



Neural networks predict response biases of female túngara frogs

Steven M. Phelps and Michael J. Ryan

Department of Zoology, University of Texas at Austin, Austin, TX 78712, USA

Artificial neural networks have become useful tools for probing the origins of perceptual biases in the absence of explicit information on underlying neuronal substrates. Preceding studies have shown that neural networks selected to recognize or discriminate simple patterns may possess emergent biases toward pattern size or symmetry—preferences often exhibited by real females—and have investigated how these biases shape signal evolution. We asked whether simple neural networks could evolve to respond to an actual mate recognition signal, the call of the túngara frog, *Physalaemus pustulosus*. We found that not only were networks capable of recognizing the call of the túngara frog, but that they made remarkably accurate quantitative predictions about how well females generalized to many novel calls, and that these predictions were stable over several architectures. The data suggest that the degree to which *P. pustulosus* females respond to a call may often be an incidental by-product of a sensory system selected simply for species recognition.

Keywords: neural networks; animal communication; *Physalaemus pustulosus*; sexual selection; mate choice

1. INTRODUCTION

Research in animal communication has long sought an explanation for the extraordinary diversity of signal form. Strategic analyses employing game theory have been successful in predicting the contexts in which an animal will convey particular information, and in predicting the general forms signals ought to take, but ultimately say more about the content of signals than about their structure (Krebs & Dawkins 1984; Dawkins & Krebs 1978; Cheney & Seyfarth 1991; Hauser 1996). Signal detection theory describes the sorts of signals that may propagate successfully through a noisy world, predicting general attributes of signal structure, like the colours and contrasts of pigments (Endler 1992), or some temporal and frequency attributes of songs and calls (Brenowitz 1982; Littlejohn 1977; Römer & Bailey 1986). Nevertheless, the question remains: why so many intricate forms? Is the diversity and complexity of signals solely a function of intent and interference?

In the last decade, researchers have placed a strong emphasis on the role receiver biases may play in driving the diversity of signal form (Guilford & Dawkins 1991; Endler 1992; Ryan 1990; West Eberhard 1979). One recent development has been the introduction of connectionist methods for exploring ‘receiver psychology’ (Enquist & Arak 1993; Johnstone 1994). The models known as artificial neural networks are in many ways ideal tools for formulating hypotheses about the nature and origins of biases in the nervous systems of receivers. An artificial neural network consists of interacting neuron-like computational units that sometimes behave remarkably like a real nervous system. This convenient attribute, combined with the relative simplicity of the networks, permits

researchers to investigate how tasks such as signal recognition might be performed. Once networks have been trained to recognize a particular signal, for example, one can probe successful networks to determine whether they recapitulate the errors and generalizations made by test subjects, to produce novel predictions about the responses of receivers, and ultimately to generate hypotheses concerning the neural substrates of signal recognition. Researchers in animal communication may augment this approach by replacing standard training algorithms, such as ‘back-propagation’, with a genetic algorithm. The genetic algorithm uses selection, mutation and recombination to train networks, and so enables one to model how receiver responses change as a function of selection pressures and population parameters.

To date, researchers in animal communication have successfully employed neural networks to investigate how selecting networks to perform tasks analogous to species recognition might lead to more general biases (reviewed in Enquist & Arak 1998). Several of these studies have demonstrated that networks selected to recognize one pattern, or to discriminate between two patterns, show emergent preferences for traits—such as pattern length and symmetry—without having been selected to cue in on the ‘good genes’ these traits are often thought to advertise (Enquist & Arak 1993, 1994; Arak & Enquist 1993; Johnstone 1994). The main strength of this approach is that the simulations can be used to make qualitative predictions about many communication systems (Arak & Enquist 1993; Hurd *et al.* 1995; Enquist & Arak 1998; Johnstone 1994). Unfortunately, in doing so one sacrifices the capacity to make quantitative predictions about the behaviour of particular animals. In the absence of data on the external validity of the models, signal and receiver

simplicity is often construed as an impediment rather than an asset (Kirkpatrick & Rosenthal 1994; Dawkins & Guilford 1995).

We asked if simple neural networks could evolve to recognize a real mate recognition signal—the call of the túngara frog, *Physalaemus pustulosus*. We then tested whether neural networks that were able to recognize the túngara frog call would generalize to novel calls, as do real females. To assess the biological validity of the networks, we compared the responses of females in a two-way choice test (novel call versus noise) to the responses of neural networks. And last, we asked if the ability of the networks to predict female responses was robust to manipulations in the network architecture.

2. METHODS

(a) *Network architecture*

The architecture consisted of four groups, or layers, of neurons. The input layer neurons respond to select frequency ranges and pass this information on to the neurons of the feature detector layer. This layer processes the stimuli and influences the activity of the context layer. The context layer feeds back on the neurons of the feature detector layer, enabling the feature detector neurons to make responses to current sound stimuli contingent on preceding stimuli. This recurrence permits the extraction and recognition of temporal features in the sound stimulus. Neurons of the feature detector layer influence the activity of the output neuron, and it is the activity of the output neuron that ultimately defines a network's response to a stimulus.

This recurrent network architecture is essentially an Elman network (1990) modified to relax the assumption that context neurons have perfect information about the activities of feature detector neurons in a preceding time-step. This was accomplished by replacing the context layer of an Elman network—in which the activity of an individual context neuron is identical to the activity of a corresponding feature detector neuron during the preceding time interval—with context neurons that receive input from all neurons of the feature detector layer. The context neurons are now allowed to weight these inputs differently and compute their activations based on the sum of these inputs.

The input layer consisted of 15 neurons, each responding selectively to a frequency range of 86.93 Hz, together spanning from 261 Hz to 1565 Hz. Each of these neurons projected to every neuron in the feature detector layer. In the standard architecture, the feature detector layer consisted of 12 neurons: each received input from every neuron in the input layer and the context layer; each sent efferent projections to the 12 neurons of the context layer and to the output neuron (figure 1). Each neuron in the network also received input from a bias neuron (not shown in the schematic) that was tonically active (activity=1.0).

The activity of individual neurons was defined by the sigmoid activation function

$$A \equiv \frac{1}{1 + e \left[-2.5945 * \left(\sum_{i=1}^n \omega_i * \alpha_i - 1.5 \right) \right]},$$

where ω_i and α_i represent the weight and activity of the input neuron i . The coefficients -2.5945 and -1.5 were chosen *a priori* to provide the neurons with low basal levels of activation in the absence of input, and to saturate activation at high input.




Figure 1. Schematic of standard neural network architecture. Each network consisted of neurons connected as shown above, with large arrows indicating that each neuron of the target layer receives information from every neuron in the source layer. On the upper left hand corner of the figure there are three spectrograms presented one at a time to the networks, with each stimulus presented at a different time in the stimulus window. Each pixel of a spectrogram represents a particular time and frequency interval. The shade of a pixel reflects the value of one FFT coefficient of the sound—darker pixels correspond to larger coefficients. To the left of the input layer is a row of pixels representing such coefficients for 15 frequency intervals and a single time interval. In each time interval, the activity of each neuron is calculated by taking the weighted sum of the activities of its inputs in the preceding time interval, and entering this number into the activation function (see text). The activity of the output neuron at the end of the stimulus window is the output of the network. Additional simulations were run using architectures with varied numbers of neurons in the feature detector and context layers.

As in the standard implementation of the Elman network (1990), the activation of a neuron did not accumulate over time (though current studies are exploring the ramifications of modifying this parameter). Weights varied from -0.6 to $+0.6$, coded as six-bit numbers, with a discrete uniform distribution ($\mu=0$). (Neurons of the input layer did not weight sound stimuli, though they did weight input from a bias neuron. The bias neuron enabled input layer neurons to differ in threshold responses to their respective frequencies; variation in the frequency tuning of ascending auditory fibres is known to influence anuran auditory processing and mate choice (Ryan *et al.* 1990).)

A neuron of the feature detector layer, for example, possesses 28 weights for each of its 28 inputs: one tonically active bias neuron, 15 neurons of the input layer, and 12 neurons of the context layer. This neuron multiplies each weight times the activity of its respective input, sums these values, and enters the weighted sum into the activity function. Simultaneously, all other neurons in the network calculate their activities based on their inputs and weights. This yields new activity levels for the neurons of the network, which are then used to recalculate activities in the subsequent time step. The activity of the output

neuron at the end of the stimulus window is recorded as a network's response to a stimulus.

To assess whether the responses of networks were stable over various network architectures, we varied the size of the feature detector and context layers. 'Small' networks consisted of eight neurons in each of these layers. 'Large' networks consisted of 16 neurons in each of these layers. These architectures were otherwise identical to the 'standard' neural network depicted in figure 1.

(b) Stimulus synthesis

Calls were synthesized at a sampling rate of 20 kHz using an AMIGA 3000 computer and a sound synthesis program developed by J. Schwartz for use with the software package *FutureSound*, as previously described (Ryan & Rand 1993). Synthetic calls were then recorded on a Macintosh using Canary software (Bioacoustics Research Program, Cornell Laboratory of Ornithology). We generated spectrograms of these calls using a filter bandwidth of 352.94 Hz, a framelength of 256 points, grid resolution of 86.93 Hz by 11.5 ms, logarithmic amplitude, a sampling frequency of 22 kHz, and a Hamming window function. Because calls were synthetic, extraneous noise could be easily identified and clipping level adjusted appropriately. Three spectrograms for each call were then extracted from the file and fast Fourier transform (FFT) coefficients were rescaled such that the smallest coefficients in each spectrogram were equal to zero, and the largest coefficient was equal to 4.0. This ensured that all stimuli had identical peak coefficients, little noise, and fell into a range appropriate for use with the activation function chosen for the neural network model. Spectrograms generated from the same calls but with different clipping levels produced very similar responses from networks. The túngara frog call resulted in a set of three matrices, each comprised of 15 rows corresponding to each frequency range, and 32 columns corresponding to each time interval; novel stimuli also had 15 frequency rows, but varied in the number of time columns that were produced.

During each generation of the evolutionary simulations, each túngara frog call was placed at a random time in a window of 70 time bins, resulting in three 15×70 matrices (figure 1). Noise stimuli were constructed *de novo* at every stimulus presentation by randomly reassigning the frequency coefficients within each time interval, producing noise in a matching amplitude envelope. Selecting for networks that ignored this stimulus ensured that networks would not recognize calls by acting as simple high-pass filters—a possible but trivial solution to the task of discriminating a call from white noise. For calls and corresponding noises, 2% of the 15×70 matrix entries were randomly selected and assigned arbitrary values between zero and four. Thus, during the evolutionary simulation there were two sources of noise: networks were selected to respond to a call in various positions of the time window, with some noise overlaid; and networks were selected to ignore a noise stimulus containing sound in a similar amplitude envelope. Although we have begun studies examining the consequences of using other noise stimuli, such as environmental background noise or the calls of sympatric species, these stimuli would confound our investigation of species recognition by adding selection to overcome transmission constraints and discriminate between species.

The two sources of noise used during the evolutionary simulation provide the variation in the 'training set' necessary for network generalization. The variation introduced by the overlaid noise alone produces approximately 1.14×10^{63} potential signals, of which only three are used in each generation. Because of the

vast number of potential signals, the repetition of patterns is negligible, and the size of the training set used in a run is approximately three times the number of generations in a run of the simulation (small networks, $\mu=1577$ generations; standard networks, $\mu=997$; large networks, $\mu=840$).

When testing the responses of networks to novel calls (see §2d below), we sought to use a noise stimulus comparable to the 'white noise' used in female phonotaxis experiments. The testing algorithm first computed the average of the 15 frequency coefficients in each time interval of the call. Each frequency coefficient of a noise stimulus was drawn from a uniform distribution between 0 and twice the call's average frequency coefficient in the corresponding time bin (note $\mu_{\text{noise}} = \mu_{\text{call}}$). This results in a noise stimulus whose coefficients are more evenly distributed than either the true call or the noise used during the evolutionary simulation, but retains the appropriate amplitude envelope. During the testing phase, no noise was added to the calls themselves.

(c) Genetic algorithm

Each evolutionary simulation began with the generation of 100 networks at random. Each network was represented by a single binary string called a 'chromosome.' Chromosomes of the standard architecture consisted of 3282 bits, representing 547 network weights, each coded as a six-bit string. We evolved populations of networks using a modified version of Goldberg's simple genetic algorithm in C (Smith *et al.* 1994). Networks encoded as chromosomes were selected for a capacity to discriminate calls from noise in the same amplitude envelope, with the fitness function defined as

$$F \equiv \sqrt{\sum_{i=1}^n (C_i - N_i)^2 / n + 0.01}$$

where F is fitness, C_i is the response of the network to call i , N_i is the response of the network to noise i , and n is the number of calls (3). (The response of the network was defined as the activity of the output neuron at the end of the stimulus window. If the response of the network to noise exceeds the response to a call, the response $C_i - N_i$ was recorded as zero.) Because we used a proportional selection method, described below, minute differences in initial fitness could cause loss of most genetic diversity very early in the simulation (data not shown); a small external fitness component of 0.01 was added to slow this loss and thereby speed the algorithm.

The algorithm drew 100 chromosomes with replacement from the starting population. For each draw, the probability that a particular chromosome would be chosen was equal to its fitness relative to the rest of the population (a proportional selection regime known as roulette wheel selection). These 100 chromosomes were selected in 50 pairs; pairs were recombined with a probability of 0.5. Recombination was equally likely at any point along the chromosome. Each bit of the resulting 100 chromosomes was mutated with a probability of 0.001. This yielded a population of daughter chromosomes that were again screened for their capacity to recognize the túngara frog call. The simulation ran until the population recognized the *P. pustulosus* call at predetermined criteria (fitness of single most fit individual in the population > 0.90 , average fitness > 0.75) for two consecutive generations.

The simulation was run 20 different times for the standard architecture, ten times for the small architecture, and ten times for the large architecture. At the end of each run, the

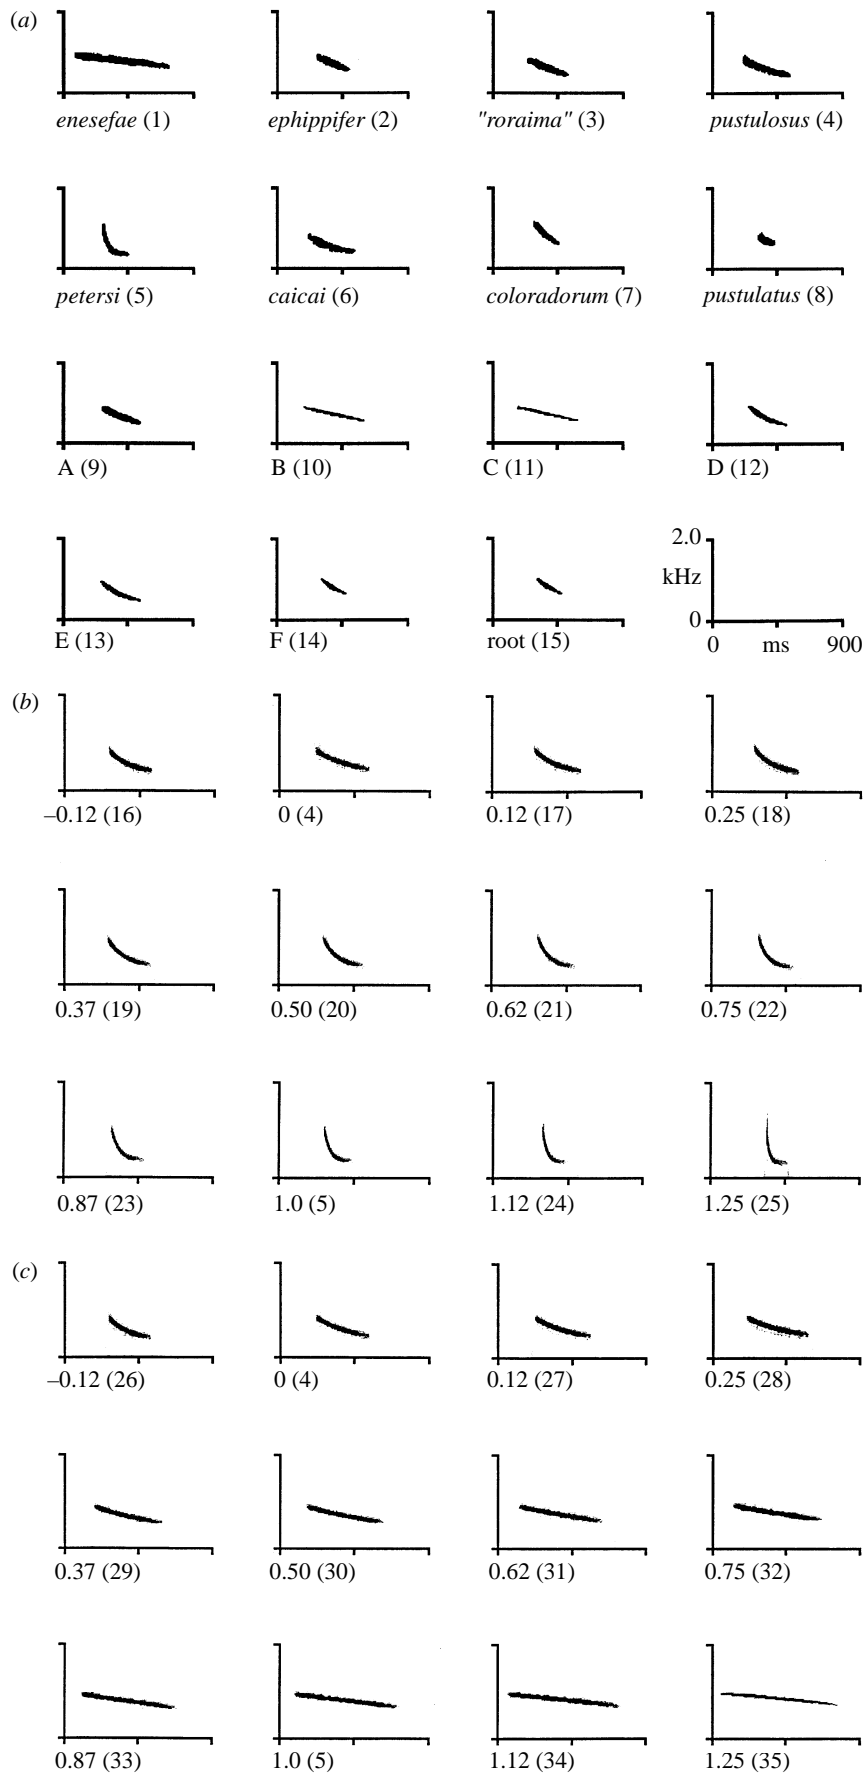


Figure 2. Stimuli presented to the neural networks during testing. (a) Sonograms of the extant and ancestral calls of the *P. pustulosus* species group, as reported in Ryan & Rand (1995). The empty frame on the right depicts the scale of sonograms used in panels a, b, and c. (b) A call transect representing intermediate states between *P. pustulosus* and *P. petersi*, with 'caricatures' (calls 16, 24 and 25) formed by overshooting the axis of variation bounded by the two calls. Intermediate and caricature calls were constructed by

weights of the single best network in the last generation were frozen, recorded, and tested for their responses to a suite of novel calls.

The simulations were implemented in an AIX 4.0 environment on IBM RS/6000 computers provided by the computer Science Department at the University of Texas at Austin.

(d) *Network testing*

We measured how well each of the 40 networks, representing separate runs of the simulation, could discriminate test stimuli from noise in a corresponding amplitude envelope. Stimuli were presented at the beginning, centre, and end of the stimulus window. Network responses were calculated using the fitness function defined in §2c above, in which C now represents the response to the novel call being tested, and N the response to 'white noise' in a matching amplitude envelope. Test stimuli consisted of 15 calls representing the extant and hypothetical ancestral calls of the *P. pustulosus* clade (fig. 2a, Ryan & Rand 1995), and 19 calls representing various intermediate states between the túngara frog call and the calls of *P. enesefae* and *P. petersi* (figure 2b,c) for which female response data were available. For each of the three network types, the average network response was considered an estimate of the probability that an individual female would approach the stimulus.

(e) *Female responses*

Female response data were obtained from phonotaxis experiments in which each call was paired with white noise in the same amplitude envelope, as previously described (Ryan & Rand 1995). Recent data suggest that the proportion of females approaching a speaker in a choice test represents the probability that any individual female will approach the speaker on repeated tests (Kime *et al.* 1998). Female responses to the 15 extant and hypothetical ancestral calls of the *P. pustulosus* clade have been published elsewhere (Ryan & Rand 1995).

3. RESULTS

Of the 40 runs of the simulation, every random initial population was able to meet the call recognition criteria within 2200 generations (small $\mu=1577$, s.d.=414; standard $\mu=997$, s.d.=168; large $\mu=840$, s.d.=292). The results from a run representative of the standard networks are shown in figure 3, which depicts the single best and average fitness of each generation.

Testing of network responses to novel stimuli versus noise revealed that the networks are able to generalize to many novel calls (figure 4). We compared the average response of the standard 20 networks for each of the 34 test stimuli (shown in figure 2) to the proportion of females responding to the corresponding calls in phonotaxis experiments. Standard networks were able to predict the preferences of females remarkably well. The resulting regression showed that the responses of the networks could explain 65% of the variation in the responses of real females (figure 4; $y=0.62x+0.10$, $R^2=0.65$, $p\ll 0.0001$). Thus selection for

mate recognition is sufficient to predict female responses to many novel stimuli.

This result proved to be extremely robust to manipulations of architecture size. Small networks yielded a regression ($y=0.60x+0.10$, $R^2=0.65$, $p\ll 0.0001$) remarkably similar to the standard network regression, as did large networks ($y=0.64x+0.13$, $R^2=0.67$, $p\ll 0.0001$). Neither of these regressions were significantly different from the regression equation produced by the standard networks ($p>0.05$) and the responses of both were highly correlated with the responses of the standard networks (small networks, $r=0.99$, $p\ll 0.001$; large networks, $r=0.97$, $p\ll 0.001$). The responses of small and large networks have not been shown because of their extreme redundancy with standard network responses.

4. DISCUSSION

Simple neural networks are reliably able to recognize the call of the túngara frog, and to discriminate it from noise in the same amplitude envelope. This is the first time, to the authors' knowledge, that a simple neural network has been shown to recognize a mate recognition signal capable of eliciting responses from females. These data are consistent, however, with reports indicating that recurrent neural networks are capable of recognizing time-varying signals from a variety of domains, including human speech (Tank & Hopfield 1987; Elman 1990; Lippman 1989).

Neural networks selected to recognize the call of the túngara frog generalize to novel stimuli with varying strengths. This is particularly encouraging, since a receiver's capacity to generalize to some stimuli while ignoring others defines the dimensions along which a signal may vary. If receiver biases influence signal evolution, signal diversity ought to be partly explained by these dimensions of receiver preference.

Interestingly, there are data suggesting that the relaxation of receiver perceptual constraints permitted the diversification of calls in different groups of anurans. One would predict that lineages of receivers with larger perceptual landscapes would provide more opportunity for signal variation and divergence. Given that changes in mating call are an important behavioural isolating mechanism (Blair 1964), increases in the size of the perceptual landscape should translate into higher speciation rates (Ryan 1986). Ryan (1986) reported exactly this correlation: anuran lineages with amphibian papillae that respond to a larger frequency range—and hence place less rigid restrictions on signal divergence—are reliably more speciose. Similarly, a number of researchers have argued that when receiver preferences extend beyond existing signal variation, novel traits able to exploit these biases may be added to the sender's repertoire (West Eberhard 1979; Basolo 1990; Ryan 1990; Endler 1992).

Arguably the most exciting find in this study is the remarkable correlation between the responses of neural

Figure 2 (*continued*) simply changing each of the call variables used in call synthesis by 1/8 the difference between the *P. pustulosus* and *P. petersi* calls. (The calls of panels (b) and (c) are part of an ongoing study. Call 16 was excluded from the regression analysis because female choice data are not yet available, but is presented in this figure for completeness.) (c) A similar call transect for the axis of variation defined by *P. pustulosus* and *P. enesefae*. The numbers in parentheses correspond to the numbers given in the regression provided in figure 4.




Figure 3. A representative run of the evolutionary simulation using the standard network architecture. The y -axis depicts the fitness of a network in the population, the x -axis the generation. A grey circle represents the average fitness of all 100 networks in a generation, a black circle the single best fitness in a generation. Twenty such simulations were run with this architecture. Ten additional simulations were run with the large architecture, and ten more with the small architecture.




Figure 4. Regression of standard neural network responses and female responses. The x -axis represents the average response of 20 standard neural networks to a call when compared to noise in the same amplitude envelope. The y -axis describes the proportion of females approaching a speaker playing the call when it is paired with another speaker playing noise in the same amplitude envelope. The numbers indicate which stimulus of those depicted in figure 2 was used to obtain each datum. $y = 0.62x + 0.10$, $r = 0.80$, $R^2 = 0.65$, $p \ll 0.0001$.

networks and the responses of females in phonotaxis experiments. This demonstrates an external validity that satisfies the criticisms levied against earlier studies, 'that future models should show that they can adequately mimic animal perception before they are used to simulate the evolutionary forces produced by animal perception in the design of animal signals,' (Dawkins & Guilford 1995). Moreover, this predictive power is not compromised by manipulating details of the architecture.

The overall similarity between the responses of networks and females suggests that selection for species recognition is sufficient to explain many of the perceptual biases exhibited by real females. The dimensions of female preference in this species appear to have little to do with the information the investigated calls might convey, or with the capacity of these calls to penetrate a noisy environment. While both game theory and signal detection theory are viable explanations for general attributes of signal structure and receiver preference (Hauser 1996), they seem less capable of explaining certain differences in signal form among related taxa.

Last, the similarity between networks and females suggests an underlying similarity in information processing. Current studies are aimed at understanding how networks are solving the task of mate recognition, and how varying selection pressures affect the properties of networks. The neural network models proposed here, with their small number of assumptions about neuronal properties and selection pressures, may serve as null models for understanding biases in call recognition. Discrepancies between the responses of networks and females may suggest selection on the female auditory system to cue in on particular traits, selection to overcome particular environmental transmission constraints, or the absence of an important proximate constraint in the design of the model. In short, neural networks that accurately predict receiver responses may help to reveal how particular signals are processed, and how this processing changes as a function of selection. These questions are central to any deep understanding of animal communication.

We thank M. S. Dawkins, M. Enquist, W. Goetz, M. Kirkpatrick, R. Miiikkulainen, A. S. Rand, W. Wilczynski, and two anonymous reviewers for their useful comments and suggestions, and J. Burgess for technical advice. This research was supported by NIMH F31 MH11194 to S.P., and a Guggenheim Fellowship to M.J.R.

REFERENCES

- Arak, A. & Enquist, M. 1993 Hidden preferences and the evolution of signals. *Phil. Trans. R. Soc. Lond. B* **340**, 207–213.
- Basolo, A. 1990 Female preference predates the evolution of the sword in swordtails. *Science* **250**, 808–810.
- Blair, W. F. 1964 Isolating mechanisms and interspecies interactions in anuran amphibians. *Q. Rev. Biol.* **39**, 334–344.
- Brenowitz, E. A. 1982 The active space of red-winged blackbird song. *J. Comp. Physiol. A*, **147**, 511–522.
- Cheney, D. L., and Seyfarth, R. M. 1991 Truth and deception in animal communication. In *Cognitive ethology* (ed. C. A. Ristau), pp. 127–152. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Dawkins, M. S. & Guilford, T. 1995 An exaggerated preference for simple neural network models of signal evolution? *Proc. R. Soc. Lond. B* **261**, 357–360.

- Dawkins, R. & Krebs, J. R. 1978 Animal signals: information or manipulation? In *Behavioural ecology: an evolutionary approach* (ed. J. R. Krebs & N. B. Davies), pp. 282–309. Oxford University Press.
- Elman, J. L. 1990 Finding structure in time. *Cog. Sci.* **14**, 179–211.
- Endler, J. 1992 Signals, signal conditions, and the direction of evolution. *Am. Nat.* **139**, S125–S153.
- 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.
- Enquist, M. & Arak, A. 1998 Neural representation and the evolution of signal form. In *Cognitive ecology* (ed. Reuven Dukas). Chicago University Press.
- Guilford, T. & Dawkins, M.S. 1991 Receiver psychology and the evolution of animal signals. *Anim. Behav.* **42**, 1–14.
- Hauser, M. D. 1996 *The evolution of communication*. Cambridge, MA: MIT Press.
- Hurd, P. L., Wachtmeister, C.-A. & Enquist, M. 1995 Darwin's principle of thesis and 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.
- Kime, N., Rand, A. S., Kapfer, M. & Ryan, M. J. 1998 Consistency of female choice in the túngara frog: a permissive preference for complex characters. *Anim. Behav.* (In the press.)
- Kirkpatrick, M. & Rosenthal, G. 1994 Symmetry without fear. *Nature* **372**, 134–135.
- Krebs, J. R. & Dawkins, R. 1984 Animal signals: mind reading and manipulation. In *Behavioural ecology: an evolutionary approach* (ed. J. R. Krebs & N. B. Davies), pp. 380–402. Oxford University Press.
- Lippmann, R. P. 1989 Review of neural networks for speech recognition. *Neural Computation* **1**, 1–38.
- Littlejohn, M. J. 1977 Long range acoustic communication in anurans: an integrated evolutionary approach. In *The reproductive biology of amphibians* (ed. D. H. Taylor & S. I. Guttman), pp. 263–294. Cambridge University Press.
- Römer, H. & Bailey, W. J. 1986 Insect hearing in the field: II. Male spacing behavior and correlated acoustic cues in the bushcricket *Mygalopsis marki*. *J. Comp. Physiol. A* **159**, 627–638.
- Ryan, M. J. 1986 Neuroanatomy influences speciation rates among anurans. *Proc. Natn. Acad. Sci. USA* **83**, 1379–1382.
- Ryan, M. J. 1990 Sexual selection, sensory systems, and sensory exploitation. *Oxford Surv. Evol. Biol.* **7**, 157–195.
- Ryan, M. J. & Rand, A. S. 1993 Species recognition and sexual selection as a unitary problem in animal communication. *Evolution* **47**, 647–657.
- Ryan, M. J. & Rand, A. S. 1995 Female responses to ancestral advertisement calls in túngara frogs. *Science* **269**, 390–392.
- Ryan, M. J., Fox, J. H., Wilczynski, W. W. & Rand, A. S. 1990 Sexual selection for sensory exploitation in the frog *Physalaemus pustulosus*. *Nature* **343**, 66–67.
- Smith, R. E., Goldberg, D. E. & Earickson, J. A. 1994 SGA-C: a C-language implementation of a simple genetic algorithm. <ftp://www.aic.nrl.navy.mil/pub/galist/src/sca-c.tar.z>.
- Tank, D. W. & Hopfield, J. J. 1987 Neural computation by concentrating information in time. *Proc. Natn. Acad. Sci. USA* **84**, 1896–1900.
- West Eberhard, M. J. 1979 Sexual selection, social competition, and evolution. *Proc. Am. Phil. Soc.* **123**, 222–234.

