and only generally constrains the fault depth. Integration of seismological data on the hypocentre, mainshock focal mechanism and distribution of aftershocks establishes the exact dip and depth of the Pico thrust.

Our interpretation indicates a pairing of the Santa Susana Mountains anticlinorium uplift and the deeply buried Pico thrust. The lateral extent of the latter can therefore be defined by the lateral extent of the Santa Susana Mountains anticlinorium which is recorded by surface geology and subsurface well data to be 30–40 km in length. More work is needed to elucidate the detailed geometry of the anticlinorium to define more precisely the regional geometry and extent of the seismically active Pico thrust.

The Pico thrust is shown to be a backthrust off the north-dipping Elysian Park thrust ramp. The Elysian Park thrust is rooted in a mid-crustal detachment at ~22 km depth. Mid-crustal detachments have been proposed by numerous workers[3,10,19–21] to underlie the western Transverse Ranges and are a basic component of fold-and-thrust belts[8,9,13]. The restored cross-section (Fig. 2*c*) shows the structural geometry before folding and faulting, and provides a check that the cross-section is balanced.

The balanced cross-section provides information about the seismic potential of the Pico and Elysian Park thrusts. This information is especially important for the latter thrust because it underlies the most urbanized parts of the Los Angeles basin (Fig. 1). Previous balanced cross-section analysis and subsurface mapping of the Elysian Park thrust[3,22] indicate that it is over 170 km long and is a fundamental thrust fault of southern California. The Pico thrust has 3.3 km of displacement and the Elysian Park thrust has 11.8 km of displacement (Fig. 2*b*). Compressive deformation probably began 2–3 Myr ago[3] within the Transverse Range fold-and-thrust belt. Initiation of the Santa Susana Mountains anticlinorium and Pico thrust is recorded by deformation of the youngest unit, the Saugus Formation (QTu, Fig. 2), which is no older than 2.3 Myr (ref. 23). The displacement and 2.3–2.0 Myr age of fault initiation yields an average slip rate of 1.4–1.7 mm yr^{-1} for the Pico thrust. A 2.0–3.0 Myr (ref. 3) age of initiation of the Santa Monica Mountains anticlinorium yields an average slip rate of 3.9–5.9 mm yr^{-1} for the Elysian Park thrust.

Geodetic and seismic modelling suggest the Pico thrust moved ~2.5 m during the Northridge earthquake[15,24] which, divided by our slip rates, yield an average repeat time of 1,500–1,800 years. This repeat-time estimate applies to the segment of the Pico thrust involved in the earthquake, which is ~15 km long. Similar recurrence calculations can be made for the Elysian Park thrust (Fig. 1). Assuming that a Northridge-size earthquake is the characteristic event for the Elysian Park thrust trend, then our slip rate and time-of-initiation estimates yield average earthquake repeat times of 420–640 years for any 15 km segment of the thrust, or an event every 39–58 years along the trend. This repeat time is not supported by the 220 years of recorded history. This discrepancy is probably due to one or more of the following; our long-term slip rates differ from the short-term rates, some crustal shortening is taken up aseismically, the earthquake repeat time is variable, or the characteristic event is even larger and less frequent than that predicted by a Northridge-size earthquake.

Fold-and-thrust belts are tectonic systems with complicated fold and fault patterns that develop over a scale of kilometres (ref. 13). In seismically active belts, the integration of earthquake data with balanced cross-sections provides a broader view of seismic risk than traditional seismic-hazard methods. Application of the balanced cross-section technique should be the first step in evaluating the seismic risk in fold-and-thrust belts and would have predicted the presence of a young south-dipping thrust beneath the Santa Susana Mountains and San Fernando Valley before the Northridge earthquake. □

Received 23 February; accepted 6 October 1994.

1. Whitcomb, J. H., Allen, C. R., Garmany, J. D. & Hileman, J. A. *Rev. Geophys. Space Phys.* **11,** 693–730 (1973).
2. Namson, J. S. & Davis, T. L. *Bull. geol. Soc. Am.* **100,** 257–273 (1988).
3. Davis, T. L., Namson, J. S. & Yerkes, R. F. *J. geophys. Res.* **94,** 9644–9664 (1989).
4. Hauksson, E. *et al. Science* **239,** 1409–1412 (1988).
5. Hauksson, E. *et al. Seism. Soc. Am. Abstr.* **65,** Northridge Progm no. 4 (1994).
6. Allen, C. R. *Bull. geol. Soc. Am.* **86,** 1041–1067 (1975).
7. Sieh, K. E. *J. geophys. Res.* **83,** 3907–3939 (1978).
8. Dahlstrom, C. D. A. *Can. J. Earth Sci.* **6,** 743–757 (1969).
9. Bally, A. W., Gordy, P. L. & Stewart, G. A. *Bull. Can. Petrol. Geol.* **14,** 337–381 (1966).
10. Namson, J. S. & Davis, T. L. *Geology* **16,** 675–679 (1988).
11. Buxdorf, A. *Naturf. Geos. Basel Verh.* **27,** 184–205 (1916).
12. Hobson, D. M. *Aust. Petrol. Explor. Ass. J.* **26,** 214–224 (1986).
13. Mitra, S. *Structural Geology of Fold and Thrust Belts* 53–77 (Johns Hopkins Univ., Baltimore, Maryland, 1992).
14. US Geological Survey staff *Seism. Soc. Am. Abstr.* **65,** Northridge Progm nos 31, 32 (1994).
15. Hudnut, K. W. *Seism. Soc. Am. Abstr.* **65,** Northridge Progm no. 40 (1994).
16. Dibblee, T. W. Newhall and Oat Mountain map (Map DF36, Diblee Foundation, Santa Barbara, California, 1992).
17. Treiman, J. A. *Seism. Soc. Am. Abstr.* **65,** Northridge Progm no. 34 (1994).
18. Mitra, S. *Bull. Am. Ass. Petrol. Geol.* **74,** 921–945 (1990).
19. Namson, J. S. & Davis, T. L. *Bull. Am. Ass. Petrol. Geol.* **74,** 467–492 (1990).
20. Shaw, J. W. & Suppe, J. *Bull. Am. Ass. Petrol. Geol.* **78,** 700–721 (1994).
21. Shaw, J. W. & Suppe, J. *Bull. geol. Soc. Am.* **106,** 607–626 (1994).
22. Hauksson, E. *J. geophys. Res.* **95,** 15365–15394 (1990).
23. Levi, S. & Yeats, R. S. *Tectonics* **12,** 688–702 (1993).
24. Wald, D. J. & Heaton, T. H. *Seism. Soc. Am. Abstr.* **65,** Northridge Progm no. 9 (1994).

ACKNOWLEDGEMENTS. We thank E. Duebendorfer and D. Schwartz for comments on the manuscript.

Symmetry, beauty and evolution

Magnus Enquist* & Anthony Arak†

* Department of Zoology, University of Stockholm,
S-106 91 Stockholm, Sweden
† Archway Engineering (UK) Ltd, Ainleys Industrial Estate,
Elland HX5 9JP, UK

HUMANS and certain other species find symmetrical patterns more attractive than asymmetrical ones. These preferences may appear in response to biological signals[1–3], or in situations where there is no obvious signalling context, such as exploratory behaviour[4,5] and human aesthetic response to pattern[6–8]. It has been proposed[9,10] that preferences for symmetry have evolved in animals because the degree of symmetry in signals indicates the signaller's quality. By contrast, we show here that symmetry preferences may arise as a by-product of the need to recognize objects irrespective of their position and orientation in the visual field. The existence of sensory biases for symmetry may have been exploited independently by natural selection acting on biological signals and by human artistic innovation. This may account for the observed convergence on symmetrical forms in nature and decorative art[11].

It is striking that many signals used for communication by organisms are judged to be beautiful by humans. Examples include the colours and symmetries of flowers, the patterns on butterflies' wings and coral reef fish, and the elaborate courtship displays of birds. The almost universal appeal of such signals to humans is surprising because they have been selected to influence the behaviour of animals whose visual systems differ in significant ways from our own, and which in some taxa (for example insects and cephalopods) have evolved independently[12]. This raises the possibility that human aesthetic sense is based on general principles of perception that have been important during the evolution of biological signals.

Studies in early ethology[13] and experimental psychology[14,15], have shown that sensory systems react strongly to certain novel stimuli. Such biases impose a selection pressure on signals[16–18], causing them to become more extreme in form and polarized during evolution[19,20] (typically larger, louder, brighter and otherwise different from stimuli that the receiver should not react to in the same way). However, it is less obvious that the striking degree of patterning seen in signals, and preferences for such patterns, are also the consequences of sensory bias.

One problem faced by animals is the need to recognize objects in different positions and orientations in the visual field. An object viewed from a particular location is focused on the retina as an 'image', which is a geometrical transformation of the object itself. An intriguing idea is that the need to generalize many such transformations of the same object may lead to preferences for symmetry[21] and the evolution of symmetrical signals. To explore this idea we have simulated the coevolution of signals and receiver preferences, using artificial neural networks as models of recognition systems.

The technique used in simulations was similar to that described in ref. 19 (see also Fig. 1). Images of patterns ('signals'), consisting of coloured squares in a grid, were 'pasted' onto an artificial retina in various positions and orientations, mimicking some of the ways in which animals may see objects in the visual field. Starting with an arbitrary signal and arbitrary network (randomly assigned colours and connection weights), coevolution proceeded, in each generation, by producing a series of mutations of the signal and the network and retaining the best one of each (artificial selection). The criterion for selecting the signal from a series of mutations was that it should elicit the highest output from the network; the criterion for selecting the network was that it should discriminate between the correct sig-

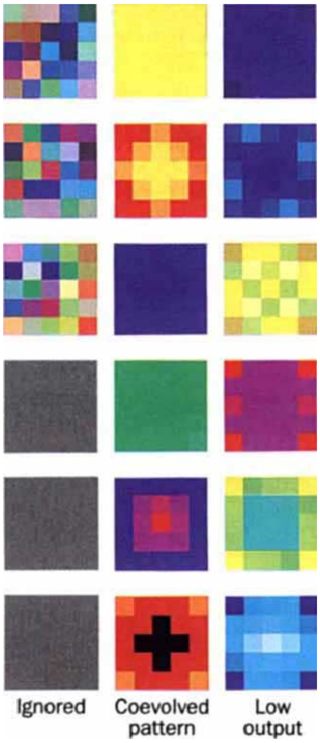

FIG. 1 Symmetry generated by coevolution between a visual signal and a recognition mechanism. Results of six independent simulations are shown, one in each row. Signals consisted of coloured squares in a 5×5 grid, each square containing a single colour. The recognition mechanism was an artificial neural network with 147 input cells, 15 hidden cells and one output cell. Each cell in a given layer was connected to all cells in the next layer, but cells within the same layer were not interconnected. The input cells were arranged on a square retina with 7×7 locations, each location containing 3 input cells sensitive to a different part of the colour spectrum according to the RGB system[22]. Each input cell receives stimulation between 0 and 1 corresponding to the degree of saturation of the respective colour element. A random pattern was generated and designated as the 'correct' signal, and an alternative pattern (either random or uniform grey) was designated as the pattern that the network should ignore. The patterns were pasted onto the retina in various positions by translation and rotation, and the response of the network measured. A network was said to react to a signal when the activity in its output cell plus a random internal factor (normally distributed with zero mean, s.d. 0.02) exceeded a threshold value of 0.5. Initially, connection weights in each network were randomly assigned, varying between -0.3 and 0.3 in a rectangular distribution. A simulation consisted of 500 generations, each involving 100 mutations of both the signal and the network. Mutations of a network were produced letting each connection weight mutate with probability 0.01 and adding a normally distributed increment (zero mean, s.d. 0.02) to each mutated weight. Mutations of the signal were produced by mutating each colour element with probability 0.05, then adding (where possible) a normally distributed increment (zero mean, s.d. 0.3) to each mutated element. In each generation the best signal and best network were kept. The low output signals are those signals which tended to elicit the lowest responses from the coevolved networks when the connection weights held constant.

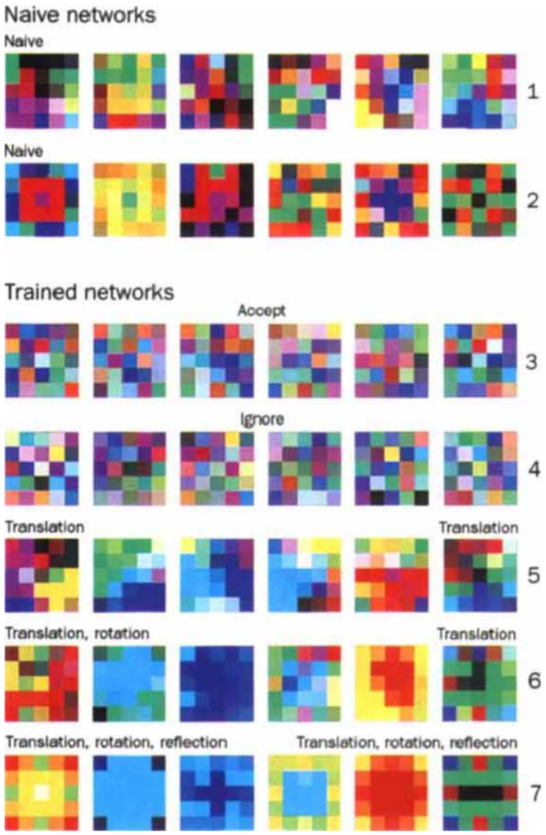

FIG. 2 The effect of different training and test procedures. The network training procedure is indicated on the left-hand side and the test procedure on the right-hand side above each row of signals. In all examples, the networks were first prepared and then connection weights held constant to explore their preferences. For six different 'naive' networks (randomly assigned connection weights), preferences are shown for signals projected onto the retina either by translation (top row, 1) or by a combination of translation and rotation (row 2). Six trained networks were prepared by teaching them to distinguish between two sets of random patterns that lack symmetry (rows 3 and 4). Training was done by projecting these patterns under translation only (row 5), translation + rotation (row 6), or translation + rotation + reflection (row 7) and continued until recognition errors were negligible ($<10^{-5}$). After training, preferences were tested under translation only (rows 5 and 6), or using all transformation types (row 7). In all cases, the preference test involved generating a signal that elicited a strong response in the network by repeated mutation and selection of signals for 500 generations (for details, see legend to Fig. 1). Signals shown directly under one another in each column are derived from the same starting network.

nal and a series of alternative patterns (Fig. 1: Ignored) with least error. This procedure eventually leads to a quasi-stable situation in which the network discriminates almost perfectly between correct and incorrect signal patterns and the signal itself (Fig. 1: Coevolved pattern) changes only slowly. After completing the simulation, the connection weights of the network were held constant, and signals were found that tended to minimize output (Fig. 1: Low output).

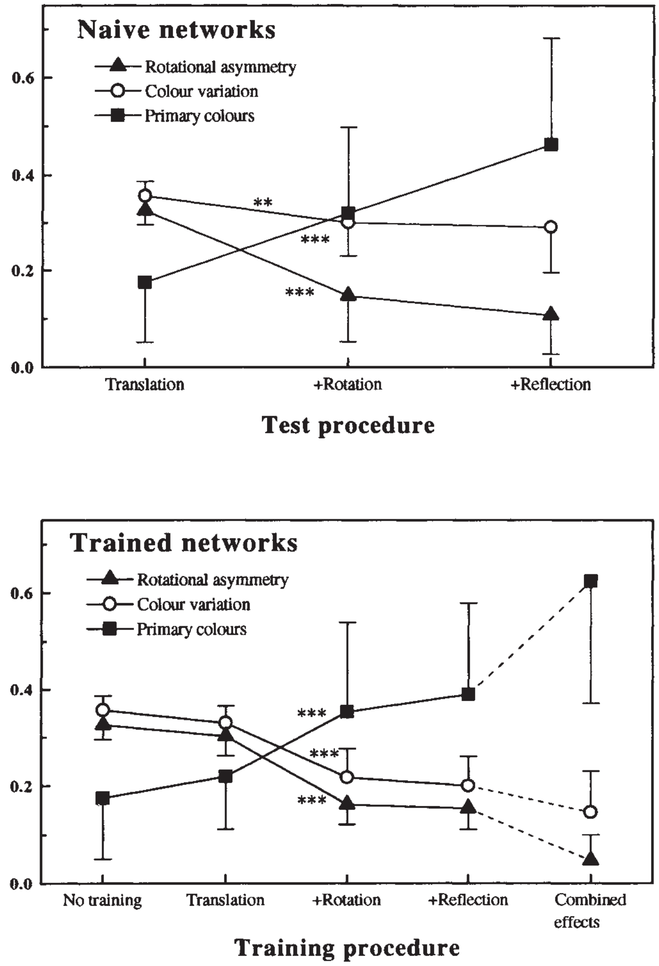

FIG. 3 Characteristics of preferred signals as a function of signal selection and network training procedures. Each point is based on 50 independent calculations. The bars represent standard deviations and the asterisks denote significant differences in comparison with preferred patterns generated by translation only. ANOVA was used except for test on primary colours which used Mann–Whitney *U*-test ($^{***}P<0.001$, $^{**}P<0.01$). Colour variation and rotational asymmetry were measured as the root mean square within groups. For colour variation, the groups were made up of the three RGB colour elements. For rotational asymmetry, groups were based on both colour element and position (3 × 6 groups). Under 90° rotation around the centre, a perfectly symmetrical pattern requires the colour to be identical at four locations (for example the four corners) within 6 different groups (excluding the central point). Primary colour is a measure of the relative frequency of such colours in the pattern. Any RGB triplet was considered to be a primary colour if none of its three elements deviated by more than 0.1 from zero or one. The *x*-axis shows differences in the method of projecting signals onto the retina during the test procedure (for naive networks) or during the training procedure (for trained networks). Preferred signals of trained networks were generated by selection of signals under translation only, except for one treatment (far right) which shows the combined effect of training networks and testing their preferences using all possible transformations permitted by the experimental design (translation + rotation + reflection).

The coevolved signals consisted of purer, brighter colours than random patterns; they also displayed significantly lower variation in colours, and marked symmetries. Signals producing a low output also displayed strong symmetries, but consisted of the chromatic opponents[22] of those colours present in the coevolved signals. These results demonstrate polarization towards brighter colours in the signal, that is, a divergence from the 'ignored' patterns. They also show that signals least likely to provoke a response in a given situation are 'opposite' in appearance to those that cause a strong response (Darwin's principle of signal antithesis[23]).

There are two possible causes for the symmetries that appeared during our simulations. Symmetry may result directly from selection acting on signals, irrespective of the state of the network. This is because a signal that, on average, is best at stimulating the network in all possible positions and orientations may take a symmetrical form. Alternatively, the network itself may become more sensitive to symmetry as a result of the evolutionary process.

To distinguish between these possibilities, we did two further tests (Figs 2 and 3). First, a random network was tested to find the signals that were most stimulating under different projection regimens (connection weights of the network were held constant). The degree of symmetry that developed in signals was found to depend strongly on the method of projecting the signals onto the retina during tests. For example, when signals were projected in all possible orientations, stronger rotational symmetries developed than when signals were projected by translation alone.

Second, networks were trained to recognize random, asymmetrical signals projected onto the retina in different ways (the signals were held constant). The preferences which then emerged depended strongly on the method of training the network. Networks trained to recognize random signals, irrespective of their orientation, developed preferences for signals with strong rotational symmetries. By contrast, preferences for symmetry were weak when networks were trained to recognize signals projected onto the retina in different locations, but in a fixed orientation.

These results show that preference for symmetry is a consequence of the need to recognize signals irrespective of their position and orientation in the visual field. In more detail, this is due to the combination of two effects that reinforce each other during coevolution: one effect is independent of the state of the network, and the other becomes important when the network is trained or evolves.

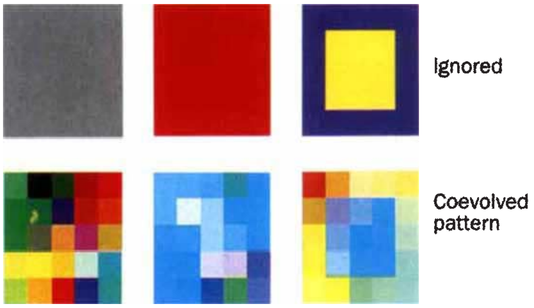

FIG. 4 The influence of the ignored pattern on the outcome of signal coevolution. Signals were selected under translation only. When the pattern to ignore was uniformly grey there are six equally likely directions of bias (red, green, blue, cyan, magenta, yellow) for each element in the coevolved signal: the precise outcome cannot be predicted. When the ignored pattern is uniformly red, the signal becomes polarized towards the blue end of the spectrum. When the ignored pattern contains more complex symmetries, the coevolved signal displays similar symmetry but in opponent colours.

Several alternative explanations exist for the occurrence of symmetrical signals and symmetry preferences in nature. It has been suggested that some morphological symmetries arise inevitably from developmental processes[24]. However, as Wallace[25] observed, the symmetrical body markings of wild animals are often lost or degraded in their domesticated descendants. This suggests that certain symmetries are not inescapable consequences of development, but are maintained by other selection pressures in nature.

Preferences for symmetry observed in animals might reflect symmetries in the pattern of connection between nerve cells in the brain. Such preferences may also arise because many objects in the world that animals should distinguish are themselves symmetrical (see Fig. 4). However, whether symmetry is common enough in nature to make this important as a general explanation for symmetry preference remains unclear. Finally, it has been suggested that preferences for symmetry may have evolved for adaptive reasons connected with mate choice[1,10]. Recent experiments on birds and insects have shown that females prefer to mate with males possessing the most symmetrical sexual ornaments[1–3]. Because deviations from perfect symmetry are negatively correlated with fitness in some species[9], the degree of symmetry in ornaments may provide females with information about male quality. But evidence that the amount of symmetry in ornaments provides females with information does not prove that preferences for symmetry evolved for that reason. Because such preferences emerge as a by-product of selection for recognition, any benefits that females gain by mating with symmetrical males may be best considered a fortuitous effect of sensory bias.

What implications, if any, do these findings have for aesthetics? Both humans[6–8], and certain other species[4,5], find symmetrical patterns attractive in contexts unrelated to signalling. Such general preferences for symmetry serve no obvious function, but may result from the universal need among organisms to recognize objects irrespective of the manner in which they are encountered in the outside world. Our results suggest that in the process of learning to recognize objects, preferences can develop for particular forms that have no objective existence in nature. Such hidden preferences[20] are revealed only when the actual forms corresponding to them appear in nature, either through biological evolution or artistic innovation. □

Received 31 May; accepted 18 August 1994.

1. Møller, A. P. *Nature* **357,** 238–240 (1992).
2. Swaddle, J. P. & Cuthill, I. C. *Nature* **367,** 165–166 (1994).
3. Radesäter, T. & Halldórsdóttir, H. *Anim. Behav.* **45,** 626–628 (1993).
4. Rensch, B. *Z. Tierpsychol.* **14,** 71–99 (1957).
5. Rensch, B. *Z. Tierpsychol.* **15,** 447–461 (1958).
6. Berlyne, D. E. *Aesthetics and Psychobiology* (Appleton, New York, 1971).
7. Hubbell, M. B. *Am. J. Psychol.* **53,** 46–69 (1940).
8. Musinger, H. & Kessen, W. *Psychol. Monogr.* **9,** 1–24 (1964).
9. Watson, P. J. & Thornhill, R. *Trends Ecol. Evol.* **9,** 21–25 (1994).
10. Brookes, M. & Pomiankowski, A. *Trends Ecol. Evol.* **9,** 201–202 (1994).
11. Gombrich, E. H. *The Sense of Order* 2nd edn (Phaidon, London, 1984).
12. Cronly-Dillon, J. R. & Gregory, R. L. (eds) *Evolution of the Eye and Visual System* (Macmillan, Basingstoke, 1991).
13. Tinbergen, N. *The Study of Instinct* (Oxford Univ. Press, London, 1951).
14. Mackintosh, N. J. *Psychology of Animal Learning* (Academic, London, 1974).
15. Doll, T. J. & Thomas, D. R. *J. exp. Psychol.* **75,** 508–512 (1967).
16. Staddon, N. E. R. *Am. Nat.* **109,** 541–545 (1975).
17. Leimar, O., Enquist, M. & Sillén-Tullberg, B. *Am. Nat.* **128,** 469–490 (1986).
18. Ryan, M. J. *Oxford Surv. evol. Biol.* **7,** 156–195 (1991).
19. Enquist, M. & Arak, A. *Nature* **361,** 446–448 (1993).
20. Arak, A. & Enquist, M. *Phil. Trans. R. Soc.* B**340,** 207–213 (1993).
21. Stewart, I. & Golubitsky, M. *Fearful Symmetry: Is God A Geometer?* (Penguin, London, 1993).
22. Derefeldt, G. in *The Perception of Colour* (ed. Gouras, P.) 218–261 (Macmillan, Basingstoke, 1991).
23. Darwin, C. *The Expression of the Emotions in Man and Animals* (Murray, London, 1872).
24. Turing, A. M. *Phil. Trans. R. Soc.* B**237,** 37–72 (1952).
25. Wallace, A. R. *Darwinism, an Exposition of the Theory of Natural Selection* (Macmillan, London, 1889).

ACKNOWLEDGEMENTS. We thank H. Cressey, O. Leimar, P. Sällström and R. Rosenberg for advice and comments on an earlier version of the manuscript. Research supported by Swedish Natural Science Research Council (M.E.) and Archway Engineering (UK) Ltd (A.A.).

Female preference for symmetrical males as a by-product of selection for mate recognition

Rufus A. Johnstone

Department of Zoology, University of Cambridge, Downing Street, Cambridge CB2 3EJ, UK

FLUCTUATING asymmetry (FA) refers to the random, stress-induced deviations from perfect symmetry that develop during the growth of bilaterally symmetrical traits[1,2]. Individual differences in the level of FA may influence mate choice[3]: in a number of species, females prefer to mate with males that have more symmetrical sexual ornaments[4–7]. As the degree of FA has been shown to reflect the ability of individuals to cope with a wide variety of environmental stresses[2,8,9], it has been suggested that mating preferences for symmetry evolve for adaptive reasons, because the degree of FA provides honest information about male quality[10,11]. Here I use simple, artificial neural networks to show that such preferences are likely to arise in the absence of any link between symmetry and quality, as a by-product of selection for mate recognition.

Female preferences for symmetrical male displays have been demonstrated in a number of species, including swallows[4,5], zebra finches[6] and earwigs[7]. These preferences have commonly been explained by reference to the honest advertisement hypothesis or handicap principle[12]. Many sexual ornaments are thought to be expensive to produce, and are under strong directional selection. Both of these factors lead to stresses that reduce the effectiveness of developmental homeostatic mechanisms. Consequently, the level of FA in male secondary sexual traits may provide females with honest information about the developmental competence of a potential mate[3,10,11]. In support of this suggestion, there is evidence that the level of FA in some species is heritable and/or negatively correlated with one or more fitness measures such as viability, fecundity, and growth rate[2,8,9,13,14]. However, evidence that FA currently provides females with useful information regarding male quality does not prove that preferences for symmetry initially evolved for that reason.

Here I test the hypothesis that preferences for low levels of FA in males may evolve in the absence of any link between symmetry and quality, as a by-product of selection for mate recognition. I investigate the evolution of symmetry biases by examining the properties of simple, artificial neural networks trained (by a process of artificial selection) to recognize a set of patterns that exhibit varying degrees of FA. Such networks, when used in the study of sexual selection and signal evolution[5,16], provide a way to examine general properties of recognition systems.

A network, representing the recognition system of a female bird, was 'trained' by artificial selection to recognize a suite of images representing a bilaterally symmetrical tail that exhibits varying degrees of FA (Fig. 1). This involved repeated mutation of the network to generate variants that differed slightly in their response to different patterns, and selection of those variants that tended to respond more strongly to the images in the training set than to random patterns. The training process thus simulates the effects of selection for mate recognition, in that it favours the ability to distinguish a specific (bilaterally symmetrical) display from other stimuli. Two different sets of training images were used: the first comprised all five of the patterns shown in Fig. 1, and the second comprised only patterns 1, 2, 4 and 5. Networks trained using the first set were thus exposed to the perfectly symmetrical tail pattern 3 during the selection procedure, whereas networks trained using the second set were not. Neither of the two training procedures gave any explicit