figure 12.2e), they learned all the pattern—and to avoid the unrewarded pattern. In their training pattern, however, there were radial versus tangential edges and also differing average orientations on the two sides, which the authors did not mention, providing obvious parameters for the bees. Either one or both of these parameters was also displayed in their test patterns, so the results threw no light at all on global vision.

More recently, bees were trained with similar patterns but with shuffled thickness and positions of the bars, versus a similar unrewarded group with a different pattern of orientations (Stach et al. 2004; Stach and Giurfa 2005). This time, the targets subtended 37° at the point of choice. Discrimination depended on green contrast and therefore edges were involved. In the training targets and tests, there were opposite average orientation cues on one side of the targets (Figures 12.2f and 12.2g) and on the other side there were radial versus tangential edges that the authors did not mention. In some tests, the pattern was reduced to one bar in each quadrant while retaining the difference in average orientation on the two sides (as in Figure 12.2d); in others, the details were

shuffled within each quadrant, but the parameters remained for all to see. In tests, the trained bees discriminated with black and white reversed, as would be expected because the feature detectors for edges were symmetrical (Figure 9.4).

The ability of the bees to discriminate the unfamiliar patterns was described as a generalisation and 'after a long training with a single pair of patterns, bees built a simplified holistic pattern representation that included all four edge orientations in their appropriate spatial relationship and that allowed transfer to novel stimuli preserving such a positive layout' (Stach and Giurfa 2005). This conclusion was a guess for which there was no evidence and no test of global vision. Moreover, obvious radial parameters were displayed. There was no evidence that the bees also generalise their response to patterns with fewer correct orientations, depending on their match with the trained layout because the parameters in the training remained in the tests. There were no tests of what the bees really detected. There was certainly no evidence for the claim that the bees responded to 'the perceived lay-out' in patterns of this size.

The same data supposedly demonstrated 'categorization based on sets of multiple features' and the bees 'were shown to assemble different features to build a generic pattern representation which could be used to respond appropriately to novel stimuli sharing the same basic layout' (Benard et al. 2006), revealing the persistence of unsubstantiated ideas about spatial reassembly in the brains of targets that subtended 37° at the point of choice. The authors say the 'results show that honeybees can recognize visual patterns on the basis of the global layout made from four different orientations, common to a series of different patterns'. In fact, there were no tests of whether the layout of quadrants was noticed at all by the bees and the test data were compatible with the recognition of the obvious rad/tan or orientation cues. Moreover, there were abundant published data to show that the individual bars, the separate quadrants and the whole patterns could not be discriminated if the rad/tan and orientation cues cancelled out in patterns subtending 37° (for example, in Figures 12.2h and 9.14j).

Illusory contours

By 1993, it was possible to 'suggest, perhaps for the first time, the existence of feature-extracting mechanisms in the insect visual system that might be comparable, functionally, to those known to exist in the mammalian cortex' (Srinivasan et al. 1993). This lyric was inspired by an inference that insects perceived illusory contours. When they had been trained to discriminate between the orientations of shuffled orthogonal gratings, bees apparently saw the contours of the Kanizsa rectangle illusion (van Hateren et al. 1990:Fig. 4). It was supposed that, as in the human cortex, lines of edge detectors with similar orientation were strung together. Bees, like humans, also responded as though they saw an illusory orientation at a fault line across a regularly striped pattern

(Horridge et al. 1992). There seemed to be nothing wrong with the idea of illusory contours, but at the time we did not know that different edge orientations in close proximity cancelled each other or that edge detectors did not span across gaps that were resolved.

When the experiments were repeated, they failed. There had been two changes to the design of the experiments. Until 1996, there were no baffles in the apparatus so the bees could enter at full speed and make a fast decision from further away. Also, they were allowed 10 minutes and two visits on each side in the tests, which allowed them to improve their success rate. After 1996, however, the baffles halted them and they took longer to peruse the targets from a fixed distance. They also had only five minutes on each side and many varied tests were intercalated, so they saw the same test at long intervals between other tests. With these precautions, the bees did not detect illusory edges or the edges at fault lines (Horridge 2003a). Also, David O'Carroll told me that he could not repeat the detection of illusory contours when recording from single neurons of the dragonfly lobula.

Transfer of shape between green and blue channels

In a brief paper, Zhang et al. (1995) trained bees to discriminate between a wide horizontal bar (rewarded) and a similar vertical bar (subtending about 36° by 8° at the choice point). To prevent input via the motion-detecting system, the edges of the bars displayed contrast against background only to the green receptors. The trained bees could immediately distinguish between such bars when they were presented in blue contrast.

The observations were not in doubt, but there was no evidence for the conclusion that 'shape is memorized in a generic form regardless of whether it is initially sensed by green-contrast, blue-contrast, luminance-contrast or motion-contrast signals' Zhang et al. (1995). The shapes of the bars or the orientations at the edges were not even probable cues for stationary bars. In the light of later findings (Giger and Srinivasan 1996), it was impossible for the bees to detect orientation with the blue channel alone, and in any case, the probable cue was the modulation difference. In another experiment, the authors in fact showed that the cue was the difference in modulation between horizontal and vertical bars and modulation was detected by both green and blue receptors. When the bars were oblique, the bees learned the orientation cue and could transfer to similar targets with green contrast but not to ones with blue contrast, because there was no difference in modulation with the oblique bars. In the first experiment, the bees did not transfer between green and blue channels; they had learned the modulation cue, which was not colourblind. The conclusion, however, has been frequently quoted as evidence of cognitive transfer of shape discrimination.

The Dalmatian dog; shape from parallax

Continuing the same saga:

To investigate whether bees encode shape in a generic form, regardless of input channel, we began by asking if bees that have learned a shape defined in terms of luminance contrast can recognize the same shape when it is defined in terms of motion contrast. (Zhang et al. 1995)

Accordingly, bees were trained to discriminate a thick black ring (rewarded) from a large black spot of the same area, both centred on the central reward hole. The ‘trained bees can immediately distinguish between the same shapes when they are presented as black-and-white textures, of pixel size 4 mm square, 6 cm in front of a similarly textured background’. The same trained bees then learned to discriminate between black and white random-pixel textured oblique bars, a task that they could not do before they learned ‘that motion contrast is the relevant cue’ (Zhang et al. 1995). Even more remarkable, having learned to discriminate the two shapes with motion cues, the trained bees recognised them in blue contrast. In the earlier version (Zhang and Srinivasan 1994), a textured Dalmatian dog on a textured background was illustrated, upside-down, to make recognition of it more difficult. Miriam Lehrer used a textured elephant in one of her illustrations.

First, let us look at the internal evidence for misplaced conclusions. Pixels of 4mm square on the background would subtend 0.8° and even the pixels raised 6cm in front would subtend 1° at the point of choice, and would not be resolved. Second, the discovery that equal lengths of edges at right angles cancelled out the orientation cue (Srinivasan et al. 1994) implied that when the pixels were large enough to be resolved, the orientation cues were cancelled. Third, discrimination of orientation required only edge detectors (Figure 9.4), not motion detectors (Srinivasan et al. 1993). Fourth, the bees were allowed 10 minutes at each arm of the Y-maze before the patterns were changed, giving an average of two choices at each test, so they could more easily reach the relatively weak borderline scores that were recorded. Finally, there was no test for whether the bees saw the shapes at all. Furthermore, when I repeated the original training with exactly the same patterns, the bees learned to avoid the spot, the cue was the absence of black near the reward hole and the trained bees had no memory at all of the shapes (Figure 11.8).

When this experiment was repeated with larger pixels that were resolved, the bees failed to discriminate. The bees trained on plain black patterns would not discriminate textured patterns raised 6cm above a textured background. Furthermore, the same bees readily discriminated between two orthogonal bars of plain white paper that were raised 6cm above plain white targets, showing that weak shadows provided sufficient cues (Horridge 2003a).

In this example, the experiment was based on a good idea and the experimental data were compatible with the premise. The patterns, however, were inappropriate, the data were suspect, alternative explanations were available, there were several reasons to reject shape perception in general and shape from parallax in particular and the results could not be repeated.

Bilateral symmetry about an axis

Untrained bees have a preference for patterns with a vertical axis of bilateral symmetry, irrespective of the pattern (Lehrer et al. 1995). Bees learned to discriminate the vertical axis in patterns of two pairs of bars at right angles to each other (Figure 9.15). In agreement with the contemporary ideas about global templates, it was proposed that 'bilateral symmetry assists discrimination' and 'if there is a global filter for this pattern, it has broad angular tuning' (Horridge 1996a).

Then, to demonstrate the cue of symmetry about an axis, with the newly introduced baffles in place, bees were trained all day on seven quite different bilaterally symmetrical patterns that were taken successively for 10 minutes each. The patterns all displayed the same four black bars in various arrangements. The rewarded ones had a vertical axis of symmetry and the unrewarded one in each pair was the same pattern rotated through 90° (Figure 9.20). Bees readily detected the orientation of the axis in tests with unfamiliar patterns.

From these results, I inferred global filters that

perhaps work in the same way as the face detectors in human vision... It is difficult for us to appreciate that the bees are sensitive to the pattern as a whole and discriminate a global feature of it without remembering the locations or orientations of individual bars, but in our own vision we are familiar with our discrimination of colours without being able to identify their constituent wavelengths. In this respect, bee vision of form resembles our vision of colour; the components of it are not separately discriminated.

Like the smile of the Cheshire cat in Lewis Carroll's *Alice Through the Looking Glass*, the abstract feature, the smile, persists although the cat is no longer distinguished. Generalization of this type is the essence of vision, in that whole objects and complex relationships are recognized irrespective of local variables. (Horridge 1996a)

With science like this, who needs poets?

In fact, low-level cues must have been detected in all the symmetrical patterns, but the tests were never done. Scanning in flight of bilateral symmetry yields the same sequence of feature detection in either direction. Alternatively, bees can discriminate the average orientation and the averaged positions of the centres

of black, colour or other cues, separately on the two sides of the target. When these averages, together with radial and tangential cues, are at equal heights on the two sides, they could be sufficient to distinguish bilateral symmetry about a vertical axis in many patterns. The topic needs further investigation.

Topology as a cue

Having found that bees discriminate between a rewarded black O and an unrewarded black S (Figure 11.10a), Chen et al. (2003) proposed that the ‘topological properties constitute a formal description of fundamental perceptual organizations, such as distinguishing [a] figure from [the] background, parsing visual scenes into potential objects, and performing other global, Gestalt-like operations’. This was typical gobbledegook borrowed from the cognitive sciences.

Chen et al. made four tests of the trained bees that neither proved nor disproved topology as a cue, but they tried none of the possible tests that would have disproved it, nor did they demonstrate what cues the bees had really learned. They also found that discrimination of the S (rewarded) from the O was learned extremely rapidly, which is now explained by the innate avoidance of the O, so probably it was not learned at all.

The choice of the broad black ring as the rewarded target was most unfortunate. In a repetition of the same experiments, the cues were the presence of black near the reward hole *on the unrewarded target* and the orientation of the middle section of the S (Figure 11.10). The topology was irrelevant.

Even without the critical tests, the discrimination between a closed and an open shape in no way demonstrated that bees recognised the topology, any more than a discrimination between two pictures of human faces showed that bees recognised faces as faces or as individuals.

Preference for radial symmetry irrespective of pattern

About 1994, attention was drawn to the evolutionary advantages of symmetry in a variety of animals and plants. Animals are intrinsically asymmetrical but symmetry has been perfected by sexual selection and forward locomotion. Bees preferred to forage from symmetrical flowers (Møller 1995) and it was supposed that flowers adopted and rewarded symmetry to attract bees.

In a circular apparatus, bees were trained to come to neutral targets placed vertically at the back of four out of 12 compartments (Figure 9.11). The training patterns were then replaced with 12 patterns with different levels and kinds of symmetry in equal numbers and the bees’ choices were recorded. When the test patterns were of the same kind, the bees preferred larger periods and broader

bars or sectors. When the patterns differed in type, the bees preferred radial and avoided circular patterns. Bilaterally symmetrical patterns were preferred when their axes were vertical.

In line with the ideas of the time, it was proposed that ‘filters tuned to radiating and circular shape elements...would enable the bee to use global parameters to discriminate numerous patterns with only a small number of specialized neurons’ (Lehrer et al. 1995). The discrimination of edges as radial or tangential when they lay at different angles to each other was, however, at odds with the discovery that edges at large angles to each other reduced the orientation cue (Srinivasan et al. 1994). To resolve the discrepancy, it was proposed that ‘bees have additional filters, of which the minimum number is two types in polar coordinates that resemble radial sectors and concentric circles’, and that the ‘large field or global detectors of polar symmetry inhibit the orientation detectors’ (Horridge 1994). Later, edge detectors in radial or circular directions on the eye were grouped into ‘innate global filters for radial and tangential contours in the pattern as a whole’ (Horridge 1996c). None of these proposals was tested. They became firmer as time passed and they were quoted by others. In hindsight, it was an illustration of science in progress, naive moonshine or a misleading catastrophe—depending on your standpoint.

Training on radial symmetry

Before the work on preferences, bees were trained simultaneously with radial patterns (rewarded) versus tangential ones (Figure 9.12a), with the positions of black shuffled at intervals. The bees transferred their training to quite different patterns displaying the same cues. Unfortunately, it was found later that bees innately preferred radial patterns and avoided circles, so they might have learned nothing. This error was later corrected (Horridge 2006b, 2007).

A pattern of three or six equally spaced radial bars was readily discriminated from the same target rotated by half the angle between the bars, but rotation of a target with four, five or seven radial bars was poorly discriminated (Horridge 2000b). This result was ‘consistent with the proposal that there is a family of global filters at small angles to each other with 3 arms and another family with 6 arms’. ‘The early visual processing retains the resolution of the retina, but at a higher level the memory has available only the outputs of large-field filters’ (Horridge 1997c). By 1998, ‘generalization over a range of certain related images...can be explained by...coarsely tuned filters but not by an eidetic image, or universal learning mechanism’ (Horridge 1998a).

Despite these observations, there were other examples where different orientations were remembered separately in the same region of the target. No-one discussed the discrimination of edges at angles in the same pattern

(Zhang et al. 1992), different orientations in concentric circles (Horridge and Zhang 1995) or different numbers of sectors in radial patterns (Wehner 1981:Figs 59, 67). There was obviously more to be found, *dessous des cartes*.

No global filter for radial symmetry

The discriminations of symmetry presented a problem because they were independent of the scale and layout, so that many templates would be required. It was impossible to imagine global filters that fitted the data. The usual tests showed that radial or circular edges were not reassembled. The bees remembered only the radial or tangential character and the position of the centre of symmetry irrespective of the position of black, pattern or scale. Later, it was discovered that bees could be trained to discriminate between the left and the right halves of a symmetrical pattern, either radial or circular. Therefore, symmetrical global filters with a single output were ruled out because either half of the pattern would excite the same filter (Horridge 2006b).

The feature detectors proposed for edge orientation in other experiments did not string together to span across gaps (Figure 9.8). They were short, independent and about 3° long (Figure 9.9). The feature detectors for radial and tangential patterns were demonstrated by the same tests as used for detectors of edge orientation and turned out to be the same (Figure 9.19).

There were now sufficient data for an explanation of symmetry detection by local feature detectors feeding into larger fields, like all sensory processing. Edges anywhere in the pattern were treated as radial when they converged towards a hub or as tangential when lines at right angles to them converged towards a hub (Figure 9.19). The position of the hub and its radial or circular character were remembered, but the original layout of the feature detector responses was lost. This was a distributed, local and flexible mechanism that would find an average centre and identify the pattern as radial or circular by a distributed administration, irrespective of the size or pattern. Hypothetical global filters were excluded and replaced by an evidence-based explanation.

How the nexus between patterns, landmarks and place was broken

In the nineteenth century, many efforts were made to understand how bees returned exactly to the rewarded place. A common technique was to give a reward on a flower, then change the flower for another of a different shape or colour or hide it with a few leaves. Felix Plateau, for example, correctly concluded that the bees ignored the altered shapes and the colours of artificial flowers. Tedious exact repetition of Plateau's experiments showed that the bees went unerringly to the place where they had found the reward irrespective of the shapes, but were lost if the place was moved (Forel 1908:170).

When Turner (1911) trained bees to distinguish between two boxes—one with horizontal black stripes and a reward of sugar inside, the other with vertical stripes and no reward—he interchanged the positions of the boxes at intervals to make the bees look for the rewarded pattern, *irrespective of the place*. By shuffling the positions while keeping the cue constant, Turner had broken the nexus between the recognition of the label on the box, which was the horizontal edges, and the recognition of the place, which required several landmarks at wide angles to each other.

The technique was adopted by von Frisch (1914) and used with various modifications by all later investigators. Hertz (1933) placed black patterns flat on a white table and placed a reward of sugar solution next to one of them. She broke the connection between reward and place by shuffling the patterns on the table, so the bees ignored everything except the cues in the correct pattern, *irrespective of the place*.

For the first time, it was noticed that the bees took a much longer time to learn. Also, they either became tolerant or liable to mistakes when trained in this way and accepted unfamiliar patterns, which was called generalisation. In contrast, when rewarded at a fixed place, they returned after a single visit and never made an error.

Let's explain. When the positions of the patterns were shuffled or two targets were interchanged, the bees were obliged to look for the familiar cues on the patterns and they were trained to ignore everything outside the patterns. Vision was restricted to one or two forward-looking local regions of the eye by the shuffling or alternation of the patterns because the rest of the eye learned to ignore the surrounding place. So blinkered, they could no longer use the coincidences of landmarks between different eye regions (Figure 10.7). Alternating or shuffling the targets exposed the bees to errors by restricting the memory to a local region, which processed only one of each type of cue. The number and variety of cues that could be learned was inadequate to distinguish every pattern, so recognition was easily fooled. The observed ambiguity, or confusion of the bees, was called *generalisation*.

This bit of history shows how bee trainers were fooled by their own training technique, combined with the small repertoire of cues. For almost a century, they believed that bees generalised patterns because they saw them as similar or they belonged to the same bee category. The error of thought was established in the literature, heels were dug in, territories were defended, referees unjustly rejected papers and contention seriously slowed the advance of understanding.

At the same time, the coincidences of cues, the total area, the position of the centre of area, total modulation, average local orientation, the tangential or radial nature of edges and positions and types of hubs were each summed over a local region. This removed the detailed distributions of contrasts within the

experimental patterns. It was called *global vision*, but it was in fact an artefact of the training and testing technique that was restricted to one local region of the eye.

On the other hand, bees were not interested in the training pattern in an unfamiliar position and generalisation was not observed in the identification of a place. The recognition of a very large target, or a pattern that the bees were allowed to examine closely or a natural situation, involved the retinotopic detection of a variety of features over very wide angles in the whole scene (Figure 10.7) and the bees made the best fit with cues in their expected places at large angles to each other (Thorpe 1956; Collett et al. 2002). In a similar way, blind people identify a place by sound, smell and touch all around. Despite the great number of publications with ‘pattern perception’ in the title, bees detect coincidences of cues, not patterns.

Coincidences in neuron responses and learning

The explanations of visual recognition offered here have been in terms of the coincidences between the cues and expected positions of landmark labels. This explanation has a long history. Sherrington (1906) called it ‘integration’. Hebb (1949) wrote an influential book with the idea that the coincidences of inputs, including those from reward channels, would strengthen synaptic contacts on a key neuron and trigger the growth of new synapses when learning occurred. Eccles (1957) described in detail the summation or inhibition of coincidences of the inputs at synapses as the key to understanding all nervous systems. Moreover, the immense, new topic of adaptive neural nets in artificial learning systems relies on the idea that the coincidences of different inputs allow the neural net to learn.

Whether or not there is a range of bee behaviour that makes use of something more thoughtful than the learning of rewarded coincidences, or the avoidance of punished ones, seems now to be a matter of opinion.

Cognitive visual behaviour in route finding and navigation

Much of this discussion depends on the education and life experiences of the contestants. An education in the Napoleonic system of Continental Europe, or as an ethologist, will lean you towards accepting intuitive explanations of performance and reliance on definitions of terms that are usually simply taken from cognitive psychology. In contrast, English empiricists or American comparative physiologists will lean towards mechanistic analysis.

The least justified, most dogmatic or fundamentalist opinion that I can find comes from Professor Randy Gallistel, of Rutgers University, who does not work primarily on bees but has just written a book on cognition. Gallistel would say that the word ‘cognitive’ implies computation, so if the bee computes, it

has cognition—simple as that! The bee does path integration and optic flow summation, therefore it computes, therefore it has cognition. The problem with this is easily discerned when I point out that a bindweed stem describes an excellent spiral as it winds up a stick and my slide rule computes without needing a battery, therefore they have cognition. Even my watch computes.

A more reasonable view comes from Adrian Dyer, who in fact works on bees (but allows the bees to land on the targets and thereby blocks his own progress). Dyer would say that bees will reveal cognition if they can use memory to solve a novel or abstract task, and he can point to several published accounts of performance. For example, bees that have learned mazes are faster at solving an unfamiliar maze. Trained bees accept some unfamiliar targets in place of the learned one. Bees familiar with a foraging ground adjust their foraging method according to the place (and time of day). Moreover, there is some evidence that bees count. Again, the problem is that these accounts are of the *performance*, and the bees have not been tested thoroughly to see what they have really learned. The idea of cognition was an intuitive inference that was not deduced from experimental results, but was a word taken from the cognitive sciences, put into the title of the paper and then claimed to be a causative agent.

Even worse, we might already have a mechanistic explanation available for these performances. For example, Hertz showed that bees discriminated between targets on the basis of more or less modulation (Figure 1.4), so four objects would generate more modulation than three of the same objects, enabling the bees to pass the test of counting. Perhaps ‘cognition’ equates to that which we do not yet understand.

It is clear from this example that empirical experimental data about parameters or landmark labels can replace cognition as a causative agent. It is not so clear that cognition can ever be demonstrated as a causative necessity, because there might always be an undiscovered mechanistic explanation.

Coming closer to home, Srinivasan would grant ‘cognition’ to any animal that can do something, such as an ability to categorise, navigate complex mazes or other tasks that might require thought in a human. Srinivasan would say, ‘If it looks like cognition, sounds like cognition, acts like cognition, then it is cognition.’ There are more problems here than the requirement to test unsupported inferences of cognition. For example, robots perform tasks more difficult than bees and there are distinguished psychologists who would allow cognition for robots but not for bees. Performance that looks like cognition is a feature of computer programs that look ahead and predict moves in chess games and also of systems with feedback loops that counteract unexpected forces and stabilise our posture. The performance is just the beginning. We look for the mechanism, not for a word that tells us that the bee does something interesting.

Finally, another hardworking experimentalist, Randolph Menzel, has been involved many times in the discussion of cognition in bees, mainly over the question of whether bees remember maps of their territory. Menzel has again shown that bees remember the territory that they have explored, either voluntarily or in search of the randomised position of a food source. He would probably allow the term 'cognitive' for novel behaviour that emerges from a combination of memories that creates expectations of outcomes and (contra Dyer and Srinivasan) he would exclude those examples where behaviour is directly controlled by the sequence of stimuli.

From quite a different standpoint, Tye (1997) reviews the most remarkable bee performances and argues that the bees are not aware of what they are doing, and are therefore not cognitive, and also that the localisation of a light by a subject with blind-sight is a response to a stimulus and not cognition.

Need I say that the analysis of the mind of the bee cannot be based on performance alone. Before a book on 'What do bees think?' can be written, there must be some experimental analysis of several kinds, followed by detailed tests and validation, otherwise 'cognition' is just a word in an arbitrary definition.

After all this, what does the bee see?

Of course, we can never know what bees really sense when they see. In human terms, they see nothing. To the experimentalist, the expression 'What do bees see?' is a query about what stimulation they detect, not about the sensations of the bee when the visual system is in action. They detect cues and direction of movement in each local region of the eye, but these stimuli are mixed in the optic lobes with other modalities from other parts of the animal. Their appreciation of their surroundings must be like that of a blind man who uses all available inputs to control his movements.

We can *guess* what bees really detect. For example, some disturbed bees chased me away from their hive, so I am not going back there—they might SEE ME. Alternatively, we can *propose* that the disturbed bees detect and follow any large moving object, even against a textured background. Then we can devise experiments to *test* this proposal. We might *conclude* that the bees follow the largest moving object through a forest of trees and bushes when there is an odour trail generated by spilt honey, a bear or bee pheromone. For every question, we follow the steps: guess, proposal, tests, conclusion, belief and unwarranted extrapolation, then rejection. Given a sensitive imagination, assiduous observation, efficient experimentation and much thought, we slowly analyse the behaviour that the bee presents. *This* is the way that small science advances.

To make an analysis at all, we depend on the repertoire of the animal. If the bee does not respond to the training or the tests, we can go no further. The bee might have detected the stimulus but was not aroused by it. For this reason, there might be a lot that we will never suspect. This is not proof that the bee will never be fully understood. If, however, the bees respond to one group of patterns but not to a related group of patterns that differs in a defined way from the first group, we are on the way to discovering a cue.

It is easy to show that bees detect edges and areas separately, that shapes are not reassembled in their memory and that orientation is cancelled by edges at angles to each other. With that visual mechanism, we can only suggest that bee vision is similar to detecting the separate tasty molecules in coffee or hearing sounds from an orchestra.

Detection and perception

The human visual system has several kinds of lapses from conscious vision that could help us imagine the vision of the bee. One of these is the ability to be aware of our surroundings although not particularly conscious of them.

Subliminal perception is the ability to take in brief or weak signals that are not consciously detected at the time. In humans, they can be recorded by brain imaging or correlated with electrical potentials, so there is no doubt of their existence, even if nothing is reported. One example is subconscious priming, when a word is flashed so briefly on a screen that it is not seen but can still be correctly reported. Other examples are masked perception, inattention blindness and diverted attention, all of which block conscious vision but the stimulus can be correctly reported later. That is all that is required by a bee that remembers a route and a place, but is not interested in pattern perception. Classically, subliminal perception was regarded as an automatic process that was independent of consciousness, and perhaps that is the way we might think about bee vision.

In humans, some brain lesions (not retinal lesions) cause a situation called 'blind-sight', in which the subject has no conscious vision in a part or whole of an eye, but is able to report correctly a strong stimulus such as a colour, a black spot or a large familiar object and its position. Perhaps it means something to suggest that bee vision is all blind-sight and therefore not cognitive by Pye's definition!

Retrospect

The idea, which persisted for 100 years, that pattern perception was based on the reassembly of a central image laid out in the brain served the bee badly. The inferences of cognitive analysis of visual images by bees were compatible with the original data and in line with the general theories of the time, but

some results were explained by merely describing the performance in different words and then guessing higher processes, so causing years of confusion. As the experimental testing of trained bees progressed, the local cues offered a low-level mechanistic explanation but the detection of the spatial layout of the pattern was not ruled out because the patterns were huge and overlapped several local eye regions. Contemporary publications, moreover, added new conclusions without cancelling the old ideas of cognition.

We can now infer that the training procedure limited the bee's vision. For the task in hand only, they learned to ignore all except a few cues in a local area of the eye. They generalised because they recognised the few cues they had learned and no unexpected cues were detected. Further discoveries, however, are never ruled out. After all, humans have a sensory processing system that depends entirely on peripheral arrays of simple feature detectors and a distributed administration.

Endnotes

1. To understand the depth of the divide between intuition and empirical methods, or between ethology and the mechanistic analysis of behaviour, this chapter should be read in conjunction with Chapter 2.