



Research paper

Implementing behaviour in individual-based models using neural networks and genetic algorithms

GEIR HUSE*, ESPEN STRAND and JARL GISKE

Department of Fisheries and Marine Biology, University of Bergen, Norway

*(*author for correspondence, e-mail: geir.huse@ifm.uib.no)*

Received 13 July 1999; accepted 3 May 2000

Co-ordinating editor: O. Leimar

Abstract. Even though individual-based models (IBMs) have become very popular in ecology during the last decade, there have been few attempts to implement behavioural aspects in IBMs. This is partly due to lack of appropriate techniques. Behavioural and life history aspects can be implemented in IBMs through adaptive models based on genetic algorithms and neural networks (individual-based-neural network-genetic algorithm, ING). To investigate the precision of the adaptation process, we present three cases where solutions can be found by optimisation. These cases include a state-dependent patch selection problem, a simple game between predators and prey, and a more complex vertical migration scenario for a planktivorous fish. In all cases, the optimal solution is calculated and compared with the solution achieved using ING. The results show that the ING method finds optimal or close to optimal solutions for the problems presented. In addition it has a wider range of potential application areas than conventional techniques in behavioural modelling. Especially the method is well suited for complex problems where other methods fail to provide answers.

Key words: adaptation, artificial neural networks, behaviour, genetic algorithms, habitat choice, individual-based model, state dependence, stochastic dynamic programming

Introduction

The dynamics of natural populations result from interaction between individual organisms. Most traditional models in ecology, however, define populations in terms of their abundance and ignore these local interactions. This has been recognised for some decades now (Łmonicki, 1988; Huston *et al.*, 1988), and has led to an increasing popularity of individual-based models (IBMs) in ecology. This approach allows a population to be specified in terms of its individuals, and the recent increase in computing power has made the approach feasible despite the considerable computational cost. Since IBMs trace individuals it is easy to compare model predictions with individual observations in experiments and field studies. Such studies have illustrated that individual variability in many cases can be important to population dynamics

(e.g. Crowder *et al.*, 1992). However, most IBMs simply follow the fates of individuals in a population without taking into account behavioural decisions. This limits IBMs to studies where will directed behaviour is unimportant (eggs, larvae, plants) or to such simple situations that motivation can be enforced by the modeller. The traditional models for describing behavioural trade-offs generally deal with relatively narrow aspects of the organism's life. Several such model frameworks prevail, including life-history theory (e.g. Roff, 1992), optimal foraging theory (MacArthur and Pianka, 1966), the ideal free distribution (Fretwell and Lucas, 1970) and evolutionary stable strategies (Maynard Smith, 1982). These models are largely analytical, which ensures an excellent understanding of their underlying logic and generality (Houston and McNamara, 1999). However, this approach suffers from a lack of flexibility, and soon gets intractable when more aspects of the organism's life are added (Taylor and Jefferson, 1995). Behavioural decisions in IBMs should therefore be implemented in a different way than by using these analytical tools.

Here we show how this can be achieved using an adaptive approach based on the genetic algorithm (GA, Holland, 1975) and artificial neural networks (ANNs, Rummelhart *et al.*, 1986). These methods are common in artificial life (Langton, 1989) and related fields, but have rarely been applied in ecological studies (Goldberg, 1989; Toquenaga and Wade, 1996). There are, however, some exceptions, and we have earlier applied these techniques in a complex scenario of fish migration, which includes both individual and environmental variability in a large-scale spatial setting (Huse and Giske, 1998). Despite the complexity of this scenario, the behavioural model provided credible solutions. However, in order to investigate the precision of the adaptive model, we investigate three simple cases where the optimal solutions can be found using conventional models. These examples include a state-dependent patch selection problem, a simple game between predators and prey, and a more complex vertical migration scenario for a planktivorous fish.

The individual-based-neural network-genetic algorithm (ING) concept

The genetic algorithm (GA) is a technique that applies the principle of evolution by natural selection to search for optimal solutions to a problem (Holland, 1975). Solutions are represented as arrays of numbers, analogous to genes (here characters) on chromosomes (here strings) in biology. The best solution is the combination of numbers that gives the highest score in a test of the particular problem. The algorithm works by starting with random solutions in generation 1. Thereafter the best individuals are selected for reproduction in each generation. As shown in Figure 1, variability in the strings is produced

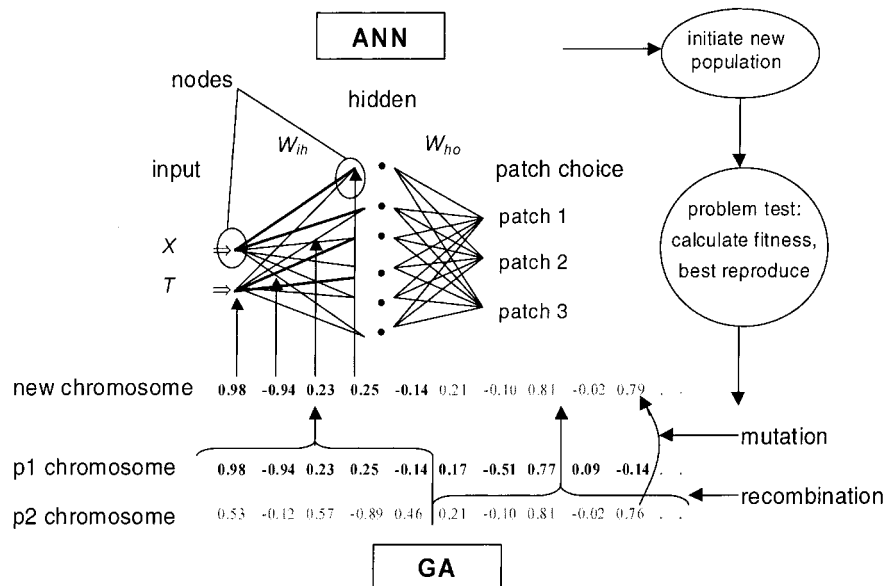


Figure 1. The ING modelling concept with ANN based behaviour and GA based evolution in an individual-based setting. Individuals with randomly set 'genetic' strings that code for the weighting of the ANN, are initiated in the first time-step. The problem cycle (19 time steps in the patch model) is simulated. The 10% of individuals with the highest fitness (p1) are then selected to reproduce with a partner (p2) randomly selected among 50% of the best individuals. New individuals are produced through recombinations and mutations. Differences between the 'genetic' strings are expressed through differences in weights of the ANN, which can lead to behavioural differences. Patch choice is calculated based on the state (X) and time-step (T) as input. W_{ih} and W_{ho} are the weight matrices of connection strengths between the layers. Behaviour is determined by choosing the patch with the highest output value.

through recombination and mutation. This procedure is repeated over many generations.

Artificial neural networks (ANNs) is a method that applies neurobiological principles of synaptic brain-activity to model behavioural output by differential weighting of input variables (Rummelhart *et al.*, 1986; Montana and Davis, 1989). The way they function is that weights and biases are trained through an adaptive process, by repeated calculations of output from input data and weights. During this process the weights and biases are altered so that the output of the ANN accommodates a fitness criteria. Depending on the problem, there exist several methods for training an ANN, but here the GA does this. After training, the ANN can be used in problem-solving. Since ANNs are inspired by brain function they are well suited to model behaviour. The networks applied here are feed forward networks with an input layer, one hidden layer and an output layer. Each of the layers consist of a number of

nodes, values that can be input data (information presented to the network), sum of weights and input data (at the hidden layer) or output (behavioural decisions, Figure 1). The input data (I_i) are multiplied by a weight matrix that specifies the strength of connection between nodes in the input layer and nodes in the hidden layer. At each hidden node the sum of the input data multiplied by the connection weights (W_{ij}) plus a value referred to as bias (B_j) are added together (Equation (1)). This sum is then transformed (F_j) by a sigmoid activation function. Transformation makes the ANN non-linear, which is a key aspect of its performance, and standardises the output values between 0 and 1:

$$F_j = \sum_{i=1}^n \frac{1}{1 + e^{-(W_{ij}I_i + B_j)}} \quad (1)$$

where n is the number of input nodes. While the weights (W_{ij}) are multiplied by the input data, the bias is not (Equation (1)) and provides the threshold value out of the hidden node. The value coming out of the hidden node (F_j) is then weighted by the connection strength between the hidden and output nodes. At each output node the values are again added and behaviour is determined. This is the general procedure of the ANNs applied here. ANNs have many similarities to statistical methods and can be considered as a special case of multiple regression.

By extending IBMs with the ANN and the GA we get an ING. Here individuals, in addition to being characterised by state variables, are equipped with a 'genetic string' that specifies the inherited basis for behavioural traits. This method removes much of the dichotomy between proximate stimuli-response models and ultimate survival-value derived models in ecology (Giske *et al.*, 1998). The ultimate GA modifies the proximate ANN and enables individuals to respond in an adaptive but not necessarily strict optimal manner to complex, unforeseen events, more like real organisms do. The more familiar the situation is and the closer the behaviour is related to fitness, the stronger will the selection pressure be in the GA to evolve a truly optimal response in the ANN.

The models

The ING method has been used to model spatial distribution and life history strategies of the Barents Sea capelin (Huse and Giske, 1998). However, there does not exist data with a spatial or temporal resolution that allows more than a qualitative validation of this ING model. The modelled situation is also so complex that other approaches (such as SDP) cannot be used without strong simplifying assumptions (Fiksen *et al.*, 1995). To test the ability of the ING, to

arrive at a true optimal solution, we therefore now investigate simpler ecological models.

The patch model

The patch selection problem, originally presented by Mangel and Clark (1988), is simply at each time step to choose the patch that maximises the probability of survival to the end of a simulation period. There are three patches, each with a different value for predation risk (β), probability of finding food (λ), and food reward if food is found (Y). There is also patch independent respiration cost α (Table 1). Patch choice is state-dependent, and in general satiated animals are expected to choose safe patches while hungry animals choose riskier patches, if such give a higher energy gain. Energetic state X and time T were used as state variables in the original model and are used as input data to the ANN here. Mortality is inflicted either through predation (β) or starvation if the energetic state falls below the threshold at $X = 4$. Since the current results are compared with those from a stochastic dynamic programming (SDP) model, emphasis is put on making our model as similar to the SDP setting as possible. First we used the SDP model of Mangel and Clark (1988) to find optimal habitat choices and maximal survival. Then we ran an ING-adaptation in the same environment. At the start of the simulation the individual has state 4. Patch choice is then determined for this state using the ANN. The proportion of individuals dying is then subtracted. Similarly the proportion of individuals finding food and increasing their state is determined as well as the proportion that do not find food. If a state gets higher than 10 it is set equal to 10. At the end of the 19th time-step simulation the proportion of the individual strategies alive in state 4–10 is calculated and used as a fitness criterion in the GA.

After a minor test period an ANN architecture with two input nodes, ten hidden nodes, and three output nodes was found to be profitable for the current problem. Individuals are born with a string of 63 characters including 20 input weights, 10 hidden biases, 30 output weights and 3 output biases. The

Table 1. Parameter values used in the patch model where β is predation risk, λ is probability of finding food, Y is food reward when food is found, and α is energetic cost (From Mangel and Clark (1988))

Patch number	β	λ	Y	α
1	0.000	0.0	0	1
2	0.004	0.4	3	1
3	0.020	0.6	5	1

two state variables, energy level and time are applied as input variables in the ANN (Figure 1), and the layers are fully connected. The three output nodes function in a “winner takes all” mode where the patch corresponding to the output node with the highest output value is chosen as each time-step. The weights of the ANN are initiated randomly with values between -1 and 1 . A population of 2000 individuals is tested in each generation of the GA, and the 10% best individuals are used as parents (p1, Figure 1) for the next generation with 10 offspring each. Partners (p2) are selected randomly among the 50% best individuals, and a single recombination of strings between these parents at a random site occurs with probability ($p = 0.5$). In case of no recombination, new individuals are a clone of p1. Mutations are carried out by choosing a node in the hidden and output layers randomly (with $p = 0.1$) and mutate all weights and biases that affect the output of this node (Montana, 1991). Mutations add random values $[-1,1]$ to the weight and bias values. The model was run over 1000 generations with four replicates altogether.

Predator–prey interaction

The second case is a habitat selection game between a zooplankton prey and its fish predator, and was originally introduced by Iwasa (1982). Iwasa imagined a water column divided into two habitats, a shallow and a deep, that differ in phytoplankton density and light intensity. The model is solved mathematically by assuming that the individuals will distribute among the habitats so that no individual is able to increase its fitness by moving to the other habitat. Iwasa referred to this as an evolutionary stable strategy (ESS, Maynard Smith, 1982) although Gabriel and Thomas (1988) have shown that not all of the requirements for an ESS are fulfilled. We constructed an ING model where the individual fish and zooplankton choose habitat based on information about the current densities of predators and prey in each habitat (Figure 2). The




Figure 2. The ANN architecture used to solve Iwasa’s game. The predator and prey are provided with similar information. Habitats 1 and 2 are the shallow and deep ones respectively.

individuals in the two populations choose habitat sequentially, alternating between zooplankton and fish. When one individual has picked a habitat, the densities are updated. The ANN is constructed with four input nodes, fifteen hidden nodes, and one output node (Figure 2). Habitat 1 is chosen if the output value is less than 0.5. The 50% most fit individuals are selected as parents, while a partner is selected randomly from the population. Since the relative abundance of predators and prey will affect solution, the model was run for different predator:prey ratios. The model was run over 700 generations with five replica simulations for each predator:prey ratio. Mutations and recombinations are implemented as in the previous case.

Vertical migration

The final case is a vertical migration scenario for the mesopelagic fish *Maurollicus muelleri* in a fjord of Western Norway. Diel vertical migration (DVM) is a classical theme in aquatic ecology, and although many different hypotheses have been proposed regarding its ultimate function, DVM is now generally considered to be a trade off between feeding and predator avoidance (Clark and Levy, 1988; Fiksen, 1998). Mesopelagic fishes perform DVM, and we present a model environment similar to that of Rosland and Giske (1994). Whereas the two previous cases only consider behaviour for its survival value (ultimate), behaviour is here determined proximately as a response to stimuli. Due to its great diel variation, light intensity is the most important physical parameter in triggering vertical migration (Aksnes and Giske, 1993). We used models developed by Aksnes and Giske (1993) and Aksnes and Utne (1997) to calculate visual range of the fish. Visual range depends upon surface light intensity, which again is dependent upon sun declination, latitude, day of the year, time of day, cloud cover, and attenuation in the water column. Rosland and Giske (1994) provide a detailed account on model implementation and parameter values of the visual models.

The model is run over 5 days with a 5 min time-step. The vertical distribution of zooplankton biomass is taken from a fjord in Western Norway (Figure 3). The predation risk of *M. muelleri* is assumed to be dependent upon the predator's visual range (Rosland and Giske, 1994). For simplicity, the visual range of the predator is set proportional to that of the modelled fish. In addition we introduced a predator density variable, which varies between days to increase the complexity of the environment. Feeding rate is calculated as the product of visual range and zooplankton density.

Similar to predator density, the zooplankton density was set to vary randomly between the five days, with a vertical profile proportional to one in Figure 3. It is assumed that the fish can determine the density of zooplankton and predators (e.g. by olfaction).

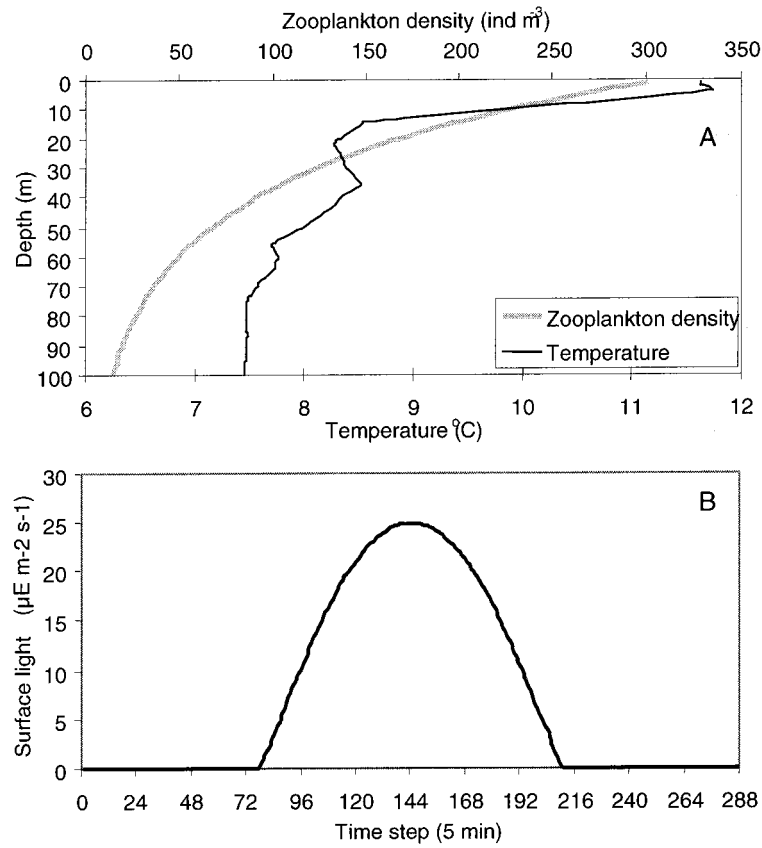


Figure 3. Vertical gradients in zooplankton density and temperature (A), and variation in surface light over the day (B). The data are for October.

Werner and Gilliam (1984) showed that juvenile individuals should choose habitats that minimised the risk of predation (μ_t) divided by the growth rate. If one assumes that growth rate is proportional to feeding rate (f_t), Gilliam's rule can be expressed on the inverse form as:

$$H_t = \frac{f_t}{\mu_t} \quad (2)$$

where H_t is habitat profitability at time-step t . Optimal habitat choice is calculated for each time-step. After 5 days of simulation the habitat profitability sums give the terminal fitness (φ):

$$\varphi = \sum_{\text{Day}=1}^5 \left(\sum_{t=1}^{288} H_t \right) \quad (3)$$




Figure 4. The ANN architecture used to solve the vertical migration scenario. The input variables used are visual range, zooplankton density, density of predators, food consumption last time-step, and temperature.

An ING model with five input nodes, thirty hidden nodes, and one output node was constructed to solve this problem, with input data shown in Figure 4. The 5% most fit individuals (Equation 3) were selected as parents (20 offspring per parent), while a random partner was selected among the 50% best. The model was run over 300 generations, and mutations and recombinations were implemented as in the previous cases.

Results

The patch model

The ANN is able to solve the patch selection problem with a high precision (Table 2). The proportion of correct patch choices is determined by comparing with the optimal choices found by SDP, and the maximally possible survival probability is found by using this optimal strategy. Although the ING model does not find the optimal patch in about 3% of the choices (Table 2), one should keep in mind that the fitness criterion in the GA is survival rather than the number of correct patch choices. The canonical costs (McNamara and

Table 2. Average patch choice, survival, and similarity with the patch choices (\pm SE) in the ING patch model with the original SDP solution of Mangel and Clark (1988). The survival in the SDP column is the maximal survival by using the optimal strategy when individuals are initiated at state 4. SE is standard error of mean ($n = 4$)

	SDP	ING
Average patch value	2.28	2.30 \pm 0.02
Survival	0.51	0.51 \pm 0.00
Patch choice similarity (%)	100.0	96.8 \pm 1.95

Houston 1986) of some of the sub-optimal patch choices are marginal. It is therefore more relevant to compare the survival data (Table 2) in which case the score of the GA is >99% of the survival rate achieved using the optimal strategy.

Predator-prey interaction

Although Iwasa (1982) presented solutions both for the day and night phases, we concentrate on the day phase, which is the most interesting phase in terms of vertical positioning. The ING model has the same qualitative predictions for all F/Z ratios (Figure 5), although it does not always find the exact numerical solution to the game situation. For $F/Z > 0.12$, Iwasa's model predicts a constant distribution of zooplankton among habitats 1 and 2, unaffected by the distribution of predators. These results are reproduced in the ING run as well. In some cases the ING model predicts a more even distribution of the fish among the habitats than the analytical solution. This is seen for F/Z ratios of 1 and 2.5 (Figure 5). As a result of the simulation method, the solution tends to oscillate around the equilibrium. When $1.00 > F/Z > 0.12$ the zooplankton will exert a greater selection pressure on the fish than vice versa. Conversely,

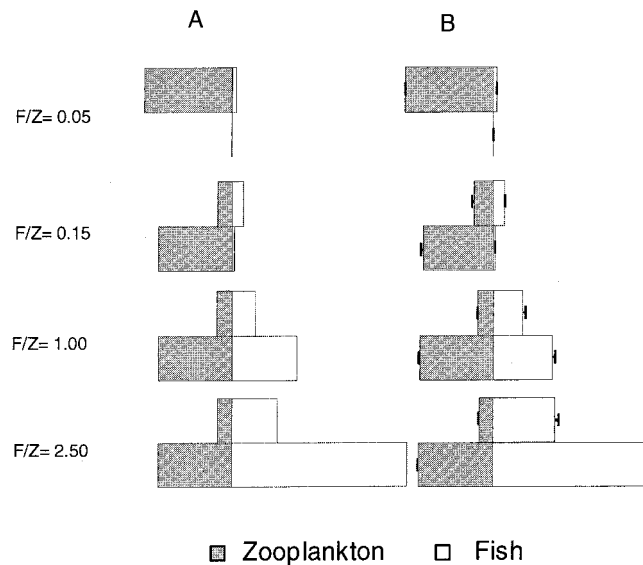


Figure 5. The results of the habitat selection game between zooplankton and fish with Iwasa's (1982) original solutions (A) and the ING model results (B). Each figure shows zooplankton and fish distribution in habitat 1 (upper) and habitat 2 (lower). The results for four different predator/prey (F/Z) ratios are shown. In each replica the average density of the last 100 generations are calculated. The error bars are standard error of the five replicates.

when $F > Z$ the fish will exert the greatest selection pressure. This leads to an oscillation around the ESS, where the population with fewer individuals will more easily be forced into their ESS, as indicated by the error bars.

Vertical migration

The vertical distribution seen in Figure 6 is a typical example of how *M. muelleri* is distributed during fall, with distinct vertical movement at dawn and dusk. Although it is not an exact match, the evolved behaviour of the ING model resembles the optimal trajectory to a high degree (Figure 6). Also the ING model responds close to optimally to the variability in zooplankton and predator densities between days. The behaviour of the best individuals in each of the replicate runs is shown (Figure 6), although there is only small differences in behaviour between individuals in each computer run. As for the previous model there are some oscillations between depths during some periods of the day.

Discussion

We have shown how genetic algorithms and neural networks can be used to implement behaviour in IBMs. Despite the different nature of the problems

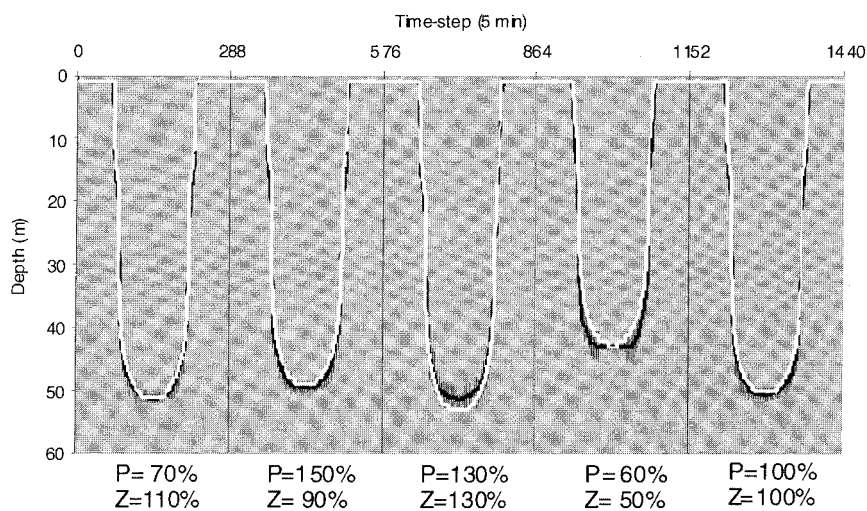


Figure 6. The results of the vertical migration scenario. The predator density parameter (P) and zooplankton biomass parameter (Z) values vary between days. The white line is the global optimum solution and the black line (\pm SE) is the adapted behaviour of the ING model.

presented, the ING model was able to solve them with a high precision. The advantages of the individual-based approach can thus be ported to behavioural ecology using the ING concept. Many basic and applied problems in ecology are complex, and there is a need for tools that can solve such problems. In theory, the adaptive approach used here can be applied to any problem regardless of its complexity, although the probability that the ING method finds good solutions is likely to be lower in highly complex environments. However, the good results achieved for the simple cases shown here are encouraging regarding its performance in more complex environments as well.

Proximate and ultimate models

As seen in the two first cases, ANNs can rely on derived information about habitat profitability when calculating behaviour in the same way that an SDP model operates. In addition, ANNs can integrate proximate responses with the ultimate functioning as seen in the last case. This allows behaviour to be modelled as a response to stimuli (Ghirlanda and Enquist, 1998), but at the same time also for its fitness value, thus providing a more realistic approach to individual behaviour. The proximate stimulus-response approach was the foundation for the classic behavioural study of ethology (Tinbergen, 1951). The optimisation approach in ecology can be criticised for neglecting this aspect of behaviour. For example in SDP models of vertical migration (e.g. Rosland and Giske, 1994), light is used to generate both expected feeding rate and the anticipated predation risk from visual predators. This gain and risk are presented to the fish as 'information' while the light intensity as such is not an input variable. In the ING model, however, individuals position themselves according to light and other stimuli. This makes the modelling process using ANNs analogous to the way that animals position themselves, which gives the modeller more flexibility with regard to the kind of problems that can be addressed. For example it can provide more information about constraints acting on the behavioural process. Also, trained ANNs can provide adapted responses towards non-adapted stimuli. This can be fruitful if one is interested in providing predictions about behaviour in response to human influences. During acoustic abundance estimation of pelagic fish for example, avoidance reactions, due to the noise produced by the vessel are common. In this case an ANN trained for a natural situation could be used to predict the avoidance behaviour of the fish, which again may be used to correct the acoustic estimate.

The 'genetic' adaptation that takes place during training of the ANN, is the process of modifying the weights so that profitable responses to environmental input are achieved. An alternative way to adapt the weights is to let individuals

learn during their lifetime. This can be achieved by using a different technique such as reinforcement learning (Ballard, 1997) instead of or in combination with the GA (Ackley and Littman, 1992). Based on the problem at hand the need to apply learning or genetic adaptation or both should be determined. In some cases learning during the life time will be the more important adaptive process and in such cases a learning routine rather than a GA could be used to adapt the weights of the ANN.

Potential pitfalls of adaptive models

A problem with the ING method is that it may get stuck on local maxima in solution space and never find the optimum. The probability of getting stuck is especially pronounced in cases where fitness gradients are weak and/or the fitness landscape is rugged. This problem is inherent in all processes governed by adaptation. A problem with ANN based models is oscillatory behaviour, as seen in both the predator–prey and vertical migration cases. This is caused by the continuous update of input data as a result of behavioural actions. Oscillations may be a real property of animal positioning in cases where the perceptive ability is poor, or where habitat choices are discrete. The oscillations are probably an inherent property of ANNs, but an optimal network and short time-steps would have fine-tuned the behaviour and reduced the degree of oscillation to the minimum. Increasing the probability that the optimal solution will be found therefore, reduces both the problems of getting stuck on local maxima and the oscillation problem. This can be achieved by using the shifting balance theory of Wright (1977) where instead of having a single population, a metapopulation with several sub populations is applied. An example of how this technique can be implemented in GAs is provided by Sumida *et al.* (1990).

Although ANNs perform well, it can be difficult to understand why particular behaviour is produced. ANNs are therefore sometimes referred to as black box models. However, by analysing the different weights and how they relate to the input data, it is possible to improve the understanding of ANN behaviour (see Aoki and Kamatsu, 1997). Such analyses are useful in providing more information about how the input data affects the output. Another way of analysing trained ANNs is to perform sensitivity analyses where the behaviour of the model over a range of different input values is studied.

Despite these challenges we think that the ING concept is an important extension of individual-based models, that can make IBMs applicable to ecological problems where behavioural aspects are important. Especially, the method is suited for complex problems where most other methods fail to provide answers.

Acknowledgements

We thank Dag L. Aksnes, Øyvind Fiksen, Paul Hart, Alasdair Houston, David Kirby, Marc Mangel and Esa Ranta for constructive comments on former drafts of this manuscript. GH was supported by The Research Council of Norway through the Marine Resource Management Programme and the Programme for Supercomputing through a grant of computing time.

References

- Ackley, D. and Littman, M. (1992) Interactions between learning and evolution. In C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds) *Artificial Life II*, pp. 487–509. Addison-Wesley, Reading, MA.
- Aksnes, D.L. and Giske, J. (1993) A theoretical model of aquatic visual feeding. *Ecological Modelling* **67**, 233–250.
- Aksnes, D.L. and Utne, A.C.W. (1997) A revised model of visual range in fish. *Sarsia* **82**, 137–147.
- Aoki, I. and Komatsu, T. (1997) Analysis and prediction of the fluctuation of sardine abundance using a neural network. *Oceanol. Acta* **20**, 81–88.
- Ballard, D. (1997) *Natural Computation*. MIT Press, Cambridge.
- Clark, C.W. and Levy, D.A. (1988) Diel vertical migrations by sockeye salmon and the anti-predation window. *Am. Nat.* **131**, 271–290.
- Crowder, L.B., Rice, J.A., Miller, T.J. and Marschall, E.A. (1992) Empirical and theoretical approaches to size-based interactions and recruitment variability in fishes. In D.L. De Angelis and L.J. Gross (eds.) *Individual-based Models and Approaches in Ecology*, pp. 237–255. Chapman & Hall, NY.
- Fiksen, Ø. (1998) Allocation patterns and diel vertical migration: modeling the optimal Daphnia. *Ecology* **78**, 1446–1456.
- Fiksen, Ø., Giske, J. and Slagstad, D. (1995) A spatially explicit fitness-based model of capelin migrations in the Barents Sea. *Fish. Oceanogr.* **4**, 193–208.
- Fretwell, S.D. and Lucas, H.J. Jr. (1970) On territorial behavior and other factors influencing habitat distributions in birds. 1. Theoretical development. *Acta Biother.* **19**, 16–36.
- Gabriel, W. and Thomas, B. (1988) Vertical migration of zooplankton as an evolutionary stable strategy. *Am. Nat.* **132**, 199–216.
- Ghirlanda, S. and Enquist, M. (1998) Artificial neural networks as models of stimulus control. *Anim. Behav.* **56**, 1383–1389.
- 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*. Addison-Wesley, Reading, MA.
- Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor.
- Huse, G. and Giske, J. (1998) Ecology in Mare Pentium: an individual based spatio-temporal model for fish with adapted behaviour. *Fish. Res.* **37**, 163–178.
- Houston, A.I. and McNamara, J.M. (1999) *Models of Adaptive Behaviour*. Cambridge University Press, Cambridge.
- Huston, M., DeAngelis, D.L. and Post, W. (1998) New computer models unify ecological theory. *BioScience* **38**, 682–691.
- Iwasa, Y. (1982) Vertical migration of zooplankton: a game between predators and prey. *Am. Nat.* **120**, 171–180.

- Langton, C.G. (1989) *Artificial Life*. Santa Fe Institute Studies in the Sciences of Complexity, Vol. 6. Addison-Wesley, Reading, MA.
- Łomnicki, A. (1988) *Population Ecology of Individuals*. Princeton University Press, Princeton, NJ.
- MacArthur, R.H. and Pianka, E.R. (1966) On optimal use of a patchy environment. *Am. Nat.* **100**, 603–609.
- Mangel, M. and Clark, C.W. (1988) *Dynamic Modeling in Behavioral Ecology*. Princeton University Press, Princeton, NJ.
- Maynard Smith, J. (1982) *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- McNamara, J.M. and Houston, A.I. (1986) The common currency for behavioral decisions. *Am. Nat.* **127**, 358–378.
- Montana, D.J. (1991) Automated parameter tuning for interpretation of synthetic images. In L. David (ed) *Handbook of Genetic Algorithms*, pp. 282–311. Van Nostrand-Reinhold, NY.
- Montana, D.J. and Davis, L.D. (1989) Training feedforward neural networks using genetic algorithms. In N.S. Sridharan (ed.) *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 762–767. Morgan Kaufmann, San Francisco.
- Roff, D.A. (1992) *The Evolution of Life Histories*. Chapman & Hall, NY.
- Rosland, R. and Giske, J. (1994) A dynamic optimization model of the diel vertical distribution of a pelagic planktivorous fish. *Prog. Oceanog.* **34**, 1–43.
- Rummelhart, D.E., Hinton, G.E. and Williams, R.J. (1986) Learning representations by back propagating errors. *Nature* **323**, 533–536.
- Sumida, B.H., Houston, A.I., McNamara, J.M. and Hamilton, W.D. (1990) Genetic algorithms and evolution. *J. Theor. Biol.* **147**, 59–84.
- Taylor, C. and Jefferson, D. (1995) Artificial life as a tool for biological inquiry. In C. Langton (ed.) *Artificial Life an Overview*, pp. 1–13. MIT Press, Cambridge.
- Tinbergen, N. (1951) *The Study of Instinct*. Oxford University Press, Oxford.
- Toquenaga, Y. and Wade, M.J. (1996) Sewall Wright meets artificial life: the origin and maintenance of evolutionary novelty. *Tr. Ecol. Evol.* **11**, 478–482.
- Werner, E.E. and Gilliam, J.F. (1984) The ontogenetic niche and species interactions in size-structured populations. *Ann. Rev. Ecol. Syst.* **15**, 393–425.
- Wright, S. (1977) *Evolution and Genetics of Evolution*. Vol. 3. *Experimental Results and Evolutionary Deductions*. University of Chicago Press, IL.