

Explicit trade-off rules in proximate adaptive agents

Jarl Giske,^{1*} Marc Mangel,³ Per Jakobsen,² Geir Huse,^{1‡}
Chris Wilcox^{4§} and Espen Strand¹

¹Department of Fisheries and Marine Biology, ²Department of Zoology, University of Bergen, Postboks 7800, 5020 Bergen, Norway, ³Department of Applied Mathematics and Statistics, Jack Baskin School of Engineering and ⁴Department of Environmental Studies, University of California, Santa Cruz, CA 95064, USA

ABSTRACT

Organisms in nature are both proximate (operating by rules of thumb) and adapted (the rules influenced by natural selection). We introduce new methods that can be used to study in silico versions of organisms behaving according to proximate adapted rules. Our approach goes beyond neural networks and offers an alternative to optimization methods. It is based on the idea that organisms receive signals from the environment, that the signals are modified by internal (state-dependent) factors to create feelings (which we refer to as hedonic tones), and that behavioural processes (decisions) are a response to the hedonic tones. We illustrate these ideas through a model of a fish moving in a vertically structured environment, subject to predation and competition from conspecifics. The fish in our model responds to food, light, temperature and conspecifics, without any reference to current or future fitness. We use a combination of hedonic modelling to process the response, and genetic algorithms to modify the response via natural selection, according to internal needs and evolutionary history. We show that many different combinations of genes can lead to similar fitnesses, so that this approach generates genetic diversity. We compare our results with those of a variety of empirical studies and show that our approach can lead to new links between empirical and simulation studies.

Keywords: adaptation, affect, behaviour, individual-based modelling, perception, proximate, sensing, trade-off.

INTRODUCTION

Individual-based modelling (*sensu* Huston *et al.*, 1988) is now widely used by modellers in population ecology (Grimm, 1999; Huse *et al.*, 2002). The structure of individual-based models allows richly specified simulations of populations in natural ecosystems. The approach has been most successful when the intra-population variation in physiology, behaviour or spatial distribution has been decisive for understanding population processes.

* Author to whom all correspondence should be addressed. e-mail: jarl.giske@ifm.uib.no

‡ Present address: Institute of Marine Research, Postboks 1870 Nordnes, 5817 Bergen, Norway.

§ Present address: Department of Zoology and Entomology, University of Queensland, St. Lucia, Queensland 4072, Australia.

Consult the copyright statement on the inside front cover for non-commercial copying policies.

However, the optimism of Huston *et al.* (1988) has only partly been met and, for example, individual-based models are still rarely used in management (Bart, 1995; Grimm, 1999; Railsback, 2001). One of the main reasons may be that individual-based models often rely on *ad hoc* formulations of behavioural and other responses to the environment (Railsback, 2001; Huse *et al.*, 2002). The use of these *ad hoc* formulations is also quite understandable, since there exists no quantitative theory on how organisms are expected to respond to the multitude of data streams that they simultaneously receive (Tyler and Rose, 1994; Giske *et al.*, 1998).

Behaviour can be implemented in individual-based models by rules derived from a fitness measure. For example, one may assume that organisms are predicted to live in a way that maximizes growth or minimizes mortality risk per growth rate. These and similar rules can be derived explicitly from life-history theory (e.g. Roff, 2002) or other evolutionary approaches. However, all the major quantitative branches of evolutionary ecology depend on major simplifications in the descriptions of both the organisms and their environment, which make their behavioural rules less applicable in more realistic, complex situations (Giske *et al.*, 1998; Railsback, 2001). Life-history theory usually requires that conditions be the same for all generations, and that trade-offs be age- or stage-dependent. Game theory (Dugatkin and Reeve, 1998) eliminates time and finds steady-state equilibrium solutions, and stochastic dynamic programming (e.g. Houston and McNamara, 1999; Clark and Mangel, 2000) relies on a knowledge or anticipation of the future events over the model horizon. In many cases, the environment is so variable that none of these assumptions holds. This is probably one of the main reasons why practitioners of individual-based modelling stick to *ad hoc* formulations of behaviour and life history. However, *ad hoc* formulations limit the validity of the results (Railsback *et al.*, 1999).

Modification of behaviour through adaptation can be an alternative to both global optimization and *ad hoc* approaches in individual-based modelling. Rather than deriving rules from equations of Darwinian fitness, models based on adaptation evolve improved solutions by natural selection (Holland, 1975; Mitchell and Taylor, 1999). The adapted random walk (Huse, 2001) and individual-based neural network genetic algorithm (ING; Giske *et al.*, 1998; Huse and Giske, 1998) allow adaptive behavioural decisions to emerge in populations inhabiting complex environments. In ING models, the behaviour of an organism is determined by an artificial neural network (Hopfield, 1982; Kohonen, 1984), where different types of sensory information are weighed to make decisions. The weighing factors in the artificial neural network are coded as genes. A genetic algorithm (Holland, 1975; Goldberg, 1989) passes the genes of the successful reproducers to the next generation, with modifications based on recombination and mutation (van Rooij *et al.*, 1996). The ING method is thus both ultimate (its genetic algorithm) and proximate (its artificial neural network) in design, which allows studies of complex scenarios and richly specified agents. The ING method provides a straightforward implementation of behaviour in individual-based models. It has been successful in the sense that ING models are both able to solve pure state- or density-dependent problems (Huse *et al.*, 1999) and complex problems where optimization techniques are inappropriate (Huse and Giske, 1998; Strand *et al.*, 2002).

The main problem with ING is that the artificial neural network is hard to visualize and analyse. The synaptic connections between (sensory) information and the hidden and output layers of neurons have no direct ecological or physiological interpretation. The artificial neural network is an imitation of how the human brain is thought to function, and both are black boxes from the perspective of the outside observer. Hence, the ING method

can solve more complex problems than can the evolutionary optimization tools, but the ecological and behavioural trade-offs remain invisible to the observers. This impairs the interpretive value of ING and other approaches based on the artificial neural network.

Here, we continue in the philosophy of the ING: search for a way to combine state, age and density in the same model. However, we will also require the solution to be understandable. The goal is a tool that can handle complex decisions and communicate the major trade-offs in use by the organisms to the researcher. As for artificial neural networks, we will base this approach on a metaphor of the brain. But rather than focus on the myriad of neurons, we mimic the overall brain architecture and the pathway of information from perception to decision.

Hedonic modelling: models based on feelings

Sensing and perception are the core of a proximate approach to behaviour. Examples abound. The flatworm has evolved a nervous system of peripheral sensors and a central nervous system where decisions are made. These simple animals are able to make complex foraging decisions depending on density of conspecifics (Cash *et al.*, 1993) and to learn to associate predation risk with novel cues (Wisenden and Millard, 2001). Fish are able to make complex decisions based on observations of predators, hunger and risk dilution by conspecifics (Magurran *et al.*, 1985). The brains of bony fish can be divided into the hindbrain, the midbrain and the tiny forebrain (MacLean, 1990). Most sensory signals enter the fish brain in the hindbrain regions, but the perception of information takes place in specialized centres for vision, hearing and olfaction in the midbrain and forebrain (Atema *et al.*, 1988; MacLean, 1990).

At present, there is no consensus on terminology and process understanding in the behavioural sciences (e.g. Ekman and Davidson, 1994). We adopt the following:

- The sensory system and the brain filter the perceived information and send relevant information to other parts of the brain, responsible for adding a ‘feeling’, through a neural response, to the perception.
- This feeling is necessary for an organism (the ‘agent’) to react to the primary stimulus.
- Feelings based on sensory information (e.g. sounds, smells, tastes, pressure and temperature) are called ‘affects’.
- Behavioural processes (‘decisions’) are associated with ‘motivations’ based on affects. That is, affects give the agent its motivation for behaviour.

The affect consists of (1) the primary stimulus, (2) a secondary perception of the stimulus and (3) a tertiary ‘hedonic tone’, derived from the organism’s genetically and physiologically determined disposition, that evaluates the ‘feeling’ associated with the stimulus (Scherer, 1984; LeDoux, 1994). The hedonic tone – positive (attraction), negative (avoidance) or neutral (no affect raised) – determines behaviour.

Affects from separate perceptions combine to create a single hedonic tone, but the physiological mechanism of this combination is currently unknown. Our objective is to evolve behaviourally interpretable artificial genes for responses to stimuli so that the resulting behaviour of the adapted individuals can be understood simultaneously from the perspectives of ultimate fitness maximization and proximate responses to stimuli (cf. Thorpe *et al.*, 1998).

MODEL

Environment

The population of proximate adaptive agents that we model inhabits an aquatic environment with gradients only along the vertical and temporal axes. Environmental variation is linked to the diel cycle. There are 75 time steps in one diel cycle. Underwater light intensity varies with the diel cycle, and the prey organisms of the agent population perform a diel vertical migration according to changes in light level. At any time, prey are found in half the water column of 30 layers, with concentrations in the centre of the bell-shaped distribution seven times higher than at the outer edges (Fig. 1). Temperature does not change over time. It is uniform (15°C) in the upper third of the water column and then decreases linearly towards the bottom (2°C).

Model organisms

In each of the potential 150 time steps in its life, the organism can make a single decision: to remain where it is or to move one cell up or down. The environmental cues are food, light, temperature and conspecifics (other agents) in each of the three depth cells. The focal agent is further able to sum the feelings for each of the three possible decisions, and will move to the cell that gives the most positive hedonic tone. It can also measure its own degree of stomach fullness, age and body mass. We assume that most aquatic animals are capable of sensing and processing these types of information (Atema *et al.*, 1988). The functional relationship between these variables and the ultimate fecundity and survival probability of the organism is not included in their decision processes. The agents do not know which depth they are at, where they were born or mortality risk at the depth, but they can get clues from light intensity and from the presence of other agents. The current movement decision is made without future anticipation.

The agents can live for two diel cycles. If the organism is considered as a fish or another pelagic visual predator, this time scale may seem unrealistically low. We chose this as a

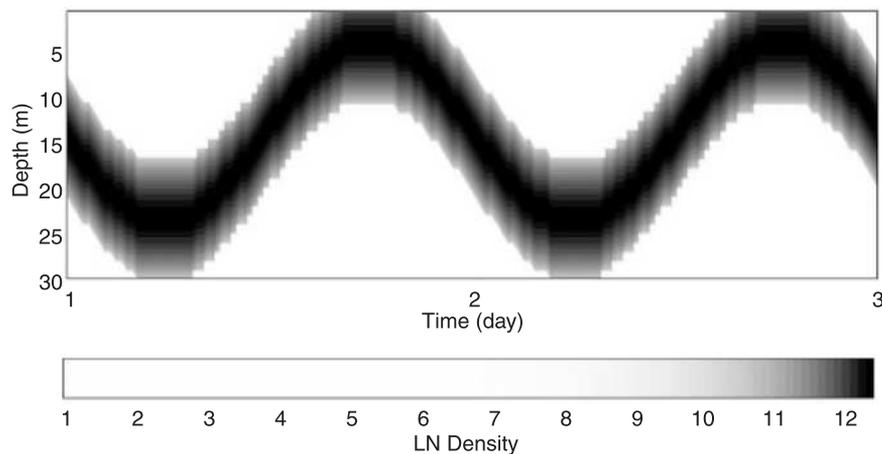


Fig. 1. Vertical distribution of the prey population during the two diel cycles of each generation.

minimum life span for age-dependent behaviour to emerge – that is, the ability of the organisms to live differently during the two diel cycles. If it has survived and has sufficient resources, an individual can reproduce in the final time step. All individuals are potential mothers, in the sense they may produce offspring; 25% of a survivor's final body mass can be converted to eggs. The weight of an egg is set to 250 times the mass of a prey item.

Each survivor may also become a father. In some simulations, individuals are born with two additional genes, specifying the probabilities of point-mutations in the genome and the probability of reproducing sexually, respectively. In the other simulations, these probabilities are fixed and equal for all individuals. In the case of asexual reproduction, all offspring are genetically identical to their mother, except for mutations. In the case of sexual reproduction, the genes of the offspring will consist of a mix of maternal and paternal genes.

The feeding rate of the agents depends on visual range, prey density, local food competition from other agents in the same depth cell, as well as stomach capacity. (Full equations for physiology and mortality of the agents are given in the Appendix.) Digestion rate and growth are temperature-dependent. The potential feeding rate doubles for every 10°C. In the absence of competition and satiation, the feeding rate is a function of the prey encounter rate. To model density-dependent processes without having to model the food resource dynamically, feeding rate is reduced proportionally to the number of competitors in the same cell. Finally, the feeding rate of an individual is constrained by its current body size and its recent feeding history, so that the stomach content cannot surpass 25% of its body mass. Stomach content is evacuated by an exponential and temperature-dependent function.

The organism also has metabolic costs, depending on temperature and body mass. The weight gain or loss of the organism is the difference between its feeding gain and its metabolic costs, depending on prey concentration, light intensity, temperature, number of competitors, feeding history and body size.

We assume that the predator of the modelled organism is a visual feeder. Then, the mortality risk can be found from light intensity and risk dilution when conspecifics are present (Mangel and Clark, 1988: chapter 5; Giske *et al.*, 1994, 1997). To avoid cancellation of the importance of light or other agents in feeding and mortality risk, dilution of risk falls off at a different density-dependent rate ($N^{-0.5}$) than dilution of feeding gain (N^{-1}). The rationale for this is that, because of their larger body size, the adaptive agents may be better than their prey in detecting and escaping an approaching predator.

Decision making

The study of perception and reaction to stimuli has a long history in ethology (Curio, 1976; Ewert, 1980; McFarland, 1985). However, modelling decisions is difficult without a fitness measure. For example, McFarland and co-workers used a utility-based approach to ethology (McFarland and Sibly, 1975; Sibly and McFarland, 1976; McFarland, 1977; McFarland and Houston, 1981; McNamara and Houston, 1986). Other traditions in research on sensory-based decisions also rely on some anticipation of the utility of each possible choice (e.g. Cabanac, 1992; Gallistel, 1994; Leven and Levine, 1996; Doya, 1999, 2000). However, these methods are hard to utilize in forward individual-based models, or require low complexity in the description of the environment or the agent (Tyler and Rose, 1994; Giske *et al.*, 1998; Railsback, 2001).

We model decisions based on the present perceptions, without evaluation of future utility. We also link to fitness, but this is provided by the adaptation of the affect genes in the genetic algorithm. We assume there is no cognitive evaluation in the decision process of our modelled agents. The affects are determined both by internal and external factors (Scherer, 1984; LeDoux, 1994). We define the internal factor by the genetic disposition M and a possible modification by one or more states S of the organism:

$$M \cdot S^m \quad (1a)$$

Genetic disposition determines the direction of the affects, which are attraction (>0), avoidance (<0) or neglect (0). The strength of the internal factor determining behaviour, however, is also influenced by the internal states of the organism. A negative value of a gene (m) for a state-dependent disposition regulator means that the response decreases with increasing state, whereas a positive value means that the response increases with increasing state. A value of zero means that the state does not impact disposition.

The external factor determining behaviour is the perception P of the environmental cue X , which can be modified by internal state in a similar manner:

$$P \cdot S^m \quad (1b)$$

We assume that the hedonic tone $H_{i,X}$ of individual i towards the perceived cue X , when m_{1i} and m_{2i} are the numeric impacts of the gene products, takes the form

$$H_{i,X} = M_{i,X} S_1^{m_{1i}} S_2^{m_{2i}} P_X \quad (2)$$

In our model, the state S_1 is always the degree of stomach fullness; S_2 may be either body mass or age. The internal states of the organism are scaled in the $[1, 2]$ range. This means that organisms with low state will have a neutral state-dependent response ($1^m = 1$), while organisms with higher state will have strengthened ($m > 0$) or weakened ($m < 0$) response.

We assume that the agent can perceive information relevant for several affects simultaneously, and that its movement depends on the combined hedonic tone of all affects. To begin, we define the net attraction or avoidance for depth z for individual i to be a composite of four affects:

$$H_{i,z} = H_{i,\text{food}} + H_{i,\text{light}} + H_{i,\text{temperature}} + H_{i,\text{other_agents}} \quad (3)$$

where each hedonic tone is of the form in equation (2).

The individual chooses the depth among $z + 1$, z , and $z - 1$ that gives the most positive (or least negative) hedonic tone. Since the aquatic environment is continuous, we assume that the agent, placed in the centre of its cell, can measure gradients that hint to the conditions in the neighbouring cells without actually visiting them. At each interval during its lifetime, an individual performs this spatial choice. Growth, mortality risk, body size and fecundity emerge directly from these choices. Survivors at the end of the final time step in a generation produce offspring in proportion to their accumulated resources.

The process of adaptation

The genotype of an agent thus consists of n (ranging from 4 to 12 in the experiments conducted) artificial genes

$$\mathbf{G}_i = [h_{1i}, h_{2i}, \dots, h_{ni}] \quad (4)$$

labelled M and m in equation (2). The n genes of the genotype form the ‘strategy’ vector (Huse *et al.*, 2002) of the agent, describing its decision rules. Its phenotype, or ‘attribute’ vector (Chambers, 1993), is a function of its birth depth and the conditions encountered as a consequence of past decisions based on G_i . The attribute vector of individual i at time t is, therefore,

$$A_{i,t} = (\text{body mass}_{i,t}, \text{ingestion}_{i,t}, \text{stomach mass}_{i,t}, \text{offspring}_{i,t}, \text{depth}_{i,t}, \text{age}_{i,t}) \quad (5)$$

where body mass, ingestion and stomach mass are measured in the same arbitrary weight units. Reproduction is 0 until the final time step, and then a step function of the final body mass, as explained above.

When three genes (equation 2) control each of four affects, 12 hedonic genes are needed. All genes are free to evolve within $[-5, 5]$ at allele steps of 0.1. The boundaries of ± 5 eliminate a co-evolutionary race of increasing strength of affect genes, and the 101 possible alleles of each gene still allow a fine resolution of the genetic disposition for behaviour.

The alleles in the gene pool evolve by a genetic algorithm so that the evolutionary reward for appropriate feelings is the likelihood of survival until reproduction and many descendants. We used a discrete generation individual-based model (Huston *et al.*, 1988; DeAngelis *et al.*, 1992) to track the life cycles in a semelparous cohort of about 10,000 newborn individuals in each generation (constrained by food and mortality risk). Each simulation runs over 1000 generations. In the long run, agents that obtain feelings leading them to profitable depths will produce more offspring than the others and their alleles will increase in abundance in the gene pool. Since the generation time is fixed in the model, the appropriate measure of fitness is expected reproductive success (R_0). But fitness is also an emergent by-product of the genetic algorithm, in that no fitness measure is specified (Strand *et al.*, 2002).

The genetic algorithm moves upwards in the fitness landscape while simultaneously searching the neighbourhood (by mutations and recombinations) for alternative peaks (Holland, 1975). However, we cannot be completely sure that the very highest tops in the fitness landscape are found after searching a tiny fraction of less than 10^{-12} of all possible genomes. As with real populations, we also risk the possibility of genetic drift, maladaptations and even extinctions. Finally, the hill-climbing ability of a genetic algorithm in a density-dependent landscape is poorly studied. We performed five independent experiments with each equation set, each starting with different random number seeds, so that all factors that have a random component (alleles of individuals in the first generation, and Monte-Carlo-driven mutations, recombinations, partner choice and individual survival from predators in all time steps) will differ among runs.

Since the trade-off factors in the decision equations are coded as genes, individuals that survive until reproduction will pass on a set of their genes to each of their offspring. In the first generation, the alleles are set randomly and some individuals feel the strangest attractions and aversions. Alleles that combine in a genome to give feelings that make the organism survive until reproduction with a large body size will tend to increase in abundance in the gene pool.

The major factor in the adaptive evolution of the gene pool is natural selection acting on the phenotypes, just as for real biological gene pools. Mutations come in three grades: 90% of the mutations are one-step mutations (changing the allele value by ± 0.1), 9% are four-step mutations and the last percentage is a 16-step mutation. If the mutation rate is x , then each of an agent’s alleles will on average mutate $1.42x$ single-allele steps, and a 12-gene

descendant genome will on average deviate $17.4x$ steps from its mother's. In addition to this mode of asexual reproduction, there is also a chance that an offspring will inherit some of its genes from its father. The agents are haploid. Technically, recombination is done by inserting parts of the father's genome into the egg. The crossover point is chosen randomly in the genome. Fathers are thus hermaphrodites, and the recombination does not affect the eggs produced by the father. Three qualities are required to become a father: to have survived until reproduction, to be spatially close to the mother and to be the largest of the 15 candidates she encounters in the mate search process. As for the case of mutation, the probability of a recombination can either be controlled by an individually inherited and evolvable gene [0.01–1] or be preset at a fixed probability for all individuals.

The aquatic world we have constructed is, like the real one, a tough place to live. To avoid extinction in the first generation of randomly created agents, we initiated all simulations with a population that was 3–4 times higher than the usual population fecundity in the adapted state. This was usually enough to allow a few thousand non-random parents for the second generation. But in some simulation experiments, this would not be enough to prevent extinction, and we then started the evolution in a benign environment where survival was easy and offspring were cheap. We then gradually increased costs of living and mortality risk and decreased growth efficiency over the first 12 generations to reach the same level as in the other simulations.

Experiments

We conducted computer experiments based on: (1) adaptive or fixed mutation and recombination rates for the speed and outcome of the adaptation process; (2) the effect of internal states in a full 2-state and 4-affect model; and (3) the effect of reducing the number of affects and states. The comparative measure was population fecundity, as a measure of the average fitness found in the experiment. In addition, the speed, in number of generations needed to find the adapted gene pool, was also measured in some experiments. From preliminary experiments, we knew that a specific combination of states usually produced more offspring than other combinations. This combination was used in Experiments 1 and 3 and tested in Experiment 2.

RESULTS

Density dependence

In all experiments, there was strong fluctuation in population fecundity between parent and offspring generations (Figs 2, 3). In very small populations, individual fecundity was high due to low competition. In very large populations, risk dilution more than compensated for reduced feeding rate, and the population produced many offspring in spite of density-dependent competition. The period for the density dependence was usually a few generations. Population sizes of 12,000 could be attained after one intermediate generation at 11,000. Population sizes of 13,000 or 14,000 required several intermediate steps, since only population sizes below 11,000 produced more than 13,000 offspring, and only population sizes around 12,000 could produce less than 11,000 offspring (Fig. 2). Since population size fluctuated so heavily between generations, estimates of average values would have to cover many generations. We therefore compared different simulation

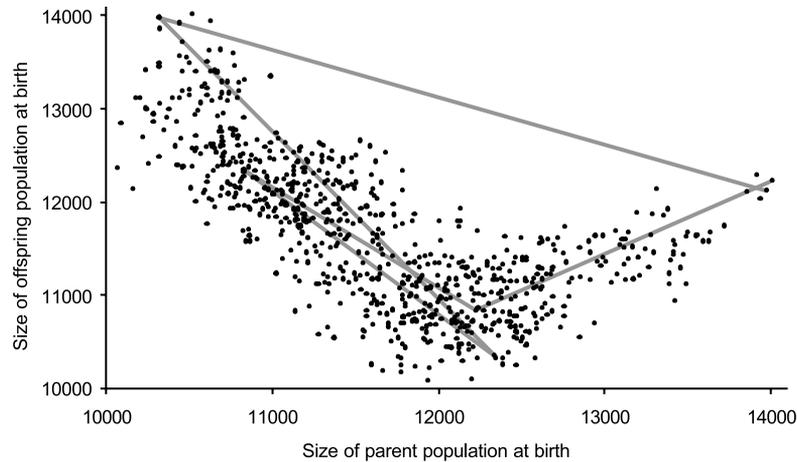


Fig. 2. Density-dependent population dynamics. Population offspring production for each of the last 800 generations in the first simulation with the highest ranked equation set. The straight lines show an example of a five-step path starting and ending at approximately 14,100 agents.

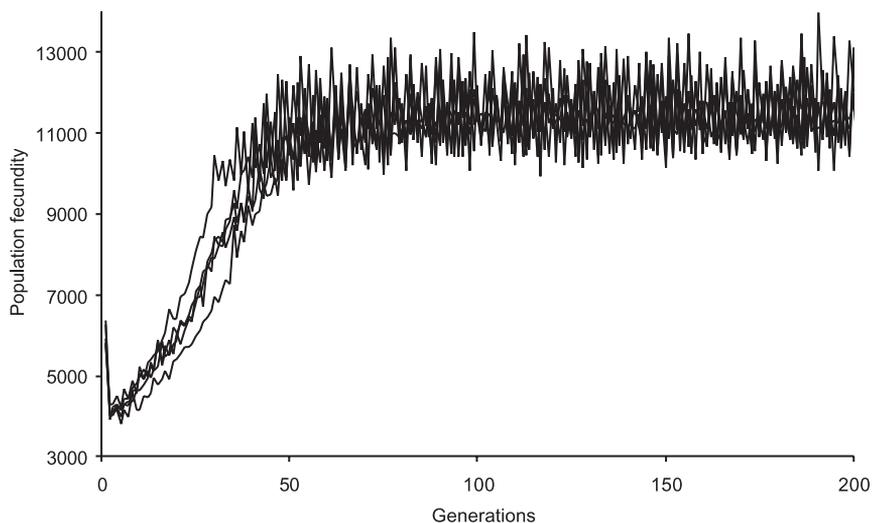


Fig. 3. Evolution of population fecundity in the five replicate simulations with mutation and evolution rates at 10%. The states chosen are the same as those in the highest ranked equation set in Table 3. The initial fall in fecundity from generation 1 to 2 is due to an initiation of 40,000 random individuals at the beginning of the first generation. This random population generated 4000–7000 offspring.

experiments by first averaging the offspring production in groups of 500 data sets (100 generations, 5 replicate simulations of the same experiment). Finally, we calculated the average population offspring production (± 1 standard deviation) for all except the first two of the 100-generation periods of each equation set. We assumed the final averages to be independent.

Evolution and evolutionary rates

All experiments started with random gene pools and random vertical distribution of the population. Differences in population fecundity between the five runs of the same experiment were always small and were not visible after about 50 generations (Fig. 3). In simulations of the effect of mutation and recombination rates on evolution, there were only minor differences in population fecundity between experiments run with adaptive and fixed probabilities at 1–15% (Table 1, Fig. 4), while runs with lower rates (1%) ended with less fecund populations. In the simulations with adaptive evolution, both mutation rate and recombination rates fell to very low levels after a few hundred generations (Fig. 5). In all experiments, except in those with the lowest evolution rates, population fecundity levelled after 60–80 generations. No significant increase in offspring production was seen after 200 generations. Also in simulations over 10,000 generations, no improvement in population performance was seen after 1000 generations (Fig. 6). Based on these results, further experiments were run over 1000 generations, with fixed point-mutation and genome recombination probabilities at 10%.

Table 1. Population adaptation under different evolution regimes

Mutation and recombination rates	Mean fecundity (± 1 SD)
Adaptive, individually inherited	11558 \pm 15
0.01	11221 \pm 58
0.05	11539 \pm 50
0.075	11579 \pm 37
0.1	11620 \pm 43
0.15	11582 \pm 20

Note: Mean fecundity (± 1 standard deviation) is given for generations 200–1000 in 100-generation intervals. The standard deviation is measured between these eight 100-generation averages. All simulations are run with the same choice of second state modifiers as in the highest ranked equation set in Table 3.

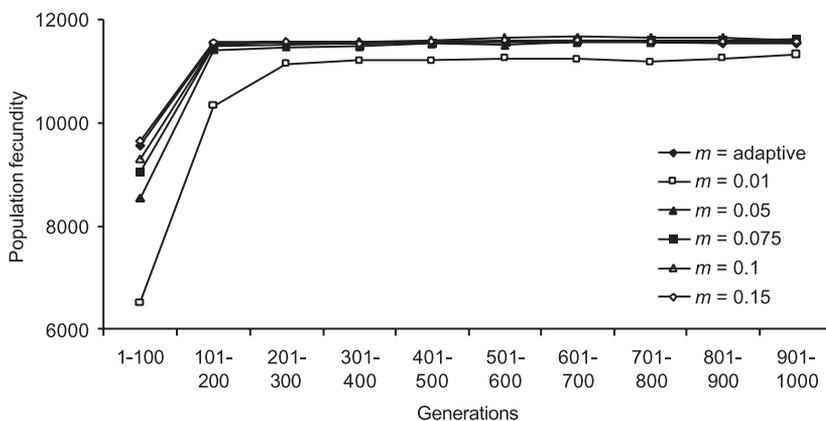


Fig. 4. Evolution of population fecundity under different recombination and mutation rates. Each point represents the average population fecundity over 100 generations in five replicate simulations.

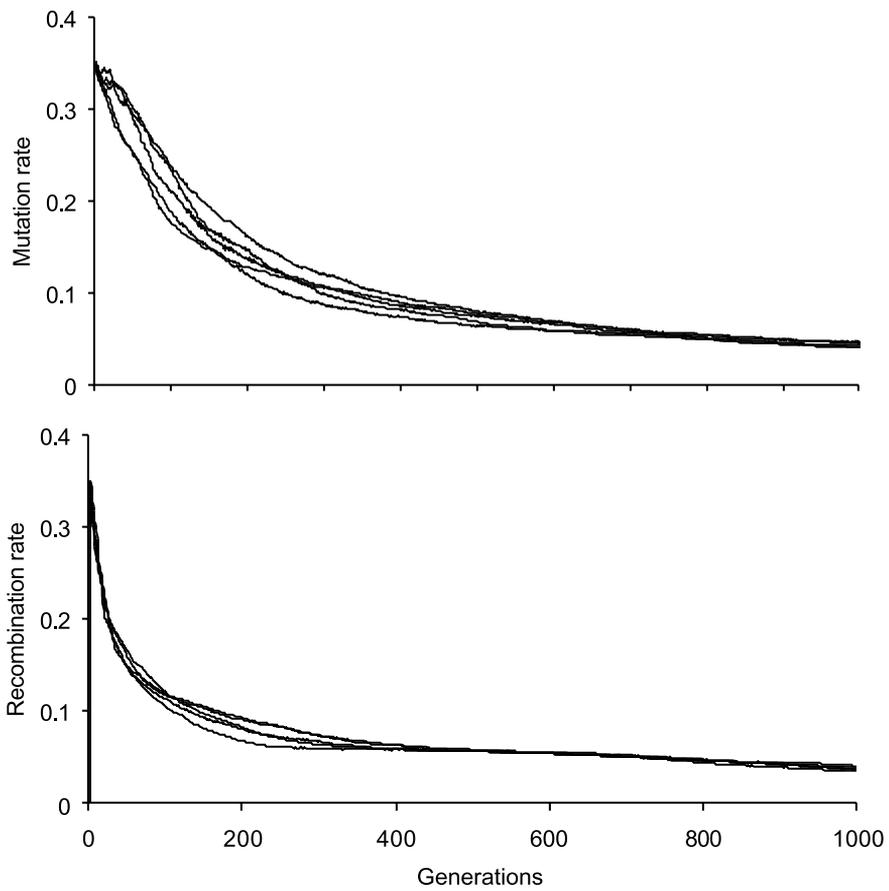


Fig. 5. Evolution of average mutation rate (top) and recombination rate (bottom) in the five simulations of adaptive evolution rates.

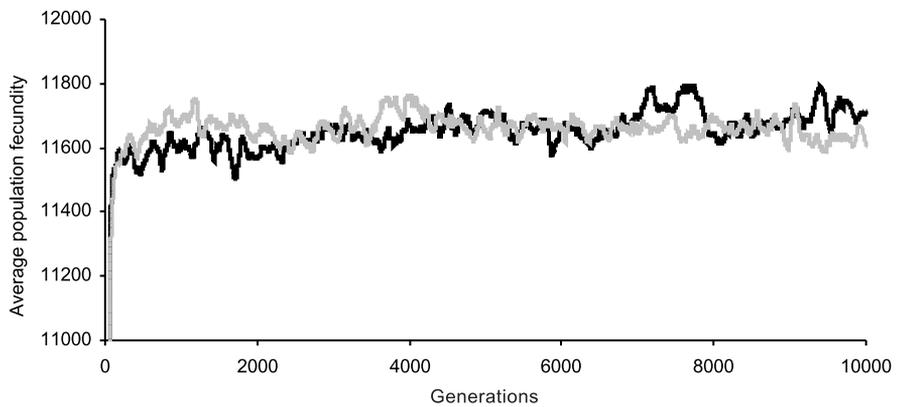


Fig. 6. The 100-generation running mean of population fecundity in two simulations over 10,000 generations using the highest ranked equation set in Table 3.

Table 2. Effect on population mean fecundity (± 1 standard deviation) of reducing complexity of sensing and feeling

Component removed relative to equation (2, 3)	Equation	New formulation	Mean fecundity (± 1 SD) for the last eight 100-generation periods
Nothing removed		Best ranked set in Table 3	11620 \pm 43
Fixed genetic disposition for all perceptions	2	$H_{i,X} = S_1^{m_1} S_2^{m_2} P_X$	4751 \pm 117
First state (stomach fullness)	2	$H_{i,X} = M_{i,X} S_2^{m_2} P_X$	11621 \pm 17
Second state (age or body mass)	2	$H_{i,X} = M_{i,X} S_1^{m_1} P_X$	11645 \pm 19
Both states (stomach and age or body mass)	2	$H_{i,X} = M_{i,X} P_X$	0 \pm 0
One affect (other agents)	3	$H_{i,z} = H_{i,food} + H_{i,light} + H_{i,temperature}$	8202 \pm 39
One affect (temperature)	3	$H_{i,z} = H_{i,food} + H_{i,light} + H_{i,agents}$	11457 \pm 23
Two affects (temperature and other agents)	3	$H_{i,z} = H_{i,food} + H_{i,light}$	10877 \pm 22
Three affects (light, temperature and other agents)	3	$H_{i,z} = H_{i,food}$	0 \pm 0

Note: The state variables in the basic equation set, which is modified, is the best performing set in Table 3.

Affect structure

The fecundity of the population strongly depends on the structure of the affect equation (Table 2). Simulations with either one or both states have high fecundities. If both state variables are removed, populations always go extinct in a few generations. Temperature is relatively unimportant to the outcome. Simulations in which the perception of other agents is removed require weak initial selection to persist. When food is the only source of information, agent populations go extinct quickly, even with weak natural selection over the first 12 generations.

State dependence

In the model with two states for each of four affects, where state 1 is always stomach content and state 2 can be either age or body mass, there are 16 different permutations of states. Each permutation, called an 'equation set', was simulated for 1000 generations in five replicates. The ranks and average fecundity of these sets over time were quite consistent during the time evolution, and four of the sets performed consistently better than the rest (Fig. 7, Table 3). The 12 poorest sets produced on average 402 fewer offspring than the best set during the final 800 generations. The effect of the second state modifier is most clearly seen for the feeling of light, where equation sets using body mass were ranked 1 to 7. For food, the value of body mass as second modifying state is clear. Here, the four best performing equation sets utilized body mass. For temperature and other agents, the average ranking of sets containing age or body mass are the same (Table 3).

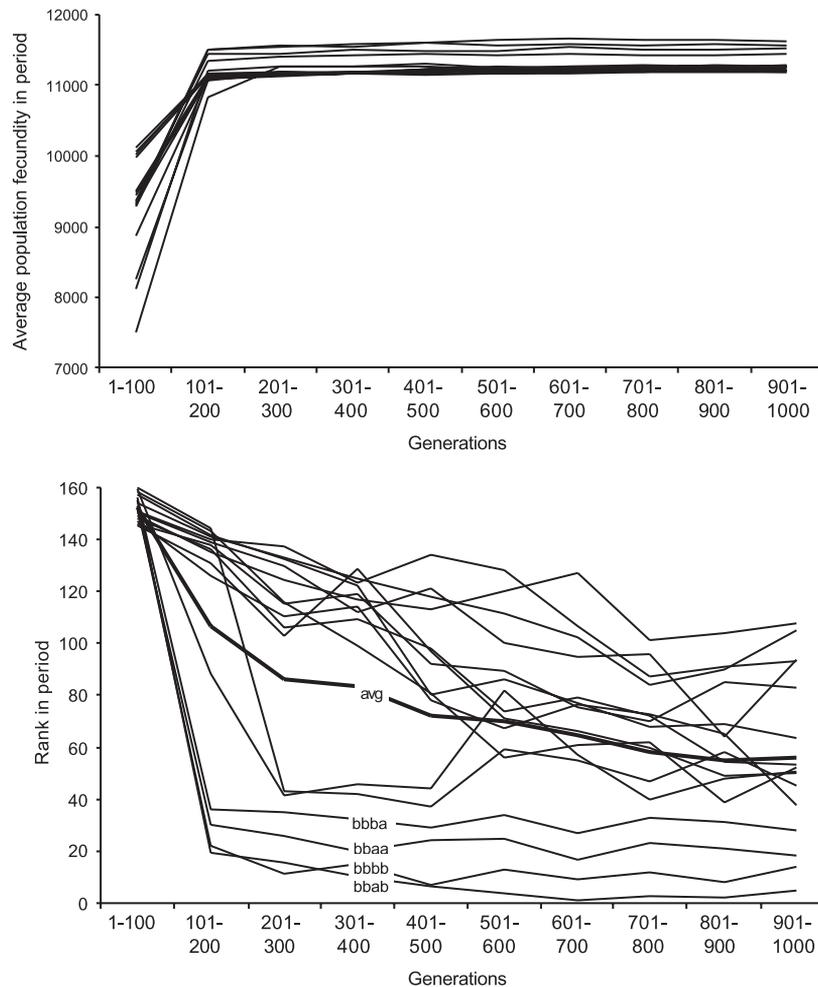


Fig. 7. Time-development in average performance of the 16 equation sets, differing in the second state modifying affect for the four perceptions. (Top) Average fecundity in each 100-generation period of the five replicate simulations. (Bottom) Rank of the equation sets in terms of average fecundity. For the four best performing equation sets, the second state ($a = \text{age}$, $b = \text{body mass}$) modifying the feeling towards food, light, temperature and other agents, respectively, is given.

The gene pool

From a start with random allele values $[-5, 5]$ of all 12 hedonic genes, the gene pool evolved rapidly towards stabilized fecundity (Fig. 3). After the first generation, alleles that made their bearers move away from food concentrations were eliminated. After about 200 generations, the gene pool changed very slowly, and generations 200 and 1000 were very similar. Density-dependent selection between generations (Fig. 2) impacts the strength of the response for light most strongly, but also the three genes controlling temperature affects (Table 4). However, these fluctuations have little impact on the short-term evolution of the total gene pool, as seen for generations 999 and 1000 (Fig. 8). Still, even after 10,000

Table 3. The impact of the second state for population fecundity in the 16 equation sets

Second state modifying the feeling from the perception of:				Mean score (± 1 SD) in each of the eight 100-generation periods	
Food	Light	Temperature	Other agents	Rank	Fecundity
body mass	body mass	age	body mass	5.9 \pm 4.9	11620 \pm 43
body mass	body mass	body mass	body mass	11.1 \pm 2.9	11573 \pm 15
body mass	body mass	age	age	21.8 \pm 3.3	11496 \pm 43
body mass	body mass	body mass	age	31.1 \pm 2.9	11426 \pm 16
age	body mass	body mass	body mass	48.3 \pm 8.1	11273 \pm 14
age	body mass	age	body mass	51.1 \pm 13.6	11267 \pm 18
age	body mass	age	age	70.8 \pm 25.9	11241 \pm 40
body mass	age	age	age	77.6 \pm 24.6	11233 \pm 37
body mass	age	body mass	body mass	78.1 \pm 28.6	11229 \pm 42
body mass	age	body mass	age	80.6 \pm 21.9	11225 \pm 35
age	body mass	body mass	age	87.1 \pm 25.8	11218 \pm 41
body mass	age	age	body mass	91.0 \pm 17.6	11208 \pm 28
age	age	body mass	body mass	101.5 \pm 20.1	11192 \pm 30
age	age	body mass	age	108.5 \pm 16.8	11182 \pm 21
age	age	age	age	112.5 \pm 20.5	11174 \pm 29
age	age	age	body mass	114.3 \pm 9.4	11176 \pm 11

Note: Mean fecundity (± 1 standard deviation) in each of eight time periods of each of 16 sets are ranked together, yielding 128 numbers to compare. Due to the strong increase in fecundity during the first 200 generations, only the last 800 are used in this comparison.

generations, there are visible differences between the gene pools of independent simulations of the same equation (Fig. 9).

Individual decisions

At a broad scale, the population follows the spatial distribution of its food resource (Figs 1, 10). Individuals born at the outskirts of the population spatial distribution experience less resource competition and lower risk dilution than individuals born at intermediate depths (Fig. 11). As a result, the population is near or at a situation where all phenotypes have the same expected offspring production (Fig. 11, Table 5). In this situation, the remaining genetic variation is conserved by frequency-dependent selection.

We collected behavioural data for 100 individuals in the last generation and we use this data set to present the feelings of a few of those individuals. We focus on the smallest (agent 39) and two of the largest (43 and 93) survivors of these 100 (Figs 12–14). In many cases, the feelings for the preferred movement choice is only slightly better than for the next alternative (Fig. 14), compared with the average values of the feelings (Fig. 13). This means that although the feeling towards other agents is usually rather weak, it may have a decisive impact on the movement decision.

Individual 39 is the slowest growing of the survivors among the 100 agents studied. It still has a higher disposition for food and a lower fear of light than the fastest growing agent 43.

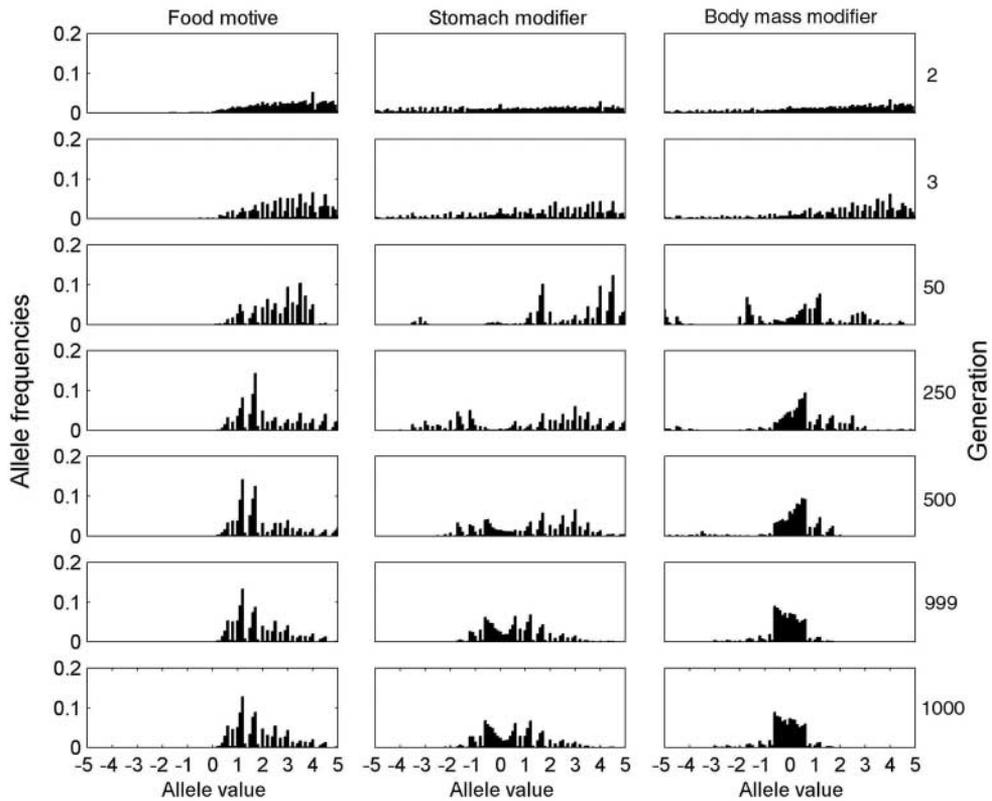


Fig. 8. Adaptation of the gene pool. Allele frequencies at selected generations for the three genes impacting the feeling for food.

However, both these differences, as well as the attraction to other agents, drive agent 39 towards the centre of the food distribution, where competition from other agents is more severe. Due to its strong attraction to light, it is often further up in the water column than the other agents. Fear of light increases with age and body mass, whereas the attraction to food, temperature and other agents fades.

Individual 43 is an agent with substantial fear. It has a slightly lower than average genetic disposition for feeding, and strong fear for light and heat (equation 6). The fear for light diminishes with increasing stomach contents, and is strongly increased with body mass. [Its stomach is never full (Fig. 12), so it utilizes stomach fullness as a signal for feeding opportunity.] It takes no notice of other agents. With these genes, agent 43 usually positions deeper than the others, following the deepest segment of the migrating food resource. By being at the spatial extreme, competition is low, and agent 43 is the fastest growing of the 100 individuals investigated. Hence, the fecundity of agent 43 is an example of frequency-dependent fitness of its alleles, since many offspring produced with identical genetic dispositions would encounter high competition in this low-productive environment.

Agent 93, the other large individual, is at the opposite extreme of 43, most of the time

Table 4. Mean (± 1 standard deviation) allele values of population

Gene		Mean allele value (± 1 SD) in generation:			Standard deviation in mean allele value between generations 950 and 1000
		250	500	1000	
Food	M	2.2 ± 1.2	1.9 ± 1.1	1.7 ± 1.0	0.04
	m_1	1.2 ± 2.4	1.1 ± 1.7	0.5 ± 1.1	0.04
	m_2	0.5 ± 1.4	0.1 ± 1	-0.2 ± 0.6	0.01
Light	M	-0.8 ± 3	-0.4 ± 2.8	-0.6 ± 2.5	0.11
	m_1	-0.3 ± 1.7	-0.4 ± 1	-0.5 ± 0.7	0.03
	m_2	1.4 ± 1.1	1.1 ± 0.7	0.6 ± 0.6	0.03
Temperature	M	0.2 ± 1.6	0.4 ± 1.9	-0.1 ± 1.4	0.06
	m_1	-0.4 ± 1.4	-1.0 ± 1.9	-1.2 ± 2.1	0.07
	m_2	-0.6 ± 1.7	-1.0 ± 1.2	-1.4 ± 1.7	0.08
Others	M	-0.4 ± 1.2	-0.3 ± 1.2	-0.1 ± 0.6	0.02
	m_1	1.2 ± 1.9	0.3 ± 1.2	-0.8 ± 1.5	0.06
	m_2	-3.3 ± 1.4	-3.3 ± 1.2	-3.5 ± 1.3	0.04

being close to the surface. However, this is also a risky behaviour, and many of the others that took the same decision died before reproduction (Fig. 11). Its risk-prone behaviour is due to a strong attraction to light, balanced by a tendency to prefer colder water. Since temperature is uniform in the upper 15 m, the heat avoidance of agent 93 will not impact spatial decisions during the night. At this time, it is drawn downwards by attraction to food and other agents. With increasing body mass, the attraction to light and other agents is weakened, while the downward pulling attraction to food increases.

The explicit decision equations for these three agents are

$$\begin{aligned}
 H_{39,F} &= 1.3S^{2.2}B^{-0.9}P_F & H_{43,F} &= 1.2S^{1.7}P_F & H_{93,F} &= 0.5S^{-0.8}B^{0.6}P_F \\
 H_{39,L} &= -1.2B^{4.4}P_L & H_{43,L} &= -1.7S^{-1.0}B^{4.5}P_L & H_{93,L} &= 1.8S^{0.5}B^{0.2}P_L \\
 H_{39,T} &= 4.9S^{0.6}A^{-1.1}P_T & H_{43,T} &= -1.2S^{-3.0}A^{-2.0}P_T & H_{93,T} &= -0.6S^{2.2}A^{-2.2}P_T \\
 H_{39,O} &= 0.3S^{-0.6}B^{-0.7}P_O & H_{43,O} &= 0.0 & H_{93,O} &= 0.4S^{2.4}B^{-2.5}P_O
 \end{aligned} \tag{6}$$

where S is the degree of stomach fullness, B is body mass and A is age, all scaled in the [1, 2] range. Here we see that the second state has no impact on the feeding affect in agent 43 (since the allele value of m_2 is 0), and that other agents have no impact on its decisions (since $M_O = 0$). A similar effect leads to loss of information about stomach fullness for the fear of light in agent 39.

DISCUSSION

After 200 generations, the population of proximate agents appears to be adapted to its environment. Although there still is some genetic change underway, major adaptation is almost nearly completed after less than 100 generations. We are aware that we have only

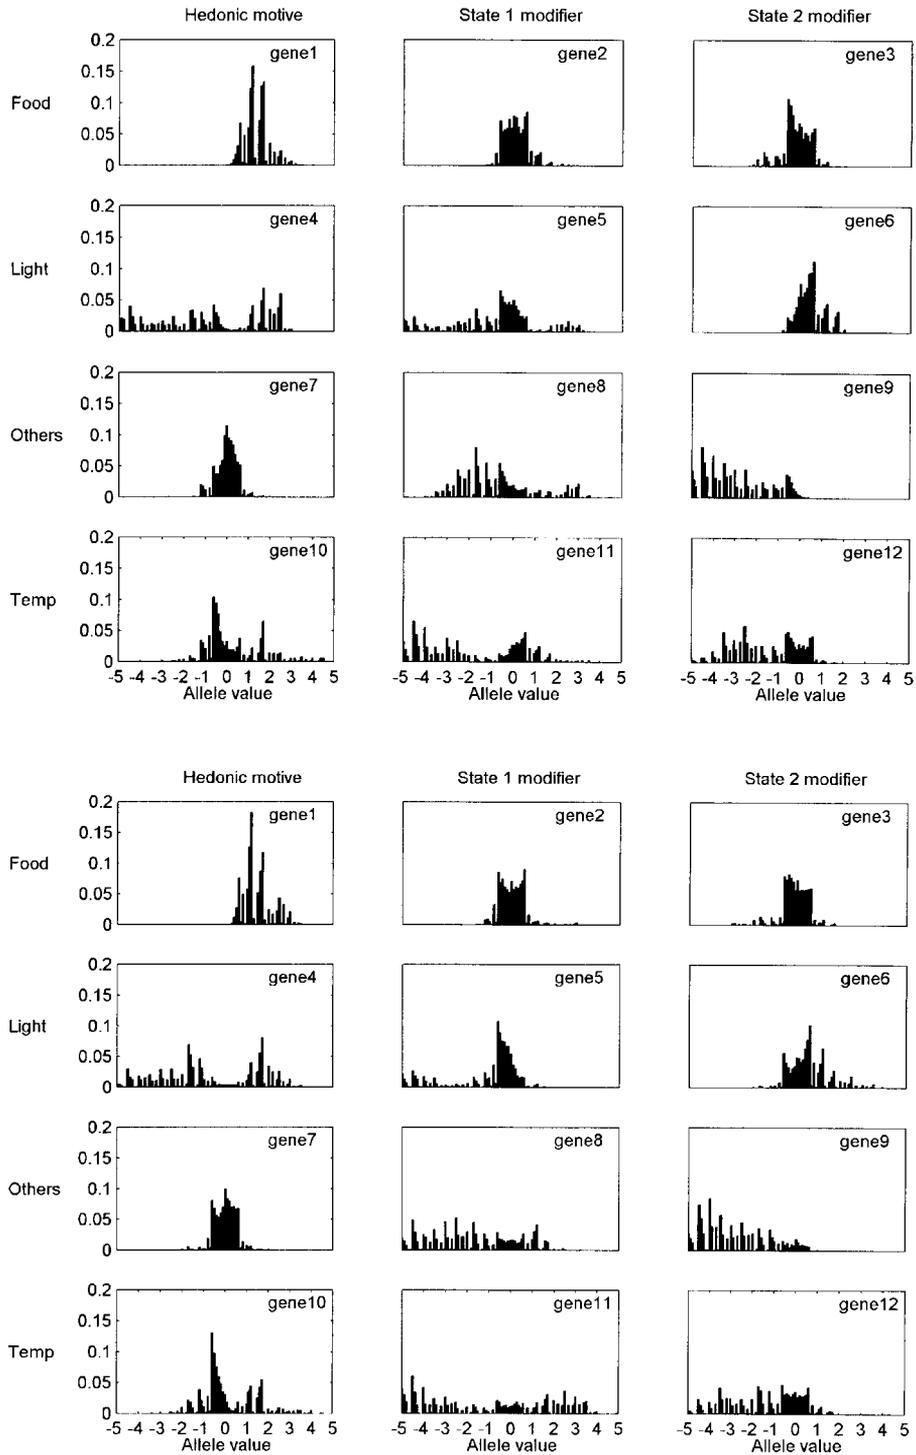


Fig. 9. Final allele frequency distribution in two replicate simulations over 10,000 generations.

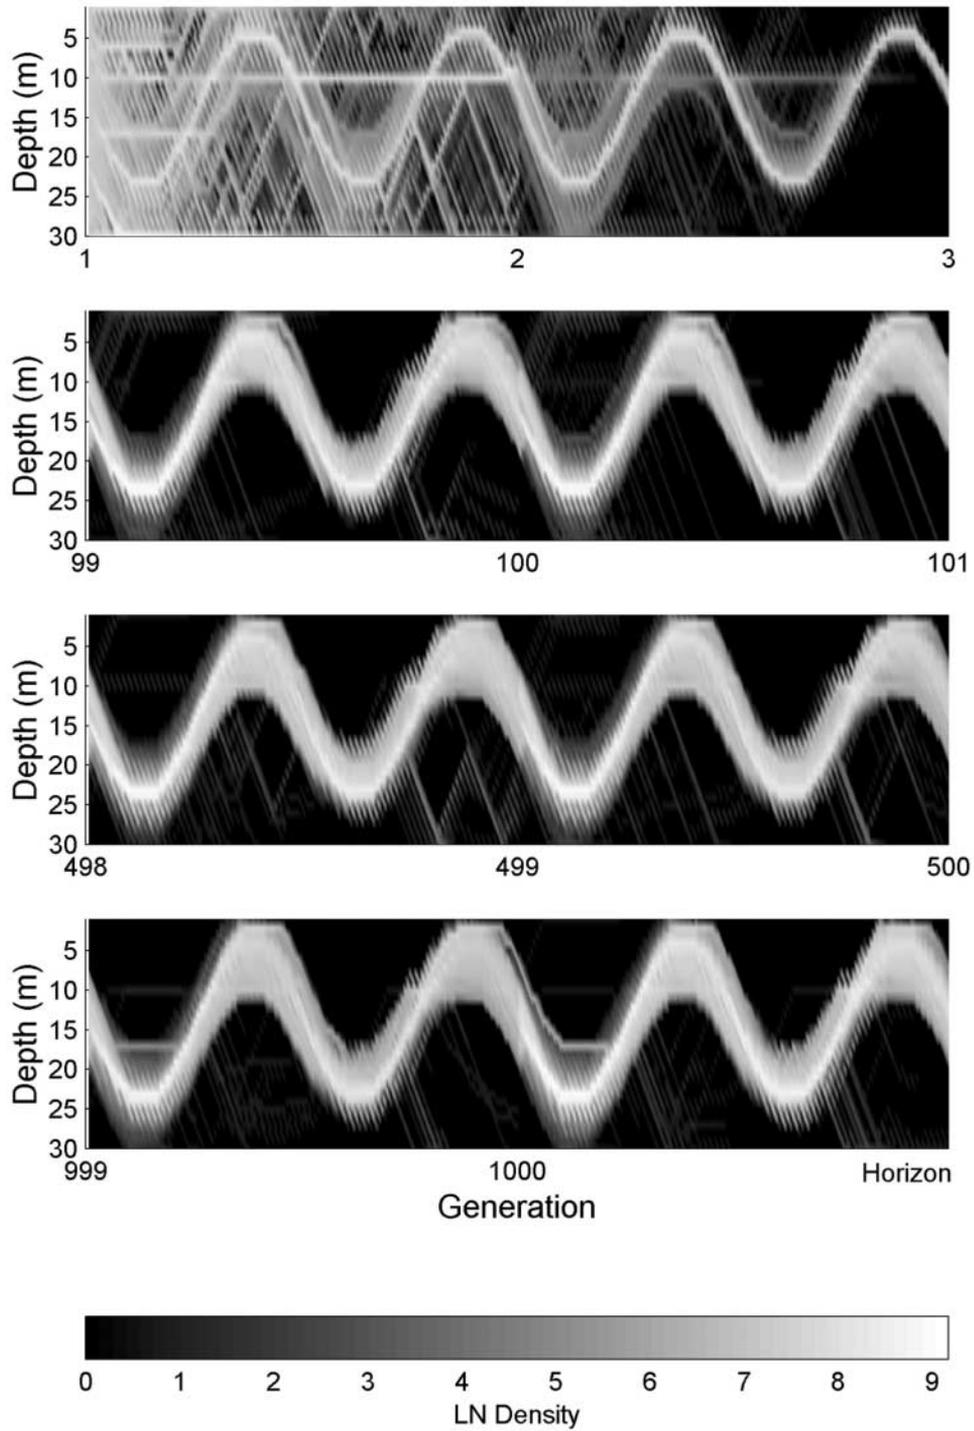


Fig. 10. Vertical distribution of the population during generations 1–2, 99–100, 498–499 and 999–1000 in the first simulation of the best equation set (Table 3).

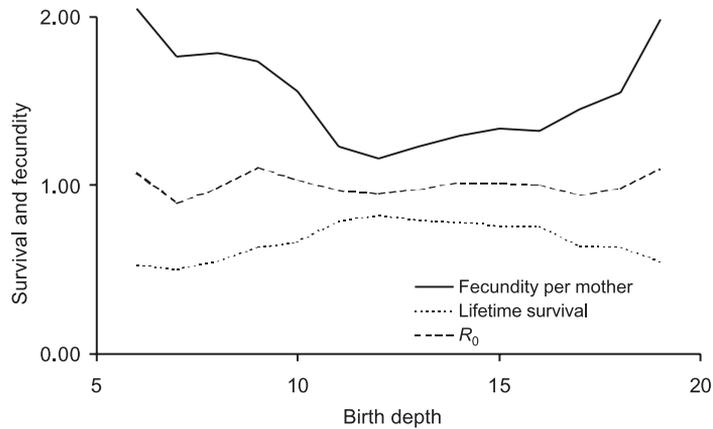


Fig. 11. Depth-dependent fitness. Survival from birth until end of last time step, fecundity of mothers, and net reproductive rate (R_0) of all individuals born at same depth. Data from all individuals born in the last five generations.

Table 5. Offspring production through fatherhood, based on all individuals born in the last five generations (only fatherly offspring produced as donor in recombination counted)

Birth depth	Individuals born	Probability of becoming a father (%)	Expected number of fatherly offspring at birth (n)	Offspring per father	
				Average	Maximum
7	321	4.7	0.14	3.0	13
8	1952	4.5	0.10	2.2	9
9	3375	4.5	0.10	2.3	12
10	5959	4.1	0.09	2.2	12
11	6019	4.3	0.09	2.2	17
12	8350	7.6	0.17	2.2	18
13	10438	4.2	0.11	2.5	18
14	4854	4.6	0.11	2.4	22
15	5593	4.6	0.12	2.5	16
16	5075	4.0	0.09	2.2	12
17	3609	4.6	0.11	2.3	17
18	1935	4.6	0.10	2.2	11
19	972	3.6	0.07	2.0	6

touched on the fringe of the phenotype and genotype spaces, but the consistency of the solutions implies that the hill-climbing ability of the genetic algorithm makes it a suitable tool for searches in such landscapes (Huse *et al.*, 1999). The current fitness landscape is not very rugged; it is bell-shaped. We have no means to find out whether the solutions at generation 100 or 10,000 are at, or close to, the ultimate peaks in the fitness landscape, but the genetic similarity between runs indicates that the peaks cannot be far away. However, since the speed of evolution decreases when fitness differences are small (Fisher, 1930), it might take a long time to arrive at significantly better solutions. Adaptation is not very good

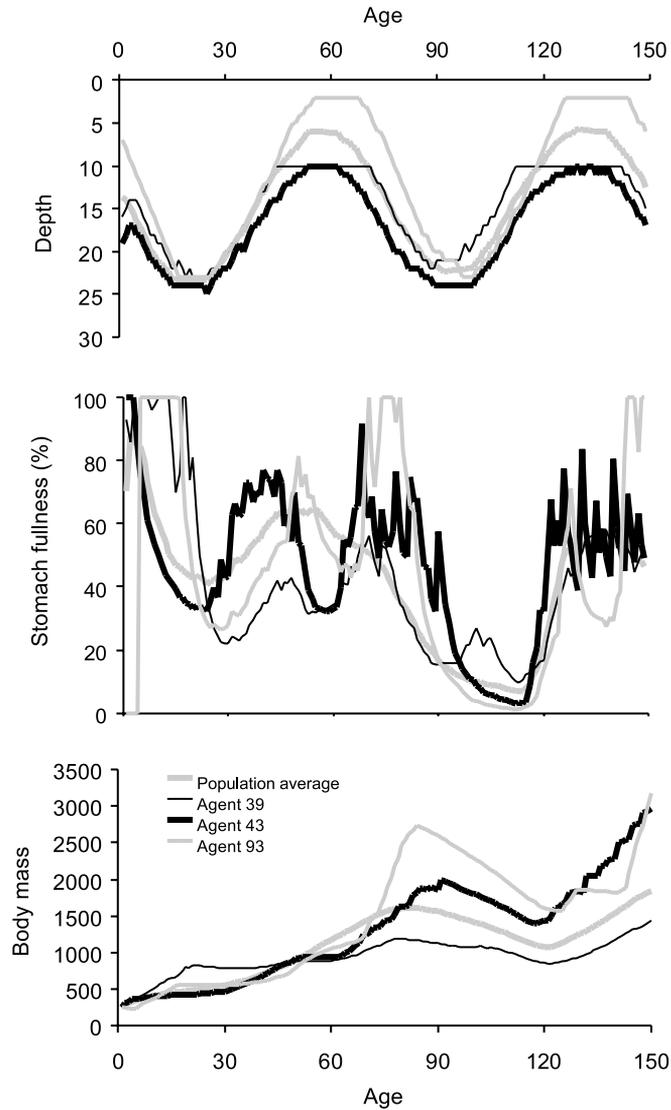


Fig. 12. Time trajectories of depth choice, stomach fullness and body mass for three selected agents. Population average is also given.

at walking the many steps of very small fitness increments, after the few very valuable steps are taken (Orr, 2002). The final range and frequency of alleles thus also depends on the mutation rate in the adapted gene pool.

Inter-generational variation in population fecundity is caused by a combination of genetic and ecological effects. A population above its carrying capacity (K) has lower fecundity and will decline in numbers, while a smaller population may grow bigger than K . Since growth, survival and fecundity are functions of the vertical behaviour of the population, and ultimately of the genetic composition of the individuals, density-dependent

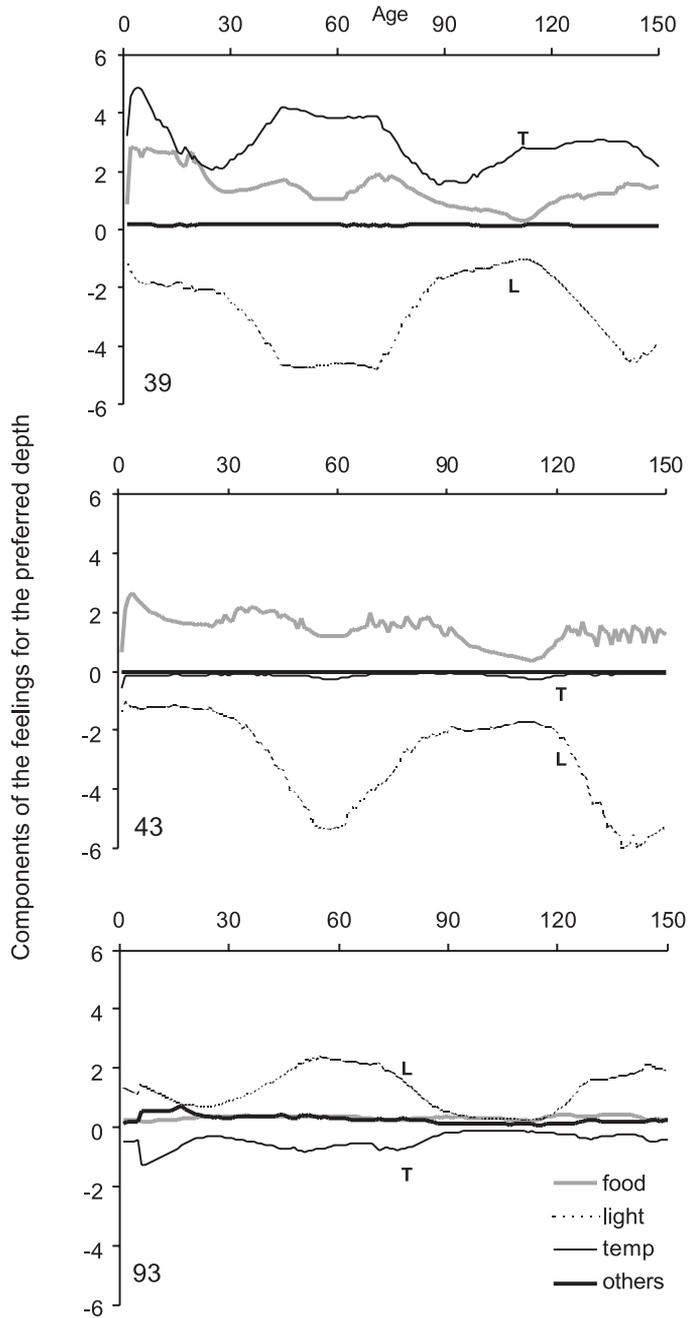


Fig. 13. The components of feelings for the preferred depth choice for three selected agents through life. Agent 39 was small, agent 43 grew large in deep water and agent 93 grew large in shallow water.

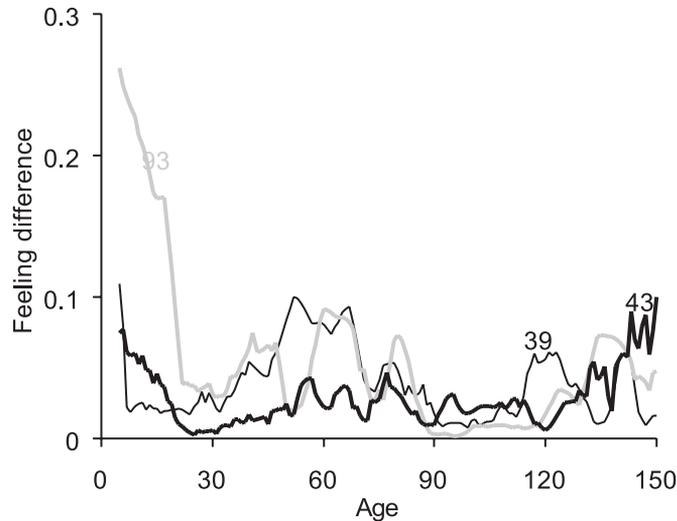


Fig. 14. Difference in strength of feelings for the best and next best of the three movement choices. Curves shown are moving averages for five ages, to increase readability. Agent 39 is small and is usually found in the centre of the population vertical distribution; agent 43 is large and prefers deeper water; and agent 93 is large and usually selects shallow water.

regulation becomes a function of the unique combination of phenotypes in each generation. A simple linear density-dependent response should not be expected, as the precise level of K depends on the present gene pool.

Interpretable individual variation

The approach of proximate adaptive agents gives interpretable rules for individual responses to sensory information (Figs 12–14) and the model output is open for experimental tests by real organisms. Such a dialogue between laboratory and simulation has been difficult with pure optimization tools, in part because fitness-derived rules of conduct should be changed as the result of new theory rather than new data (Mangel *et al.*, 2001). Experiments with real animals cannot test *whether* the organisms are adapted (since neither evolved animals nor evolved agents can be perfectly adapted to anything), but can test the assumptions in the model of how the proximate reaction to perception operates. Tests with real organisms can also discover deficiencies in model structure leading to compensations in emotions and response.

There is a lesson from our modelling exercise for experimental ecology. Substantial individual variation in behaviour, growth and survival is to be expected from variation in affect genes. In our model, this can be due to random mutations, near-neutral selection of competing alternatives, or the frequency-dependent fitness of alleles in density-dependent environments. The evolution of real populations may also yield competing alleles for things other than behaviour, for example for physiological efficiency or abiotic factors. Salvanes and Hart (2000) found genotypes coding for oxygen affinities in cod haemoglobin to be linked to individual competitive abilities.

Adaptive individuals utilize state, density and environment

The vertical distribution in the population is the result of all individual decisions, which are influenced by three independent factors: (1) the genome controls the unique feelings each individual obtains; (2) the historic path and present position decide which depth cells an individual may choose among; (3) the presence of other individuals (and hence their historic paths and, ultimately, their genomes) impact the feelings each agent has towards the three possible choices.

The path and choices of each individual are to varying extents influenced by individual state, other agents and external factors. The model (equation 2) allows for the use of state information during the decision process. Of 65,500 agents checked in the last few generations of the simulation in Fig. 8, no agents had evolved to utilize state components less than four times. Only two utilized state to modify feelings in four ways, while 1600 individuals utilized state information five or six times, 13,000 utilized such information seven times and more than 50,000 agents used all eight possibilities to accommodate their feelings by internal state. Overall, the agents followed the path of the prey population through the light cycle. Some agents avoided competition but most agents did not. Only 6% of the adapted agents did not let other agents influence their feelings and choices.

Instead of arriving at one optimal allele for each gene, the gene pool evolved to a range of alleles for all genes. The similarity in allele frequencies in different simulations with the same equation set shows that the final allele distribution is the result of adaptive evolution. Hence, these alleles are present because in previous generations they have given their genomes high fecundity. One could be tempted to call the final spatial distribution a density-dependent ideal free distribution. However, this will be misleading. The ideal free distribution concentrates on phenotypes, assumes perfect information and no travel costs, and does not include individual state variation. Our model of proximate adaptive agents deviates in all these respects. We have not shown that each individual could not benefit from a slightly different path (Fig. 10), and the very character of natural selection makes it increasingly harder to improve a well-performing genome (Fisher, 1930; Orr, 2002).

Although most organisms utilized internal state in many ways in their decisions, some simulations indicate that removing one of the two states did not impact fitness (Table 2). However, this does not mean that state is not important for behaviour, and removal of both states caused rapid extinctions of the populations. Our simulations thus shed light on the explanatory success of state-based models, such as stochastic dynamic programming (e.g. Mangel and Clark, 1988; Houston and McNamara, 1999; Clark and Mangel, 2000). In this respect, it is interesting to note that those equation sets that utilized body mass (i.e. an internal state in dynamic programming jargon) fared better than those using age (as in life-history theory). However, while optimization models based on ultimate considerations often consider state, there is little emphasis on proximate mechanisms for state perceptions in the ecological literature (Thorpe *et al.*, 1998). The general impact of physiological state on hormone levels and motivation has long been recognized (e.g. Woods *et al.*, 1998); a more detailed description of the physiology of our proximate adaptive agents could possibly lead to increased importance of state for regulating behaviour.

Comparison with real populations

The polymorphism in behavioural responses to density seen here resembles the polymorphism in the *foraging* locus at chromosome 2 in *Drosophila melanogaster* (Sokolowski

et al., 1997). Larvae with the *rover* allele move more than larvae with the *sitter* allele. Adult rovers also move farther from the food source than sitters. Sitters were selected for low density-situations allowing rapid population growth, while rovers were adapted to high-density situations with low resource availability. Hence, the two alleles and phenotypes were maintained by density-dependent selection.

The clonal evolution of water fleas in the Oud Heverlee Pond, Belgium (Cousyn *et al.*, 2001) is a particularly interesting case of behavioural genetics. At any time, the resident population of *Daphnia magna* consists of a series of clones, differing in their propensity for phototactic behaviour. De Meester (1993) crossed positively phototactic *D. magna* with negatively phototactic ones and obtained offspring with intermediate phototactic behaviour. The co-occurrence of genetically distinct clones in Oud Heverlee Pond resembles the situation in the density-dependent world we have simulated. Over the past 30 years, fish predation pressure on these cladocerans has changed systematically. Initially, the shallow man-made lake was stocked with benthivorous fish. Later, it was also stocked with a high number of planktivorous fish. This increased predation intensity selected for negatively phototactic clones (Cousyn *et al.*, 2001). A related case is the co-existence of two species of *Daphnia* in Lake Constance (Stich and Lampert, 1981). Here, *D. galeata* is found in the warm and food-rich upper parts of the water column, while *D. hyalina* is found in deeper, darker and colder waters below. Both species grow better and produce more offspring under experimental conditions of high food and high temperature (Stich and Lampert, 1984). One could then assume that the affects of *D. galeata* in the lake were less risk-sensitive and more focused on food and temperature than those of the more fearful *D. hyalina* below. Daphnids also show clonal differences in sensitivity to fish kairomones, which are proximate cues for mortality risk (Boersma *et al.*, 1998). Clones from fishless lakes were less responsive to fish kairomones, indicating local adaptation.

Small genetic changes may also have great effects on real fishes. Both brown trout and rainbow trout treated with growth hormone (i.e. a gene product from an 'affect gene') are more willing to risk exposure to a predator than unmanipulated fish (Johnsson *et al.*, 1996; Jönsson *et al.*, 1996). As a result, these fish grow faster than their unmanipulated conspecifics. The same lack of fear is seen in Atlantic salmon genetically modified to contain and transmit a growth hormone transgene (Abrahams and Sutterlin, 1999). Our model is not meant to represent the real genetics of an organism. While our proximate adaptive agents possess three artificial genes impacting their feeding motivation, more than 400 genes impact fat storage in the nematode *Caenorhabditis elegans* (Ashrafi *et al.*, 2003). Hopefully, not all of them are needed to model the animal's feeding decisions. But the studies on *C. elegans* indicate that each behavioural decision in multicellular organisms is probably under the influence of several genes. There do, though, exist a few cases where single-gene effects on behaviour are evident. In *C. elegans*, the two behaviourally distinct phenotypes *social* and *solitary* only differ in a single nucleotide in a single gene (de Bono and Bargmann, 1998; Thomas, 1998). Individuals of the social strain differ from others behaviourally in two respects: they are less likely to remain on a food spot and they are far more likely to remain close to others. This attraction to other *C. elegans* is only seen in the presence of food, implying that some unidentified odorant is related to food or to feeding. Compared with our model, it seems that food availability impacts the hedonic tone towards other *C. elegans* in the *social* phenotype. Thus *C. elegans* may obtain a single general arousal, rather than a series of discrete feelings. The final formulation of the general arousal (equation 3) obviously requires additional experimental investigation.

Level of complexity

The many runs of the different equations gave similar final population reproductive output, showing that many similar equation structures may evolve to yield the same fitness. This is often seen in other adaptation tools such as ING (Huse and Giske, 1998; Huse *et al.*, 1999), where repeated runs yield similar population behaviour, but based on quite different weightings in the artificial neural network (Strand *et al.*, 2002). In our model, the allele frequency distributions are very similar between repeated runs of the same equation set, indicating that the remaining variability is maintained by frequency-dependent selection. Thus, in a real population of proximate adaptive agents, there may be substantial individual differences in tolerances (e.g. to light) and preferences (e.g. between food attraction and light avoidance).

While the precise state representation in the model had little impact, the overall model structure was important (Table 3). Experiments with changes in model structure show that all four affects, represented by an overall genetic disposition and one or two states for each affect, yield the highest adaptation to the local conditions. It appears possible to utilize one rather than two states, but all other simplifications in model structure lead to lowered population fecundity.

The best performing equation set found here could probably be improved, since the affect equations are not derived from optimizing a fitness measure. The equations are pre-defined, but the parameter values are adapted. An alternative would be to use genetic programming (Koza, 1992) to evolve the equation structure. Both the choices of functional relationships and of state-modifications require more attention in further research on affect-based decisions. Since the benefit extracted from each decision is density dependent, we can assume that the greatest problem for the modelled fish was to anticipate the density in the cell they were moving into. They only knew the densities at the three cells in the current time step. A realistic improvement could be to divide each time step into smaller partitions, with a fraction of the population moving at a time. This will have a computational cost, but this may be offset by the shorter adaptation time needed. Another unrealistic assumption in the model is that the organisms do not plan ahead. Although that may not be necessary in the cyclic and one-dimensional world they live in, it could be needed in other models. Huse and Giske (1998) allowed for planning in an adaptive model by giving the fish two separate brain hemispheres. One was used for navigation in local gradients, the other for large-scale migrations (i.e. for reaching long-term goals). A separate gene was evolved to decide under which condition each hemisphere (artificial neural network) should be used. Johnston (2000) equipped his dog with an artificial neural network that could learn to recognize smells. Contrary to ING models, the artificial neural network in Johnston (2000) was reset at birth, and the weighting resulted from individual learning and experience. Such a system of basic affect genes and an individually learning neural network could allow both long-term adaptation of affect genes and short-term habituation to new perceptions.

Reasons for modelling by proximate adaptive agents

There are two key reasons for modelling by proximate adaptive agents. First, with the advent of individual-based modelling (DeAngelis *et al.*, 1979; Huston *et al.*, 1988), ecologists started to explore dynamics of populations at a far more complex level than those for which optimization tools are developed. With a few affect genes we can maintain much of the

transparency of a simple model and yet avoid the validity problems that *ad hoc* formulations encounter. For the purpose of improvements in individual-based modelling, a deep understanding of perception, affect or arousal is not necessary. Rather, the concept of proximate adaptive agents can be used as a calculation tool for complex choices. Second, organisms are proximate, and combined ultimate–proximate models allow for a new level of realism in ecological explanations. Everyone who has tried to tell friends that animals behave in order to maximize Darwinian fitness knows the scepticism this evokes. But try then to explain that organisms with such feelings that make them avoid danger while exploiting resources are more likely to survive and make offspring, and thereby pass the genetic basis for these feelings on to future generations. This is not merely a difference in explanation, it is a difference in concept. Adaptive proximate agents are not necessarily optimal. For problems at low complexity, where the global optimum can be found, adaptation models are able to identify optimal behavioural solutions (Huse *et al.*, 1999). For more complex cases, one simply does not know the optimal solution and thus how near perfection adaptive agents behave. However, optimization theory is only practical for low-dimensional artificial organisms in low-dimensional artificial worlds. Optimization tools should not be abandoned, since their simplicity yields a huge learning potential about the real world. Still, real organisms are adapted and proximate (Thorpe *et al.*, 1998), and the present and other adaptive modelling tools can show us how adapted organisms behave and the difficulty of approaching the optimal solution from the near-optimal (Orr, 2002).

CONCLUSION

An approach based on proximate adaptive agents differs from classical approaches to animal behaviour in a number of ways. First, the animals are recognized as feeling their choices. This brings us one step closer to how real animals make decisions, compared with previous proximate models. It thus allows for new communication between scientists interested in adaptive behaviour and animal psychology. In a model of a simple environment, there may be few or no differences between modelling the decisions based on feelings and those based on rational utility, since the genetic algorithm will have secured a tight link between feelings and fitness. However, if a proximate agent is transported or migrates to a different environment, its feelings may yield sub-optimal decisions from a fitness-maximization perspective (cf. Clark and Mangel, 2000). Second, modelling proximate adaptive agents will benefit from detailed information on sensing and brain structures in the species studied, thus binding ecology closer to neurology and molecular biology. In the present model, where we have studied a purely artificial agent, a genetic algorithm has been allowed to find the best possible set of feelings. However, in studies of real organisms, the role of the genetic algorithm will be reduced to find those aspects of the affective system that are not yet understood. Third, the behaviour seen is the result of selection over previous generations, and no information on the near or distant future needs to be offered, other than what can be guessed from the present conditions. This allows for modelling organisms with more complex trade-offs. The optimality demand in classical ultimate theories strongly restricts the level of complexity that can be studied (Parker and Maynard Smith, 1990). The organism modelled here can simultaneously consider the value of other agents (as in game theory), its own age (as in life-history theory), its state (as in stochastic dynamic programming) and the environment (as in non-evolutionary biogeochemical models and in community ecology). Decisions can also be made in highly variable or

non-repetitive environments (as often is the case in individual-based modelling; Bartsch *et al.*, 1989; Grimm, 1999), and individual learning can be included (Bourgine and Snyers, 1995; Johnston, 2000). However, although proximate adaptive agents are capable of performing more complex choices, there is no guarantee that any decision will be optimal. Proximate adaptive agents are selected for lifetime performance, while the optimization tools secure optimal decisions at every decision point, although only in simpler fitness landscapes. Affect genes may therefore yield more static trade-offs, since the genes do not change during life. With relatively few genes, one can have a population in which individuals make many suboptimal decisions during their lives. An alternative could then be the ING method (Huse *et al.*, 1999; Strand *et al.*, 2002), where the affect genes are replaced by genes coding for an artificial neural network. This could, in some cases, improve short-term decisions, but at the cost of loss of an intuitive interpretation of the trade-offs.

ACKNOWLEDGEMENTS

We thank Dan Depp, Alex Kacelnik, Victor S. Johnston, Suzanne Henson Alonzo, Angie Shelton, Nick Wolf, Øyvind Fiksen, Dag L. Aksnes, Thomas Torgersen, Sigrunn Eliassen and Jarle Berntsen for stimulating discussions. Most of this study was undertaken during J.G.'s sabbatical visit to M.M. in 2000–2001, funded by a grant from the Research Council of Norway. The work of M.M. was partly supported by the National Marine Fisheries Service.

REFERENCES

- Abrahams, M.V. and Sutterlin, A. 1999. The foraging and antipredator behaviour of growth-enhanced transgenic Atlantic salmon. *Anim. Behav.*, **58**: 933–942.
- Aksnes, D.L. and Giske, J. 1993. A theoretical model of aquatic visual feeding. *Ecol. Model.*, **67**: 233–250.
- Aksnes, D.L. and Utne, A.C.W. 1997. A revised model of visual range in fish. *Sarsia*, **82**: 137–147.
- Ashrafi, K., Chang, F.Y., Watts, J.L., Fraser, A.G., Kamath, R.S., Ahringer, J. and Ruvkun, G. 2003. Genome-wide RNAi analysis of *Caenorhabditis elegans* fat regulatory genes. *Nature*, **421**: 268–272.
- Atema, J., Fay, R.R., Popper, A.N. and Tavolga, W.N., eds. 1988. *Sensory Biology of Aquatic Animals*. New York: Springer-Verlag.
- Bart, J. 1995. Acceptance criteria for using individual-based models to make management decisions. *Ecol. Appl.*, **5**: 411–420.
- Bartsch, J., Brander, K., Heath, M., Munk, P., Richardson, K. and Svendsen, E. 1989. Modelling the advection of herring larvae in the North Sea. *Nature*, **340**: 632–636.
- Boersma, M., Spaak, P. and De Meester, L. 1998. Predator-mediated plasticity in morphology, life history, and behaviour of *Daphnia*: the uncoupling of responses. *Am. Nat.*, **152**: 237–248.
- Bourgine, P. and Snyers, D. 1995. The hedonic agent: a constructivist approach to abductive capacities. *Lect. Notes Artif. Intel.*, **957**: 199–216.
- Cabanac, M. 1992. Pleasure: the common currency. *J. Theor. Biol.*, **155**: 173–200.
- Cash, K.J., McKee, M.H. and Wrona, F.J. 1993. Short-term and long-term consequences of grouping and group foraging in the free-living flatworm *Dugesia tigrina*. *J. Anim. Ecol.*, **62**: 529–535.
- Chambers, C.R. 1993. Phenotypic variability in fish populations and its representation in Individual-Based Models. *Trans. Am. Fish. Soc.*, **122**: 404–414.
- Clark, C.W. and Mangel, M. 2000. *Dynamic State Variable Models in Ecology: Methods and Applications*. New York: Oxford University Press.

- Cousyn, C., De Meester, L., Colburne, J.K., Brendonck, L., Verschuren, D. and Volckaert, F. 2001. Rapid, local adaptation of zooplankton behaviour to changes in predation pressure in the absence of neutral genetic changes. *Proc. Natl. Acad. Sci. USA*, **98**: 6256–6260.
- Curio, E. 1976. *The Ethology of Predation*. Heidelberg: Springer-Verlag.
- DeAngelis, D.L., Cox, D.C. and Coutant, C.C. 1979. Cannibalism and size dispersal in young-of-the-year largemouth bass: experiments and model. *Ecol. Model.*, **8**: 133–148.
- DeAngelis, D.L. and Gross, L.J., eds. 1992. *Individual-Based Models and Approaches in Ecology*. New York: Chapman & Hall.
- de Bono, M. and Bargmann, C.I. 1998. Natural variation in a neuropeptide Y receptor homolog modifies social behaviour and food response in *C. elegans*. *Cell*, **94**: 679–689.
- De Meester, L. 1993. On the genetical ecology of phototactic behaviour in *Daphnia magna* (Crustacea: Cladocera). *Academia Analecta*, **55**: 145–188.
- Doya, K. 1999. What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex? *Neural Networks*, **12**: 961–974.
- Doya, K. 2000. Metalearning, neuromodulation, and emotion. In *Affective Minds* (G. Hatano, N. Okada and H. Tanabe, eds), pp. 101–104. Amsterdam: Elsevier.
- Dugatkin, L.A. and Reeve, H.K., eds. 1998. *Game Theory and Animal Behavior*. New York: Oxford University Press.
- Ekman, P. and Davidson, R.J. 1994. *The Nature of Emotion, Fundamental Questions*. New York: Oxford University Press.
- Ewert, J.P. 1980. *Neuroethology*. Heidelberg: Springer-Verlag.
- Fisher, R.A. 1930. *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.
- Gallistel, C.R. 1994. Foraging for brain stimulation: toward a neurobiology of computation. *Cognition*, **50**: 151–170.
- Giske, J., Aksnes, D.L. and Fiksen, Ø. 1994. Visual predators, environmental variables and zooplankton mortality risk. *Vie Milieu*, **44**: 1–9.
- Giske, J., Rosland, R., Berntsen, J. and Fiksen, Ø. 1997. Ideal free distribution of copepods under predation risk. *Ecol. Model.*, **95**: 45–59.
- Giske, J., Huse, G. and Fiksen, Ø. 1998. Modelling spatial dynamics of fish. *Rev. Fish Biol. Fish.*, **8**: 57–91.
- Goldberg, D.E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA: Addison-Wesley.
- Grimm, V. 1999. Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future? *Ecol. Model.*, **115**: 129–148.
- Holland, J.H. 1975. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.
- Hopfield, J.J. 1982. Neural networks and physical systems with emergent collective properties. *Proc. Natl. Acad. Sci. USA*, **79**: 2554–2558.
- Houston, A. and McNamara, J. 1999. *Models of Adaptive Behaviour: An Approach Based on State*. Cambridge: Cambridge University Press.
- Huse, G. 2001. Modelling habitat choice in fish using adapted random walk. *Sarsia*, **86**: 477–483.
- Huse, G. and Giske, J. 1998. Ecology in the Sea of Pentium: an individual-based spatio-temporal model for fish with adapted behaviour. *Fish. Res.*, **37**: 163–178.
- Huse, G., Strand, E. and Giske, J. 1999. Implementing behaviour in individual-based models using neural networks and genetic algorithms. *Evol. Ecol.*, **13**: 469–483.
- Huse, G., Giske, J. and Salvanes, A.G.V. 2002. Individual-based modelling. In *Handbook of Fish and Fisheries*, Vol. 2 (P.J.B. Hart and J.D. Reynolds, eds), pp. 228–248. Malden, MA: Blackwell Science.
- Huston, M., DeAngelis, D. and Post, W. 1988. New computer models unify ecological theory. *BioScience*, **38**: 682–691.

- Johnsson, J.I., Petersson, E., Björnsson, B.Th. and Järvi, T. 1996. Domestication and growth hormone alter antipredator behaviour and growth patterns in juvenile brown trout, *Salmo trutta*. *Can. J. Fish. Aquat. Sci.*, **53**: 1546–1554.
- Johnston, V.S. 2000. *Why We Feel: The Science of Human Emotions*. Reading, MA: Helix Books.
- Jönsson, E., Johnsson, J.I. and Björnsson, B.Th. 1996. Growth hormone increases predation exposure of rainbow trout. *Phil. Trans. R. Soc. Lond. B*, **341**: 375–397.
- Kohonen, T. 1984. *Self Organisation and Associative Memory*. Heidelberg: Springer-Verlag.
- Koza, J.R. 1992. *Genetic Programming: On the Programming of Computers by means of Natural Selection*. Cambridge, MA: MIT Press.
- LeDoux, J.E. 1994. Cognitive–emotional interactions in the brain. In *The Nature of Emotion: Fundamental Questions* (P. Ekman and R.J. Davidson, eds), pp. 216–223. Oxford: Oxford University Press.
- Leven, S.J. and Levine, D.S. 1996. Multiattribute decision making in context: a dynamic neural network methodology. *Cognitive Sci.*, **20**: 271–299.
- MacLean, P.D. 1990. *The Triune Brain in Evolution: Role in Paleocerebral Functions*. New York: Plenum Press.
- Magurran, A.E., Oulton, W. and Pitcher, T.J. 1985. Vigilance behaviour and shoal size in minnows. *Z. Tierpsychol.*, **67**: 167–178.
- Mangel, M. and Clark, C.W. 1988. *Dynamic Modeling in Behavioral Ecology*. Princeton, NJ: Princeton University Press.
- Mangel, M., Fiksen, Ø. and Giske, J. 2001. Theoretical and statistical models in natural resource management and research. In *Modeling in Natural Resource Management* (T.M. Shenk and A.B. Franklin, eds), pp. 57–72. Washington, DC: Island Press.
- McFarland, D.J. 1977. Decision making in animals. *Nature*, **269**: 15–21.
- McFarland, D.J. 1985. *Animal Behaviour*. Oxford: Pitman.
- McFarland, D.J. and Houston, A.I. 1981. *Quantitative Ethology: The State Space Approach*. Oxford: Pitman.
- McFarland, D.J. and Sibly, R.M. 1975. The behavioural common final path. *Phil. Trans. R. Soc. Lond. B*, **270**: 265–293.
- McNamara, J.M. and Houston, A.I. 1986. The common currency for behavioural decisions. *Am. Nat.*, **127**: 358–378.
- Mitchell, M. and Taylor, C.E. 1999. Evolutionary computation: an overview. *Annu. Rev. Ecol. Syst.*, **30**: 593–616.
- Orr, H.A. 2002. The population genetics of adaptation: the adaptation of DNA sequences. *Evolution*, **56**: 1317–1330.
- Parker, G.A. and Maynard Smith, J. 1990. Optimality theory in evolutionary biology. *Nature*, **348**: 27–33.
- Railsback, S.F. 2001. Concepts from complex adaptive systems as a framework for individual-based modelling. *Ecol. Model.*, **139**: 47–62.
- Railsback, S.F., Lamberson, R.H., Harvey, B.C. and Duffy, W.E. 1999. Movement rules for individual-based models of stream fish. *Ecol. Model.*, **123**: 73–89.
- Roff, D.A. 2002. *Life History Evolution*. Sunderland, MA: Sinauer Associates.
- Salvanes, A.G.V. and Hart, P.J.B. 2000. Is individual variation in competitive performance of reared juvenile cod influenced by haemoglobin genotype? *Sarsia*, **85**: 265–274.
- Scherer, K.R. 1984. On the nature and function of emotion: a component process approach. In *Approaches to Emotion* (K.R. Scherer and P. Ekman, eds), pp. 293–317. Hillsdale, NJ: Lawrence Erlbaum.
- Sibly, R. and McFarland, D. 1976. On the fitness of behaviour sequences. *Am. Nat.*, **110**: 601–617.
- Sokolowski, M.B., Pereira, H.S. and Hughes, K. 1997. Evolution of foraging behaviour in *Drosophila* by density-dependent selection. *Proc. Natl. Acad. Sci. USA*, **94**: 7373–7377.

- Stich, H.-B. and Lampert, W. 1981. Predator-evasion as an explanation of diurnal vertical migration by zooplankton. *Nature*, **293**: 396–398.
- Stich, H.-B. and Lampert, W. 1984. Growth and reproduction of migrating and non-migrating *Daphnia* species under simulated food and temperature conditions of diurnal vertical migration. *Oecologia*, **61**: 191–196.
- Strand, E., Huse, G. and Giske, J. 2002. Artificial evolution of life history and behaviour. *Am. Nat.*, **159**: 624–644.
- Thomas, J.H. 1998. Social life and the single nucleotide: foraging behaviour in *C. elegans*. *Cell*, **94**: 549–550.
- Thorpe, J.E., Mangel, M., Metcalfe, N.B. and Huntingford, F.A. 1998. Modelling the proximate basis of salmonid life-history variation, with application to Atlantic salmon, *Salmo salar* L. *Evol. Ecol.*, **12**: 581–599.
- Tyler, J.A. and Rose, K.A. 1994. Individual variability and spatial heterogeneity in fish population models. *Rev. Fish Biol. Fish.*, **4**: 91–123
- van Rooij, A.J.F., Jain, L.C. and Johnson, R.P. 1996. *Neural Network Training Using Genetic Algorithms*. Series in Machine Perception and Artificial Intelligence, Vol. 26. Singapore: World Scientific.
- Wisenden, B.D. and Millard, M.C. 2001. Aquatic flatworms use chemical cues from injured conspecifics to assess predation risk and to associate risk with novel cues. *Anim. Behav.*, **62**: 761–766.
- Woods, S.C., Seeley, R.J., Porte, D. and Schwartz, M.W. 1998. Signals that regulate food intake and energy homeostasis. *Science*, **280**: 1378–1383.

APPENDIX

Environment

Surface light intensity follows a diel cycle, with maximum daylight of $L_{\max} = 500 \mu\text{E} \cdot \text{m}^{-2} \cdot \text{s}^{-1}$. Light intensity just beneath the surface is modelled by assuming a 50% loss through scattering at the surface

$$s(t) = L_{\max} 0.5 \sin(3.14 D t \Omega^{-1})$$

where Ω is the potential life span of the organism (in time steps), D is the number of diel cycles during a life span and t is the current time step. In the present model, $D = 2$ cycles and $\Omega = 150$ time steps. Underwater light is attenuated following Baer's law,

$$E_b(z, t) = s(t) \exp(-kz)$$

where $E_b(z, t)$ is background irradiance at depth z at time t and k is the diffuse light attenuation coefficient.

Physiology

The feeding rate of the proximate agent depends on visual range $r(z, t)$, prey density $C(z, t)$ and temperature. The visual range of a predator at (z, t) is

$$r(z, t) = \sqrt{e^{-(cr+kz)} |C_0| A E \frac{E_b}{K_c + E_b}}$$

where c and k are the beam and diffuse light attenuation coefficients, C_0 and A are prey contrast and area, E and K_c are the visual capacity and the eye saturation parameter of the fish, and E_b is the background irradiance (Aksnes and Giske, 1993; Aksnes and Utne, 1997). The same equation is used to find r for the agents and for their predators. Hence, variation in body size of the agents is not included in this part of their mortality risk. This is done separately, below. The temperature-dependence is a Q_{10} function, doubling the potential feeding rate for every 10 degrees.

$$f_T = f_0 \exp(T(z)q_f)$$

Feeding rate is a function of the prey encounter rate and the number of conspecifics in the same cell:

$$f_c(z, t) = \pi r(z, t)^2 C(z, t) f_T N(z, t)^{-\beta}$$

Feeding rate is also constrained by its current body size and its recent feeding history, so that the stomach contents cannot surpass 25% of its body mass. Stomach contents $s(t)$ are evacuated by an exponential and temperature-dependent function d_T :

$$d_T = d_0 \exp(T(z)q_d)$$

$$s(t) = s(t-1) \exp(-d_T)$$

Thus, the feeding gain of an organism is either a function of the prey encounter rate reduced by competition or the difference between 25% of body mass and the stomach contents, whichever is smallest.

The organism also has metabolic costs, depending on temperature and body mass B :

$$c_T = c_0 \exp(T(z)q_d)$$

$$C(z, t) = B c_T$$

The weight gain or loss of the organism is the difference between its feeding gain and its living costs, depending on prey concentration, light intensity, temperature, number of competitors, feeding history and body size.

Mortality

Assuming that the mortality agent of the modelled organism also is a visual predator, the mortality risk can be found from light intensity, body size and risk dilution due to other agents (Giske *et al.*, 1994, 1997). To avoid cancellation of the importance of light or other agents in feeding and mortality risk, we formulated the latter slightly differently:

$$M(z, t) = m r B^{1.5} N(z, t)^{-a}$$

Here, m is the overall mortality risk (per time step) of the environment. The rationale for $a < \beta$ is that a fish may be better than its prey in detecting and escaping an approaching predator. According to the surface area to body mass ratios, the visible surface will fall off by a factor of 0.67. When we have used a much higher factor, we have also assumed that the predator would prefer to eat the largest of the agents it sees. Individual body size is, for computational reasons, not a variable in the calculation of visual range.

