

Modelling Biases and Biasing Models: 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, U.K.
sethb@cogs.susx.ac.uk, davec@cogs.susx.ac.uk

March 1996

Abstract

Recently, within the biology literature, there has been considerable interest in exploring the evolutionary function of animal displays through computer simulations of evolutionary processes (Arak & Enquist, 1993, 1995a; Enquist & Arak, 1993, 1994; Johnstone, 1994; Hurd, Wachtmeister, & Enquist, 1995; Krakauer & Johnstone, 1995). Whilst we applaud biologists' adoption of the simulation techniques pioneered within the artificial sciences (see, for example, Meyer & Wilson, 1991; Meyer, Roitblat, & Wilson, 1993; Cliff, Husbands, Meyer, & Wilson, 1994, for collections of such research), and feel that *bi-directional* cross-fertilisation between natural and artificial sciences has a bright future, we suggest that the application of such techniques to evolutionary modelling may prove to be problematic. Some debate has accompanied the work (Cook, 1995; Johnstone, 1995; Arak & Enquist, 1995b; Stamp Dawkins & Guildford, 1995) but attention to the methodology employed within this embryonic research paradigm has been cursory. Here we provide a critique of this methodology, concentrating on Enquist and Arak's (1994) exploration of the evolutionary function of complex symmetrical displays. We investigate their hypothesis that complex signal form, rather than being the product of evolutionary pressure for information exchange, is the product of 'hidden preferences' inherent in sensory systems (i.e. sensory biases). Through extending their work and relaxing their assumptions we reveal that the 'hidden preference' for symmetry proffered by Enquist and Arak (1994) is in reality a preference for homogeneity. We show that the flaws present in Enquist and Arak's (1994) study are immanent in any such evolutionary simulation model, and must be challenged if research within this paradigm is to prove worthwhile.

1 Introduction

Traditionally much of the research falling under the auspices of the 'International Conference on Simulation of Adaptive Behaviour' (now entering its sixth year) derives inspiration from, or emulates, results in the biological sciences (Meyer & Wilson, 1991; Meyer et al., 1993; Cliff et al., 1994). This phenomenon is not without its detractors (Miller, 1995).

Recently a small but rapidly increasing number of studies have marked the inception of an analogous trend; techniques pioneered within the artificial sciences are being adopted by evolutionary biologists in an effort to model natural evolutionary scenarios, and test hypotheses within evolutionary biology (Arak & Enquist, 1993, 1995a; Enquist & Arak, 1993, 1994; Johnstone, 1994; Hurd et al., 1995; Krakauer & Johnstone, 1995). Perhaps it should not be surprising that such research (termed *evolutionary simulation modelling* throughout this paper) is (potentially) problematic.

Here we replicate, discuss, and extend an evolutionary simulation model concerned with the evolutionary function of naturally occurring symmetrical visual displays (Enquist & Arak, 1994). We conclude that the problems inherent in its design and execution can be expected to be typical of evolutionary simulation models implemented without an appreciation of the methodological issues surrounding the simulation of adaptive behaviour.

2 Symmetry

That naturally occurring symmetry is both attractive and intriguing is attested to by at once the prevalence of artistic work aimed at replicating or eulogising natural symmetry, and the body of academic literature attempting

*Submitted to the Fourth International Conference on Simulation of Adaptive Behavior (SAB96), Cape Cod, Massachusetts, USA, September 9-13, 1996

to explain it (e.g. Stewart & Golubitsky, 1993). Within the discipline of evolutionary biology, recent speculation concerning the possible role of animal symmetry as an honest indicator of viability has prompted an explosion of interest in the degree to which females (and to a lesser extent males) are interested in the symmetry of their suitors, and the degree to which such symmetry is correlated with factors deemed advantageous to a prospective mate (see Watson & Thornhill, 1994, for a recent review).

The degree of Fluctuating Asymmetry (FA) evidenced by a prospective suitor, defined as “random deviation from perfect bilateral symmetry in a morphological trait for which differences between the right and left sides have a mean of zero and are normally distributed” (*ibid.*, p.21), is thought to reveal the fragility of the developmental mechanism which translated the suitor’s genotype into its phenotype. This speculation is founded on the observation that “since both sides of any bilateral trait are produced by the same genome, the degree of symmetry reveals an individual’s ability to canalize development in the face of stress” (*ibid.*, p.21).

The hypothesis that the symmetry of secondary sexual ornaments or colouration functions as a signal, imparting information about the state of the signaller, is one of a family of related hypotheses which consider the primary role of signals to be one of information exchange (see, for example, Zahavi, 1975, 1977). Recently animal behaviour theorists have taken pains to demonstrate that signal form may be the product of evolutionary forces which, far from involving pressures for information exchange, may simply drive signals to exploit ‘hidden preferences’, or sensory biases, inherent in the signal receivers (e.g. Arak & Enquist, 1993, 1995a; Enquist & Arak, 1993, 1994; Hurd et al., 1995). That the symmetry of natural signals may be amenable to an explanation in terms of such ‘hidden preferences’ is the thesis of Enquist and Arak’s (1994) study.

3 A Sensory Bias for Symmetry

Enquist and 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 perceptual invariance of such patterns across such situations.

They model an artificial visual system using an array of receptive elements (a seven-by-seven grid of idealised 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 idealised 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 of these ‘hidden’ neurons, when summed, and perturbed by a random internal factor, rises above a certain threshold.

The visual system’s task was to accurately discriminate between a signal pattern and various distractor patterns (each represented as a five-by-five array 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, 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 ‘reproduce’, creating new patterns and networks that were similar to their parents 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” (*ibid.*, p.171).

Prima facie these results are strong support for the hypothesis that symmetry “may arise as a by-product of the need to recognise objects irrespective of their position and orientation in the visual field” (*ibid.*, p.169).

4 Replication

In replicating this study we re-implemented the stochastic hill-climbing algorithm employed by (Enquist & Arak, 1994)¹. 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

¹Copies of the code, and a version of this paper with colour figures, are available on request from the authors, or from the world-wide-web page URL: <http://www.cogs.susx.ac.uk/users/sethb/sym.html>

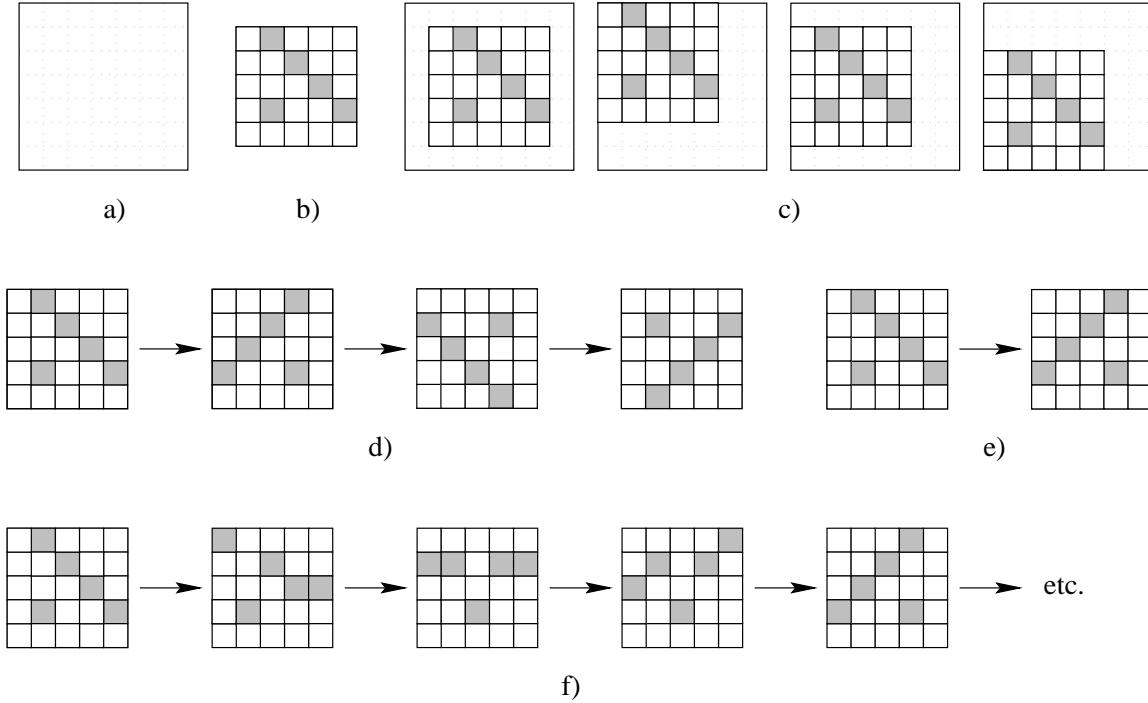


Figure 1: Depicting (a) the retina, (b) an example signal, Enquist and Arak's (1994) (c) translation, (d) 90° rotation, and (e) reflection transformations, and (f) the higher resolution rotation transformation we employ in this paper (n.b. in (c) four of the nine possible translations are depicted, and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

the space of retinal networks. In order to avoid biasing the simulation through preferentially evolving either the network or the signal first, the order in which the evolutionary 'steps' were carried out was counter-balanced across simulations.

Hill-climbing is carried out within the space of retinal networks 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 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)². If a mutant is assessed as fitter than the current best signal, it takes the best signal's place and acts as model for the next mutant. After n mutants have been generated attention reverts to the best retinal network and the process repeats.

Enquist and Arak (1994) 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:

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

²Perturbations which produced colour component values outside the legal range [0.0, 1.0] were truncated to the nearest extreme value. This effectively means that the mutation rate is not uniform over the range of colour component values, but decreases as values approach either 0 or 1 due to the more frequent generation of perturbations which must be truncated. This could account for the "purer, brighter colours" observed by Enquist and Arak (1994, p.169).

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, whilst that of a network neuron was calculated as a sigmoid function of y , the sum of its weighted inputs, thus:

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

The function $\xi(w)$ is an error function which smoothes the network output (w) by modelling a stochastic internal variable (Θ) which perturbs the network output by a value normally distributed with zero mean ($\mu=0$), and standard deviation 0.02 ($\sigma=0.02$).

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

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

Where τ is a scaling constant used to normalise the function's output to between 0 and 1, and 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 and 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 and 1d). The results of 20 simulation runs are displayed in Figure 2. 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)³. Under the conditions employed by Enquist and Arak (1994), symmetrical signals do indeed evolve.

5 Discussion

The predictions made by the hypothesis that object discrimination inherently favours symmetrical patterns were upheld by our replication of Enquist and Arak's (1994) study. However, we were interested in exploring the possibility that the preferences that the networks were exhibiting might not tell the whole story. Were there 'hidden' preferences that could be revealed through further simulation?

The transformations that Enquist and Arak's patterns undergo are intended to mimic the different ways in which the light reflected from a real-world object falls onto a retina. Objects are not always seen at the same orientation, or from the same distance, etc., but visual systems have evolved to cope with such variation in the retinal image. As Enquist and Arak's retinae and patterns are square, reflections in the vertical and horizontal axes seem natural transformations to perform. However, what situations are these transformations analogous to? Natural visual systems are very rarely exposed to a pattern and then, subsequently, a reflection of that pattern. The only situations we can envisage involve still, mirror-like, bodies of water, which surely comprise a vanishingly small proportion of most visual systems' experience⁴.

Rotations of 90 degrees 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 (not a frequently occurring pattern in nature – although strongly resembling the bilateral symmetry that *is* prevalent). However, in naturally occurring situations, the rotation transformations that images undergo are *continuously* distributed between 0° and 360°. Presumably (extrapolating from Enquist and Arak's results) patterns exhibiting continuous rotational symmetry (i.e. concentric circles) should arise under such conditions.

³Bilateral 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.

⁴Enquist (pers. comm.) has suggested that the reflection transformation might correspond to situations in which the visual system is presented first with one (patterned) side of a creature (e.g. a thin tropical fish), and subsequently with the other side, on which there exists a symmetrical version of the first side.

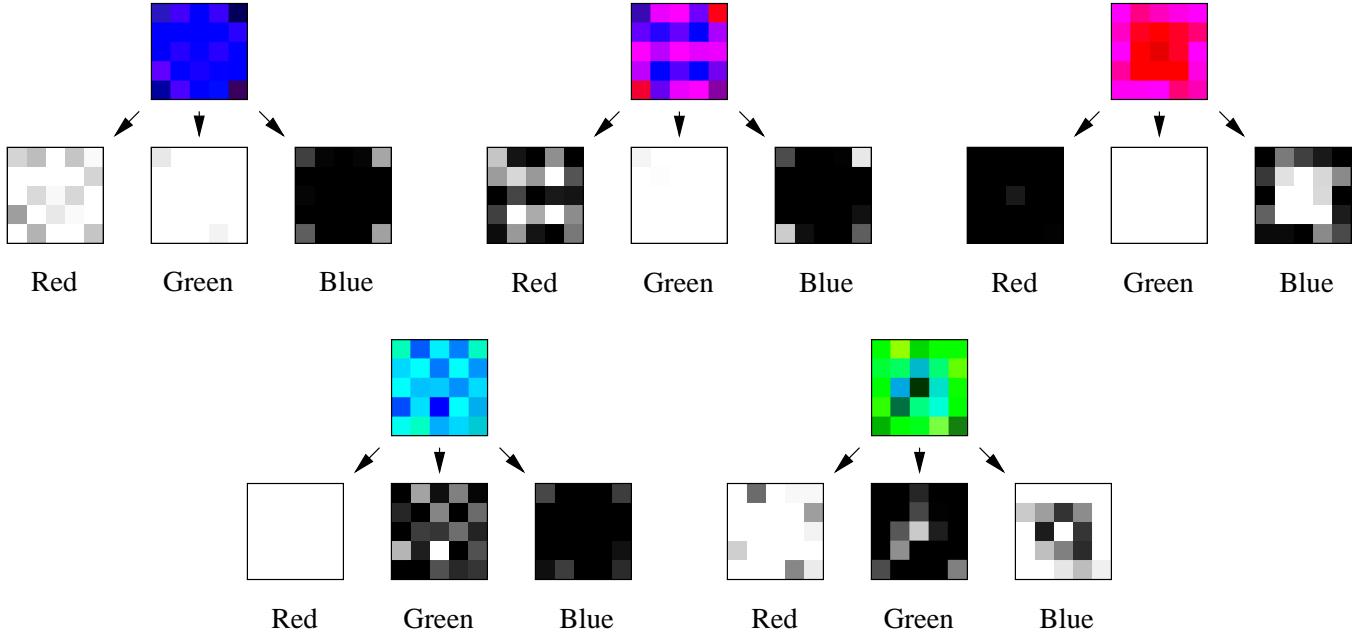


Figure 2: Five signals typical of the 20 evolved under Enquist and Arak's (1994) presentation regime. Each signal is shown as a colour composite and three colour separates in which the increasing intensity of each colour component is represented as increasingly heavy shading.

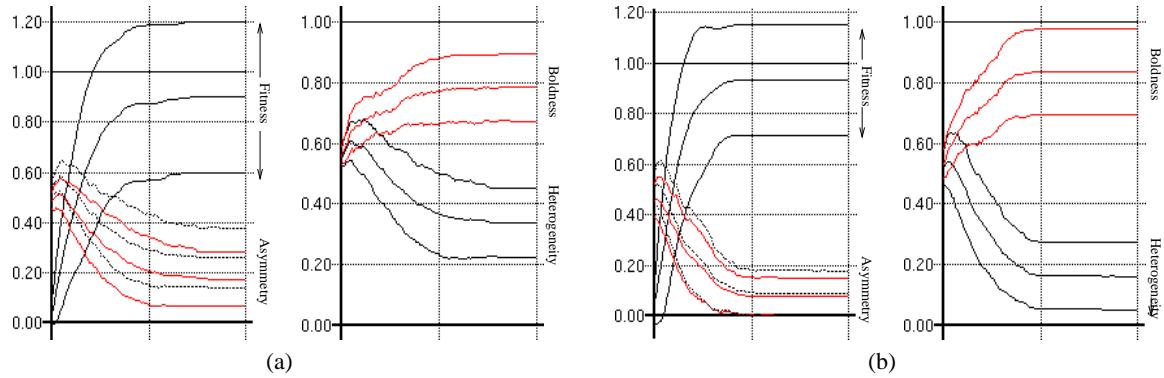


Figure 3: Initially random signals were subjected to (a) Enquist and Arak's (1994) presentation regime comprising nine translation, and four 90° rotation transformations, and (b) a presentation regime comprising nine translation, and sixteen 22.5° rotation transformations. Graphs depict (left) plots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

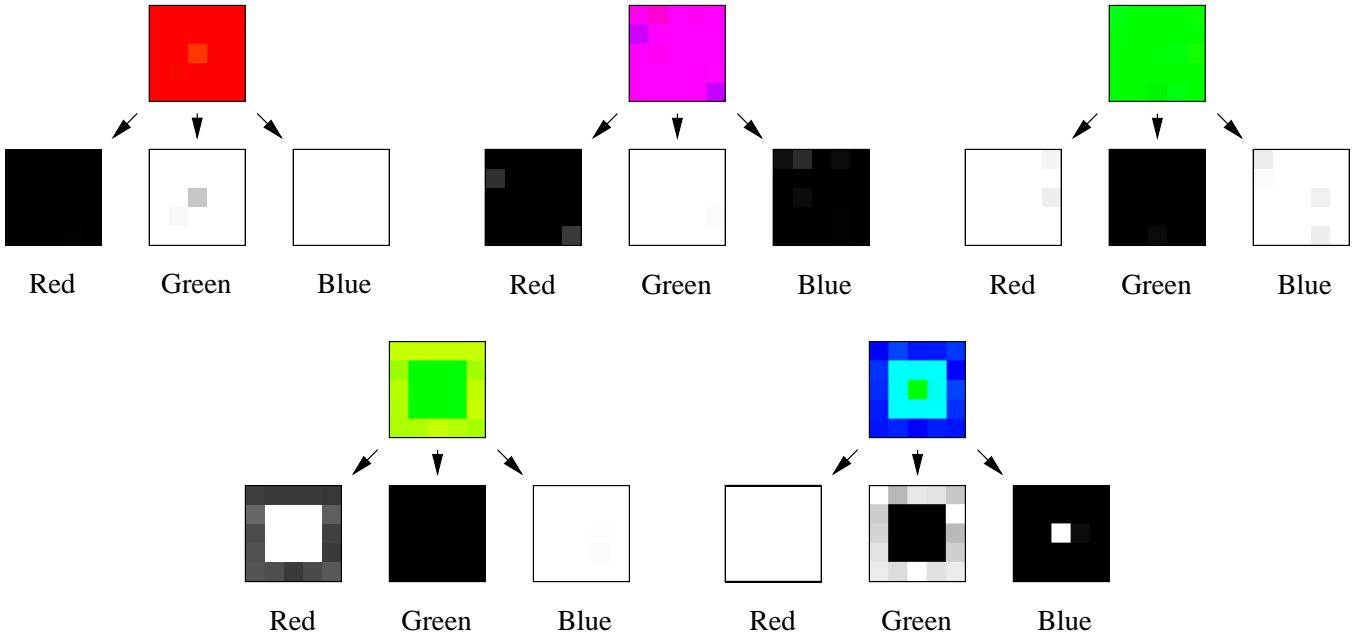


Figure 4: Five signals typical of the 20 evolved under a presentation regime comprising nine translation, and sixteen 22.5° rotation transformations. Each signal is shown as a colour composite and three colour separates in which the increasing intensity of each colour component is represented as increasingly heavy shading.

If invariance under transformations is what makes a signal easy to discriminate, then biological signals which are selected for conspicuity should exhibit maximum symmetry. Under definitions of symmetry provided by Stewart and Golubitsky (1993) such patterns should be one solid colour (i.e. invariant under the most transformations). Enquist and Arak's patterns are occasionally uniform, but the naturally occurring patterns that they are attempting to account for exhibit a high degree of *symmetry breaking*, i.e. they are both symmetrical *and* complex.

Could it be the case that the networks evolved under these conditions had preferences for structure (e.g. complex symmetry), only insofar as they were subjected to unnaturally structured sets of transformations during their evolution? As the set of transformations was naturalised, would the 'hidden preference' for homogeneity emerge? In order to test these hypotheses we extended Enquist and Arak's (1994) paradigm to incorporate more natural transformations and explored the system's sensitivity to initial conditions.

6 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 and 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 sixteen rather than four different orientations. All other parameter values were held constant. The results of 20 simulation runs are displayed in Figures 4 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 and 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 5a).

These results demonstrate that the networks evolve a preference for bold homogeneous signals. The sensory bias that Enquist and Arak propose as a pressure for complex symmetry cannot be responsible for natural complex bilateral symmetries. Furthermore, seeding 20 simulations with initially perfectly bilaterally symmetrical signals also

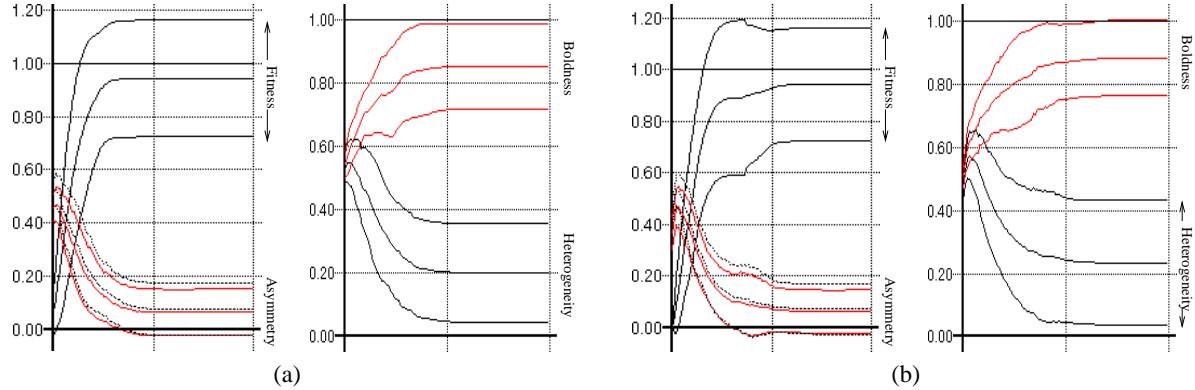


Figure 5: Initially (a) random, and (b) bilaterally symmetrical signals were subjected to a presentation regime comprising nine translation, and sixteen 22.5° rotation transformations. In addition, signals depicted in graphs (a) were subjected to a bilateral reflection transformation. Graphs depict metrics as per Figure 3.

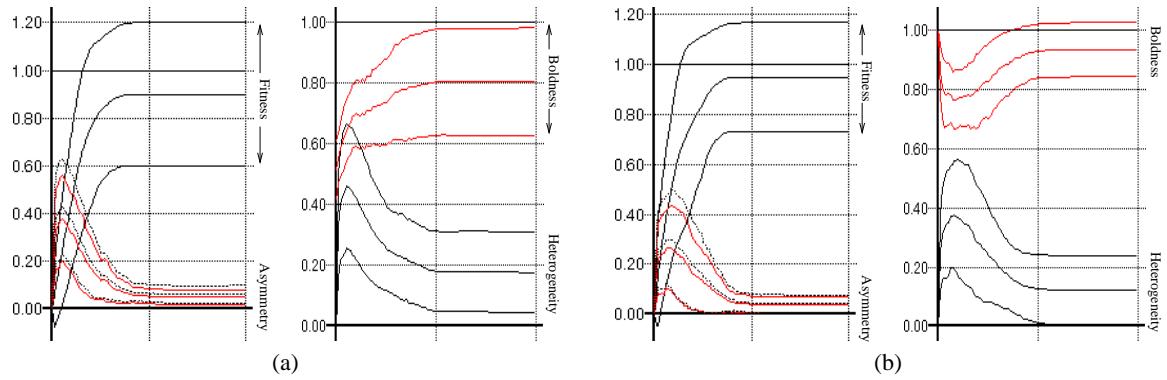


Figure 6: Initially (a) homogeneous signals and (b) maximally bold homogeneous signals were subjected to a presentation regime comprising nine translation, and sixteen 22.5° rotation transformations. Graphs depict metrics as per Figure 3.

resulted in evolved signals not significantly different from those generated from random initial conditions, indicating that, as well as failing to *generate* symmetrical structure, the sensory bias presented by Enquist and Arak (1994) could not *maintain* pre-existing bilateral symmetries (see Figure 5b).

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 6a).

Seeding 20 simulations with signals initially homogeneous in one colour chosen randomly from the six maximally bold colours that the networks were not penalised for favouring (i.e. one of green, red, blue, magenta, cyan, or yellow, but not white or black) also resulted in evolved signals not significantly different from those evolved from initially random signals (see Figure 6b). 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 and Arak (1994) propose.

7 Exploitation, Biases, and ‘Hidden Preferences’ in Modelling Paradigms

The debate concerning the status of biological signals *vis a vis* their role as informative indicators has focused in part on alternative explanations for signal form based on the observation that evolution will exploit biases in receiver systems, resulting in signals which, although carrying no information, effectively manipulate their recipients to the advantage of the signallers.

Gulls brooding basketballs, fish tempted by angler fish lures, and bees fooled by bee-orchid stamen, are all examples of sensory systems being ‘deceived’ by both natural and artificial (supernormal) stimuli. In this section we explore the possibility that evolutionary simulation models are susceptible to the very same exploitation.

Evolution’s penchant for exploiting simplistic mechanisms is not limited to the natural world. The forces responsible for naturally occurring exploitation are just as effective within artificial evolutionary models. This leads us to ask of simulation models the same questions that advocates of sensory bias explanations ask of signal form. Are such simulations informative or exploitative? Can they usefully support or refute evolutionary hypotheses, or do they merely exploit weaknesses, loop-holes, biases, or ‘hidden preferences’ in the simple mechanisms and methodology that modellers currently employ?⁵

Unfortunately, whereas the debate over the informative vs. exploitative role of signal form is (literally) academic, the consequences of the analogous debate over the theoretical status of evolutionary simulation modelling paradigms are much more serious.

If models are to provide theoretically worthwhile results, the methodology employed in their design and execution must be rigorous enough to withstand critical accusations of artefactual results due to poor simplifying assumptions. Although simplifying assumptions are a necessary part of modelling (whether through simulation, or formal analytic methods), they must be explicit, justified and, any conclusions drawn from such models must be qualified with respect to those simplifying assumptions.

The assumptions made by Enquist and Arak (1994) involving their presentation regime, their implementation of noise, their use of colour, and their model evolutive process, all fail to survive close scrutiny and in drawing conclusions from their results the authors fail to appreciate the considerable part played by these assumptions.

We have shown herein, that the form of the signals that Enquist and Arak (1994) evolve is dependent on the presentation regime that they undergo. As Enquist and Arak’s (1994) unnatural regime is relaxed so as to more closely approximate that experienced by natural signals, the interesting structure of the artificial patterns, which marks them as possible candidate exemplars in a theory of symmetry as sensory exploitation, dissolves, until bland homogeneity is reached. We have shown that Enquist and Arak’s (1994) signals exploited the structure inherent in the *presentation regime* rather than any ‘hidden preferences’ inherent in their artificial retinae and, by extension, visual systems in general.

The random internal factor implemented by Enquist and Arak (1994) serves to smooth the fitness landscape and, in combination with rank-based selection, enables evolutionary progress over areas of the landscape which have a very low fitness gradient. Modelling the ‘noise’ more traditionally, as a random value drawn independently from the normal distribution specified by Enquist and Arak (1994) each time the internal stochastic variable is invoked, results in poorer evolutionary progress and, as a result, evolved signals only slightly more symmetrical than random ones. Progress is impeded because slight improvements in discriminatory ability on the part of the retinal networks or discriminability on the part of signal patterns are not consistently scored fitter than marginally poorer competitors.

For example, a network responding slightly more positively to a particular signal presentation, but still failing to overcome the output unit’s threshold, will enjoy an increase in fitness only if a random noise value perturbs the output such that the threshold is now exceeded. The fact that noise values are typically very small ensures that such occurrences will be vanishingly infrequent for any output not extremely close to the threshold. This contrasts starkly with Enquist and Arak’s (1994) deterministic model which rewards *any* improvement in performance with an improvement in fitness on the basis that over an infinite number of trials any improvement, however small, *must* result in more correct discriminations. Whether natural evolutionary systems enjoy a similar selection dynamic is not discussed, and the (poorer) results of differing models of noise are not presented.

Enquist and Arak (1994) mention the bright colours of their evolved signals but wisely provide no adaptive explanation for them. The fact that colours in their simulation are represented as vectors in a simple three-dimensional colour space, and evolve by means of large mutations, typically along a single dimension of this space, render any

⁵ Anecdotal evidence abounds suggesting that evolutionary simulations will mercilessly exploit modellers’ errors. For example, Karl Sims tells of creatures evolved in a realistically simulated world (e.g. Sims, 1995) which exploit his model’s failure to conserve momentum correctly by banging bits of their body together in order to get around.

conclusions drawn from their results at best tenuous. We feel that the colours present in the evolved signals are the result of their non-uniform mutation operator (which effectively favours mutations towards the extremes of the colour space and suppresses mutations away from those extremes⁶) and the predispositions of simple artificial neural networks to favour extreme valued inputs, rather than any evolutionary force analogous to that responsible for the generation of naturally occurring brightly coloured signals.

The hill-climbing algorithm employed by Enquist and Arak (1994) suffers from problems typical of local search algorithms. Roughly 1 in 20 simulations failed to make any progress, as no mutant retina could discriminate the initial signal at above chance, and no mutant signal was discriminated at above chance by the initial retina. Arak and Enquist (1995a, p.340) seem predisposed to attribute evolutionary, functional explanations for this type of phenomenon rather than explain it as a consequence of the simplicity of their model evolutive process. A population size of greater than one (i.e. a parallel search algorithm more typical of evolutionary simulation models) reduces the incidence rate of such ‘sterile’ initial conditions, which rapidly falls to near zero as the size of each population increases beyond 100.

The two-step nature of Enquist and Arak’s (1994) algorithm also proves problematic. As has already been stated, precautions must be taken in order to avoid preferentially evolving either network or signal. Furthermore the notion of successive ‘generations’ in Enquist and Arak’s (1994) description of their algorithm is suspect, as individuals do not exist concurrently. This approach precludes the appreciation of frequency-dependent selection effects, or other effects of the interaction between individuals. In addition, the arbitrary length of each evolutionary ‘step’ (i.e. the value given to n in this paper) is a free parameter, the effect of which is difficult to predict or explore effectively.

In a wider sense these problems are symptomatic of a failure on the part of modellers within the evolutionary simulation modelling paradigm to appreciate the methodological issues pertaining to the use of neural networks and genetic algorithms in the modelling of adaptive evolutionary processes. As has been shown here, evolutionary simulation models are not tolerant of poor simplifying assumptions, or simple adaptive mechanisms. The results of such models are always potential artefacts, the products of exploitation, bias and what (Arak & Enquist, 1993) have dubbed ‘hidden preferences’.

8 Conclusion

In conclusion, Enquist and Arak’s (1994) model of the evolution of symmetrical patterns was found to be seriously flawed in a manner which suggests problems for evolutionary simulation models in general. In light of this, although the hypothesis that symmetrical patterns may be the result of sensory exploitation is still healthy (see Osorio, 1996, for an alternative formulation of this hypothesis), the prospect of evolutionary simulation models suffering exploitative evolutionary dynamics is a less welcome challenge for researchers working within this paradigm.

Therefore, in order that evolutionary biologists can usefully employ the techniques being developed within the simulation of adaptive behaviour community, two commitments must be undertaken. First, a commitment to more tightly-coupled interdisciplinary collaboration between scientists studying natural and artificial systems. Secondly, a commitment to the wider dissemination of studies explicating the methodological issues pertaining to research within what we have referred to here as the evolutionary simulation modelling paradigm.

Acknowledgements

We thank Magnus Enquist, Geoffrey Miller, and Daniel Osorio for discussions, and Henrietta Wilson for her meticulous proof reading.

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.

Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.). (1994). *From Animals to Animats 3:Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*. MIT Press/Bradford Books.

⁶See footnote 2.

Cook, N. D. (1995). Artefact or network evolution. *Nature*, 374, 313.

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.

Krakauer, D. C., & Johnstone, R. A. (1995). The evolution of exploitation and honesty in animal communication: a model using artificial neural networks. *Phil. Trans. R. Soc. Lond. B*, 348, 355 – 361.

Meyer, J.-A., Roitblat, H. L., & Wilson, S. W. (Eds.). (1993). *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behaviour*. MIT Press/Bradford Books.

Meyer, J.-A., & Wilson, S. W. (Eds.). (1991). *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*. MIT Press/Bradford Books.

Miller, G. F. (1995). Artificial life as theoretical biology: How to do real science with computer simulation. Tech. rep., University of Sussex Cognitive and Computing Sciences Technical Report CSRP 378.

Osorio, D. (1996). Symmetry detection by categorisation of spatial phase, a model. *Proc. R. Soc. Lond. B*, 263, 105 – 110.

Sims, K. (1995). Evolving 3D morphology and behavior by competition. *Artificial Life*, 1, 353 – 372.

Stamp Dawkins, M., & Guildford, T. (1995). An exaggerated preference for simple neural network models of signal evolution?. *Proc. R. Soc. Lond. B*, 261, 357 – 360.

Stewart, I., & Golubitsky, M. (1993). *Fearful Symmetry: Is God A Geometer?* Penguin, London.

Watson, P. J., & Thornhill, R. (1994). Fluctuating asymmetry and sexual selection. *Trends in Ecology and Evolution*, 9(1), 21 – 25.

Zahavi, A. (1975). Mate selection – a selection for a handicap. *Journal of Theoretical Biology*, 53, 205–214.

Zahavi, A. (1977). The cost of honesty (further remarks on the handicap principle). *Journal of Theoretical Biology*, 67, 603–605.