

Evolutionary optimization and neural network models of behavior

Marc Mangel

Zoology Department and Center for Population Biology, University of California,
Davis, CA 95616, USA

Abstract. One of the main challenges to the adaptationist program in general and the use of optimization models in behavioral and evolutionary ecology, in particular, is that organisms are so constrained by ontogeny and phylogeny that they may not be able to attain optimal solutions, however those are defined. This paper responds to the challenge through the comparison of optimality and neural network models for the behavior of an individual polychaete worm. The evolutionary optimization model is used to compute behaviors (movement in and out of a tube) that maximize a measure of Darwinian fitness based on individual survival and reproduction. The neural network involves motor, sensory, energetic reserve and clock neuronal groups. Ontogeny of the neural network is the change of connections of a single individual in response to its experiences in the environment. Evolution of the neural network is the natural selection of initial values of connections between groups and learning rules for changing connections. Taken together, these can be viewed as “design parameters”. The best neural networks have fitnesses between 85% and 99% of the fitness of the evolutionary optimization model. More complicated models for polychaete worms are discussed. Formulation of a neural network model for host acceptance decisions by tephritid fruit flies leads to predictions about the neurobiology of the flies. The general conclusion is that neural networks appear to be sufficiently rich and plastic that even weak evolution of design parameters may be sufficient for organisms to achieve behaviors that give fitnesses close to the evolutionary optimal fitness, particularly if the behaviors are relatively simple.

Key words: Behavior — Evolution — Neural network — Dynamic optimization

1. Introduction

Among the main challenges to the "adaptionist program" (Gould and Lewontin 1979) and to the use of optimality models in biology are the assertions that (i) natural selection does not optimize and (ii) organisms are so constrained by structure that they may not be able to obtain optimal solutions, however such solutions are defined. This paper and an accompanying one (Mangel 1989) are attempts to answer these challenges, using behavioral modeling as an example. Mangel (1989) considers the very general problem of habitat selection and asks when optimal behavior will arise as a result of natural selection. Here optimal is interpreted to mean the sequence of behaviors that maximize individual fitness. Thus, neither ESSs nor co-evolution are considered.

This paper treats two specific problems—the movement by a tubular worm and learning by tephritid fruit flies—and asks if the evolutionary optimum can be achieved by an organism given its biological constraints. Two modeling techniques are used. The first is the "evolutionary optimization" approach, based on dynamic state variable modeling of behavior (Mangel and Clark 1988). In this approach Darwinian fitness is defined in terms of survival and reproduction. State variables are used to link physiology, ecology and fitness. I consider an interval of length T and assume that evolution has selected behaviors in the interval that maximize the "terminal" fitness at time T (Mangel and Clark 1988); this provides a means for computing both fitness and optimal behaviors.

The second modeling approach is based on "neural networks". (See, e.g. Brady 1985; Bounds 1987; Dehane et al. 1987; Edelman 1987 or Cowan and Sharp 1988 for a discussion of neural networks; the paper by Cowan and Sharp is especially biological.) In this approach, a neural network that solves the biological problem is developed. The strength of connections between different neural groups and "learning rules" for changing those connections provide a "template" for behavior. Natural selection causes learning rules to evolve. Once the template is given, ontogeny determines the detailed nature of connections between different neural groups and the feedback that different groups provide to each other (see Aoki and Siekevitz 1988 for a discussion of plasticity in brain development). The model based on evolutionary optimization provides a standard against which the model based on neural networks can be assessed. In particular, we ask: can an organism, modeled by a neural network, achieve fitness that is close to optimal, as defined by the evolutionary optimization model?

In Sect. 2, the evolutionary optimization model is described for a polychaete worm. This animal is very simple and there exists abundant information about it (Evans 1966, 1971, 1981; Evans and Downie 1986). The neural network model for this worm is described in Sect. 3 and the "evolution" of the neural network model is described in Sect. 4, using Monte Carlo simulation to mimic natural selection. In Sect. 5, the ontogeny of the evolved neural network is modeled, and the fitness of the neural network model is compared to the fitness of the evolutionary optimization model. In Sect. 6, I describe more complex evolutionary optimization models for polychaetes. Section 7 describes learning problems

of tephritid fruit flies such as the apple maggot *Rhagoletis pomonella*. Finally, Sect. 8 contains discussion and conclusions.

2. Polychaete behavior: evolutionary optimization model

Polychaete (tubular) worms often have a life pattern (e.g. Evans 1966, 1971, 1981; Evans and Downie 1986) that consists of moving in and out of a tube in response to hunger and to visual stimulation. I assume that the worm can sense shadows, but does not have the ability to detect images. A shadow can be caused by a piece of food or a piece of debris drifting by the tube or a predator moving by the tube. If the worm exits its tube in response to the shadow and the shadow is caused by a predator, the worm is killed. If the shadow was caused by a piece of food, then food is ingested and the worm's reserves are increased by the network energetic value of the food. This simple life cycle continues from period $t = 1$ until period $t = T$, at which time reproduction occurs.

The evolutionary optimization model is used to compute the set of behaviors (rules for exiting from the tube) that maximize expected lifetime reproduction at time T (Mangel and Clark 1988). Introduce a state variable $X(t)$ defined by

$$X(t) = \text{level of reserves of the worm at the start of period } t \quad (1)$$

The reserves are bounded with $x_c < X(t) \leq x_m$, where x_c denotes a critical, starvation level of reserves such that the worm dies if its reserves fall below this level and x_m denotes a maximum level of reserves (e.g. limited by the size of the tube). In Sect. 5, the effects of letting $x_m \rightarrow \infty$ are discussed.

At the start of period T , terminal fitness $R(X(T))$ is assessed. In general, this function will be some measure of resources available for reproduction or reproduction itself. For the computations used in this paper, I chose a linear function of reserves with a threshold:

$$R(X(T)) = \begin{cases} r(X(T) - x_r) & \text{if } X(T) > x_r \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where x_r denotes a minimum reserve level needed for reproduction. Here $x_m > x_r > x_c$, so that an individual can be alive, but still unable to reproduce. The choice of the terminal fitness function $R(z)$ is not particularly important, but one wants $R(0) = 0$ and $R(z)$ non-decreasing in z . Here T is fixed and known; Mangel and Clark (1988), p. 71, show how to treat uncertain end times. The overall objective of the evolutionary optimization method is to be able to relate short term behaviors to a measure of long term fitness; this is accomplished by specifying the terminal fitness function and solving for the short term behaviors using stochastic dynamic programming.

The environment of the worm is characterized by two probabilities associated with the visual stimulus:

$$\lambda_s = \text{Prob}\{\text{a visual stimulus occurs in period } t\} \quad (3)$$

$$\beta = \text{Prob}\{\text{stimulus is caused by a predator} \mid \text{visual stimulus occurs}\} \quad (4)$$

and by a set of probabilities for the energetic value of possible food types

$$p_i = \text{Prob}\{\text{stimulus is caused by food particle with energetic value } Y_i \mid \text{visual stimulus occurs}\} \quad (5)$$

The food types are indexed by $i = 0, \dots, K$. Here $Y_0 = 0$, corresponding to a piece of debris, which can be thought of as food with no energetic value.

The energetic consequences of the worm's behavior are determined by two "costs" that measure decrement in the reserve levels

$$\alpha_h = \text{cost of staying in the tube for a single period} \quad (6)$$

$$\alpha_g = \text{cost of exiting the tube and returning in a single period} \quad (7)$$

Since it should be energetically more costly to exit the tube than to stay in it, $\alpha_g > \alpha_h$.

Fitness is defined in terms of expected reproduction and survival. In particular,

$$F(x, t, T) = \max E\{R(X(T)) \mid X(t) = x\} \quad (8)$$

is the Darwinian fitness, where the maximization is taken over feasible behavioral decisions (i.e. movement out of the tube) between period t and T and E denotes the expectation over the stochastic environment.

In order to determine $F(x, t, T)$ we derive an iterative equation (Mangel and Clark 1988) characterizing the time evolution of $F(x, t, T)$. This equation is

$$F(x, t, T) = (1 - \lambda_s)F(x - \alpha_h, t + 1, T) + \lambda_s \max \left[F(x - \alpha_h, t + 1, T); (1 - \beta) \sum_{i=0}^K p_i F(x - \alpha_g + Y_i, t + 1, T) \right] \quad (9)$$

This equation is derived by use of the law of total expectation (Mangel and Clark 1988, Chap. 2). The left hand side of Eq. (9) is the maximum expected lifetime fitness from period t to period T , assuming that the worm is alive at the start of period t . The first term on the right hand side is the future expected fitness if no stimulus occurs during period t (this happens with probability $1 - \lambda_s$). If no stimulus occurs, the worm remains in its tube and the state variable is decremented by α_h ; its expected future fitness from period $t + 1$ to period T is given by $F(x - \alpha_h, t + 1, T)$. A stimulus occurs with probability λ_s . In that case, a decision is made (corresponding to the "max" on the right hand side of Eq. (9)). If the worm stays in its tube, its expected future fitness is still $F(x - \alpha_h, t + 1, T)$. If the worm exits the tube and the stimulus was caused by a predator (which occurs with probability β), the worm is killed and its expected future fitness is 0. If the stimulus was caused by a piece of food or debris of type i (which happens with probability $(1 - \beta)p_i$) the state variable is augmented by the network energetic value of the food $Y_i - \alpha_g$ so that expected future fitness is $F(x - \alpha_g + Y_i, t + 1, T)$. If $x - \alpha_g + Y_i$ either exceeds x_m or falls below x_c , then $x - \alpha_g + Y_i$ is replaced by x_m or x_c in the formula for the future fitness on the right hand side of (9). Thus, the condition that the fitness is 0 at $x = x_c$ acts as a boundary condition.




Fig. 1. Switching curve for the following parameters described in the text. For values of $X(t)$ above the switching curve, the worm should remain in its tube if a visual stimulus occurs; otherwise it should exit

The fitness also satisfies the end condition $F(x, T, T) = R(x)$. Equation (9) is solved "backwards" in time, starting at $t = T - 1$ and decreasing in t until $t = 1$; thus the condition on $F(x, T, T)$ acts as an initial condition in the iterative solution of (9). As Eq. (9) is solved, the optimal decisions are obtained. For this model, the optimal decisions are $d^*(x, t)$ indicating whether the worm should exit its tube when a stimulus appears in period t and $X(t) = x$. These decisions can be summarized by a "switching curve" $X^*(t | \lambda_s, \beta, \{p_i, Y_i\})$ in the (t, x) plane. Figure 1 shows such a switching curve for reproduction given by Eq. (2) and the following parameters: $x_m = 20$, $x_r = 4$, $x_c = 2$, $T = 50$, $\lambda_s = 1$, $\alpha_h = 1$, $\alpha_g = 3$, $\beta = .05$, $p_0 = .9$, $p_i = .025$ for $i = 1$ to 4, and $Y_i = i + 3$ for $i = 1$ to 4.

For values of $X(t)$ above the switching curve, the worm should remain in its tube if a visual stimulus occurs; otherwise it should exit. Such switching curves determine the optimal behavior of the worm. Except for values of t near T , the decision threshold for moving out of the tube in response to a visual stimulus depends only on the physiological variable and not on time. As t approaches T , the worm becomes more "risk averse", only exiting the tube if physiological reserves are low. Similar kinds of switching curves are obtained for other combinations of parameters.

3. Polychaete behavior: neural network model

Neural network models have become very popular in the last few years (Bounds 1987; Cowan and Sharp 1988). The key features for a neural network model are: (i) nonlinearity of interactions leading to a rich and wide variety of behaviors, (ii) plasticity of the interactions so that the neural network can change in

response to the environment and (iii) lack of knowledge of the specific problem being solved so that the neural network “works in the dark” about the objective functional; behaviors are selected and reinforced by the environment. In particular, recent work on models for central pattern generators (Getting 1988; Mulloney and Perkel 1988), olfactory detection (Getz 1989), drosophila motor neurons (Friesen and Wyman 1980), crayfish motoneurons (Edwards and Mulloney 1987) and swimming (Getting 1983a,b) have shown how the “neural groups” used in this paper can actually be identified for relatively simple invertebrates.

The neural network model for the polychaete consists of different neural groups chosen to mimic the biology of the polychaete. The different neural groups are linked by connections. The strengths of these connections are modified by “learning rules” that describe how the connections change in response to experience of the individual. For example, the connections will change in one way when the worm exits the tube and finds food and in a different way when food is not found. In addition to connections, the states of the neural groups follow non-linear dynamics. The dynamics chosen below are typical logistic-like functions. These mimic the generally observed neural physiology (Edleman 1987).

The neural network thus has a set of “design parameters” (Williams 1966; Grafen 1988) that includes the initial values of the connections, parameters for learning rules describing how the connections change upon experience, and the parameters of the nonlinear dynamics of the neural groups themselves; this set of design parameters is denoted by \mathcal{P} . Ontogeny of the network is manifest in the change of connections, given initial values and learning rule parameters, according to individual experience. Evolution of the network is the result of the natural selection of the initial values and the learning rules. Darwinian fitness for a set of design parameters of the neural network is measured in terms of the expected terminal fitness of individuals with this set of design parameters. Natural selection does not “optimize” design parameters. Rather, for a given set of environmental conditions, natural selection provides a “ranking” for the quality of different sets of design parameters, with the better parameter sets being more prevalent in the subsequent generation.

The evolution of the design parameters can be mimicked by Monte Carlo simulation. We begin by specifying a set of environmental parameters $\{\lambda_s, \beta, \{p_i, Y_i\}\}$, which generate various realizations of the environment. We then consider a large number of worms with different design parameter sets \mathcal{P} . This corresponds to the current “generation”. The environmental parameters are used to generate realizations of the environment (within the generation) and the worms with design set \mathcal{P} respond according to the rules of the neural network. At $t = T$, fitness is determined by $R(X(T))$. This gives a fitness $F\{\mathcal{P}\}$ for the design parameters \mathcal{P} . Evolution consists of considering another set of design parameters \mathcal{P}' , computing the fitness $F\{\mathcal{P}'\}$ of that set and comparing the two fitnesses; the precise rules for choosing the variants of the design parameter sets will be given below. Briefly, the design parameters will be chosen uniformly over given fixed intervals. If $F\{\mathcal{P}'\} > F\{\mathcal{P}\}$ then the design parameters \mathcal{P} are replaced

by the new set \mathcal{P}' . This process continues over many generations; with one set of design parameters being replaced by a superior set as natural selection occurs. This entire process can be repeated over many different “worlds” or “patches” (essentially enlarging the number of different designs studied).

In this paper, no explicit genetics are assumed, so that design parameters replicate perfectly from one generation to the next. Thus, from one generation to the next, the best designs are maintained and additionally many variants are constructed, for comparison with the current best design. In the actual computations reported below, the best design was maintained from one generation to the next, and in each generation 500 variants were compared to it.

The neural network model involves four “neural groups”: (i) a sensory group that responds to the visual stimulus, (ii) a motor group that moves the worm in and out of the tube, (iii) a “reserves group” (stretch receptor) that measures the level of energetic reserves of the worm and (iv) a clock group that indicates when t is approaching T . Figure 2 shows a schematic drawing of the neural groups and their interactions. In principle, the motor, sensory and reserve groups all interact with each other; the clock group interacts in a one-way fashion with the motor group. The purpose of the clock group is to “shut down” the motor group when $t = T$ so that reproduction occurs.

The interactions between the different neural groups are determined by connections

$$\beta_{ij} = \text{strength of the interaction from neural group } i \text{ to neural group } j \quad (10)$$

Here i, j will range over the indices m (motor), s (sensory), r (reserve) and c (clock). In particular, I assume that $\beta_{ms}, \beta_{sm}, \beta_{rm}, \beta_{rs}$ and β_{cm} are non-zero and that β_{mr} and β_{sr} are identically 0.

We must also specify the nonlinear dynamics of the neural groups and learning rules for the connections between the neural groups. The states of the neural groups are denoted by X_s, X_m, X_r and X_{c1} . A discrete time formulation is used, so that the group dynamics connect a state X_i at time t with its value at

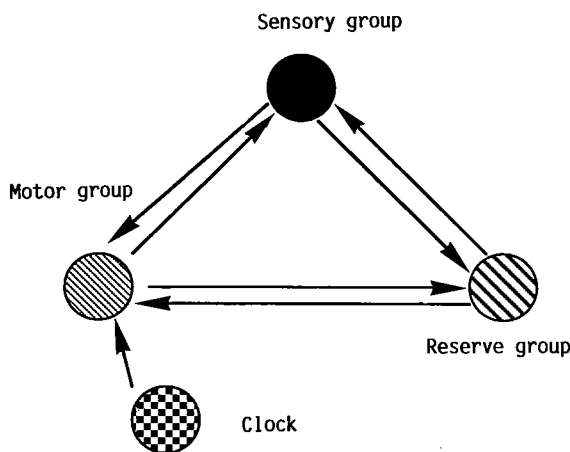


Fig. 2. The neural groups and their interactions. In principle, the motor, sensory and reserve groups all interact with each other; the clock group interacts in a one-way fashion with the motor group

time $t + 1$. All neural groups should have bounded state variables (see Cowan and Sharp 1988 for a general exposition of the modeling approach).

For the computations reported above and below $\lambda_s = 1$; this is done to simplify the study in terms of the number of parameters involved. This means that in each period the sensory group is stimulated. Following Dehane et al. (1987), Eqs. (11, 12), assume that the updating of the sensory group follows a sigmoidal response. In particular, $X_s(t)$ is updated according to

$$X_s(t + 1) = [1 + \exp(-(Q_s - \mu_s)/\varepsilon)]^{-1} \quad (11)$$

In this equation, μ_s centers the response, ε determines the sensitivity of the response to the exponent and Q_s summarizes the interactions between the sensory, motor and reserve groups. When choosing variants for the neural network, μ_s and ε are uniformly distributed as $[0, 1]$. The function Q_s is used to relate the inputs from the other neural groups to the updating of the sensory group. The form of this function is

$$Q_s = \gamma X_s(t) + Z_s \exp[-\beta_{ms} X_m(t) - \beta_{rs}(X_r(t)/x_m)] \quad (12)$$

In this equation, γ and Z_s are fixed parameters. When choosing variants, γ is uniformly distributed on $[0, 1]$ and Z_s is uniformly distributed on $[0, 2]$. The β_{ij} are the strengths of the neural connections between groups i and j . These change according to individual experience, but even so initial values need to be specified. When choosing variants of the neural network, the initial value of each β_{ij} is uniformly distributed on $[0, 1]$.

The important aspects of (12) are that Q_s increases with X_s and decreases with X_m and X_r . That is, $X_s(t + 1)$ increases as the sensory group is stimulated, the motor group rests, or reserves become low. Other functional forms would probably work as well (see, e.g. Finkel and Edelman 1985). The dynamics of X_s thus involves the states of the motor, sensory and reserve groups at time t ; γ , Z_s , μ_s , and ε are parameters determined by evolution through the natural selection of the neural network.

The dynamics of the motor and reserve groups, and the connections between the groups, depend upon the "decision" at period t . If a combination of inputs, denoted by X^* , is below a threshold value \hat{x} , then the worm remains in its tube in response to the visual stimulus. When generating variants of the neural network, \hat{x} is uniformly distributed on $[0, 1]$.

If $X^* < \hat{x}$, the dynamics of the motor group are

$$X_m(t + 1) = \gamma X_m(t) \quad (13)$$

Thus, according to (13), the state of the motor group decreases exponentially when the worm remains in its tube. When generating variants of the neural network, γ is uniformly distributed on $[0, 1]$.

The combination of inputs that determines exit is

$$X^* = X_{c1} \{X_m(t) + X_s(t)\beta_{sm} + \beta_{rm}/[X_r(t)/x_m]\} \quad (14)$$

The clock state is basically 0 or 1, so that the worm does not exit when t is close to T unless the term in brackets is very large. The crucial aspects of the term in

brackets are the increase in X^* with the state of the sensory group X_s and with a reduction in the state of the reserve group X_r . Recall that X^* needs to be above the threshold \hat{x} for exit to occur. The functions in (12) and (14) are chosen because they are mathematically convenient; they are not based on neurobiological considerations (but see Getting 1983a,b, 1988 for a way of getting such forms from neurobiological data).

The dynamics of the reserve group depend upon the dynamics of the motor group. If the worm does not exit its tube during period t , then

$$X_r(t+1) = \max\{x_e, X_r(t) - \alpha_h\} \quad (15)$$

and the connections change according to the rule

$$\beta_{ij} \rightarrow \beta_{ij} + \delta_{ij}(s) \quad (16)$$

Here the set $\{\delta_{ij}(s)\}$ is the set of learning rules that the connections follow when the worm stays in its tube.

If $X^* \geq \hat{x}$, then the worm makes an excursive cycle from the tube and returns to the tube. If the worm is not killed by a predator, then it may find food (with probability p_i of finding food with energetic value Y_i) or may not find food (with probability p_0). In either case, the motor group is "reset" to a resting value

$$X_m(t+1) = x_{m, \text{rest}} \quad (17)$$

If the visual stimulus is caused by debris, then $X_r(t+1)$ is given by Eq. (15) with α_h replaced by α_g and the connections are modified by the learning rule $\beta_{ij} \rightarrow \beta_{ij} + \delta_{ij}(nf)$ where the set of learning parameters $\{\delta_{ij}(nf)\}$ is used when the worm exits but does not find food.

If the worm makes an excursion and does find food, then the reserve group dynamics are

$$X_r(t+1) = \min(x_m, X_r(t) - \alpha_g, Y_i) \text{ with probability } p_i \quad (18)$$

and the connections are modified by $\beta_{ij} \rightarrow \beta_{ij} + \delta_{ij}(f)$ where the set of learning parameters $\{\delta_{ij}(f)\}$ is used when the worm exits its tube and food is found. In this model, I do not presume that the worm "knows" the Y_i or p_i ; rather these simply describe what happens to the neural group.

When generating variants of the neural network, all of the learning rules $\{\delta_{ij}(k)\}$ are assumed to be uniformly distributed on $[-\delta, 2 + \delta]$, where δ itself is uniformly distributed on $[0, 1]$.

The clock is simply used to insure that the motor group is turned off when $t = T$. Many chemical relaxation oscillators could be used to model the clock dynamics.

4. Results of evolution of the neural network model by Monte Carlo simulation

The neural network described in the previous section thus has a set of "design parameters" (Williams 1966; Grafen 1988) given by $\mathcal{P} = \{\gamma, Z_s, \mu_s, \varepsilon, \{\delta_{ij}(k) \mid k = s, nf, f\}, \{\text{initial values of } \beta_{ij}\}\}$. Evolution by Monte Carlo simulation proceeds

Table 1. Rate of evolution of the neural group for four different patches

Case	Increase in fitness at generation	Size of increase in fitness
1(277) ^a	1	9
	2	12
	73	11
	112	20
2(429)	3	14
	5	2
	9	9
	25	39
	216	2
	225	13
3(500)	1	14
	32	3
	56	10
	66	7
	83	7
	223	2
	263	12
398	2	
4(500)	1	24
	2	10
	5	5
	7	2
	80	13
	87	5
	163	13

^a Numbers in paranthesis in the first column indicate the total number of generations simulated

as described in the previous section. Table 1 shows the results of evolution of the neural network parameters and learning rules for four different cases. In each case, there is a rapid initial increase in fitness, followed by a slower improvement in the fitness of the neural network model. This result is similar to that of Kauffman and Levin (1987), Figs. 5 and 6, who also found that, in general by 200–500 generations the waiting times for the next increase in fitness become very large. That is, after many generations the waiting time for the next improvement of design parameters is long. Another way to view this is that after many generations of natural selection, the currently best neural network parameters are drawn from the tail of a distribution. When sampling from the entire distribution, the probability of obtaining a “better” set of parameters is small.

There is another way of viewing the evolution of the neural network. Suppose that instead of specifying a single design parameter set and a variety of realizations of the environment, we specify a single realization of the environment and a large number of different design parameter sets. The fitness of each of these sets can be compared with the fitness determined by evolutionary optimization. Most of the randomly selected design parameter sets will have




Fig. 3. Distributions of fitness of design parameters, for two different realizations of the environment (■, case 1; ▨, case 2). The cumulative number of individuals with neural fitness as a percentage of evolutionary optimal fitness is shown for the situation in which a single realization of the environment and 500 different design parameter sets are specified. The fitness of each of these sets can be compared with the optimal fitness determined by evolutionary optimization. Most of the randomly selected design parameter sets will have fitness far below the optimal fitness, but some will be close. Note that in each case there are one or two designs that actually have neural network fitness higher than evolutionary optimal fitness. This occurs because a single environmental realization is selected; by definition the evolutionary optimal is only maximized as an expectation over many different realizations of the environment

fitness far below the optimal fitness, but some neural networks will have fitness that is close to the evolutionary optimal fitness. These would correspond to sub-optimal peaks on rugged landscapes (cf. Kaufmann and Levin 1987). Mangel (1989) describes explicitly how such behavioral problems can be viewed as walks on rugged landscapes. Figure 3 gives two examples of distributions of fitness of design parameters. The cumulative number of individuals with fitness of the neural network as a percentage of fitness of the evolutionary optimization model is shown. Note that in each case there are one or two designs that actually have fitness of the neural network greater than evolutionarily optimal fitness. This occurs because a single environmental realization is selected; by definition the evolutionary optimal fitness is only maximized as an expectation over many different realizations of the environment. The environmental processes are generalizations of the Poisson process, so that the variances will be of the same order as means. Thus, the variance of the evolutionarily optimal fitness will be high; there will be many individual realizations in which the neural network can outperform the evolutionarily optimal behaviors.

5. Ontogeny of the neural network and comparison with evolutionary optimization models

Suppose now evolution leads to a set of “very good” design parameters \mathcal{P}^* , in the sense that many generations have passed without improvement in the fitness

of the neural network (i.e. that no better design parameters have been found for many generations). The design set specifies initial values of connections β_{ij} and learning rules $\delta_{ij}(k)$; these have been determined by the evolution of the neural network. Consider a large number of worms, each with the naturally selected design parameter set \mathcal{P}^* . Each of these worms experiences a different random environment during its ontogeny, determined by the environmental parameters. Although the initial values of the connections the learning rules are the same, the different realizations of the environment will lead to different ultimate values of connections and states of neural groups. The ontogeny and plasticity of the neural network can thus be modeled.

Each realization of the environment experienced by a worm using decisions based on the evolved neural network can also be applied to a worm making decisions based on evolutionary optimization. Doing this allows a comparison of the relative fitness of the neural network model to the evolutionary optimization model. Figure 4 shows the results of comparing 8 different design parameter sets; it is a plot of fitness of the neural network versus fitness of the evolutionary optimization model. About 4000 different realizations of environment were experienced by the worms. There is a generally positive, although weak, correlation between the fitness of the neural network and the fitness based on evolutionary optimization. The grand average of the data in Fig. 4 shows that $\text{Fitness}_{\text{neural network}}/\text{Fitness}_{\text{evolutionary optimization}} = 0.85$. That is, on average the fitness of the neural network model is about 85% of the evolutionarily optimal fitness.

There are, in fact, three cases in which the fitness of the neural network is actually higher in the simulation than the fitness of the evolutionary optimization model (neural network fitness = 116%, 104% and 103% of evolutionary optimization fitness). This is caused by the relatively small number of realizations (4000) studied, compared to the large sample space over which the expectation




Fig. 4. Comparison of 8 different design parameter sets. About 4000 different realizations of environment were experienced by the worms. This figure is a plot of fitness of the neural network versus fitness of the evolutionary optimization model

in the evolutionary optimization model is taken. That is, micro fluctuations in the "local" realization of the environment allow the neural network model to actually outcompete the evolutionary optimization model when the expression of environmental stochasticity is limited.

The two neural network design parameters with the highest levels of relative fitness were then compared with the evolutionary optimization model for a much larger number of realizations of the environment, 100000 realizations and 50000 realizations respectively. The relative fitnesses in these cases were 99.8% of optimal and 85% of optimal, with an average fitness over the total 150,000 realizations of the environment equal to 94.9% of the evolutionary optimal fitness. For this simple neural network model of a polychaete worm, natural selection can attain a "design" that is almost equal in fitness to the evolutionary optimum.

The good performance of the neural network, relative to the evolutionarily optimal fitness, can be attributed to a number of factors. First, there are no deleterious mutations and design parameters are replicated perfectly, so that design parameters never get worse, but only get better. Second, the decision of the worm is a binary decision (to exit or not), but there are many parameters available for tuning this decision.

A smaller study was done for evolutionary optimum and neural network models for the case in which $x_m \rightarrow \infty$. There is essentially no change in the formulation of the evolutionary optimization model, but the dynamics of the neural groups do change since they involve x_m . These equations were changed by simply dropping x_m , but bounding the value of the connections associated with x_m . The same procedure for studying the evolution, ontogeny and comparison of the neural network and evolutionary optimization models was followed on a smaller scale and qualitatively similar results were obtained. Over a total of 55000 different realizations of the environment, the average fitness of the neural network model was 79% of the average fitness of the evolutionary optimization model.

6. Polychaete behavior: more complex models

A more complicated polychaete worm might have "visual" and "tactile" receptors, so that the worm could receive either a visual stimulus or a tactile stimulus or both in a single period. Assume that such stimuli occur independently, but may occur simultaneously. Other assumptions are, of course, possible. Let

$$\begin{aligned}\lambda_v &= \text{Prob}\{\text{visual stimulus received in a single period}\} \\ \lambda_t &= \text{Prob}\{\text{tactile stimulus received in a single period}\}\end{aligned}\quad (19)$$

The different stimuli will have different associated probabilities of predation and food:

$$\begin{aligned}\beta_v &= \text{Prob}\{\text{predator caused the visual stimulus} \mid \text{visual stimulus occurred}\} \\ \beta_t &= \text{Prob}\{\text{predator caused the tactile stimulus} \mid \text{tactile stimulus occurred}\}\end{aligned}\quad (20)$$

A stimulus not caused by a predator is, by assumption, caused by a particle of food or debris.

Finally, assume that either predators or food/debris might be present even though no stimulus occurs. This can be modeled by

$$\begin{aligned}\beta_0 &= \text{Prob}\{\text{predator is present} \mid \text{no stimulus occurs}\} \\ \omega_0 &= \text{Prob}\{\text{food/debris is present} \mid \text{no stimulus occurs}\}\end{aligned}\quad (21)$$

Defining $F(x, t, T)$ as before leads to

$$\begin{aligned}F(x, t, T) &= (1 - \lambda_v)(1 - \lambda_t) \max \left[F(x - \alpha_h, t + 1, T); \right. \\ &\quad \left. (1 - \beta_0)[(1 - \omega_0)F(x - \alpha_g, t + 1, T) + \omega_0 \sum_{i=0}^K p_i F(x - \alpha_g + Y_i, t + 1, T)] \right] \\ &\quad + \lambda_v(1 - \lambda_t) \max \left[F(x - \alpha_h, t + 1, T); (1 - \beta_v) \sum_{i=0}^K p_i F(x - \alpha_g + Y_i, t + 1, T) \right] \\ &\quad + (1 - \lambda_v)\lambda_t \max \left[F(x - \alpha_h, t + 1, T); (1 - \beta_t) \sum_{i=0}^K p_i F(x - \alpha_g + Y_i, t + 1, T) \right] \\ &\quad + \lambda_v\lambda_t \max \left[F(x - \alpha_h, t + 1, T); \right. \\ &\quad \left. (1 - \beta_t)(1 - \beta_v) \sum_{i=0}^K p_i F(x - \alpha_g + Y_i, t + 1, T) \right]\end{aligned}\quad (22)$$

This equation is solved subject to the end condition that $F(x, T, T) = R(x)$. The decision set generated during the solution of Eq. (22) will consist of rules $d_v^*(x, t)$ and $d_t^*(x, t)$ indicating whether or not the worm should exit the tube if a visual or tactile stimulus is received during the period in which $X(t) = x$.

The network corresponding to Eq. (22) would be richer in terms of neural groups and connections than the one studied in this paper. However, the problem being solved by the polychaete (to exit the tube or not) is still the same. For example, there will be many more connections and rules for changing connections. We have added just 2 environmental parameters, λ_v and ω_0 , to the evolutionary optimization model but have added many more parameters (2 sensory groups and 3 rules) to the neural network model. There are two possible arguments regarding the fitness of the neural network when compared to evolutionary optimization fitness. The first is that the neural network will still be solving the problem for a binary decision, but with many more parameters available for tuning. This suggests that it should be possible to find a neural network that has fitness nearly equal to the evolutionary optimal fitness. On the other hand, there is a tradeoff in that the parameter space being sampled is of higher dimension. This means that, in general, it will take longer to evolve the neural network and that evolution is more likely to become stuck on suboptimal peaks (cf. Kauffman and Levin 1987).

7. Speculative neurobiology of tephritid fruit flies

In this section, the ideas developed throughout the paper are applied to a different and more complicated organism, the tephritid fruit flies (see Roitberg and Prokopy 1987 for a general review). Eggs are typically laid in healthy fruit and the larvae spend essentially their entire lives in the fruit. The adults are free flying. Females often mark fruit with a pheromone after oviposition and will usually reject a marked fruit as an oviposition site. The discussion in this section is based on experiments pertaining mainly to the flies *Rhagoletis pomonella* (apple maggot fly), *Rhagoletis basiola* (rose hips fly) and *Ceratitidis capitata* (Mediterranean fruit fly).

The kinds of problems that the fruit flies must solve include:

Learning to recognize marking pheromones (Roitberg and Prokopy 1981). Naive flies do not recognize their own marking pheromone or the marking pheromones of other flies. It is only after ovipositions in fruit, and dragging her ovipositor across the fruit after oviposition, that a fly begins to recognize marking pheromone.

Responding to marked fruit and host deprivation (Roitberg and Prokopy 1983, Averill and Prokopy 1987). In general flies will not oviposit in a fruit that is marked with a pheromone. After a sufficiently long period of host deprivation, however, flies will accept marked fruit. Evolutionary optimization arguments (Mangel 1987) show how such behavior is adaptive.

Learning to recognize and accept or reject novel hosts (Prokopy et al. 1986). When flies are given known and novel hosts (e.g. hawthorn and apple), there is a time dependence of learning to accept or reject novel hosts. Figure 5 shows hypothetical time courses proposed by Prokopy et al. (1986) in which the acceptance level of known fruit remains constant, but the acceptance level of novel fruit changes over time.

These phenomena can be discussed and interpreted in the framework of neural network models. Host finding typically involves visual and olfactory




Fig. 5a,b. Hypothetical time course proposed by Prokopy et al. (1986) in which the acceptance level of known fruit remains constant, but the acceptance level of novel fruit changes over time

stimuli (Miller and Strickler 1984). Thus, it is likely that at least two sensory groups are involved in the neural network.

For *Rhagoletis* flies, at least, there is clear evidence that the neural receptors for marking pheromone are found on tarsal hairs (Crnjar and Prokopy 1982; Prokopy et al. 1982; Papaj and Prokopy 1986). When a fly lands on a fruit, she walks across the fruit, triggering the tarsal pheromone receptors if the fruit is marked and if she crosses the marking trail. (After oviposition, a fly drags her ovipositor across the fruit, leaving a thin trail of marking pheromone.) A fly may circumscribe a fruit 6 or 8 times before ovipositing or rejecting the fruit. The neural network model thus should involve a group representing tarsal hair receptors.

Walking across the fruit may also help identify the size (and thus type) of fruit. For parasitoid wasps, Schmidt and Smith (1985, 1986, 1987a) have demonstrated that walking across the fruit allows the wasp to assess the size of the fruit. A neural network model for an oligophagous or polyphagous fruit fly should thus include a "fruit size neural group".

Abdominal stretch receptors provide a means for the fly to determine egg complement, which is one of the key variables for the evolutionary optimization model. As before, a clock group will allow the fly to assess physical time. Schmidt and Smith (1987b) describe how the parasitoid's clock might work.

Finally, the neural network model should contain a neural group representing the ovipositor. This group is similar to the motor group for the polychaete model: it opens the ovipositor in the same way that the motor group causes the worm to exit its tube. After oviposition, the fly is assumed to drag her ovipositor across the fruit. Roitberg and Mangel (1988), using evolutionary optimization, describe how oviposition and marking by dragging could evolve.

Figure 6 shows a neural network model for such insects. There are about 20 connectance parameters and a much larger set of learning rules (depending upon

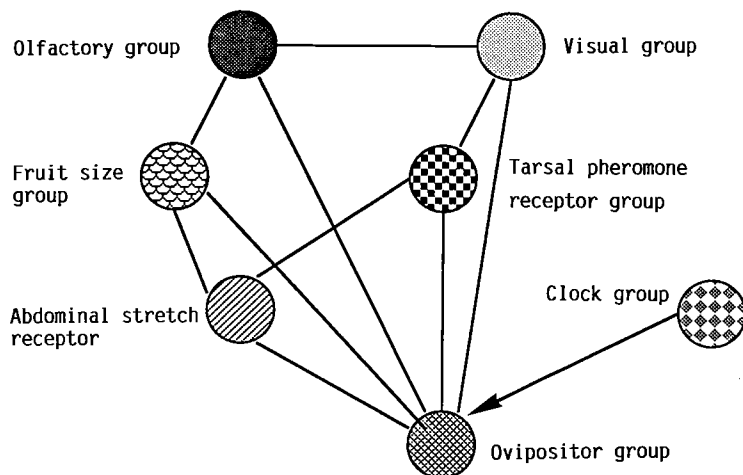


Fig. 6. A possible neural network for a tephritid fruit fly

fruit type, whether the fruit is marked or not, whether oviposition occurs or not, etc), so that simulating "natural selection" by Monte Carlo methods would be a daunting task.

On the other hand, knowing the biology of the flies and thinking about the neural network models leads to certain predictions. Consider, for example, learning to accept a novel host of different size (e.g. a small host like hawthorn compared to a large host like a Macintosh apple). Thinking about such a learning problem leads to a prediction that dynamics of the neural group $X_{size}(t)$ that is used to measure size has two stable steady states, a "small" steady (indexed by 1, say) and a "large" steady state (indexed by 2). The evolution of the neural network provides the fly with initial values of the connectance between the size group and the ovipositor group $\beta_{so}(X_{s,eq}(1))$ that depends upon the state of the size group. A sufficient delay in encountering any kind of fruit will ultimately cause the fly to oviposit in a large fruit. When this occurs, a successful oviposition causes $\beta_{so}(X_{s,eq}(2))$ to increase. The net effect will be the acceptance dynamics shown in Fig. 5a. Similarly, acceptance of a known fruit may lead to a decrease in $\beta_{so}(X_{s,eq}(2))$; a simple form of re-entry interactions (Edelman 1987). A decrease in $\beta_{so}(X_{s,eq}(2))$ caused by successful ovipositions in small fruit will lead to rejection dynamics similar to those shown in Fig. 5b. Prokopy et al. (1986) describe experiments that are consistent with this picture of the neural network dynamics. A tephritid fly may be too complex for a neurobiological study in which one tries to measure steady states of nerve groups that respond to pheromone or host size, but it might be a worthwhile project to try to determine if those groups do indeed have multiple steady states.

8. Discussion and conclusions

This paper contains a comparison of behaviors determined by evolutionary optimization methods and neural networks. The overall objective of this study was to address the challenge to the adaptationist program that "developmental constraints" will prohibit the achievement of evolutionarily optimal solutions. To answer this challenge, simple models for a polychaete worm were developed using both evolutionary optimization methods and neural networks. Allowing natural selection to act on the neural network connections, parameters, and learning rules lead to a number of neural networks with fitness nearly as high as the evolutionary optimum fitness. I propose that more complicated organisms (which must solve more complicated problems) will have neural networks that grow in complexity of connections and learning rules at a much faster rate than that difficulty of the problem that must be solved. In general, I thus propose that neurobiological mechanisms do exist for the achievement of near optimal or optimal fitness. A remaining open question concerns the general relationship between increased complexity of the neural network and increased complexity of the behavioral problem. In particular, if the complexity of the neural network increases at a rate much faster than the behavioral problem, will a neural network always lead to near evolutionarily optimal behavior? As mentioned

above, it is likely that the neural networks studied here performed so well because of no deleterious mutations, and the coupling of binary decisions to a neural network with a number of tunable parameters. In addition, since variants were selected uniformly, it was possible to escape false peaks in parameter space. Hill-climbing processes (Kauffman and Levin 1987) are more likely to be trapped at false peaks. On the other hand, because of uniform sampling for generating variants used in the paper, the rate of evolution is slower than it might be in a hill-climbing procedure that also had a method for escaping false peaks (e.g. simulated annealing, Aarts and Korst 1989). In more difficult problems, such as using a neural network to drive a system to one of many steady states of a large phase space, the success of the neural network is more problematical. Even so, there already exists encouraging evidence about the effectiveness of neural networks (Hoffmann et al. 1986). Hinton and Nowlan (1987) also provide an example of how a learning network can guide evolution (also see Maynard Smith 1987). Because of the large number of connections, nonlinearities, learning rules and plasticities, neural networks will always have rich behaviors. I suspect that one can always find a neural network that solves a complicated problem if one knows what the answer is. But it is the evolutionary optimization approach that provides a benchmark for comparison.

A model that somewhat bridges the neural network and evolutionary optimization approaches would be to assume that the polychaete evolves "knowing" that there is a switching curve, as in Fig. 1, but not knowing what the curve is. Natural selection will then act on genotypes that code different switching curves. This approach is similar to that used in Mangel (1989); it does not, however, deal with the constraints of biological design for achieving the evolutionary optimal solution.

Acknowledgements. For helpful discussions, I thank Gerald Edelman, Richard Gomulkiewicz, Alasdair Houston, Alex Kacelnik, John Krebs, Simon Levin, John McNamara, and Sara Shettleworth. Richard Gomulkiewicz, Simon Levin, and Brian Mulloney read previous versions of this manuscript and their comments were very helpful. This work was partially completed at the Department of Zoology and Centre for Mathematical Biology, University of Oxford. I thank John Krebs and Jim Murray for hospitality. Funding in partial support of this work came from NSF Grant BSR 86-1073, the John Simon Guggenheim Foundation and the Fulbright Foundation.

References

- Aarts, E., Korst, J.: Simulated annealing and Boltzmann machines. New York: Wiley 1989
- Aoki, C., Siekevitz, P.: Plasticity in brain development. *Sci. Am.* **259**, 56-67 (1988)
- Averill, A. L., Prokopy, R. J.: Residual activity of oviposition deterring pheromone in *Rhagoletis pomonella* (Diptera: Tephritidae) and female response to infested fruit. *J. Chem. Ecol.* **13**, 167-177 (1987)
- Bounds, D. G.: New optimization methods from physics and biology. *Nature* **329**, 215-219 (1987)
- Brady, R. M.: Optimization strategies gleaned from biological evolution. *Nature* **317**, 804-806 (1985)
- Charnov, E. L.: Phenotypic evolution under Fisher's fundamental theorem of natural selection. *Heredity*, **62**, 113-116 (1989)
- Cowan, J. D., Sharp, D. H.: Neural networks. *Quart. Rev. Biophys.* **21**, 365-427 (1988)

- Crnjar, R. M., Prokopy, R. J.: Morphological and electrophysiological mapping of tarsal chemoreceptors of oviposition-detering pheromone in *Rhagoletis pomonella* flies. *J. Insect Physiol.* **28**, 393-400 (1982)
- Dehaene, S., Changeux, J.-P., Nadal, J.-P.: Neural networks that learn temporal sequences by selection. *Proceedings of the National Academy of Sciences* **84**, 2727-2731 (1987)
- Edelman, G.: *Neural Darwinism*. New York: Basic Books 1987
- Edwards, D. H., Mulloney, B.: Synaptic integration in excitatory and inhibitory crayfish motoneurons. *J. Neurophysiol.* **57**, 1425-1445 (1987)
- Evans, S. M.: Non-associative behavioural modifications in nereid polychaetes. *Nature* **211**, 945-948 (1986)
- Evans, S. M. Behavior in polychaetes. *Quart. Rev. Biol.* **46**, 379-405 (1971)
- Evans, S. M.: Effects of changes in sensory input on patterns of tubicolous activity in the polychaete *Platynereis dumerilii*. *Marine Behav. Physiol.* **7**, 307-319 (1981)
- Evans, S. M., Downie, P. J.: Decision-making processes in the polychaete *Platynereis dumerilii*. *Anim. Behav.* **34**, 472-479 (1986)
- Finkel, L. H., Edelman, G. M.: Interactions of synaptic modification rules within populations of neurons. *Proc. Nat. Acad. Sci. USA* **82**, 1291-1295 (1985)
- Friesen, W. O., Wyman, R. J.: Analysis of *Drosophila* motor neuron activity patterns with neural analogs. *Biol. Cybernetics* **38**, 41-50 (1980)
- Getting, P. A.: Mechanisms of pattern generation underlying swimming in *Tritonia*. II. Network reconstruction. *J. Physiol.* **49**, 1017-1035 (1983a)
- Getting, P. A.: Mechanisms of pattern generation underlying swimming in *Tritonia*. III. Intrinsic and synaptic mechanisms for delayed excitation. *J. Physiol.* **49**, 1036-1050 (1983b)
- Getting, P. A.: Comparative analysis of invertebrate central pattern generators. In: Cohen, A. (ed.) *Neural control of rhythmic movements in vertebrates*, chap. 4, pp. 101-127. New York: Wiley 1988
- Getz, W. M.: A neural network for processing olfactory-like stimuli. *Biol. Cybernetics*, in press
- Gould, S. J., Lewontin, R. C.: The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist program. *Proc. Rl. Soc Lond. B* **205**, 581-598 (1979)
- Grafen, A.: On the uses of data on lifetime reproductive success. In: Clutton-Brock, T. H. (ed.) *Reproductive success*, pp. 454-471. Chicago: University of Chicago Press 1988
- Hinton, G. E., Nowlan, S. J.: How learning can guide evolution. *Complex Systems* **1**, 495-502 (1987)
- Hoffmann, G. W., Benson, M. W., Bree, G. M., Kinahan, P. E.: A teachable neural network based on an unorthodox neuron. *Physica* **22D**, 233-246 (1986)
- Kauffman, S., Levin, S.: Towards a general theory of adaptive walks on rugged landscapes. *J. Theor. Biol.* **128**, 11-45 (1987)
- McClelland, J. L., Rumelhart, D. E., PDP Research Group: *Parallel distributed processing. Volume 2: Psychological and biological models*. Cambridge, MA: MIT Press 1986
- Mangel, M.: Oviposition site selection and clutch size in insects. *J. Math. Biol.* **25**, 1-22 (1987)
- Mangel, M.: The evolution of optimal behavior by natural selection. Preprint, Department of Zoology, University of California, Davis
- Mangel, M., Clark, C. W.: *Dynamic modeling in behavioral ecology*. Princeton, NJ: Princeton University Press 1988
- Maynard Smith, J.: When learning guides evolution. *Nature* **329**, 761-762 (1987)
- Mulloney, B., Perkel, D. H.: The roles of synthetic models in the study of central pattern generators. In: Cohen, A. (ed.) *Neural control of rhythmic movements in vertebrates*, Chap. 11, pp. 415-453. New York: Wiley 1988
- Miller, J. R., Strickler, K. L.: Finding and accepting host plants. In: Bell, W. J., Carde, R. T. (eds.) *Chemical ecology of insects*. London: Chapman and Hall 1984
- Papaj, D. R., Prokopy, R. J.: Phytochemical basis of learning in *Rhagoletis pomonella* and other herbivorous insects. *J. Chem. Ecol.* **12**, 1125-1143 (1986)
- Prokopy, R. J., Averilli, A. L., Bardinelli, C. M., Bowdan, E. S., Cooley, S. S., Crnjar, R. M., Dundulis, E. A., Roitberg, C. A., Spatcher, P. J., Tumlinson, J. H., Weeks, B. L.: Site of production of an oviposition-detering pheromone components in *Rhagoletis pomonella* flies. *J. Insect Physiol.* **28**, 1-10 (1982)

- Prokopy, R. J., Papaj, D. R., Cooley, S. S., Kallet, C.: On the nature of learning in oviposition site acceptance by apple maggot flies. *Anim. Behav.* **34**, 98–107 (1986)
- Roitberg, B. D., Mangel, M.: On the evolutionary ecology of marking pheromones. *Evol. Ecol.* **2**, 289–315 (1988)
- Roitberg, B. D., Prokopy, R. J.: Experience required for pheromone recognition by the apple maggot fly. *Nature* **292**, 540–541 (1981)
- Roitberg, B. D., Prokopy, R. J.: Host deprivation influence on response of *Rhagoletis pomonella* to its oviposition deterring pheromone. *Physiol. Entomol.* **8**, 69–72 (1983)
- Roitberg, B. D., Prokopy, R. J.: Insects that mark host plants. *Bioscience* **37**, 400–406 (1987)
- Rumelhart, D. E., McClelland, J. L., PDP Research Group: Parallel distributed processing. Volume 1: Foundations. Cambridge, MA: MIT Press 1986
- Schmidt, J. M., Smith, J. J. B.: Host volume measurement by the parasitoid wasp *Trichogramma minutum*: The roles of curvature and surface areas. *Entomol. Exp. Appl.* **39**, 213–221 (1985)
- Schmidt, J. M., Smith, J. J. B.: Correlations between body angles and substrate curvature in the parasitoid wasp *Trichogramma minutum*: A possible mechanism of host radius measurement. *J. Exp. Biol.* **125**, 271–285 (1986)
- Schmidt, J. M., Smith, J. J. B.: Measurement of host curvature by the parasitoid wasp *Trichogramma minutum*, and its effect on host examination and progeny allocation. *J. Exp. Biol.* **129**, 151–164 (1987a)
- Schmidt, J. M., Smith, J. J. B.: Short interval time measurement by a parasitoid wasp. *Science* **237**, 903–905 (1987b)
- Williams, G. C.: Adaptation and natural selection. Princeton, NJ: Princeton University Press 1966

Received June 23, 1989; received in revised form October 8, 1989