



Neural network for female mate preference, trained by a genetic algorithm

Masashi Kamo, Takuya Kubo and Yoh Iwasa

Department of Biology, Faculty of Science, Kyushu University, Fukuoka 812-81, Japan

In some animals, males evolve exaggerated traits (e.g. the peacock's conspicuous tail and display) because of female preference. Recently Enquist and Arak presented a simple neural network model for a visual system in female birds that acquires the ability to discriminate males of the correct species from those of the wrong species by training. They reported that the trained networks were attracted by 'supernormal stimuli' where there was a greater response to an exaggerated form than to the images used as the correct species for training. They suggested that signal recognition mechanisms have an inevitable bias in response, which in turn causes selection on signal form. We here examine the Enquist and Arak model in detail. A three-layered neural network is used to represent the female's mate preference, which consists of 6×6 receptor cells arranged on a regular square lattice, ten hidden cells, and one output cell. Connection weights of the network were modified by a genetic algorithm, in which the female's fitness increases if she accepts a conspecific male but decreases if she accepts a male of a different species or a random image. We found that: (i) after the training period the evolved network was able to discriminate male images. Female preference evolves to favour unfamiliar patterns if they are similar to the images of the correct species (generalization); (ii) the speed and the final degree of learning depended critically on the choice of the random images that are rejected. The learning was much less successful if the random images were changed every generation than if 20 random images were fixed throughout the training period; (iii) the male of the same species used for training achieved the highest probability of being accepted by the trained network. Hence, contrary to Enquist and Arak, the evolved network was not attracted by supernormal stimuli.

Keywords: generalization; genetic algorithm; neural network; sensory bias; supernormal stimulus; sexual selection

1. INTRODUCTION

Males of some animals have exaggerated traits, which are apparently disadvantageous to the holder's survival. A classical example is the tail of the male peacock. Darwin (1871) suggested that males develop a structure costly to produce and maintain because females have a propensity to choose mates with such exaggerated traits. After Darwin, numerous experimental and theoretical studies have been carried out on female mate preference (reviewed by Andersson 1994; Moller 1994; Andersson & Iwasa 1996).

Fisher (1930) was the first to succeed in explaining the evolution of female preference for an exaggerated male trait. If females in the initial population have a small propensity to choose mates with a longer tail, then a female with a stronger preference than the population average can enjoy an indirect benefit of producing sexy sons who inherit their father's 'good looking' long tail. As a consequence, exaggerated male traits and strong female preference will coevolve simultaneously, resulting in Fisherian runaway selection (Fisher 1930; Lande 1981; Pomiankowski *et al.* 1991; Pomiankowski & Iwasa 1993; Iwasa & Pomiankowski 1995). According to this argument, any male trait and a female preference for it can coevolve.

In contrast, Zahavi (1977) noted that male traits used for sexual selection are often good indicators of male quality, revealing physical strength, mental ability and resistance to parasites, traits that are difficult to determine directly. A visible male trait can be used by females in their mate choice, when only a really strong male can afford to produce and maintain it. The peacock's tail is a trustworthy signal of a good and strong male simply because it is a costly handicap to the holder (Hamilton & Zuk 1982; Grafen 1990; Iwasa *et al.* 1991; Iwasa & Pomiankowski 1994).

A third theory emphasizes the inherent bias of the sensory organ (Kirkpatrick & Ryan 1991; Ryan & Keddy-Hector 1992). If females are attracted by stronger stimuli, such as a louder call or a longer tail, then males may evolve to exploit that sensory bias by exaggerating signals.

Why is there bias in sensory organs? Enquist & Arak (1993) and Arak & Enquist (1993) developed a simple neural network model to imitate a signal recognition system that is adjustable by training, and demonstrated that the trained network has an inherent bias. The neural networks had an artificial retina on which simple bird-like images were presented. They were trained to discriminate two simplified images: one represents a conspecific male bird with a long tail (figure 1c), and the other pattern with a short tail for the male of a wrong species (figure




Figure 1. Images used during the training period and the test period. During the training period, the image (c) of tail length two was used as a right male, and the image (b) of tail length one was used as a wrong male, and they were presented to the network with shifting and rotations, together with 20 random images. During the test period, various images illustrated in this figure were presented to their network without changing connecting weights.

1b). After the networks were trained to discriminate between them, some new images were presented. The network responded to some of the novel images that are similar to the stimuli used in training, which implies 'generalization' (Guttman & Kalish 1956; Hanson 1959). In some cases, a new stimulus gave a response even stronger than the image of the correct species used in the training, and hence it is a 'supernormal stimulus' (Staddon 1975). These images had a tail longer than the correct stimulus, and were thus an exaggerated form (figure 1d). The reaction for signals of different tail length had a peak shifted from the one that had been used in training. Enquist & Arak suggested that supernormal stimuli and 'peak shift' are a natural outcome of the sensory mechanism formed by training, and that a biased signal recognition system might be the basic mechanism for the evolution of exaggerated male secondary traits.

In this paper, we examine Enquist & Arak's (1993) system carefully and attempt to confirm their results by training a neural network by genetic algorithm. After training, the network was able to distinguish males of the correct species and wrong species. In addition, the trained neural networks show the greatest response to the pattern used for training, implying that there was no supernormal stimuli. We will discuss possible reasons for the discrepancy between Enquist & Arak's and our study.

2. NEURAL NETWORK MODEL

Neural networks are mathematical models imitating the function of a network of neurons (neural cells) in a computer. Neural network can learn to distinguish patterns. Each neuron, the unit of a neural network, receives n input signals x_1, x_2, \dots, x_n with weights w_1, w_2, \dots, w_n that may differ between inputs. A negative weight ($w_i < 0$) implies that the corresponding input has

an inhibitory effect. The simplest deterministic version of the model assumes that the neuron responds by outputting 1 when the weighted sum of the input signals is larger than a threshold denoted by h . If not, it outputs nothing. Due to noises from various sources, a neuron is likely to respond stochastically with a probability given by a sigmoid function of the sum of the input signals. We assume that the probability of getting the reaction 1 from the neuron is a sigmoid function:

$$\left[\begin{array}{l} \text{probability of} \\ \text{output}=1 \end{array} \right] = \frac{1}{1 + \exp \left[-a \left(\sum_{i=1}^n w_i x_i - h \right) \right]} \quad (1)$$

Figure 2 is the neural network used in Enquist & Arak (1993) and in the present paper. It consists of an artificial retina of 6×6 receptor cells, 10 cells in the hidden layer, and one output cell. Each cell in one layer is connected to every cell in the next layer, each connection being accompanied by a weight. The learning progresses by changing the weights appropriately. Male bird-like patterns were projected onto the retina (receptor cell layer), and a male is accepted by the female network as a mate when a positive response from the output cell is given.

The images used for training and testing of the networks are shown in figure 1 as filled squares on a 6×6 retina. The bird-like image figure 1c with tail length two was defined as a male of the same species, and the image of figure 1b with tail length one was considered as a male of a different or wrong species. The correct response of the female is to accept figure 1c and to reject figure 1b.

Random images were also used as reject images. Enquist & Arak did not describe clearly how to generate the random images they used. We made random images with five to eight filled squares (figure 3). We chose the number of filled squares with care to avoid the situation




Figure 2. The structure of the network. Input image was regarded as a male's pattern. A female's recognition system is modeled as a three-layered neural network. The input layer includes 6×6 receptor cells arranged on a two-dimensional square lattice, which are connected to ten cells in the hidden layer, and then to a single output cell. Each cell follows stochastic rule given by equation (1), where connection weights are modified during the training period (modified from Enquist & Arak 1993).

in which the total number of filled squares can be used as a useful cue to distinguish correct and wrong species.

When presented to the networks for training, both the patterns of tail length two (the correct species) and tail length one (wrong species) were moved and rotated on the retina, but the 20 random images were fixed. Considering the number of different positions and orientations, there are 48 possible configurations of tail length two, and 16 possible patterns of tail length one.

(a) Genetic algorithm learning

Back propagation is the learning method most frequently adopted when correct reactions (teach signals) of the network are given (Welstead 1994). Weights are appropriately changed to decrease the difference between the output by the network and the teach signal. However, we adopted an alternative learning method, called genetic algorithm (Davis 1990; Michalewicz 1994), that was used in the modelling evolution of signalling (Arak & Enquist 1993; Enquist & Arak 1993; Johnstone 1994). Recently, a neural network model learned by both genetic algorithm and back propagation was applied to a study of the artificial life (Toquenag *et al.* 1995).

We consider a string of connection weights of a neural network as one full set of genes (or genome), which can specify an individual network. We then consider a population of 20 individuals (or 20 strings of connection weights). Selective reproduction and mutation of these individuals allow the evolution of the network to make the neural network discriminate the images. The procedures are as follows.

1. Initial configuration. The initial values of the connection weights were generated by random numbers uniformly distributed between -1 and 1 .
2. Calculating scores. In every generation, various images were presented to each of the 20 networks in the population. Each network (or individual) experienced 150 input images per generation. Fifty of these images were correct species, which are chosen randomly (with resampling allowed) from 48 possible configurations, considering shift and rotation. Another 50 images are on the wrong species, which are chosen randomly from 16 possible configurations. The remaining 50 were then randomly sampled from 20 possible random images in figure 3. When the network responded (i.e. the output was 1) to the image of the correct species, it gained one score. If instead the network responded to the image of the wrong species or to random images, then its score was reduced by one.
3. Reproduction. When each of 20 networks in the population had experienced exposure to 150 images, we calculated the accumulated score of each network. The ten networks with the highest score were allowed to have two copies of themselves in the next generation. The ten networks with the lowest scores had no possibility of contributing to the next generation. In this way, the population size was kept to 20 all the time. In terms of population genetics, we adopted a 'truncation selection scheme with respect to the score.
4. Mutation then occurred on the connection weights. The probability of mutation for a particular weight was one-tenth and when mutation occurred, an increment drawn from a normal distribution (mean=0, standard deviation=0.4) was added to the weight.

These steps were repeated for 10 000 generations (training period). The networks converged to some kinds of equilibrium by then, and the average responses to learning tasks did not change so much even if the training period was doubled. The iteration of these procedures could develop a network that can discriminate males of the same and different species.

(b) Test period

After the training period, we presented various images, to the educated networks without modifying the weights, and examined their reaction to the test images. In the test period, the connection weights were fixed. The test images included the ones used in the training, but also those that have never been shown to the networks, as illustrated by figure 1a and figures 1d–g.

To count the strength of reaction of a network to an image, we chose 100 images randomly among all the patterns generated by the shifting and rotation of the original image. We then calculated the 'average response', the number of times in which the network gave output 1 among 100 trials. To examine the reaction of a particular pattern to 'random images', we examined 100 random samples (with resampling allowed) from the same 20 random reject images fixed throughout the training period.

In addition to these standard training procedures, we also examined the several different training schedules, as explained below.




Figure 3. Twenty random images used for training. These 20 images are reject images used in training period and test period. The number of filled squares ranged from five to eight. If we generate different random reject images every generation, instead of a fixed set, the speed of the learning of the network was much slower.

3. RESULTS

(a) *Average response*

In the initial population of 20 networks of randomly generated connection weights, the average response differed greatly between networks, but were similar between different images. Some networks responded to most of the input images, whereas others rejected most of the images. After the training period the average response became similar between networks in the population and different between images.

Figure 4a illustrates the average response of 100 trials for each of the 20 networks in the final population. The average response for the pattern with tail length two (correct species) was quite high (98.1%), and that for the pattern with tail length one (wrong species) was very small (2.6%).

(b) *Absence of supernormal stimuli*

In the test period, various images other than those used for training were presented to the networks (figure 1), but no image had a higher average response than the images of the tail length two (figure 4a). The average response to a pattern with a longer tail (tail length three) than the conspecific male (tail length two) was quite high, but was lower than that of the conspecific male (Wilcoxon signed rank test, $z = -3.926$, $p < 0.0001$). We did not therefore observe supernormal stimuli. This is very different from the result of Enquist & Arak (1993), in which the average response was always higher for an image with a longer tail.

Several replicates were calculated for the same set of parameters. The results described above held in all the cases. The connection weights of the networks obtained in

the final population were totally different, but they behaved in a very similar way.

As we examined only patterns with tails of integer length (e.g. tail length 1, 2, and 3), we may not be able to detect a peak shift of a small magnitude even if it exists. To overcome this difficulty, we examined responses of the trained network to patterns with a tail of non-integer length. For example, a pattern of tail length 1.75 implies a pattern with tail length 1 appended by a 'grey' unit with intensity of 0.75. Figure 4b illustrated the results, which shows the pattern with tail length two (correct species) caused the highest average response of the network, and hence there was neither peak shift nor supernormality.

(c) *Trajectory of learning procedures*

Figure 5 illustrates the time-course of the average response of a population to images of three classes: the images of correct species, those of wrong species and random images. In every generation 100 images from each of these three classes were presented to the networks, and the average responses in figure 5 are the proportion of times at which the network responded positively. The line with symbol II is for the response to the images of correct species (tail length two), which increases with time and reached a level of almost 100% at the end of the training period. The line with symbol I is for the average response of the network to the images of wrong species (tail length one), which became very low and close to zero. The reaction to random reject images with symbol R also became lower with time.

The broken line with symbol III shows the response by the network to the images of tail length three (with shift and rotation). The responses of the network to this class of




Figure 4. The average response of the trained network. (a) The acceptance rate was low (2.6%) for a wrong species image (b), and high (98.1%) for a right species image (a). (b) The reaction of the trained network to patterns with intermediate tail length. To express patterns with a non-integer valued tail length, a pattern of tail length 1.75 implies a pattern with tail length one appended by a 'grey' unit with intensity of 0.75. Horizontal axis is the tail length (1, 1.25, 1.5, ..., 3). The pattern with tail length two (used as the correct signal in training) caused the highest average response of the network, and hence there was neither peak shift nor supernormality.

images were just measured but not used for training the network (unlike the responses indicated by the other three lines).

The response to the image of tail length three wildly fluctuated with time but was consistently lower than the response to the image of tail length two (correct species). Hence there was no indication that supernormal stimuli might cause a stronger response than the stimuli used for training.

(d) *Different choice of random images*

In the standard procedures of training, we fixed 20 random images and used them as reject images throughout the training period. Alternatively, we may generate different random images in every generation. Figure 6a illustrates the average reaction of the network trained with such random reject images. The average reaction for the image of the conspecific males was 100%. However, the average response for the image of the wrong species was not small, remaining at around 20%,

even after 10 000 generations. The average reaction to random images remained at 60%. Hence we can conclude that the learning was not as effective if the random images changed in every generation.

Second, in the standard training procedures, we used 20 random images that include images of five to eight filled squares. This is so that the network cannot use the total number of black squares for the purpose of distinguishing the images.

We have also trained using 20 random images, which all have five filled squares. The results in figure 6b do not show clear differences from the standard training procedure of random images of five to eight filled squares (figure 4a). However, if random reject images are of two filled squares only, the trained networks show a similar but smaller response to some images with larger number of filled squares (figure 6c). For example, the difference between (c) and (f) in figure 6c was not significant (Wilcoxon signed rank test, $z = -0.784$, $p = 0.4331$).

Third, we have done training the networks without random images. The results are shown in figure 6d. The average response to the correct species was again significantly higher than that to the other images (Wilcoxon signed rank tests with Bonferroni adjustment, $|z| > 3.922$, $p < 0.0026$). Interestingly, the average response to the wrong species (b) was very low, but the responses to unfamiliar patterns were higher than the case random reject images were used. We suspect that this enhanced acceptance rate might be related to the reduced total number of reject images used during training.

(e) *Different number of images to present per generation*

In the standard training procedure, each network experiences 150 images every generation. We also examined cases with fewer images per generation for training. In a particular case, for example, we presented all images once per generation, so that 48 images of the correct species (including shift and rotation), 16 images of the wrong species (including shift and rotation) and 20 images of random reject images. The learning progressed somewhat slower than the standard training procedures. Interestingly, networks that accepted all the patterns increased and dominated the population temporarily, but then slowly disappeared. The final outcome of the evolution was almost the same as the standard training procedures.

4. DISCUSSION

There was a major difference between the network trained in the present paper and the one reported in Enquist & Arak (1993). We did not observe supernormal stimuli, as illustrated in figure 4a and 4b.

As the network structure is exactly the same in Enquist & Arak's study and ours, the difference must be due to the training procedures. From the comparison, we must conclude that Enquist & Arak's (1993) networks might have been trained less than perfectly. However, Enquist & Arak did not describe in detail the procedures they used for training the network. In the following we would like to list potential processes that might have caused the difference in the effectiveness of the learning processes,

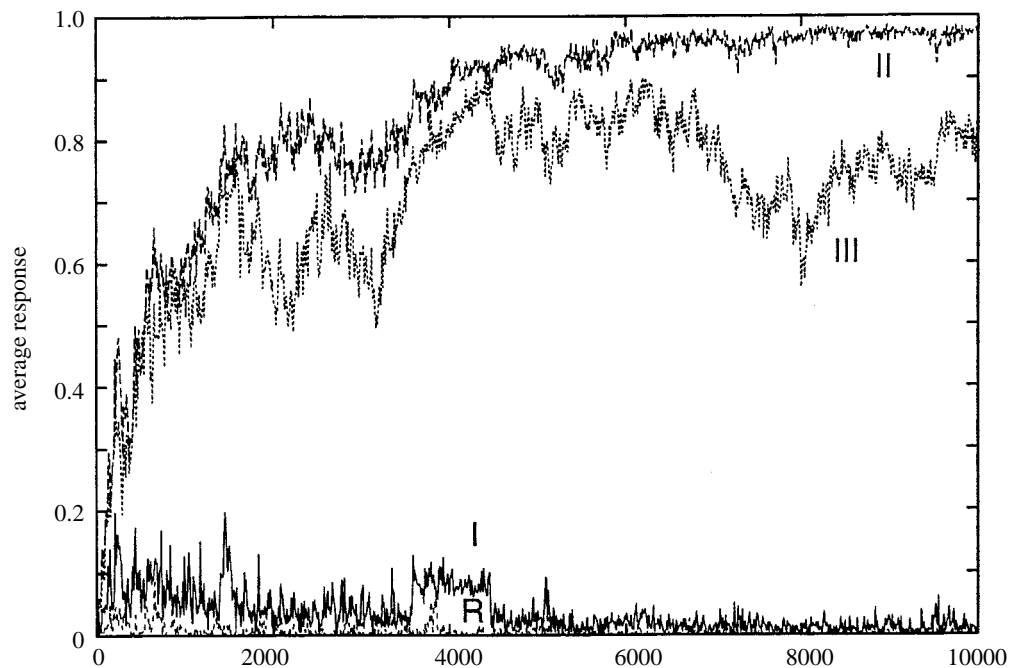


Figure 5. The average response changing with time during training period. Four lines are the response for the images of the correct species with tail length two (with symbol II), the image of a wrong species with tail length one (I), and random images (R). The average response is the fraction of times when the network output 1 among given classes of images. Connection weights were modified by genetic algorithm. The curve of a broken line (with symbol III) is the average response to the image of tail length three. The response to this class of images was not used for training.

irrespective of their applicability to Enquist & Arak's training procedures.

The first and simplest possibility is that they stopped training before the network achieved the full ability to distinguish the patterns. The second possibility is that they might have used a selection scheme that was less efficient than the truncated selection adopted in our training scheme. If we used a genetic algorithm with the fitness proportional to the contribution to the following generation, the training would be slower than in the case in which ten networks with the higher score remain and the other ten of lower score are discarded.

Third, the choice of random reject patterns used during training requires some care. In our training scheme, when we made random images we chose the number of filled squares so that the networks could not distinguish images based on the total number of filled squares. If we used random images including only a few filled squares, images with two wings longer than the correct species were accepted by the network nearly as often as the correct species (figure 6*c*). Training without using random images at all showed no indication of supernormality but generally high response to unfamiliar images (figure 6*d*).

Fourth, we fixed 20 random images that we used as reject images throughout the training and test periods (figure 3). If instead we used random images generated each generation, the learning efficiency was not very high (figure 6*a*).

Fifth, there are other relatively minor differences between Enquist & Arak and our networks. For example, the sigmoid curve for each neuron used by Enquist & Arak was an integral of a normal distribution, which is different from equation (1) used in our training, the latter being more commonly used in neural network models

(Welstead 1994). However, we do not expect this difference is the major reason of the difference in efficiency in training.

Sixth, a possible reason of the absence of supernormality is that the neural network might have been overtrained. Overtraining would reduce the generalization ability lower. To examine this possibility we ran a simulation in which four input patterns to accept (among 48 configurations considering shift and rotation) and one reject pattern (among 16 configurations) were not shown to the network during training. We then examined the average response of the networks at different times (1000, 2000, 3000, 4000 and 10 000 generations). The reaction to these images excluded in training fluctuate considerably, but the trend was clear: as training proceeds, the networks came to accept correct stimuli more often, hence generalization tendency increased rather than decreased. On the other hand they also accept one image of tail length one (wrong species) throughout the training period. In short, there was no indication of overtraining because the tendency of generalization did not decline with time.

In the present study, the average response to the image of tail length three was higher than other images except the one used for training (tail length two) (Wilcoxon signed rank tests with Bonferroni adjustment, $|z| > 3.922$, $p < 0.0004$). The image of tail length three was accepted because it resembles the conspecific male image. This illustrates that recognition by a trained neural network also shows generalization.

Generalization is a common property of recognition systems that classify unfamiliar stimuli into familiar categories. A classical example of generalization is that of a pigeon trained to perform a behaviour if it observed a light bulb of wave length 550 M μ (Guttman & Kalish




Figure 6. The average response of the trained network when different training schemes were adopted. (a) Different random images were chosen each generation. (b) All the random reject patterns have five filled squares, but are fixed throughout the training period. (c) All the random reject patterns have two filled squares, but are fixed throughout the training period.

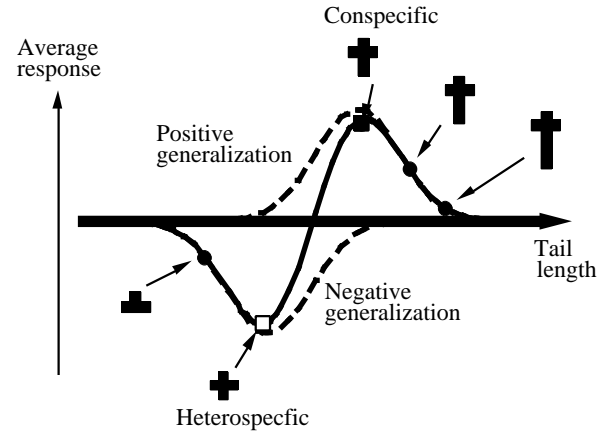


Figure 7. A scheme of generalization, summarizing the result of the learning simulation in the present paper.

1956; Hanson 1959). Naturally the response of the pigeon was highest when stimulated by the same colour. However, the pigeon also responded to light of similar wave length.

The results obtained for the present study can be understood intuitively in the following way (figure 7). The positive training by images of the correct species results in the network favouring patterns similar to them, illustrated by a broken line in figure 7. The resulting reaction of the trained network has some breadth around the peak reaction corresponding to the image of the correct species. The curve is symmetric around the peak, and hence the generalization is not biased. On the other hand, negative training by images of wrong species males would make the network avoid patterns similar to the image of wrong species, as indicated by another line in figure 7. A combination of these two tendencies would be simply the sum of the two curves, resulting in the observed pattern of the average responses to the tail images of different tail length (figure 7). Spence (1937) proposed a similar idea: a combination of positive and negative generalizations might explain an animal's response after training.

If this simplified picture holds, we may be able to infer the generalization pattern when choices of male images for correct and wrong species are different from the one adopted in the present paper. For example, suppose the images of tail length three are defined as conspecific species males and those of tail length one are defined as wrong species males. The networks can probably discriminate these two classes of images almost perfectly and the average response to the images of tail length four will be high, but probably not higher than the reaction to conspecific pattern (tail length three). Average response to the pattern of tail length two will be intermediate because it is affected by both the positive generalizations of the conspecific pattern of tail length four and the negative generalization of tail length one. Although the distances from the optimum (tail length three) of tail length two

The notations and symbols are the same as in figure 4a. The pattern with two longer wings (f) caused a similar response to the conspecific (c). (d) Training without using random reject pattern at all. The trained network tends to accept unfamiliar patterns at a higher rate than in the standard training procedure. These figures demonstrate the importance of the choice of random reject patterns.

and tail length four are the same, the acceptance rate was different. In this case, sensory bias could be more apparent. Whether the scheme in figure 7 can work in this or other situations will be an important subject of future study.

Whatever the reasons, we did not see supernormality in spite of our effort to reproduce the same procedures described in Enquist & Arak (1993), in which unfortunately the computational procedures were not explained well. In the present paper, we tried to present our work in sufficient detail that any interested reader can redo our training procedures. The case in which we have come closest to that reported in Enquist & Arak may be the case in which the network was trained with additional random reject patterns of very few filled squares, but even then the average reaction of an image with two longer wings ((f) in figure 6c) was not higher than the correct images ((c) in figure 6c). We must conclude that we have not obtained a convincing case of supernormality for various modifications.

Inspired by Enquist & Arak's pioneering work, there have been developed numerous theoretical studies on the property of sensory systems that are adjusted by training procedures (Johnstone 1994; Enquist & Arak 1994; Arak & Enquist 1995). Using the same models as Enquist & Arak (1993), Arak & Enquist (1993) trained networks to discriminate flowers that had petals of different lengths, which resulted in a bias in pollinators' preference. The evolution of symmetrical visual patterns was discussed in Enquist & Arak (1994) and Johnstone (1994). Although these works have received some criticisms as being too simplistic (Cook 1995; Dawkins & Guilford 1995), we believe that modelling preference and choice systems as a simple neural network that can be trained through experiences is a very promising approach in the study of biological signals. However, to derive robust and useful conclusions, it is important to examine the training procedures carefully, as we have done in this paper.

This work was supported in part by a Grant-in-Aid from the Ministry of Education, Science, Sports and Culture of Japan. We thank the following people for their comments on an early version of the paper: Anthony Arak, Mike Boots, Magnus Enquist, Isamu Kajitani, Hiroyuki Matsuda, Akira Sasaki, Yasuhiko Takeda, Yoshinari Tanaka and Shigeo Yachi.

REFERENCES

- Andersson, M. B. 1994 *Sexual selection*. New Jersey: Princeton University Press.
- Andersson, M. B. & Iwasa, Y. 1996 Sexual selection. *Trends Ecol. Evol.* **11**, 53–58.
- 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. 1995 Conflict, receiver bias and the evolution of signal form. *Phil. Trans. R. Soc. Lond. B* **349**, 337–344.
- Cook, N. D. 1995 Artifact or network evolution? *Nature* **374**, 313–314.
- Darwin, C. 1871 *The descent of man and selection in relation to sex*. London: John Murray.
- Davis, L. 1990 *Handbook of genetic algorithms*. New York: Van Nostrand Reinhold.
- Dawkins, M. S. & Guilford, T. 1995 An exaggerated preference for simple neural network model 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.
- Fisher, R. A. 1930 *The genetical theory of natural selection*. Oxford: Clarendon Press.
- Grafen, A. 1990 Biological signals as handicaps. *J. Theor. Biol.* **144**, 517–546.
- Guttman, N. & Kalish, H. I. 1956 Discriminability and stimulus generalization. *J. Exp. Psychol.* **51**, 79–88.
- Hamilton, W. D. & Zuk, M. 1982 Heritable true fitness and bright birds: a role for parasites? *Science* **218**, 384–387.
- Hanson, H. M. 1959 Effects of discrimination training on stimulus generalization. *J. Exp. Psychol.* **58**, 321–334.
- Iwasa, Y. & Pomiankowski, A. 1994 Evolution of mate preferences for multiple sexual ornaments. *Evolution* **48**, 853–867.
- Iwasa, Y. & Pomiankowski, A. 1995 Continual change in mate preferences. *Nature* **377**, 420–422.
- Iwasa, Y., Pomiankowski, A. & Nee, S. 1991 The evolution of costly mate preferences. II. The 'handicap' principle. *Evolution* **45**, 1431–1442.
- Johnstone, R. A. 1994 Female preference for symmetrical males as a by-product of selection for mate recognition. *Nature* **372**, 172–175.
- Kirkpatrick, M. & Ryan, M. J. 1991 The evolution of mating preferences and the paradox of the lek. *Nature* **350**, 33–38.
- Lande, R. 1981 Models of speciation by sexual selection of polygenic traits. *Proc. Natn. Acad. Sci. USA* **78**, 3721–3725.
- Michalewicz, Z. 1994 *Genetic algorithms + data structures = evolution programs*, 2nd edn. New York: Springer.
- Moller, A. P. 1994 *Sexual selection and barn swallow*. New York: Oxford University Press.
- Pomiankowski, A. & Iwasa, Y. 1993 Evolution of multiple sexual ornaments by Fisher's process of sexual selection. *Proc. R. Soc. Lond. B* **253**, 173–181.
- Pomiankowski, A., Iwasa, Y. & Nee, S. 1991 The evolution of costly mate preferences. I. Fisher and biased mutation. *Evolution* **4**, 1422–1430.
- Ryan, M. J. & Keddy-Hector, A. 1992 Directional patterns of female mate choice and the role of sensory biases. *Am. Nat.* **139**, S4–S35.
- Spence, K. W. 1937 The differential response in animals to stimuli varying within a single dimension. *Psychol. Rev.* **44**, 430–444.
- Staddon, J. E. R. 1975 A note on the evolutionary significance of 'supernormal' stimuli. *Am. Nat.* **109**, 541–545.
- Toquenag, Y., Kajitani, I. & Hoshino, T. 1995 Egrets of a feather flock together. *Artificial Life* **1**, 391–411.
- Welstead, T. 1994 *Neural network and fuzzy logic application in C/C++*. New York: Wiley.
- Zahavi, A. 1977 The cost of honesty (further remarks on the handicap principle). *J. Theor. Biol.* **67**, 603–605.