



Visual processing of the bee innately encodes higher-order image statistics when the information is consistent with natural ecology

Richard Clarke, R. Beau Lotto *

Institute of Ophthalmology, University College London, 11-43 Bath Street, London EC1V 9EL, United Kingdom

ARTICLE INFO

Article history:

Received 13 March 2008
Received in revised form 27 February 2009

Keywords:

Vision
Colour
Natural image statistics
Neuroscience
Ecology
Empirical vision

ABSTRACT

Determining the statistical relationships of images that facilitate robust visual behaviour is nontrivial. Here we ask if some spatial relationships are more easily learned by the visual brain than others. Visually naïve bumblebees were trained to recognise coloured artificial flowers in scenes of equal spatial complexity but differing patterns of stimulus intensity. When flowers of similar intensity were grouped into extended regions across the array (coincident with natural patterns of light), the accuracy of the bees' foraging behaviour was dependent on spatial context, even though this information was redundant to the task. When the same intensity information was organised into a pattern that was less consistent with natural patterns of illumination but of equal order, their behaviour was independent of spatial context and they required double the training time to solve the same conditional task. These observations suggest the brain is biased to more efficiently encode/learn ecologically 'meaningful' image correlations.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

How the visual brain encodes the patterns of light that fall onto its receptors, and turns these encoded images into useful behaviours is a question of general interest to the neuroscience and machine vision communities. The ambiguity of visual information makes this a complex process: at the most fundamental level, the spectral quality of the part of a retinal image arising from a surface in the world is determined by both the surface's reflectance and its illumination. This means the problem of vision is underspecified, since the absolute quality of each single part of an image may be the result of an infinite number of real world configurations. Thus the brain must use the relationships between elements of a visual image to generate robust visual behaviour, since it is the relationships between stimuli that remain stable. But which relationships does the visual brain innately use?

Here we consider whether the naïve visual brain is equally able to learn any contextual relationship or whether it predisposed to learn *certain spatial* relationships (perhaps those that were common or 'meaningful' during its phylogenetic history). Using the bee, an important system shown to be relevant for addressing general questions of colour vision, we first ask whether vision innately encodes spatial image correlations (when it is not required to do so) and secondly whether certain relationships are encoded more easily than other relationships.

We set bumblebees a biconditional discrimination task, where the *combination* of brightness and colour in a visually distinct object (a circular stimulus which is best conceived of as an artificial 'flower') unambiguously specified whether a stimulus was a target (i.e., rewarding), or a non-target (un-rewarding). Bees are known to solve cue-combination tasks of this complexity (Schubert, Lachnit, Francucci, & Giurfa, 2002), and in experiments of this kind bees' performance is consistent with 'unique cue' theory (whereby the bee assigns a unique identity to the compound as a result of the interaction of its components; Whitlow & Wagner, 1972; Deisig, Lachnit, Sandoz, Lober, & Giurfa, 2003). Thus the null hypothesis is that the bees in our experiments will simply associate a stimulus' specific qualities – in this case intensity and colour – with a reward, independent of the information presented in its surrounding spatial context.

On the other hand, it has been known for some time (Katz, 1935; see review by Maloney and Schirillo (2002)) that articulation – by which we mean the presence of multiple, distinct, chromatic surfaces – can alter human perception of stimuli. More recent studies have show that whilst perceptions of brightness contrast (and constancy) are indeed augmented by articulation, what matters in human perception is not articulation per se, but the specific relationships between stimulus elements in an articulated image (Lotto & Purves, 1999, 2000 and references therein). The explanation offered in this latter case was that these specific correlations provide additional statistical information about the source of the underlying scene, giving the visual brain more information on which to make a judgment as to the real world source of an inherently ambiguous stimulus.

* Corresponding author.

E-mail address: Lotto@ucl.ac.uk (R.B. Lotto).

Here then, in an array of 64 visually discrete surfaces, we organised the conditional intensity cue into large regions in space – i.e., we simply put randomly coloured surfaces of *one* relative intensity next to each other over an extended region, forming regions of lighter and darker flowers of variable colour. Note that there was significant overlap in the lightness of different regions: inherent lightness differences in the different flower colours meant that some individual flowers in the dark region were in fact lighter than some flowers in the light region. If the bees adopted a simple unique cue strategy, the spatial correlation of the brightness cue should be ignored since encoding this additional (low frequency) spatial contextual information was not necessary to be successful in the task.

We present data showing that, when presented with the arrays described, bees innately encode the spatial correlation in the brightness relationships, and indeed make it central to their learning strategy. We further show that certain contextual relationships are more directly accessible to the mechanisms of visual processing than other contextual relationships that are equivalent in terms of order or structure. We speculate that the more rapidly and robustly learned brightness/colour contextual cues are those that are most consistent with probabilistic relationships between natural images and scenes (Rubin & Richards, 1982; Cavanagh, 1991; Mullen & Kingdom, 1991; Párraga, Troscianko, & Tolhurst, 2000; Fine, MacLeod, & Boynton, 2003), and as such, this bias enables the bee, and indeed all visual systems to overcome the inherent ambiguity of sensory data quickly and robustly.

2. Results

Bees were housed within a nest-box connected to a 1 m³ Plexiglas flight arena (see Fig. 1A), and were trained to forage from stimulus arrays shown schematically in Fig. 1C (see Section 4 for full protocol). Each array to which the bees were trained was made up from stimuli, each kind of which was one of four colours.

In human terms, these approximated to orange, yellow with an admixture of UV, blue and purple (but of course were presumably perceived quite differently by the bees). These colours were then paired with one of two intensities. Flower colour was controlled with one layer of gelatine filter that altered the spectral quality of its transmitted light, and intensity was controlled using neutral density filters that reduced intensity by 25% whilst minimally altering its spectral characteristics (see Table 1). By combining each chromatic filter with each intensity filter, eight distinct stimuli were possible, and each was presented in every training session. Furthermore, all flower colours were uniformly distributed across the array (so that the average colour of the entire array was constant) and their location within the four quadrants of the array was random, thereby eliminating any possible spatial cue that the bees could have used to obtain a reward. During training, only two kinds of flower were rewarding (which we will refer to subsequently as ‘targets’). The remaining were unrewarding (which we will call ‘non-targets’). Targets were always the ‘blue’ flowers paired with the brighter intensity (i.e., light that passed through one ND filter, as well as a blue flower filter), and the ‘yellow-UV’ flowers paired with the darker intensity (i.e., light that passed through two ND filters, as well as the yellow flower filter). These rewarding stimuli remained constant for all experiments described here – only the spatial configuration of intensity filters was altered.

2.1. Experiment 1: trained to homogenous quadrant intensity, tested on homogenous quadrant intensity

In the first training condition, flowers were grouped by their intensity into quadrants of 16 flowers that were either ‘bright’ or 16 flowers that were ‘dark’. Two quadrants of bright and two quadrants of dark flowers were always present, though the location of the bright and dark quadrants was random. For a schematic representation of one of these possible arrays, see Fig. 2A. Bees were trained as a group for an average of 13 foraging cycles per bee




Fig. 1. (A) The Bee Matrix. Light from the Matrix light-box are generated by six Reptistar 5.0 fluorescent tubes placed behind a sheet of Plexiglas, and a UV-transmitting diffuser (Rosco No. 216). (B) The flowers are cylindrical in shape, measuring 20 mm in height, and are decorated with Plexiglas discs measuring 80 mm in diameter. The spectral quality of each flower is independently controlled with gel filters (Rosco, Germany). Bees feed on sugar water contained in a central chamber.

Table 1

Receptor activation of each of the five filters used in the experiments described. Spectral transmittance was measured with an Ocean Optics S2000 (Dunedin, FL) spectrometer relative to a calibrated deuterium/halogen radiation source DH 2000-CAL (Ocean Optics, Dunedin, FL). Measurements are in Watt/cm²/nm and had to be converted into quantum-based spectra. Spectral transmittance was then multiplied by the known receptor activation of UV, B & G receptors at that wavelength (Briscoe & Chittka, 2001) and summed over the whole visible range.

	UV	B	G
Y	44	39	468
O	7	7	274
B	16	250	311
P	56	222	257
NDF	53	200	365

(see Section 4 for full protocol), and then tested individually on a novel stimulus configuration, conforming to the constraints described above. The results of this control test are the following. Eighty three percent of all visits were to previously rewarding stimuli, split evenly between the bright blue and dark yellow targets – the latter proportion was consistent for all subsequent tests unless stated otherwise. Eleven percent of visits were to non-rewarding colours (purple and orange) and 6% of visits were to previously rewarding colours, but under the ‘incorrect’ intensity. Since it is well known that bees are able to use colour as a reliable foraging cue, in reporting our results we extract visits to flowers of the incorrect colour. Thus, the ‘conditional response’ reported here is: the probability that the bee selected the correct intensity given that they landed on either a blue or yellow flower. The bees generated a conditional response of 92.5% accuracy (see Fig. 2A ‘Results’; in all tests a random conditional response would result in 50% accuracy). No area of the array was favoured, a finding consistent with previous work using this experimental paradigm (Lotto & Wicklein, 2005). In order to test this, firstly stimuli were considered as “boundary stimuli” and “non-boundary stimuli”. The former were immediately adjacent to a panel of different lightness, the later were at least one stimulus removed in any direction. Note that stimuli adjacent to the edge of the array were not considered to be “boundary stimuli”. In this first test there were 41 selections of blue or yellow flowers – 38 of which were to the correct colour – lightness combination, and three of which were incorrect. In these testing arrays, there was a ratio of 1:1.91 boundary stimuli to non-boundary stimuli (all colours and lightness). Of the 41 selections, 13 were to boundary stimuli (unbiased selections would produce an expected value of 14.09), 28 were to non-boundary stimuli (expected value: 26.9). Thus the null hypothesis that the bees selections were affected by the boundary between lightness regions could not be rejected at $P > 0.01$ (binomial test gives $P(\geq 28 \text{ boundary selections}) = 0.4299$). Thus bees did not appear to be favouring selections on the boundary or otherwise. Because the bees were so accurate, the influence of examining correct selections (as opposed to just selections) was minimal (of the 38 correct selections, 26 were to boundary stimuli (1 mistake was made on a boundary), giving an expected value of 24.94 and $P(\geq 26 \text{ boundary selections}) = 0.4314$).

Secondly stimuli were considered as “edge” stimuli (i.e., those immediately adjacent to the outer edge of the arena) and “non-edge” stimuli. There was now a ratio of 1:1.29 edge to non-edge stimuli in the testing arrays. Of the 41 selections, 19 were to edge stimuli (expected value 17.94, $P(\geq 19 \text{ edge selections}) = 0.4273$ (binomial test)). The three mistakes were all made on non edge stimuli; though again, lack of data rendered this of indeterminate significance.

These data show that the bees learned to differentiate the intensity of the previously rewarding flowers from unrewarding flowers of the same colour ($P < 0.01$, see Section 4 for details of statistical

tests). In other words, the bees were able to learn to differentiate light blue (rewarding) from dark blue (un-rewarding) flowers, and dark yellow (rewarding) from light yellow (un-rewarding) flowers. The next test determined the nature of the information used to differentiate rewarding from unrewarding flowers.

2.2. Experiment 2: trained to homogenous quadrant intensity, tested on randomised intensity

In the second test, the same trained foragers were presented with an array of flowers in which the *spatial location* of the two intensities was random (see Fig. 2B). For brevity we will term this test condition as the ‘uncorrelated intensity array’, since the statistical correlation between the ‘dark’ and ‘light’ flowers no longer existed given that the two intensities were distributed randomly across the array. More formally, in this condition the higher-order spatial correlation of intensity was eliminated, leaving the absolute chromatic quality and intensity of each individual flower making up the array identical to that during training. Under these conditions, in contrast to the control test described above, the conditional accuracy of selecting the correct intensity was now only 36%. This did not constitute significant behaviour to the intensity condition (i.e., the probability that it was random behaviour was greater than 0.01). We concluded, therefore, that when the spatial organisation of flower intensity that was presented during training was removed, the bees were no longer able to discriminate the intensity of previously rewarding flowers from unrewarding flowers, even when the absolute qualities of the rewarding stimulus remained unchanged. In other words, whilst the bees selected only blue and yellow flowers, they no longer distinguished light blue flowers (rewarding) from dark blue flowers (un-rewarding), nor dark yellow flowers (rewarding) from light yellow flowers (un-rewarding). The possibility that the bees failed to find the ‘correct’ flowers simply because of the novelty of the testing array with respect to the training array was controlled for in a subsequent test, which is described later.

2.3. Experiment 3: trained to ‘chequerboard’ quadrant intensity, tested on ‘chequerboard’ intensity

It is important to note that during training it was not necessary to encode stimulus intensity *relationships* between flowers within and between panels to solve the training task. The same result could have been achieved by simply selecting the ‘light blue’ and ‘dark yellow’ flowers wherever they appeared – independent of their surrounding context. To better understand why the bees encoded instead the spatial correlations between flowers in the training arrays, a second group of visually naïve bumblebees were trained to an array in which the spatial location of intensity was structured in a pattern of alternating intensity. Described another way, if any given flower was bright, then its neighbour on all sides would be dark, and, in turn, those neighbours bright, and so on (Fig. 2C). For simplicity, we will refer to this configuration as the ‘chequerboard array’ – for obvious reasons. Note that the extent to which intensity across the array in this condition is ordered is no less than the original training condition (experiment 1). The amount of information contained in the spatial correlation of this new training array was no less than that in the first training array (the spatial intensity information was simply represented in a different way), and so could – in principle – equally serve as a cue for deciding between blue and yellow flowers.

After an average of 12 foraging cycles, as close as possible to the point of first testing in the previous experiments described above, bees were tested on a novel configuration of flowers that conformed to the spatial organisation of the new training array. The conditional probability of selecting the flower colour under the ‘correct’ intensity was 58% (colour errors were, consistent with

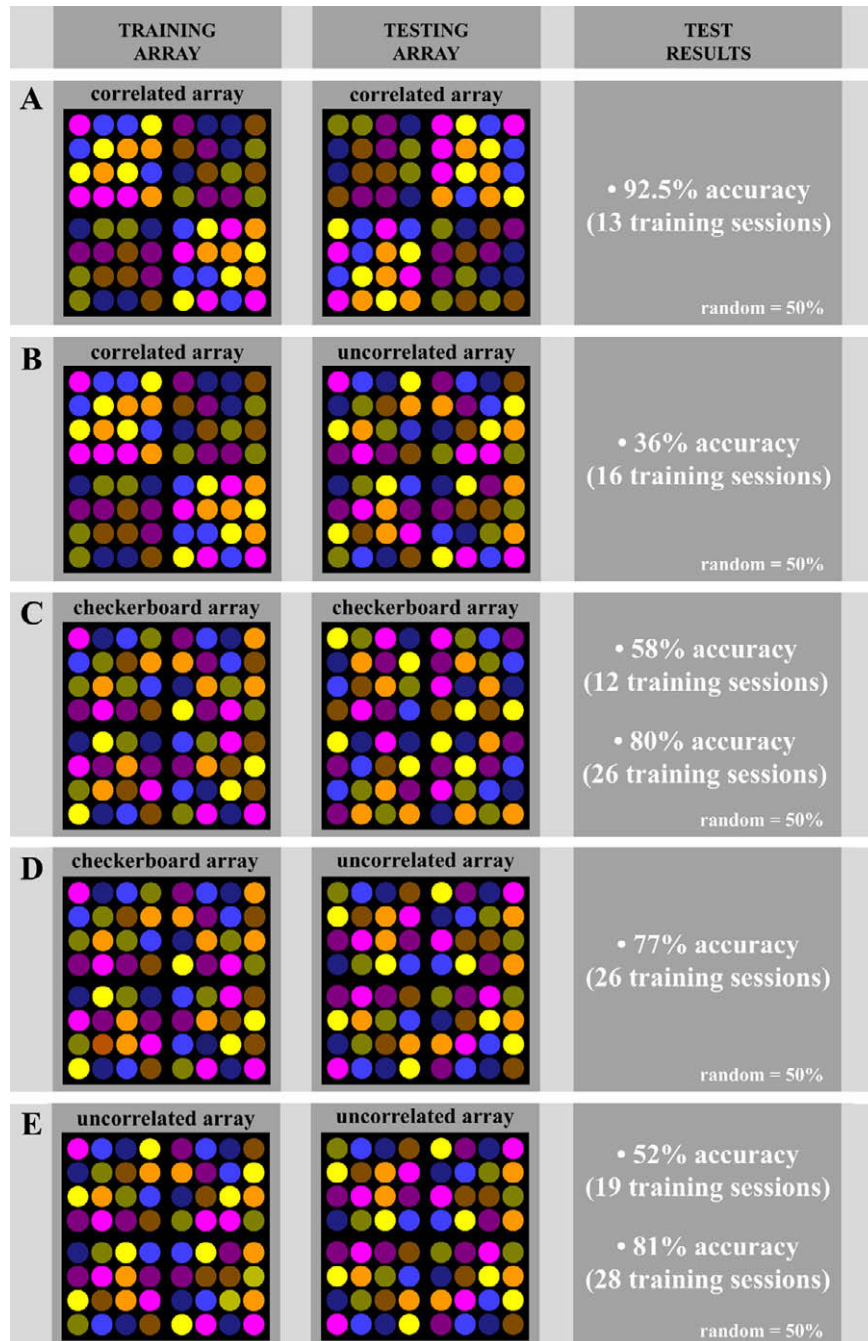


Fig. 2. Since bees are known to accurately find the target colour, results shown are conditional probabilities (i.e., given the target colour has been selected, the proportion of selections that are of the target illumination for each experiment). Bees trained to 'ecologically likely' intensity contrast: (A) Bees are able to use correlated brightness information to guide behaviour (see text for further description). (B) After approximately the same training experience, bees perform no better than chance if the brightness correlation is disrupted. Bees trained to 'ecologically unlikely' intensity contrast: (C) After 12 training cycles, bees trained to this checkerboard array of flowers are now unable to reliably solve the task (see text for further description). Increasing the training to 26 cycles results in similar levels of task performance to the ecologically likely arrays. (D) Once successful on this task, bees can generalise well to arrays containing uncorrelated intensity information, suggesting the strategy they employ does not use the correlation. (E) Similar behaviour towards the checkerboard arrays is observed when bees are trained to arrays of flowers whose intensity is organised randomly.

all previous experiments, around 12%). This does not demonstrate a significant level of response towards the intensity cue, demonstrating that after the same number of foraging bouts as the first experiment the bees were unable to solve the same training task (namely, to forage from bright blue flowers and dark yellow flowers). Training and testing continued. After 26 foraging cycles, the bees' conditional probability increased to 80%, which did not increase further with more training (Fig. 2C). This represents a significant response to the intensity stimulus ($P < 0.01$). Thus, when intensity was ordered in a checkerboard of alternating light and

dark flowers, only after twice the training experience did bees learn to use the intensity of flower as a conditional cue for behaviour. In addition, maximum selection accuracy remained lower by comparison with the previous experiment.

2.4. Experiment 4: trained to 'checkerboard' quadrant intensity, tested on randomised intensity

To determine if the bees also used the equally ordered spatial contextual cue to solve the task, as in the first experiment, this

second group of bees were tested on an array in which the spatial location of intensity was randomised across the array, which, again, eliminates the higher-order correlation of intensity across the array whilst preserving the absolute quality of each individual flower stimulus. Unlike the initial training paradigm, under these conditions the conditional probability of selecting the previously rewarding flowers remained significantly high: 77% ($P < 0.01$). Thus the strategy that the bees had adopted to solve the chequerboard arrays continued to work well even when the higher-order spatial correlations offered by the chequerboard pattern were eliminated, suggesting that the strategy was not dependent on this higher-order spatial organisation of intensity. This is exactly opposite to what was observed of the bees in training condition 1 (homogenous quadrant intensity).

2.5. Experiment 5: trained to randomised quadrant intensity, tested on randomised intensity

We next describe a series of further control experiments that confirm and extend the observations just described. In conducting these experiments a third set of foragers was trained to arrays in which the spatial location of stimulus intensity was randomised (Fig. 2E) (denoted 'randomised arrays' for brevity). Foragers were given the same number of training sessions as the two previous conditions, and tested on a novel stimulus configuration of the training parameters. After 13 foraging cycles, bees demonstrated a conditional probability of selecting the correct flower intensity, given the correct colour had been selected, 52% of the time, constituting a random response to the intensity cue. Training was again continued for a further 15 cycles per bee. Upon testing, the conditional probability now increased to 81% (further training did not increase this statistic). This constituted a significant conditional response towards the local intensity cue ($P < 0.01$). Note that the training required and accuracy achieved are nearly identical to when the bees were trained to the chequerboard arrays, and thus double that required to learn the same task when the bright and dark intensities were grouped into quadrants. Importantly, this consistency in the behaviour of two groups of foragers from different hives demonstrates the robustness of the experiments presented.

To review the findings thus far: when presented with flowers under spatially extended regions of intensity, bees used the higher-order information, even though this information was redundant to the training task. This may enable more robust behaviour to the stimuli (see control experiment 7). When presented with scenes of equal order, but in which intensity was spatially discontinuous, bees did not use the higher-order spatial information, but used instead the lower-order (local) cues. Whilst this strategy is, on the face of it, less complex, it took the bees twice as long to solve the same training task.

2.6. Experiment 6: trained to randomised quadrant intensity, tested on homogenous quadrant intensity

The foragers trained to the randomised arrays were next tested on arrays in which the spatial ordering of intensity into homogenous quadrants had been re-introduced (Fig. 3A). Under these conditions, the conditional probability of selecting the previously rewarding stimulus intensity was 88% – demonstrating that the task of finding the bright blue and dark yellow flowers remained tractable ($P < 0.01$). This means that the strategy used to solve the randomised arrays generalised to the single intensity per quadrant arrays, and makes it unlikely that the statistical novelty (testing array versus training array) in experiment two was the cause of the bees' failure to solve the task.

2.7. Experiment 7: trained to homogenous quadrant intensity, tested on globally darker homogenous quadrant intensity

To further examine the use of higher-order spatial correlations between intensities during training, bees trained to the homogenous panel arrays were presented with scenes statistically identical to those during training (i.e., lighter and darker intensity grouped into spatially extended regions), with the exception that an additional layer of neutral density filter (NDF) was added to each stimulus (Fig. 3B). Thus, the regions of higher intensity were now light passing through two NDFs (instead of one NDF as during training), and the regions of lower intensity were light passing through three NDFs (instead of two NDFs as during training). The relevance of this simple manipulation (i.e., adding a single layer of NDF to all stimuli) is that the lighter flowers during testing were physically identical (i.e., in both colour and intensity) to the darker flowers in training. If the bees encoded the absolute intensity of the rewarding flowers during training and/or did not compare the relative intensity of the groups of flowers across the entire array, then in this test they should forage only from the yellow–UV stimuli in the higher intensity panels during testing (i.e., the yellow–UV coupled with two NDFs), since this stimulus was identical to the rewarded during training. Yet of all visits to yellow flowers, 70% were to the lower intensity stimuli, whereas of all the visits to blue stimuli, 73% were to the higher intensity blues (significant behaviour in both cases to $P < 0.01$). Thus, the bees selected the blue flowers in the panels that were the lighter in the test array (though these same panels were the darker during training), and the yellow flowers in the darker two panels during testing, demonstrating that bees discriminated the relative intensities between panels across the array (otherwise there would have been no way to determine which panel was the lighter and which the darker). An important corollary of these observations is that the data are consistent with bees generating 'lightness constant' behaviour (which has never previously been demonstrated) whilst simultaneously using the intensity as a conditional cue for behaviour.

2.8. Experiments 8 and 9: trained to use orientation of a contour as the conditional cue

A final control experiment was undertaken that effectively repeated the original experiment, but using a different conditional cue: line orientation rather than intensity. In this experiment, horizontal and vertical bars – instead of stimulus intensity – provided the conditional cue for colour-mediated behaviour. That is, in direct parallel with the previous experiments, a predominantly yellow flower coupled with a horizontal dark bar, and a predominantly blue flower coupled with a vertical bar, were the only rewarding stimuli. The horizontal and vertical bars were grouped into panels, providing a spatially correlated conditional cue (exactly like the light and dark intensity cue in the first training session; Fig. 3C). Bees were trained to this format as before. After an average of 11 foraging cycles, bees were tested under conditions consistent with their training. Selection accuracy was 70% (statistically significant at $P < 0.01$). Mistakes were exclusively colour errors, with no visits to the correct colour coupled with the incorrect bar orientation (orientation errors) meaning that the comparable conditional probability was 100%.

To test whether the bees similarly encoded the higher-order spatial information provided by the spatial organisation of orientation, the bees were next tested with an array in which the orientation cue was randomized across the array (Fig. 3D). As before, this manipulation preserved the absolute qualities of each flower's stimulus, whilst eliminating the higher-order spatial correlations across the array. Despite this, selection accuracy remained statistically significant at 90%. Thus, even though the conditions were

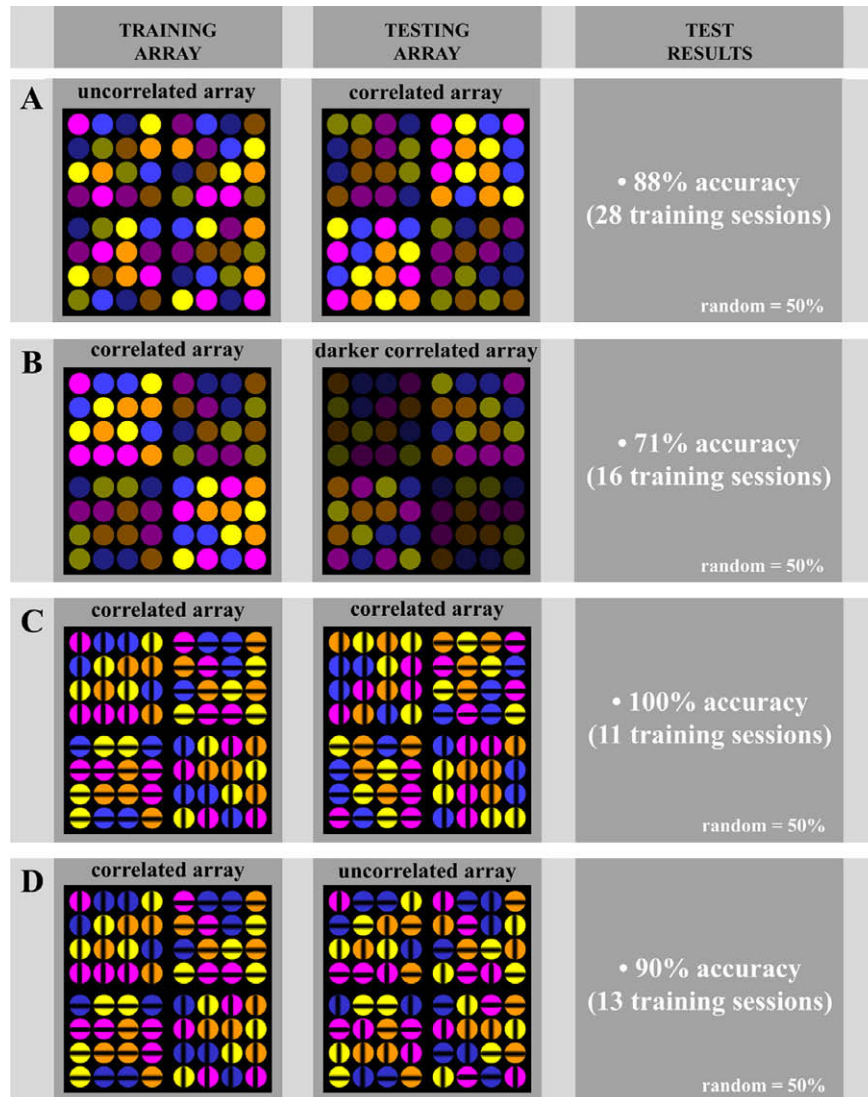


Fig. 3. Control experiments. (A) Bees trained to the flowers with random spatial intensity are tested on arrays in which the intensity cue is highly correlated, which does not disrupt their foraging behaviour (see text for further explanation). (B) Changing the overall intensity of the floral array (by adding an additional neutral density filter to each quadrant) does not alter the bee's behaviour, demonstrating a high degree of lightness constancy (see text for further explanation). Thus, the absolute physical nature of stimuli (local cues) is ignored and bees learn the global relationships between the stimuli, indicating an ability to display brightness constant behaviour. (C) and (D) Arrays in which the conditional stimulus intensity cue has been replaced with an orientated line. Thus, rather than selecting light blue and dark yellow flowers, they must selected blue flowers when associated with a vertical bar and yellow flowers when it is associated with a horizontal bar. Under these conditions, the bees do not use the spatially correlated contour information to solve the training task (see text for further explanation).

identical to the original experiment, when the conditional cue is orientation rather than intensity, the bees learned to differentiate the rewarding flowers *independent of spatial context*, i.e., without reference to the spatial correlation of contours across the array.

3. Discussion

Previous work on bee vision has shown that the task presented here – to use a subtle intensity cue (a brightness ratio of 1:1.33) as a conditional stimulus for colour behaviour – should be difficult for the bees to perform. This is because when bees are presented with surfaces that are large enough to be seen by their colour system they are unable to distinguish between stimuli on the basis of their intensity alone (Backhaus, Menzel, & Kreissl, 1987; Giurfa, Vorobyev, Kevan, & Menzel, 1996; Vorobyev & Osorio, 1998) except when the intensity difference is very large (Hempel de Ibarra, Vorobyev, Brandt, & Giurfa 2000). We examine the evidence for this conclusion in more detail below.

In honeybees, a stimulus presenting both chromatic contrast and L-receptor contrast (the L-receptor is exclusively used in the bee's achromatic visual system) needs to subtend a visual angle of more than 5° to be detectable. For a stimulus presenting *only* chromatic contrast and no L-receptor contrast, this minimum visual angle increases to 15° (Giurfa et al., 1996). Stimuli that present no chromatic contrast and only L-receptor contrast could *not* be learned (i.e., could not be 'seen') when stimuli subtended large visual angles (Giurfa et al., 1996). It was therefore concluded that the achromatic system was not used when stimuli subtended more than 15°. The idea that achromatic and chromatic systems are used independently when approaching flowers was undermined somewhat by other work: firstly the fact that the presence of *chromatic* contrast considerably enhances the detectability of small targets (subtending 5°; Giurfa et al., 1996) and secondly, honeybees were found to be able to use *achromatic* cues at visual angles > 15° (Hempel de Ibarra et al., 2000). However in the latter experiments, very high achromatic contrast levels (1:11) were used, which was

offered as the explanation for the apparent simultaneous access to both chromatic and achromatic information. Recent work has contrasted the ‘psychophysics’ of the honeybee and the bumblebee. Dyer (2008) found that the minimum visual angles for which stimuli containing both L-receptor contrast and chromatic contrast was 2.3° (compared with 5° in honeybees), and stimuli containing only chromatic contrast could be detected at 2.7° of visual angle (15° for honeybees). Thus, the spatial resolution of the chromatic system of bumblebees is nearly 5-times more acute than that of honeybees (though the trade off appears to be the bumblebee’s ability to differentiate the chromatic differences of targets: this is somewhat poorer than the honeybee’s). Like honeybees then, bumblebees also could not learn or performed poorly on stimuli that presented only (or predominantly) L-receptor contrast at large visual angles (10 – 25°). From this it was concluded that colour at these larger angles is the most important factor for stimulus detection (Dyer, 2008).

In experiment 1 flowers of similar intensity were placed in spatially correlated locations during training: i.e., quadrants of flowers that collectively subtended visual angles larger than 30° – when viewed from the back of the arena – were either light or dark (see Fig. 4). In the lighter quadrants only the blue flowers were rewarding, and in the darker quadrants only the yellow flowers were rewarding. Thus, the task presented to the bees seems simple: To differentiate light blue from dark blue flowers and dark yellow from light yellow flowers. Two other colours of flower (both light and dark) were also presented, though unrewarding. When tested under similar conditions as their training environment, each bee’s preference for the previously rewarding flowers was high, demonstrating that they rapidly learned to use the subtle intensity differences (ratio of 1:1.33) between the flowers as a conditional cue for selecting flowers of different colour. Furthermore, no area of the arena was preferred. In particular, accuracy at the boundary between quadrants was no different than elsewhere (see Section 4 – experiment 1 and Lotto & Wicklein, 2005).

Experiment 2 tested whether the bees used the intensity of individual flowers or the intensity of a local quadrant of flowers to solve the training task. This was achieved by eliminating the spatial correlation between flower intensity by placing the light and dark flowers of different colours at random locations across the array. As such, the actual physical stimuli arising from each flower was identical to the training conditions. Only the spatial correlation between flower intensity was removed. The result was that bees no longer distinguished between the previously rewarding light blue flowers from unrewarding dark blue flowers, and previously rewarding dark yellow flowers from unrewarding

light yellow flowers. Thus it appears that the bees used the spatial correlation of intensity as a conditional cue for behaviour.

However, it might be argued that the bees ‘failed’ to recognise the previously rewarding targets in experiment 2 simply because of the statistical novelty they were confronted with when presented with these randomized testing arrays (put another way, the array in experiment 2 was simply too dissimilar for the bees to ‘know what to do’). This possibility was directly tested in experiment 5 and 6, where a new set of bees were trained and tested in a manner exactly opposite to experiments 1 and 2. In training the bees experienced arrays in which stimulus intensity was randomized across the array (as in the test arrays of experiment 2). When, in experiment 5, they were tested on a randomized array, they were able to differentiate light blue from dark blue flowers, and dark yellow from light yellow flowers—though doing so took twice the number of training sessions compared to when intensity information was correlated (an important observation which is discussed further below). When these same bees were presented with intensity information organised into quadrants (experiment 6), the performance of the bees remained high—i.e., they continued to prefer light blue and dark yellow flowers. Because the physical difference between arrays in experiments 5 and 6 is exactly the same as the difference between arrays in experiments 1 and 2, we conclude that the decrease in performance in experiment 2 (as compared to experiment 6) is not because of the novelty of the stimulus configurations presented in that test.

We further tested whether bees encoded the absolute qualities of each quadrant of flowers or their relative intensities. In experiment 7 the stimuli arising from the flowers in the lighter quadrants were now physically identical to the darker quadrants of flowers during training. Despite this physical equivalence, in the now lighter panels of flowers, the bees selected the blue flowers, not the yellow as they would have done during training. Thus, in both this and the first experiment, the bees must have encoded not the absolute qualities of the panel’s intensity during training, but the relative intensity across the whole array (though we cannot rule out receptor adaptation as an explanation for the latter observation). Whilst previous studies have demonstrated that bees can generate colour constant-like behaviour (Neumeyer, 1981; Werner, Menzel, & Wehrhahn, 1988), and our own studies have shown that bees can use the colour of illumination as a contextual cue for surface-colour mediated behaviour (Lotto & Chittka, 2005; which was subsequently confirmed by Dyer, 2006), to our knowledge this is the first demonstration of lightness constancy in insects that is simultaneous for surface-colour and illumination intensity (see also Lotto




Fig. 4. The size of the flight arena is such that the angle subtended by a single flower stimulus cannot be less than 5° – the limit at which the bees can detect and discriminate stimuli. According to previous research (Giurfa et al., 1996; Hempel de Ibarra et al., 2001) chromatic cues from a single flower stimulus can only be used in the 1/3 of the flight arena that is closest to the stimulus wall. Elsewhere, only achromatic cues are used. The stimuli involved in low spatial frequency brightness contrast must subtend an angle of greater than 30° wherever the bee is in the flight arena.

& Chittka, 2005). Together, then, the above experiments suggest that bees encode the relative intensity of spatially correlated information – when it is present – to solve a colour-mediated training task, even when doing so is, in principle, unnecessary to solving the task.

We next asked whether bees could use any spatially correlated intensity data. To do so a new set of bees were trained to arrays in which the intensity component of each flower was arranged into a “chequerboard” pattern across the floral matrix (experiment 3). It is important to stress here that such chequerboard arrays are *no less structured* than the homogenous arrays presented in experiments 1 (i.e., there are no more degrees of freedom and hence, as a simple measure, the entropy of the arrays in experiment 1 and in experiment 3 are equal). Despite this equivalence, we first observe that bees required more than double the training experience to solve the problem (and solved it to a lower maximum accuracy). In fact the results are essentially identical to experiment 5 where the bees were presented with intensity information randomized across the array. We also observe that, when bees are confronted with the randomized intensity array (experiment 4) which disrupts the spatial correlations across the array without altering the absolute qualities of the stimuli themselves, the bees continued to perform well. This means the spatial correlation of the chequerboard present during training was not used by the bees, suggesting it is not *any* spatial structure, but a particular type of structure that facilitates visual (colour) learning.

Why, then, does a particular spatial configuration of intensity information facilitate visual learning of colour intensity? Here we speculate that studies on human colour perception suggest an answer, namely the ecological relevance of the articulated information presented in the first training array. Previous studies in humans show that perceptions of brightness contrast (and constancy) are augmented by articulation (stimuli composed of multiple visually distinct elements defined by a change in contrast at the

element’s boundary) (Lotto & Purves, 1999, 2000 and references therein). Importantly, however, it has been argued that the effects of articulation on humans are not due to articulation per se but to the ecological relevance of the information presented. Thus, equally articulated scenes of similar spectral average will generate different perceptions of brightness and colour depending on whether the statistical content of the information presented is more or less consistent with one combination of reflectance/illumination or another (Lotto & Purves, 1999; Lotto & Purves, 2000). The explanation offered for these observations was that each element of an articulated image provides statistical information about the source of the underlying scene.

Whereas previous studies on bees probed *behavioural* colour space using large *uniform* targets presented in either y-maze (dual choice) (Giurfa et al., 1996; Brandt & Vorobyev, 1997) or spatially uncorrelated experimental paradigms (Backhaus et al., 1987; Hempel de Ibarra, Giurfa, & Vorobyev 2002), our experiments used spatially correlated, articulated scenes. We suggest that parsing intensity stimulus information into panels of multiple chromatic stimuli results in a correlation more like that in the ‘real world’ than the chequerboard arrays, if we interpret that correlation as a change in illumination. The basis of this idea is that in natural scenes, pure or near-pure luminance variations mainly arise from inhomogeneous illumination such as shadows or shading, and tend to change more slowly than chromatic variations (Kingdom, 2003; Rubin & Richards, 1982; Cavanagh, 1991; Mullen & Kingdom, 1991; Fine et al., 2003). Put another way, illumination tends to vary at lower spatial frequencies than surface reflectance. Thus, the first training condition with a lower spatial frequency of intensity contrast (with respect to the chromatic ‘surface’ changes) was more consistent with the natural environment: e.g., a shadow cast over an extended area made up of multiple surfaces of different reflectance. The chequerboard training condition was no less ordered and contained an equal number of individual surfaces –




Fig. 5. Bees are presented with an array of flower-like stimuli in which the relationship between illumination and surface reflectance is made (A) consistent with natural images, (B) less consistent with natural images (though importantly *no less structured* than (A)) or (C) less consistent with natural images and entirely decorrelated. We show that *only* in the case where the illumination–reflectance relationship is ecologically likely (A) do the bees quickly learn to use brightness as a conditional cue for colour selection. Additionally, when trained to this array, the spatial correlation is an integral part of the bees’ innate, context-dependent learning strategy, as their behaviour is completely independent of the actual brightness of individual stimuli. When the conditional cue is changed (from illumination to orientation), no difference is seen in training times. Since this abstract relationship has no parallel in natural ecology, this latter finding further indicates that the bee’s brain is biased to encode the statistics of images that would have facilitated robust visual behaviour towards the sources of images in the past.

i.e., equally articulated. But since the spatial frequency of intensity change was equal to that of colour change across the array, the higher-order structure could be argued to be less consistent with the natural environment. Hence we suggest that, because it is behaviourally useful to retain information in stimuli that they have experienced previously (quite apart from its source) the bee's visual system is sensitive to encoding (in this case) intensity information across articulated stimuli that are spatially correlated. (see Fig. 5). Thus bees can more easily learn to use stimulus cues when those cues are organised in a way that is natural, though, of course, in the experiments above there is no evidence that they actually recognise the lightness changes to be illumination changes.

A limited test of this more general hypothesis was provided by experiment 8, in which the conditional cue was not intensity information but the orientation of a bar: bees were required to pair a colour selection with an orientation cue rather than the brightness cue. That the bees were equally able to solve arrays where the orientation cue was arranged in a chequerboard pattern and where it was parsed into panels is at least consistent with the 'ecological relevance' hypothesis discussed above, in that there is no systematic bias in the environment for spatially extended (in two dimensions) regions of identically oriented lines rather than a chequerboard arrangement. In other words when the link between conditional cue and spatial order was removed, the bees performed equally well on both training conditions.

In summary, the evidence that bees have only a 2D 'perceptual' colour space is weakened by the ease with which the bees in our experiments were able to use subtle intensity changes as a conditional cue for colour selection if and only if intensity information was (i) articulated and (ii) spatially correlated into extended regions. Thus, information about stimulus intensity and colour may not be processed separately in the bee's brain, but may be combined into a 3D perceptual colour space, as in humans. Second, the data demonstrate that the naïve bee's visual system preferentially encodes the higher-order intensity information, even when this information is redundant to a task that could have been solved using a simpler strategy of encoding local image information only. We speculate that the reason for the latter observation is the ecological relevance of the information presented, in particular whether the image information is or is not consistent the statistical relationship between images and scenes experienced in nature (Lotto & Purves, 1999, 2000).

4. Materials and methods

Methods used for the bees experiments have been described previously (Lotto & Wicklein, 2005; Lotto & Chittka, 2005). Individually marked bumblebees were housed within a nest-box connected to a 1 m³ Plexiglas flight arena. On one wall of the arena, bees were presented with 64 trans-illuminated, Plexiglas stimuli or 'flowers' (Fig. 1A). Equal numbers of four differently transmitting filters were used to colour the stimuli across the array, and the illumination intensity of each flower stimulus was controlled by neutral density filter(s) (NDF). Thirty-two of the flowers were illuminated with higher intensity light (one layer of NDF) and the other 32 with lower intensity light (two layers of NDF). The change in consequent light intensity preserved the colour of the stimuli as much as possible (see Table 1). All stimuli (S) in the array, therefore, represent the spectral transmittance characteristics of a flower's filter ($F1$) and an illuminant's filter or filters ($F2$), or $S = F1 \cdot F2$ (see Fig. 1C 'Stimuli'). In each bee experiment eight foragers were presented with eight

30 min training sessions, during which time the number of foraging bouts for each bee was recorded (one bout equating to leaving the hive, visiting 1–3 flowers, returning to the hive and disgorging the sugar solution). To eliminate spatial information during training, the locations of each flower and its illumination were changed between training sessions, within the statistical constraints of the given problem. To eliminate olfactory cues, the Plexiglas flower-stems were also washed between sessions. After training, individually marked bees, whose training histories were fully known, were tested in isolation in the arena for 5 min each (during testing, all flowers were unrewarding). Visits were recorded by hand only when a bee landed on a flower and attempted to obtain a reward by extending its proboscis into the flower's empty, central chamber, consistent with the bee's rewarded behaviour during training. Results from all foragers were combined, after first confirming that their individual responses were homogeneous across the population (at a 95% confidence level; using a Chi-Squared test for homogeneity). Significance was then determined as a normal approximation to the binomial distribution.

Acknowledgements

We would like to thank Lars Chittka, Dale Purves, Daniel Hulme and Martina Wicklein for their assistance and conversation during this study. Bee research in the lottolab is supported by the EU, Leverhulme Trust, Wellcome Trust, EPSRC and BT.

References

- Backhaus, W., Menzel, R., & Kreissl, S. (1987). Multidimensional scaling of color similarity in bees. *Biological Cybernetics*, 56, 293–304.
- Brandt, R., & Vorobyev, M. (1997). Metric analysis of threshold spectral sensitivity in the honeybee. *Vision Research*, 37, 425–437.
- Briscoe, A. D., & Chittka, L. (2001). The evolution of color vision in insects. *Annual Reviews on Entomology*, 46, 471–510.
- Cavanagh, P. (1991). Vision at equiluminance. In J. J. Kulikowski, I. J. Murray, & V. Walsh (Eds.), *Vision and visual dysfunction: Limits of vision* (Vol. 5, pp. 234–250). Boca Raton, Florida: CRC Press.
- Deisig, N., Lachnit, H., Sandoz, J. C., Lober, K., & Giurfa, M. (2003). Modified version of the unique cue theory accounts for olfactory compound processing in honeybees. *Learning & Memory*, 10, 199–208.
- Dyer, A. (2006). Bumblebees directly perceive variations in the spectral quality of illumination. *Journal of Comparative Physiology A* 192, 333–338.
- Dyer, A. G., Spaethe, J., & Prack, S. (2008). Comparative psychophysics of bumblebee and honeybee colour discrimination and object detection. *Journal of Comparative Physiology A*, 194, 617–627.
- Fine, I., MacLeod, D. L. A., & Boynton, G. M. (2003). Surface segmentation based on the luminance and colour statistics of natural scenes. *Journal of the Optical Society of America A*, 7, 1283–1291.
- Giurfa, M., Vorobyev, M. V., Kevan, P. G., & Menzel, R. (1996). Detection of coloured stimuli by honeybees: Minimum visual angles and receptor specific contrasts. *Journal of Comparative Physiology A*, 178, 699–709.
- Hempel de Ibarra, N., Giurfa, M., & Vorobyev, M. (2001). Detection of coloured patterns by honeybees through chromatic and achromatic cues. *Journal of Comparative Physiology A*, 187, 215–224.
- Hempel de Ibarra, N., Giurfa, M., & Vorobyev, M. (2002). Discrimination of coloured patterns by honeybees through chromatic and achromatic cues. *Journal of Comparative Physiology A*, 188, 503–512.
- Hempel de Ibarra, N., Vorobyev, M., Brandt, R., & Giurfa, M. (2000). Detection of bright and dim colours by honeybees. *Journal of Experimental Biology*, 203, 3289–3298.
- Katz, D. (1935). *The world of colour*. London: Kegan Paul, Trench, Trubner & Co.
- Kingdom, F. A. A. (2003). Color brings relief to human vision. *Nature Neuroscience*, 6, 641–644.
- Lotto, R., & Chittka, L. (2005). Seeing the light: Illumination as a contextual cue to colour choice behaviour in bumblebees. *Proceedings of National Academy of Sciences*, 102(10), 3852–3856.
- Lotto, R. B., & Purves, D. (1999). The effects of color on brightness. *Nature Neuroscience*, 2, 110–114.
- Lotto, R. B., & Purves, D. (2000). An empirical explanation of color contrast. *Proceedings of National Academy of Science USA*, 97, 12834–12839.
- Lotto, R. B., & Wicklein, M. (2005). Bees encode behaviorally significant spectral relationships in complex scenes to resolve stimulus ambiguity. *Proceedings of National Academy of Science USA*, 102, 16870–16874.
- Maloney, L. T., & Schirillo, J. A. (2002). Color constancy, lightness constancy, and the articulation hypothesis. *Perception*, 31, 135–139.

- Mullen, K. T., & Kingdom, F. A. A. (1991). Colour contrast in form perception. In P. Gouras & J. Cronly-Dillon (Eds.), *Vision and visual dysfunction: The perception of color* (Vol. 6, pp. 198–217). Oxford: Macmillan.
- Neumeyer, C. (1981). Chromatic adaptation in the honeybee: Successive color contrast and color constancy. *Journal of Comparative Physiology*, 144, 543–553.
- Párraga, C. A., Troscianko, T., & Tolhurst, D. J. (2000). The human visual system is optimised for processing the spatial information in natural visual images. *Current Biology*, 10, 35–38.
- Rubin, J. M., & Richards, W. A. (1982). Colour vision and image intensities: When are changes material. *Biological Cybernetics*, 45, 215–226.
- Schubert, M., Lachnit, H., Francucci, S., & Giurfa, M. (2002). Nonelemental visual learning in honeybees. *Animal Behaviour*, 64, 175–184.
- Vorobyev, M., & Osorio, D. (1998). Receptor noise as a determinant of colour thresholds. *Proceedings of the Royal Society of London B*, 265, 351–358.
- Werner, A., Menzel, R., & Wehrhahn, C. (1988). Color constancy in the honeybee. *Journal of Neuroscience*, 8, 156–159.