The role of ‘hidden preferences’ in the artificial co-evolution of symmetrical signals

SETH BULLOCK AND DAVE CLIFF

School of Cognitive and Computing Sciences, University of Sussex, Brighton BN1 9QH, UK
sethb@cogs.susx.ac.uk, davec@cogs.susx.ac.uk

SUMMARY

Recently, within the biology literature, there has been some interest in exploring the evolutionary function of animal displays through computer simulations of evolutionary processes. Here we provide a critique of an exploration of the evolutionary function of complex symmetrical displays. We investigate the hypothesis that complex symmetrical signal form is the product of a ‘hidden preference’ inherent in all sensory systems (i.e. a universal sensory bias). Through extending previous work and relaxing its assumptions we reveal that the posited ‘hidden preference’ for complex symmetry is in reality a preference for homogeneity. The resulting implications for further accounts of the evolutionary function of complex symmetrical patterning are considered.

1. INTRODUCTION

A small but rapidly increasing number of studies (Arak & Enquist 1993, 1995a; Enquist & Arak 1993, 1994; Johnstone 1994; Hurd *et al.* 1995) have undertaken to explain evolutionary phenomena through the implementation of iterative computer programs, designed to model biological systems. These models differ from more traditional iterative numerical models in that they typically rely on modelling techniques pioneered within the artificial sciences (e.g. artificial neural networks, genetic algorithms, etc.).

Some debate has accompanied this work (Cook 1995; Johnstone 1995; Arak & Enquist 1995b; Dawkins & Guilford 1995), but attention to the methodology employed within it has been cursory. Here we replicate, discuss, and extend an evolutionary simulation model (Enquist & Arak 1994) concerned with the evolutionary function of naturally occurring symmetrical visual displays (e.g. the complex symmetrical patterns found on the wings of many birds). First, we present a brief account of Enquist & Arak’s (1994) argument advocating the interpretation of symmetrical signals as having evolved to exploit a universal sensory bias. Secondly, we describe our replication of their study and provide a discussion of their model’s validity. We hypothesize that networks evolved under the conditions employed by Enquist & Arak had preferences for structure (i.e. complex symmetry), only insofar as they were presented with patterns subjected to unnaturally structured sets of transformations during their evolution. Subsequently, extensions to their study which support this hypothesis and demonstrate that a different account of their proposed sensory bias must be accepted are described. The paper concludes with some discussion of the potential contribution that a sen-

sory bias for homogeneity might make to the evolution of complex symmetry.

2. SYMMETRY

Enquist & Arak (1994) attempt to show that symmetrical patterns are inherently favoured by perceptual mechanisms which must consistently classify objects despite viewing them from different angles, at different distances, in different orientations, and so on, due to the invariance of such patterns across such situations. Following Enquist & Arak, we will sometimes refer to this postulated disposition to favour symmetrical patterns as a (universal) sensory bias, in that it is proposed as an exploitable latent preference inherent within visual systems. However, we do not intend this to imply that these visual systems’ preference for symmetry need be established prior to, or in some way independently from, the associated signal evolution.

Enquist & Arak’s (1994) model an artificial visual system using an array of receptive elements (a 7×7 grid of idealized retinal cells each containing three receptors respectively sensitive to the red, green, and blue components of the light falling on their cell) and an artificial neural network consisting of 15 idealized neurons, each receiving 147 weighted inputs (one from each of the three receptors in each of the 49 retinal cells), which in turn excite or inhibit a decision neuron which responds ‘yes’ or ‘no’ according to whether the weighted activation from these ‘hidden’ neurons, when summed, and perturbed slightly by a random internal factor, rises above a certain threshold.

This visual system’s task was to accurately discriminate between a signal pattern and various distractor patterns (each represented as a 5×5 array

Figure 1. Depicting (a) the retina, (b) an example signal, Enquist & Arak's (1994) (c) translation, (d) 90° rotation, and (e) reflection transformations, and (f) the higher resolution rotation transformation we employ in this paper (note that in (c) four of the nine possible translations are depicted, and in (f) five of the 16 orientations are depicted).

of coloured pixels; each colour being represented by a vector in the red–green–blue colour space) despite each being presented to the system in various positions on the system's 'retina' and in various orientations. These positions and orientations were arrived at through subjecting the signal (and the distractors) to reflection, rotation, and translation transformations (see figure 1).

Artificial selection pressures were applied to such visual systems and the signals they attempted to classify (both of which were initially randomly configured). Repeatedly, the best of each (i.e. the network that discriminates signal from distractor under the most transformations and the signal that is discriminated correctly under most transformations) were allowed to 'asexually reproduce', creating new patterns and networks that were each similar to their parent but not identical due to low probability 'mutation' events which arbitrarily corrupted the reproduction process.

Under this artificial evolutionary paradigm impressive results were obtained. Patterns and networks co-evolved over evolutionary time, reaching '... a quasi-stable situation in which the network discriminates almost perfectly ... and the signal itself changes slowly ... [the signals] consisted of purer brighter colours than random patterns; they also displayed ... marked symmetries' (Enquist & Arak 1994, p. 171).

Prima facie these results are strong support for the hypothesis that symmetry 'may arise as a by-product of the need to recognize objects irrespective of their position and orientation in the visual field' (Enquist & Arak 1994, p. 169).

3. REPLICATION

In replicating this study we reimplemented the stochastic hill-climbing algorithm employed by En-

quist & Arak (1994)[†]. Initially a random network (with connection weights drawn from a uniform distribution $[-0.3, 0.3]$) and a random signal (colour components drawn from a uniform distribution $[0.0, 1.0]$) are generated. The algorithm then repeatedly carries out two consecutive periods of hill-climbing, one within the space of signal patterns and one within the space of retinal networks.

Hill-climbing is carried out in the following manner. A mutant of the current best network is generated by copying each of the best retina's network weights with probability 0.01 of a weight being perturbed by a normally distributed value (zero mean, s.d. = 0.02). The mutant is assessed with respect to the current best signal, and, if discovered to be fitter than the current best network, takes the best network's place, acting as the model for the next mutant. Once n mutants have been generated and tested in this manner, attention is switched to the best signal pattern.

The signal is tested with respect to the current best network and compared to a mutant which is generated by copying each colour component of the signal with probability 0.05 of perturbing it with a normally distributed value (zero mean, s.d. = 0.3). Perturbations which produced colour component values outside the legal range $[0.0, 1.0]$ were truncated to the nearest extreme value.

If a mutant signal is assessed as fitter than the current best signal, it takes the best signal's place and acts as the model for the next mutant. After n mutants have been generated attention reverts to the best retinal network and the process repeats.

Although the retinal network evolution is described first here, in order to avoid biasing our replication through preferentially evolving either the net-

[†] Copies of the code, and a version of this paper with colour figures, are available on request from the authors, or from: <http://www.cogs.susx.ac.uk/users/sethb/sym.html>

Figure 2. Five signals typical of (a) the 20 signals evolved under Enquist & Arak's (1994) presentation regime, and (b) the 20 signals evolved under a presentation regime comprising nine translation, and sixteen 22.5° rotation transformations. Each signal is shown as a colour composite and three colour separates in which the increasing intensity of each colour component is represented as increasingly heavy shading.

work or the signal first, the order in which the evolutionary 'steps' were carried out was counter-balanced across simulations.

Enquist & Arak dub one cycle of the above process (in which, arbitrarily, $n = 100$) a 'generation', and run their simulations for 500 such 'generations'. The simulations described here employed shorter run-times ($n = 50$ for 100 'generations') in response to time pressure, and the fact that there was no significant progress made by either the networks or the signals after the first 50 ($n = 100$) generations.

Fitness was calculated thus,

$$f(s, r) = \frac{1}{p} \sum_{i=1}^p \left[\xi(\phi_r(\lambda_i(s))) - \frac{1}{q} \sum_{j=1}^q \xi(\phi_r(\lambda_i(d_j))) \right],$$

where $f(s, r)$ denotes the fitness score ascribed to

signal and retina $\{s, r\}$ under an assessment regime comprising p presentations of different transformations of both signal s and q distractor patterns (d_j denotes the j th distractor pattern). The compound function $\phi_r(\lambda_i(x))$ denotes the output of retinal network r when presented with pattern x under transformation λ_i .

This output is calculated as follows. The output of a retinal receptor was equal to the intensity of its input, while that of a network neuron was calculated as a sigmoid function of y , the sum of its weighted inputs, thus

$$\text{output} = \begin{cases} 1 - \frac{1}{2}[1/(1 + y)], & y \geq 0, \\ \frac{1}{2}[1/(1 - y)], & y < 0. \end{cases}$$

The function $\xi(w)$ is an error function which smoothes the network output (w) by modelling a

Figure 3. Initially random signals were subjected to (a) Enquist & Arak's (1994) presentation regime comprising nine translation, and four 90° rotation transformations, and (b) a presentation regime comprising nine translation, and sixteen 22.5° rotation transformations. Graphs depict (left) plots ($\pm$ s.d.) of mean signal fitness, and rotational asymmetry metrics (four-fold and sixteen-fold) averaged over 20 simulations of 100 ($n = 50$) generations each, and (right) plots ($\pm$ s.d.) of mean signal boldness and mean heterogeneity metrics for the same 20 simulation runs.

stochastic internal variable (Θ) which perturbs the network output by a value normally distributed with zero mean ($\mu = 0$), and standard deviation, $\sigma = 0.02$.

Rather than draw a random value from this distribution for each trial, the effect of this internal factor was calculated in the limit, thus

$$\xi(w) = P(w + \Theta \geq \theta) = \frac{1}{\tau} \int_{\theta}^{\infty} e^{-(z-w/(\sigma\sqrt{2}))^2} dz,$$

where τ is a scaling constant used to normalize the function's output to between 0 and 1, while θ is a threshold value above which the network is said to have responded positively to its input. The value returned by the error function was thus a deterministic measure of the proportion of trials in which a network output value (w), after perturbation by an internal stochastic variable (Θ), would exceed the arbitrarily designated threshold value (θ).

For all simulation runs considered here, $\theta = 0.5$, $q = 3$ (one white, one black, and one random distractor).

Enquist & Arak's (1994) presentation regime was replicated. Each pattern was presented to a network under 36 different transformations (nine translations by four orientations, see figures 1c, d). The results of 20 simulation runs are displayed in figure 2a.

Bilateral asymmetry was calculated as the average geometric distance in the red-green-blue colour space between corresponding pixels. Rotational asymmetry was calculated as the mean of the standard deviations of the groups of pixels invariant under the rotation transformation. Brightness or boldness was calculated as the average geometric distance in the red-green-blue colour space between each pixel and the nearest corner of the colour space.

Evolved signals exhibited significantly more bilateral symmetry ($t = 6.526$, d.f. = 19, $p < 0.01$), four-fold rotational symmetry ($t = 9.638$, d.f. = 19, $p < 0.01$) and brightness ($t = 10.825$, d.f. = 19, $p < 0.01$), than random signals (see figure 3a). Under the conditions employed by Enquist & Arak (1994), symmetrical signals do indeed evolve.

4. DISCUSSION

Objects are not always seen at the same orientation, or from the same distance, etc. The transformations that Enquist & Arak's (1994) patterns undergo are intended to mimic these different ways in which the light reflected from a real-world object falls onto a retina. As Enquist & Arak's retinæ and patterns are square, reflections in the vertical and horizontal axes seem natural transformations to perform in order to create variation in the retinal image. However, natural visual systems are very rarely exposed to a pattern and then, subsequently, a reflection of that same pattern.

Rotations of 90° also seem natural transformations to perform upon square patterns. In response to such rotations, Enquist and Arak's simulation produces patterns with four-fold rotational symmetry. However, in reality, the rotation transformations that images undergo are *continuously* (although not necessarily uniformly) distributed between 0 and 360° .

Finally, if a signal's invariance under transformations is what makes it easy to discriminate, then biological signals which are selected for conspicuousness should ideally be one solid colour (i.e. maximally symmetrical).

To summarize, could it be the case that the networks evolved under the conditions employed by Enquist & Arak had preferences for structure (i.e. complex symmetry), only insofar as they were presented with patterns subjected to unnaturally structured sets of transformations during their evolution?

5. EXTENSION

In order to explore the effect of more continuous rotational transformations, we relaxed the rotation regime to include transformations that fell at 22.5° intervals (see figure 1f) rather than the 90° intervals employed by Enquist & Arak (1994). This quadrupled the number of presentation trials per assessment (i.e. $p = 144$), as, under each translation transformation, each signal was now presented in 16, rather than

Figure 4. Initially (a) random, and (b) bilaterally symmetrical signals were subjected to a presentation regime comprising nine translation, and sixteen 22.5° rotation transformations. In addition, signals depicted in graphs (a) were subjected to a bilateral reflection transformation. Graphs depict metrics as per figure 3.

four, different orientations. All other parameter values were held constant. The results of 20 simulation runs are displayed in figures 2b and 3b.

Evolved signals display significantly greater rotational symmetry, both sixteen-fold ($t = 5.07$, d.f. = 19, $p < 0.01$), and four-fold ($t = 3.3$, d.f. = 19, $p < 0.01$), and are significantly more homogeneous ($t = 4.87$, d.f. = 19, $p < 0.01$), than those evolved under the conditions employed by Enquist & Arak. They also feature significantly brighter colours than random signals ($t = 10.135$, d.f. = 19, $p < 0.01$). In a further 20 simulation runs, introducing a bilateral reflection transformation (see figure 1e), which doubled the number of presentation trials per assessment (i.e. $p = 288$), did not result in evolved signals significantly different in any way from those generated under this sixteen-fold rotational regime (see figure 4a).

Furthermore, seeding 20 simulations with initially perfectly bilaterally symmetrical signals also resulted in evolved signals not significantly different from those generated from random initial conditions, indicating that, as well as failing to *generate* complex symmetrical structure, the sensory bias presented by Enquist & Arak (1994) could not *maintain* pre-existing bilateral symmetries (see figure 4b).

Further experimentation reveals that not any homogeneous signal suits any network. Seeding 20 simulations initially with signals homogeneous in a random colour resulted in evolved signals not significantly different from those evolved from random initial signals under the same presentation regime. However, as networks evolve to prefer signals of above average brightness, the final evolved signals were significantly more bold than their earliest ancestors ($t = 10.734$, d.f. = 19, $p < 0.01$; see figure 5a).

Seeding 20 simulations with signals initially homogeneous in one colour chosen randomly from the six maximally bold colours that the networks were not penalized for favouring (i.e. one of green, red, blue, magenta, cyan, or yellow, but not white or black) again resulted in evolved signals not significantly different from those evolved from initially random signals (see figure 5b). However, the corner of the colour space occupied by the initial signal was not necessarily the same as that occupied by the final signal.

Thus, the networks were not merely favouring any homogeneous, bold signal, but initially had innate biases which could result in signals evolving from one bold, homogeneous pattern (e.g. solid green) to another (e.g. solid blue).

In summary, signal form was shown to be a result of pressure, not for symmetry *per se*, but for homogeneity coupled with arbitrary boldness (the nature of which being determined by the initial constitution of the network), which *coincidentally* results in signals which exhibit a high degree of both rotational and bilateral symmetry. As such, complex signal form cannot be explained as a product of the sensory bias that Enquist & Arak (1994) propose.

6. CLOSING REMARKS

Several points should be made about the scope and nature of this result. First, this result should not be taken to be an explanation for the presence of homogeneous patterning in the natural world. Although certain cases of bold, homogeneous coloration may well be adaptations exploiting the conspicuousness of such patterns, no general claim to that effect is being made here. The thesis of this paper is purely negative; complex, symmetrical patterning cannot be explained as the result of evolutionary pressure to exploit the universal sensory bias proposed by Enquist & Arak (1994).

Secondly, it might be thought that a presentation regime in which the orientation of signals was biased in some way could lead to more complex symmetry. Consider a presentation regime in which the orientation of the signal is drawn from a normal distribution centered on some arbitrary angle. This can be thought of as modelling a situation in which the signal has a typical orientation, but sometimes deviates from it slightly. Under such a regime, might bilaterally symmetrical signals evolve?

Enquist & Arak's model ensures that in order that one part of a signal comes to resemble some other part (i.e. that some symmetry develops), both parts of the signal must at different times fulfil the same role for the signal receiver. That is, both signal parts

Figure 5. Initially (a) homogeneous signals and (b) maximally bold homogeneous signals were subjected to a presentation regime comprising nine translation, and 16 22.5° rotation transformations. Graphs depict metrics as per figure 3.

must be projected onto the same part of the receiving retina.

If we choose a candidate axis of symmetry passing through the centre of a signal we see that while light reflected from signal areas on either side of this axis might sometimes fall on the same part of a receiver's retina, the chance of this occurring decreases as the angular distance between such signal areas increases. A presentation regime in which signal orientation is drawn from some normal distribution centred upon some typical orientation will ensure that light reflected from signal elements at $\pm 5^\circ$ from an arbitrarily chosen axis will fall on the same area of retina far more frequently than light reflected from signal elements at $\pm 45^\circ$ from the same axis. As a result, signals will gain only very limited benefits from any degree of bilateral symmetry. Since this is true for any axis passing through the centre of the signal, no one axis of symmetry has precedence over others, and reflectional or rotational symmetry will only be favoured in that it increases homogeneity (homogeneous signals still being easier to discriminate from distractors than more complex signals).

However, consider a probability distribution function governing signal orientation which is bimodal, with the angular distance between modes equal to 180° . Such a distribution, while again not affording precedence to one particular axis of reflectional symmetry, will induce two-fold rotational symmetry (although not at the expense of homogeneity) as signals which are invariant under a 180° rotation will be favoured. The orientation aspect of Enquist & Arak's presentation regime can be considered to be a probability distribution function with four equidistant modes. Enquist & Arak (1994) demonstrated that such a presentation regime resulted in signals exhibiting four-fold rotational symmetry.

However, it cannot be claimed that such multimodal probability functions underlie many natural signalling scenarios. For example, consider a patterned fish swimming from right to left across a potential mate's visual field. In order to present the same pattern at an angular offset of 180° , the fish would have to swim upside down (i.e. rotated through 180°). For the fish to provide a four-mode

display, it would have to spend time swimming vertically up and vertically down as well as right-side up and up-side down (or some other combination of four equidistant orientations), while spending comparatively little time swimming at the intervening orientations. Note that a fish merely swimming from right to left, and subsequently left to right, would not display a pattern at two orientations, nor a pattern and subsequently a reflection of that pattern, but simply two patterns (one from each side of the fish). While the patterns on each side of such a fish might be related, such a relationship is contingent rather than necessary, and is not modelled by Enquist & Arak's scheme which does not include morphological-developmental constraints on pattern form (see below).

It is not impossible to imagine the discovery of naturally occurring signaling behaviour with a complex stereotypical structure (e.g. an animal that displays a patterned body part at four distinct orientations) coupled with an associated discovery that the structure of this signaling behaviour is reflected in the structure of the displayed signal patterns. However, the complex structure of such signals could not be said to have arisen 'as a by-product of the need to recognize objects irrespective of their position and orientation in the visual field' (Enquist & Arak 1994, p. 169). On the contrary, the complex structure of such signals would have arisen *despite* such a 'need'. In such a case, the general pressure which favours homogeneity as a response to this 'need' would have been resisted through restricting the presentation of the signal in some specific (and unusual) manner. Such complex signal patterns could not be explained as the consequence of 'general principles of perception' (Enquist & Arak 1994, p. 169), but could only be explained through reference to specific facts about their evolutionary history.

Finally, a further consideration which might influence the expression of a bias for homogeneity is the degree to which the developmental or morphological constraints of a creature might interfere with the form of its signals. It is attractive to argue that a creature's bilaterally symmetrical patterning might be the result of constraints imposed by its bilateral

Bauplan. Could it be that a sensory bias for homogeneity might be expressed across a bilaterally symmetrical body as similarly bilaterally symmetrical patterning?

Augmenting Arak & Enquist's account of complex patterning with an appeal to morphological constraints is unsatisfactory as such an appeal must posit constraints which *resist* homogeneity. For example, if morphological constraints which impose bilaterality upon a creature's morphology can be shown to influence the creature's patterning such that the patterning is similarly bilaterally symmetrical, then any 'hidden preference' for homogeneity may be regarded as superfluous to an explanation of such complex bilateral symmetry. An adequate explanation may be based entirely on the aforementioned morphological constraints. Furthermore, a sensory bias for homogeneity will actually resist the development of complex bilaterally symmetrical patterning in favour of more homogeneous bilateral patterns (i.e. uniform coloration).

7. CONCLUSION

In conclusion, we have shown herein that the form of the signals that Enquist & Arak (1994) evolved was dependent on the presentation regime that the signals underwent. The signals exploited the structure inherent in the presentation regime rather than any 'hidden preferences' inherent in the artificial retinæ. We relaxed Enquist & Arak's unnatural regime so as to more closely approximate that experienced by natural signals. As a result, the interesting structure of their artificial patterns dissolved until bland homogeneity was reached. Such homogeneous patterns fail to support the hypothesis that complex symmetry 'may arise as a by-product of the need to recognize objects irrespective of their position and orientation in the visual field' (Enquist & Arak 1994, p. 169). Furthermore, a sensory bias for homogeneity such as the one considered here would actively resist the evolution of complex patterns (whether sym-

metrical or not). Despite this negative conclusion, alternative hypotheses linking the form of complex patterns to facts about sensory architecture remain healthy (see Osorio (1996) for an example of this kind of hypothesis).

We thank Anthony Arak, Magnus Enquist, Geoffrey Miller, and Daniel Osorio for discussions, and Marian Dawkins, Henrietta Wilson, and an anonymous reviewer for help with previous drafts of this paper.

REFERENCES

- Arak, A. & Enquist, M. 1993 Hidden preferences and the evolution of signals. *Phil. Trans. R. Soc. Lond. B* **340**, 207–213.
- Arak, A. & Enquist, M. 1995a Conflict, receiver bias and the evolution of signal form. *Phil. Trans. R. Soc. Lond. B* **349**, 337–344.
- Arak, A. & Enquist, M. 1995b (reply to Cook). *Nature* **374**, 313–314.
- Cook, N. D. 1995 Artefact or network evolution. *Nature*, **374**, 313.
- Dawkins, M. S. & Guilford, T. 1995 An exaggerated preference for simple neural network models of signal evolution? *Proc. R. Soc. Lond. B* **261**, 357–360.
- Enquist, M. & Arak, A. 1993 Selection of exaggerated male traits by female aesthetic senses. *Nature* **361**, 446–448.
- Enquist, M. & Arak, A. 1994 Symmetry, beauty and evolution. *Nature* **372**, 169–172.
- Hurd, P. L., Wachtmeister, C.-A. & Enquist, M. 1995 Darwin's principle of antithesis revisited: a role for perceptual biases in the evolution of intraspecific signals. *Proc. R. Soc. Lond. B* **259**, 201–205.
- Johnstone, R. A. 1994 Female preference for symmetrical males as a by-product of selection for mate recognition. *Nature* **372**, 172–175.
- Johnstone, R. A. 1995 (reply to Cook). *Nature* **374**, 313.
- Osorio, D. 1996 Symmetry detection by categorization of spatial phase, a model. *Proc. R. Soc. Lond. B* **263**, 105–110.

Received 10 October 1996; accepted 27 November 1996