Motivation Networks – A Biological Model for Autonomous Agent Control

Content Areas: cognitive modelling, multiagent systems, simulation

Abstract

In this paper, we introduce a biologically inspired model for reasoning in autonomous agents that we call motivation networks. The main idea of this approach is to imitate fundamental mechanisms of animal behaviour. When animals interact with their environment, they constantly have to make a choice between different behavioural patterns that they could perform (such as feeding or resting). Their decisions are based on temporary priorities regarding different behavioural drives, which are determined by functional relationships with inputs from sensors and internal states. The resulting feedback interaction between decision-making, internal states, and stimuli creates behaviour in respect to time and space, which is adapted to the demands of an ecological niche. For this mechanism, we designed a formal model, which we used as an agent controller for an adaptive multiagent system. In the present investigation, we tested the use of motivation networks in a biological context by a simulation of group feeding behaviour in house sparrows. The results indicated a great potential of this approach for controlling autonomous agents and verification of biological hypotheses.

1 Introduction

Effective decision-making and problem solving are complex mechanisms that are very difficult to accomplish with machines when in a real world context. In traditional AI, one has to face problems such as uncertainty and incompleteness of the symbolic knowledge representation, which violates the closed-world assumption [Rich & Knight 1991]. In 'new' AI, this problem is avoided by a subsymbolic, adaptive representation based on biologically motivated models such as artificial neural networks (ANN) or evolutionary algorithms (EA). Indeed, ANNs and EAs turned out to be of great use for a variety of problems such as pattern classification or low-level motor-sensory control of robots. However, the use of neural network models for high-level planning and temporal reasoning is rather controversial [Rich & Knight 1991], though several recent advances are promising [Cox Hayslip *et al.* 1990], [Chappell & Taylor 1993], [Briscoe & Caelli 1997]. Interestingly, in nature, even plants and bacteria are able to behave in a non-trivial way without a central nervous system, and to evolve sophisticated mechanisms for processes like reproduction, feeding, respiration and migration. In this investigation, we have tried to improve machine reasoning by simulating mechanisms of animal decision-making.

The behaviour of animals is based on motivation and decision-making in respect to their genetic background and individual life history. Constantly shifting priorities lead the animal to 'decide' how to act in a particular situation. In fact, it has to weigh up all sorts of internal (e.g. metabolism) and external (e.g. predator avoidance, food availability) stimuli, through a decision-making process and give priority to one behaviour over another. The internal state of the animal, which is the net result of stimuli arising both inside and outside its body, constitutes its 'motivation' [Manning & Dawkins 1998].

Thus, animal behaviour is not solely determined by neural networks, but made of feedback interactions between network, organism and environment.

Typically, the motivational state decides about the selection and execution of entire action sequence patterns as for instance nest building, digging, feeding, and others. For instance, the internal stimulus of a decreasing metabolic budget over time is mapped to an increase of hunger - the motivation for feeding (figure 1a). However, foraging behaviour will not be performed until its assigned motivation dominates all other needs. Successful food intake lowers the hunger stimulus, thus lowering the motivation for further foraging until the agent stops feeding (dashed lines in figure 1a). Hence, this feedback mechanism regulates the energy budget of the animal. However, other needs can interfere with this process when they get urgent, such as daily resting routines (figure 1b) or short-term escape behaviour when an animal gets threatened by a predator (figure 1c).

In this context, time is a very important factor, because (i) behaviour is composed by sequences of subactions [Manning & Dawkins 1998], (ii) internal states and environmental conditions change, and (iii) various

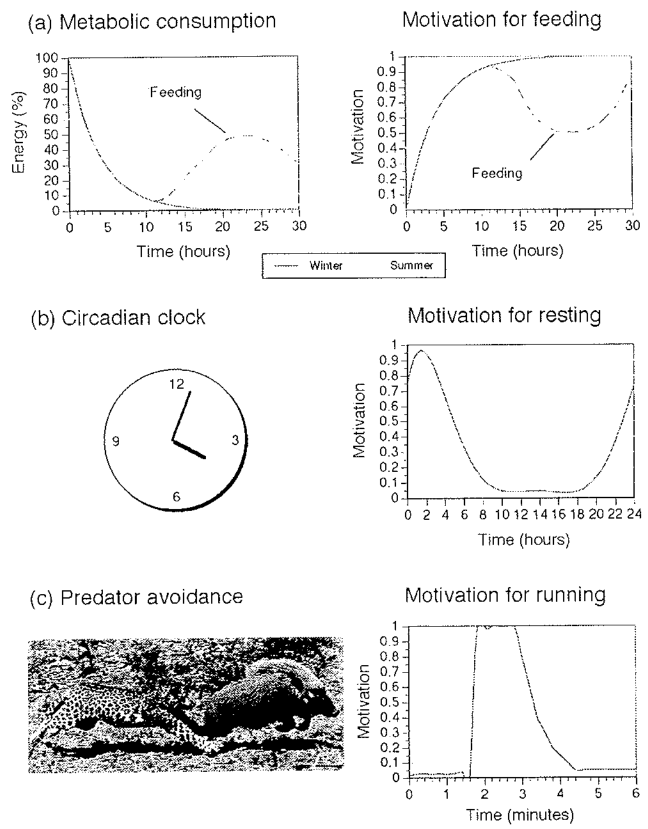

Figure 1: Decision-making by shifting-priorities
This figure shows internal and external variables which simultaneously act on motivation levels for specific sequences of action. (a) Regulation of the metabolic consumption: Mapping from the energy budget to the motivation for feeding. The dashed line indicates the effect of feeding; the light grey line shows how these curves vary due to seasonal changes; (b) periodic triggering of resting periods by circadian clocks; (c) immediate predator avoidance, which causes a shift in priorities within seconds.

proesses require synchronisation [Campbell *et al.* 1997], [Hewitt & Butlin 1997].

Many environmental factors such as temperature and light conditions change in a periodic manner and often serve as external clocks, which trigger behaviour in combination with internal periodic processes [Campbell *et al.* 1997]. Moreover, these clocks act on different scales such as lifetime (maturation), seasons (mating), hours (resting, feeding), minutes (foraging), seconds (motor control).

The idea for motivation networks is based on these considerations regarding high-level decision-making and motivation in animals.

2 Concepts of motivation networks

A motivation network is a functional unit that controls the behaviour of an autonomous agent based on the dynamic input from the agent's sensors and internal states. Its task is to take high-level decisions by selection and activation of behaviour patterns, which consist of entire action sequences.

In a biological context, the repertoire of optional behaviour patterns would include feeding, breeding, resting, etc. Each pattern is assigned to a continuous motivation variable that represents the agent's will for its execution.

The entire motivation network (figure 2) consists of (i) a set of mapping functions, (ii) motivation variables, (iii) a decision-maker, and (iv) behaviour patterns. The reasoning process is subdivided into three steps:

Firstly, the network receives information from the input variables, i.e. stimuli from *sensors* and *internal states*, which is mapped to *motivation variables*. The value of each variable ranges from 0.0 (no motivation) to 1.0 (maximum motivation) and is determined by one or more *mapping functions*, which are specified by *artificial genes* of the agent (see below). Multiple function values are combined by one of four standard operators min (acting as AND), max (acting as OR), multiplication and mean.

Secondly, the *decision-maker* determines and schedules a *behaviour pattern* according to the motivation variable with the presently highest value. Conflicts in cases of equally high priorities are resolved by a *conflict priority list*, which defines a hierarchy of the behaviour patterns. For instance, collision avoidance could be ranked higher than collecting objects in a robotics scenario.

Thirdly, the behaviour pattern is executed by activation of *physical exchange* (e.g. feeding) and *motors* (e.g. forward movement), which affects the agent's environment and its internal states.

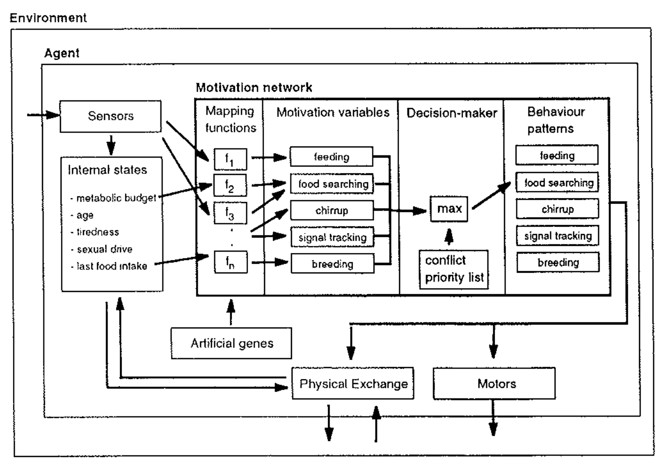

Figure 2: Conceptual model of a motivation network
Arrows indicate the information flow between the different functional units.

The repeated feedback between actions and stimuli occurring on different time scales and in varying contexts creates emergent agent behaviour.

It is important to note, that the assignment, shape and combination of the mapping functions between input and motivation variables crucially determine the behaviour of the agent. Currently, we are using a selection of standard function types (linear, logarithmic, cubic, etc.), which are fine-tuned by a set of parameters defining their slope and translation.

These parameters are encoded as artificial genes and transferred between generations by (currently asexual) reproduction and mutation, whereas the remaining mapping specifications are predefined and considered as a fixed part of the model. This set-up allows a rough design of the agent's control mechanisms, which is fine-tuned by adaptation. In addition, biological hypotheses can be falsified regarding their evolutionary plausibility by testing if agents would evolve certain motivational demands.

However, additional flexibility could be accomplished by encoding of the entire mapping specification using genetic programming, which we currently implement. Alternatively, mapping between inputs and motivation variables could be learned by ANNs. However, in this case functional relationships could not be directly implemented if required and the model would lose its transparency.

We have tested the potential of the introduced motivation network model by a simulation of the foraging behaviour of house sparrows as a biological example of decision-making in multiagent societies.

3 Simulations

3.1 Background and model specifications

An example which greatly illustrates motivation and decision-making are house sparrows (*Passer domesticus*) 'deciding' whether or not to approach a patch of food. When a sparrow discovers food, it will often not approach immediately but instead sit on a nearby perch close to cover and recruit other sparrows by a characteristic call to come and join it on the perch [Elgar 1986a]. All birds fly down to the food together, which has the advantage that they cover each other.

We modelled this biological scenario with our adaptive multiagent system SocialLab. The scenario was designed as a biotope containing autonomous agents (sparrows and predators) and passive resources (renewable food items and breeding space).

Space was represented as a 2-d grid world with a local carrier capacity, limiting the maximum number of individuals per grid cell. The goal of this design was to model direct and local interactions explicitly, and to control the effects of local spatial density on the number of concurrent local events.

Concurrency was simulated in a two step process: first, agents made decisions and scheduled actions; second, the simulation shell executed the scheduled actions and resolved temporal conflicts. For instance, when two agents decided to feed at the same location, their feeding actions were first scheduled, but not immediately executed. Depending on the local food availability and the demand of the agents, food resources were split and shares were provided for both agents instead of fulfilling their demands sequentially.

Sparrows were model as organisms with various properties (e.g. metabolic budget, minimal metabolic consumption, maximum food intake per time unit, limited lifespan, etc.), physical capabilities (movement, feeding, reproduction, attraction of other sparrows by chirruping), and simplified sensors for detection of other individuals, food items, and attraction signals. Each sparrow was controlled by a motivation network and competed for food and breeding space in order to propagate its genetic information to future generations. A sparrow was capable to reproduce when its metabolic budget exceeded a specified threshold. Its metabolic budget was increased by successful food intake and decreased due to regular metabolic maintenance costs and extra energy loss due to movement. According to their current motivation priorities, sparrows selected and executed one out of five different behaviour patterns at each time step: feeding, searching, chirruping (attracting), signal tracking, and breeding.

Input factors were given by their sensory stimuli: number of detected predators, distance to nearest predator, number of sparrows in the local neighbourhood, estimated food quality; and internal states: metabolic budget and last food intake. In the current model, most functional relationships between inputs and motivation variables were prespecified according to the biological context.

Predators tried to catch sparrows only when in range and feeding. By this constraint, we aimed to model a flightless predator and sparrows that safely sit on trees unless they fly to the ground to reach for food items.

Moreover, the population of predators neither evolved nor changed population size by birth or death. Successful catching of prey items was followed by short feeding and resting periods. Food was presented in one large food patch of renewable food resources.

3.2 Simulation results

In model 1, sparrows sent attracting signals if they were holding patches with high food availability and predators were visible, and followed attracting signals if signals were recent and their local food availability was low. Predators tried to attack sparrows if in reach.

Surprisingly, signalling and grouping lead to rapid extinction of the sparrow population. This was due to two reasons: firstly, predators did not pause after catching sparrows, but immediately continued to hunt more prey; second, attracted sparrows became signal senders when in reach of the predator while the original sender already got killed. These two effects worked like maelstroms feeding the predators with sparrows.

Conclusively, the predator must be tightly constrained in his catching abilities. In fact, safe group feeding apparently requires that (i) the mortality due to predator attacks gets significantly reduced, (ii) predators are relatively low in number compared to the number of prey items, and (iii) predators pause after successful hunting. Moreover, sparrows should only attract mates if in safe distance from the predator. These conditions are fulfilled in many similar prey-predator relationships in biology such as in hunting lions and wolves.

The motivation network turned out to be an appropriate tool for the control of an agent in this behavioural scenario and generated interesting emergent behaviour sequences.

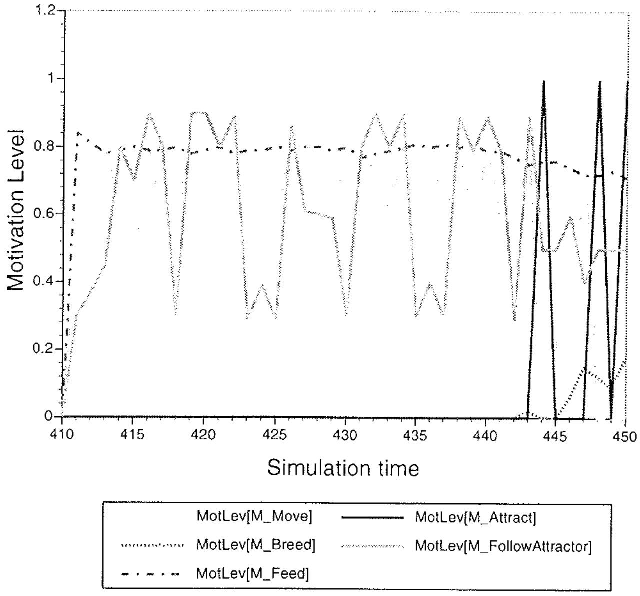

Figure 3 shows the shift of priorities over time in a motivation network of a single sparrow.

Figure 3: Shifting priorities of an agent's motivation variables

Actions were taken by the agent according to the motivation variable with the temporarily highest value. Note the sudden and clear spikes in the motivation level for sending attraction calls at simulation time t=444 and t=448 after the agent repeatedly took actions to follow an attracting signal (indicated by the two solid line).

In model 2, we changed sparrow and predator behaviour due to the conclusions from the previous model. This model resulted in flocks of straying and "grazing" sparrows where single individuals were keeping the group together by attraction signals (figure 4). These sparrow populations remained relatively stable in size. The group size of these flocks turned out to be spatially constrained by the density of predators.

However, sparrow populations always evolved genes that suppressed the sending and receiving of attraction signals. Non-signalling sparrow populations turned out to be more equally spread over the available food resource space. Though attraction calls lowered a sparrow's risk to get killed while feeding, increased food competition reduced the benefits of high food availability significantly. Moreover, attracted sparrows were taking a risk when approaching a sender and were facing poor food resources on arrival due to high local population densities caused by other attracted sparrows.

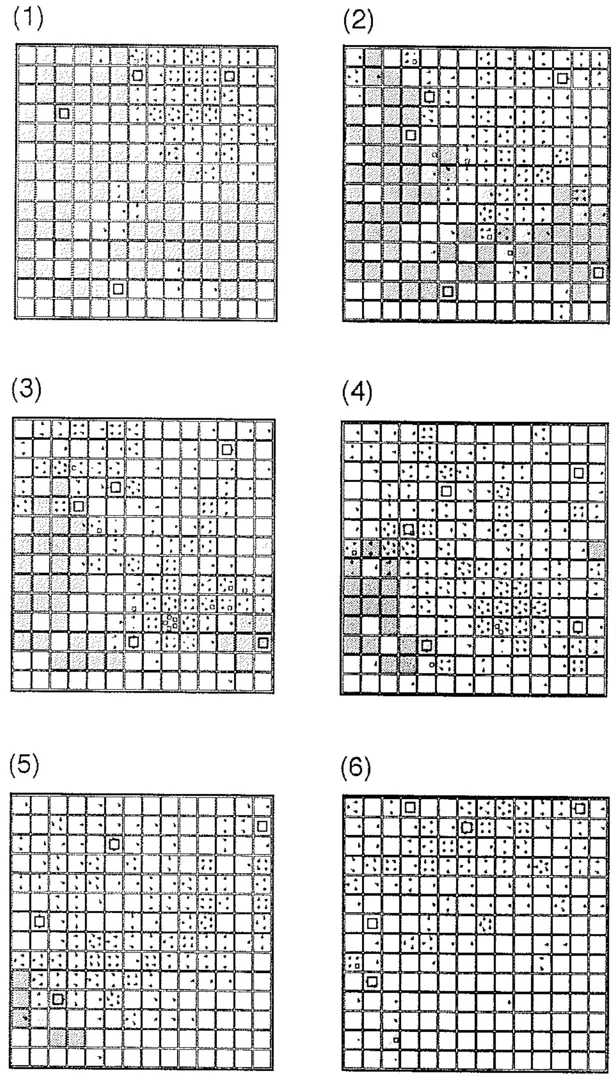

Figure 4: Flocking behaviour of group feeding agents Snapshot sequences from an arena section.

(1)-(3): attraction of several agents towards the bottom of the scence with high local population density around the sender in (3). (4)-(6) distraction and new orientation of the flock to another location.

From model 2, we can conclude that group feeding only pays when food resources are very little constrained in respect to the feeding motivation of the sparrows. Otherwise the disadvantages for signal senders (increased competition) and receivers (competition by the time when they arrive and predation risk when they approach the sender) cannot counterbalance the advantage of reduced predation risk.

4 Discussion and conclusions

In this investigation, we introduced motivation networks as a formal model for animal decision-making, which we designed as a high-level controller for autonomous agents. The application of motivation networks to the biological scenario of sparrow group feeding behaviour revealed some interesting new insights into environmental and behavioural mechanisms in respect to their evolutionary plausibility. Interestingly, group feeding could be easily modelled, though the behaviour did not turn out to be competitive with non-grouping strategies, in contrast to the biological predictions.

This simulation example demonstrates that a combination of multiagent systems and motivation networks can serve as a useful falsification tool for biological research. One great advantage is the structural similarity between motivation networks and biological theory on decision-making, which provides an easy way to model and verify hypotheses. Further, the repeated local interactions of the agents with their environment allow to mimic the dynamic impacts of individual behaviour on behavioural ecology. Another example for this potential is our recent investigation on the effect of interaction strategies on population dynamics in lizards, which uncovered a misleading result of a game theoretical model [Krink & Mayoh subm.].

Moreover, our investigation shows that the introduced motivation networks are capable to control the behaviour of autonomous agents in space and time. Regarding temporal reasoning, we are currently testing the effect of periodically varying environmental parameters which act as synchronising clocks in multiagent coordination. In this context, we will apply motivational networks to control tasks in robot coordination.

Acknowledgments

We would like to thank Prof. Bernd Schmidt for interesting and helpful discussion.

References

[Briscoe, G. & Caelli, T., 1997] Briscoe, G. & Caelli, T. Learning temporal sequences in recurrent self-organising neural nets. In: *Advanced Topics in Artificial Intelligence. 10th Australian Joint Conference on Artificial Intelligence* , pp. 427-435: Springer, 1997.

[Campbell, N. A., Mitchell, L. & Reece, J. B., 1997] Campbell, N. A., Mitchell, L. & Reece, J. B. *Biology - Concepts & Connections*. Menlo Park: Benjamin/Cummings, 1997.

[Chappell, G. J. & Taylor, J. G., 1993] Chappell, G. J. & Taylor, J. G. The Temporal Kohonen Map. 6, 441-445, 1993.

[Cox Hayslip, I., Rosenking, J. P. & Eilbert, J., 1990] Cox Hayslip, I., Rosenking, J. P. & Eilbert, J. Adaptive Planning for Air Combat Maneuvering. *Int. J. Pat. Recog. a. AI*, 4, 447-471, 1990.

[Elgar, 1986a] Elgar. The establishment of foraging flocks in house sparrows: risk of predation and daily temperature. *Behavioural Ecology and Sociobiology*, 19, 433-438, 1986a.

[Hewitt, G. M. & Butlin, R. K., 1997] Hewitt, G. M. & Butlin, R. K. Causes and Consequences of Population Structure. In *Behavioural Ecology* (ed. J. R. Krebs & N. B. Davies), pp. 350 - 372. Oxford: Blackwell, 1997.

[Krink, T. & Mayoh, B., subm.] Krink, T. & Mayoh, B. Do 'rock-scissors-paper' players continue to evolve in a cyclical manner indefinitely. *Nature*, , subm.

[Manning, A. & Dawkins, M., 1998] Manning, A. & Dawkins, M. *An introduction to animal behaviour.* Cambridge: Cambridge University Press, 1998.

[Rich, E. & Knight, K., 1991] Rich, E. & Knight, K. *Artificial Intelligence*. New York: McGraw-Hill, 1991.