Presented at the Workshop Underlying Motivation and Emotion, University of Birmingham, 1993.

Motivation and Attention in an Autonomous Agent

Christian Balkenius

Lund University Cognitive Science
Kungshuset, Lundagård, S-223 50 LUND, Sweden

e-mail: christian.balkenius@fil.lu.se

Introduction

In an attempt to construct a neural network based architecture for autonomous agents, we have been investigating motivation and attention as natural parts of a cognitive system (Balkenius 1993a, b). The role of motivation was originally suggested to us by problems encountered in behaviour selection. For agents with a large set of interacting behaviours, interference between behaviours becomes a substantial problem. The only tractable solution seems to be to include a functionally central system for behaviour selection. This system is responsible for the activation and inhibition of the behaviours of the agent. We identify this system with the motivational system of an animal where a central decision determines the motivational state which in turn determines behaviour.

The design of our motivational system suggests answers to the following questions: (1) What are the determinants of motivation? (2) In what way do motivational states interact with each other? (3) How are behaviours selected by the motivational state? (4) To what extent is perception biased by the motivational state of the agent?

The General Architecture

The motivational system is a central part of an agent architecture that is composed of three functional layers (Balkenius 1993a, Gärdenfors and Balkenius 1993). All three layers communicate reciprocally with the motivational system. The processing in each layer is directed by the motivational system and they in turn may influence the motivational state (Figure 1).

The first layer is responsible for the execution of reactive behaviour of the kind that will be more thoroughly discussed below.

The second layer is used in sequential learning tasks for chaining and chunking of action sequences. Motivational states are used both to construct and to select action chunks. The motivational system also acts a role as a mechanism for task-decomposition (cf. Tenenberg 1993).

The final layer maintains a forward model of the environment (Jordan 1992). Aided with this model, the agent can engage in planning and problem solving (Sutton, 1992). By generating motivational states that simulate future needs, anticipatory planning is being made possible (Gulz 1991). Once given the ability to control its motivational states, the agent will also engage in less fruitful activities such as daydreaming, worry and regret. In our model, these activities are necessary consequences of anticipatory planning abilities.

Below, we consider further the first layer and its interaction with the motivational system.

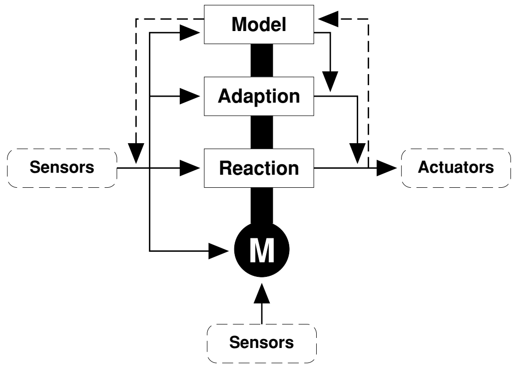

Figure 1. Outline of an agent architecture

Motivation

The determinants of motivation can be roughly divided into three categories:

(1) **External Incentive** The perception of a desired object increases the motivation for its corresponding motivational state. The perception of food increases hunger although the physical need is unchanged.

(2) **Internal Incentive** The motivational state can be changed based on various perceptual or cognitive processes. This is the case, for instance, when after a glance at the watch one realizes that the bus leaves in two minutes and is suddenly motivated to run towards the bus stop. Internal incentive can also be based on much simpler processes as when the candy store serves as an incentive to buy sweets. Internal incentive also includes high level processes such as anticipatory planning that might change the motivational state based on a complex series of inferences.

(3) **Internal Drives** The homeostatic state of the agent does also influence motivation. For example, the deprivation of food increases hunger.

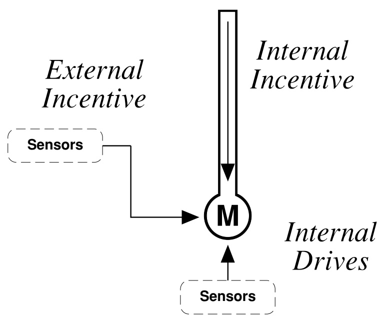

Figure 2. The determinants of motivation

These three determinants are all excitatory in that they try to increase the activation of a motivational state. However, all excited motivational states can not be allowed to direct the agent at one time since this would generate incoherent behaviour. This problem cannot be handled solely by behavioural competition but must be resolved at an earlier stage of processing. This is especially important when we consider the generation of actions that are co-ordinated in space or time which requires some sort of motivational persistence. It is unclear how such a persistence could be the result of behavioural competition only.

To select one motivation at a time we need a mechanism for motivational competition. An obvious solution is to let the motivational states inhibit each other in proportion to their current level of activation (Grossberg 1973). Such a competition allows one motivational state only to be active at any time. This does not mean that all other motivational states are simply discarded. They can be activated again at any time if the external or internal determinants change.

Motivational competition certainly exists. For example, you are less likely to be hungry when you are angry. This property of motivational competition can be exploited in more or less successful ways. For example, one may develop the strategy to eat to decrease frustration or to evoke anger to suppress tiredness.

However, studies of animals imply that the situation is more complicated. Behaviours can roughly be divided into two categories:

(1) Appetence

(2) Avoidance

An appetence behaviour can be further divided into the two components:

(1) Approach

(2) Consumption

This implies that motivational states may have a similar composition. The choice between appetence and consummation is controlled by external stimuli. When the goal situation can not be immediately obtained, an appetence behaviour is generated that causes the agent to search for or move towards the goal. Otherwise the agent will generate the appropriate consummatory behaviour. In general, these two components can be further divided into sub components that generate hierarchically organised motor patterns (Tinbergen 1989).

The type of motivational states discussed above determine what the agent should do. For example, fear should activate an avoidance behaviour, and hunger should generate eating behaviour or a search for food. Other states determine what the agent should have done, such as shock and relief. Shock tells the agent that it should have been in a state of fear but it was not. Relief tells the agent that it should not have been in a state of fear although it was. These states are caused by unconfirmed expectation and are used to change future behaviour.

The interacting motivational states involved in approach behaviour are even more complex. For every positive motivation, such as hunger, there are at least three potential sources of behavioural modification. These motivational states can be categorized as early, synchronous or late with respect to the consummatory behaviour.

(1) **Early learning** is based on expected reward and are connected to appetence behaviour. The level of perceived success or closeness to the goal act as a reinforcing motivational state for the current appetence behaviour. This type of mechanism may be used to explain latent learning (Tolman 1930). While this class of motivational states have many important merits they do critically depend on a competent evaluation of the distance to the goal. Perhaps it is the failure to perceive this distance that makes the dedicated gambler risk even more money after ‘almost winning the bet’.

(2) Synchronous learning can be either positive or negative. Positive learning is caused by a real reward, e. g. eating, and has an effect equivalent to the maximal expected reward. If early learning is used in an agent, this type of motivational state may not be necessary. Negative learning caused by disappointment can only occur synchronously with the supression of the consummatory behaviour although the determination of this moment is a far from trivial matter in most cases.

(3) Late learning would be caused by some effect of consummation on the subsequent motivational states of the agent. Learning based on the drive-reduction hypothesis would fall into this category (cf. Hull 1943).

At least early and synchronous negative motivational states must be present in an autonomous agent.

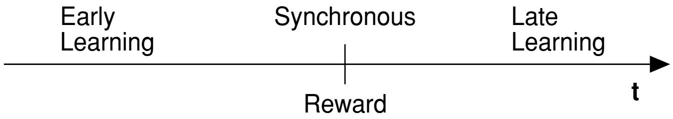

Figure 3. Different times when learning can occur.

Attention as an Emergent Property

Apart from influencing behaviour, the motivational state can also select among several simultaneous percepts. This is the role of attention. In our model, attention is an emergent property of the neural network architecture. There is no distinct attentional module, but attention results from the interaction between the perceptual and motivational systems.

In summary, there are three types of interaction between motivation, perception and action (figure 4):

(1) Interaction between motivation and perception The motivational state influences perception and the perceived stimuli influences motivation. The perceptual system can be tuned to respond more strongly to sensory patterns that are important for the current needs of the agent.

(2) Interaction between motivation and behaviour Behaviour is selected by the current motivational state. This behaviour in turn may change the motivational state.

(3) Interaction between perception and action Behaviour is obviously influenced by the current perception. The motivational system can activate an appetence behaviour, but a perceptual system is required to select a particular appetence routine. Behaviour also influences perception, both in the sense that behaviour changes the world and thus our perceptions, but also in active perception where we reach out to the world as in saccadic scanning motions.

These interactions are resposible for attention both in the sense of attending to some particular stimulus instead of some other and also in the sense that the agent attends to one of its need at a time.

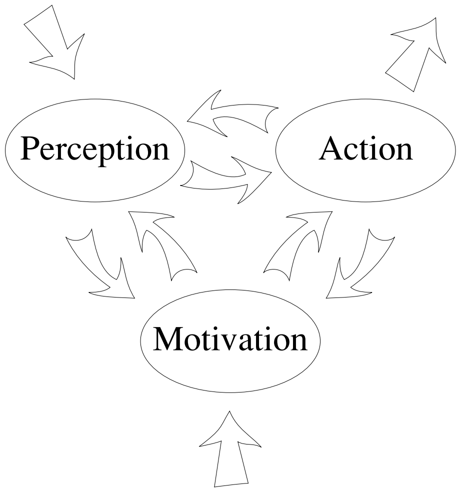

Figure 4. The interaction between perception, motivation and action.

A Computational Model

The ideas presented above have been incorporated in a working computer simulation of an autonomous agent (Figure 5). The domain we have chosen to simulate consists of a two-dimensional continuous plane together with various objects such as food and walls. The agent is allowed to move around in this world using olfactory and tactile sensors on each side of its body. Its task is to choose a suitable diet while avoiding dangerous situations. Both the sensory input and the signals that control movement in the plane are continuous. This lets us view the task of the agent more as a control problem than as a game of chess. Since there are no discrete input states and no finite set of actions to choose from, we have been led to think along lines that differ to a large extent from those of traditional AI.

The agent is controlled by a highly structured modular neural network inspired by some structures in olfactory cortex and the limbic system. The model demonstrates all the properties mentioned above. Most important, it incorporates a new type of motivationally biased auto-associative network for perceptual categorization and attention, a motivational system and a network for instrumental learning. The entire network consists of two mirror images of the same structure, each of which can be divided into five interacting stages.

(1) **Sensation** The sensory system receives olfactory signals that decreases with the distance from objects in the environment. The different smells activate random overlapping patterns at the sensors. No a priori categorization of smells is given to the agent.

(2) **Perceptual Segmentation** This stage consists of two coupled auto-associators. One for the olfactory input from each side of the body. The auto-associators are used to learn common olfactory patterns and to filter out all but the most salient pattern in the current sensory input. The coupling of the networks for each side of the body forces them to make the same choice although the same smell may not be the strongest on both sides of the body. This is essential as the differences in smell intensities between the two systems are used in later stages to direct the movement of the agent.

(3) **Perceptual Categorisation** The this processing stage constructs categories of the patterns in the previous stage using competitive learning. The orthogonal representations constructed have the dual role of both facilitating segmentation in the previous stage, to which they are sent back, and to communicate with the motivational system. The motivational system receives input from this stage and sends signals back that make some categories more likely to get activated than others. In this way, motivation can influence segmentation. The orthogonalisation process is essential when the agent senses complex smells. Objects with similar smells may have quite different affordances. The categorization process makes such representations more different and does thus avoid erroneous associations to motivational states as well as making memory formation easier.

(4) **Motivation** As already mentioned, the motivational system receives three types of inputs. The first comes from the categorisation modules and are used to bias the selection of motivational state towards one that can be satisfied by the current perceptual cues. The second input comes from the higher layers of the architecture and will not be described here. Finally, the motivational system receives input signifying the current needs of the agent. The motivational states are organised as described above and are used both to determine what actions to perform and to control learning and attention.

(5) **Motor Control** Two modules make up the final stage of the architecture. The total input to each of these modules are used to move the agent around its environment. The larger the input the faster the agent moves. When the left module receives a larger input than right, the agent is made to turn right, otherwise it turns left. When both modules receive equal input, the agent moves straight ahead.

The evaluation of the architecture is as yet at an early stage. We are currently investigating the possibility of using simulated evolution to search a space of alternative designs for the agent. This method have already been used to validate some of our design choices. Simulated evolution is also useful as a method for the characterisation of different environments in relation to invariances of the agent architecture.

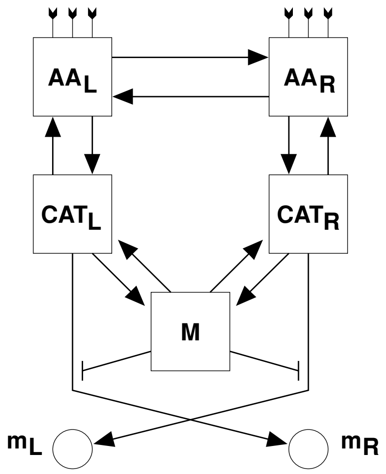

Figure 5. A specific agent architecture with attentional abilities

Figure 6 shows a preliminary computer simulation of the agent depicted in figure 5. The agent orients by olfaction only. Each of the goal objects activates random overlapping patterns at the receptors of the agent. These patterns are segmented under the influence of the current motivational state (6). The smell most relevant to the current motivation is attended to. It is also possible for the olfactory input to change the motivational state (2).

As a result of the interaction between these two processes, attention emerges.

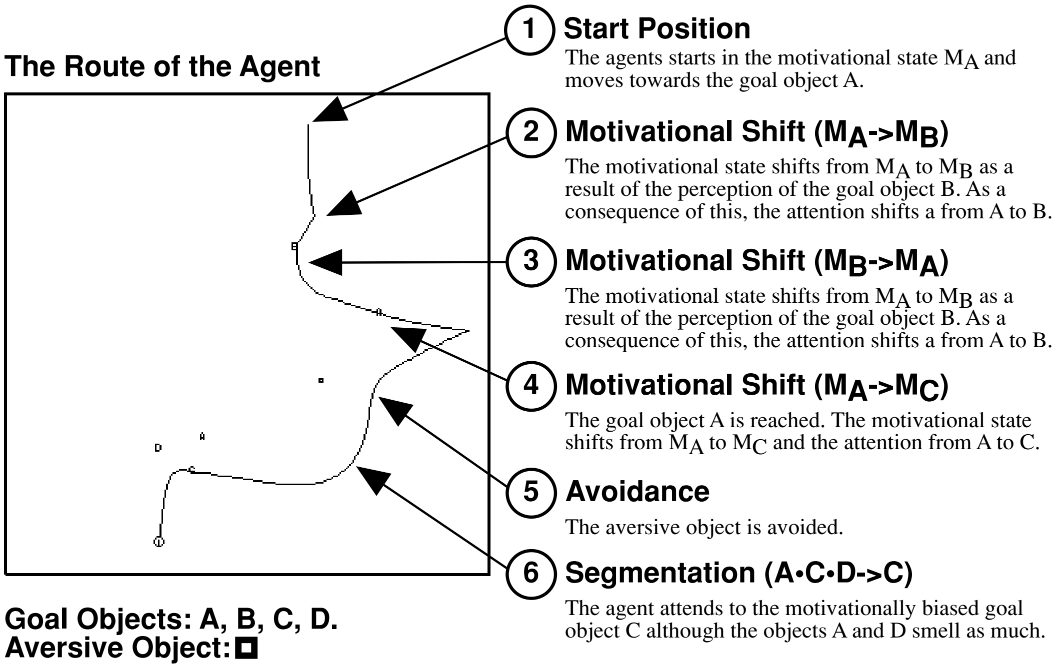

Figure 6. Computer simulation of an agent.

Conclusion

While our current model may be to simple to explain motivation and emotion in humans, it does nevertheless show how a successful control architecture for an autonomous agent can be constructed by combining ideas from psychology, biology and neurophysiology. By building a working computer model, we have been able to investigate the importance of motivation under varying environmental constraints. Our main conclusion may be that it is not possible to study motivation as an isolated system. A motivational system can only be understood once it is combined with perception and action in a well defined environment.

References

Balkenius, C., (1993a), “Natural intelligence for autonomous agents”. In B. Svensson (ed.) *Proceedings of the international workshop on mechatronical computer systems for perception and action*,

Balkenius, C., (1993b), “The roots of motivation”. In J.-A. Mayer, H. L. Roitblat and S. W. Wilson (ed.) *From animals to animats II*, MIT Press, Cambridge, MA.

Grossberg, S., (1973), “Contour enhancement, short term memory, and constancies in reverbating neural networks”, *Studies in applied Mathematics*, **2**, 49-57.

Gulz, A., (1991), “The planning of action as a cognitive and biological phenomenon”, *Lund University Cognitive Studies*, **2**.

Gärdenfors, P. & Balkenius, C., (1993), “Varför finns det inga riktiga robotar?”.

Hull, C. L., (1943), *Principles of behavior*, Appleton-Century-Crofts, New York.

Jordan, M. I. & Rumelhart, D. E., (1992), “Forward models: supervised learning with a distal teacher”, *Cognitive Science, in press*,

Sutton, R. S., (1992), “Reinforcement learning architectures for animats”. In J.-A. Meyer and S. W. Wilson (ed.) *From animals to animats*, MIT Press, Cambridge, MA.

Tenenberg, J., Karlsson, J. & Whitehead, S., (1993), “Learning via task decomposition”. In J.-A. Mayer, H. L. Roitblat and S. W. Wilson (ed.) *From animals to animats II*, MIT Press, Cambridge, MA.

Tinbergen, N., (1989), *The study of instinct*, Oxford University Press, Oxford.

Tolman, E. C. & Honzik, C. H., (1930), “"Insight" in rats”, *Univeristy of California Publications in Psychology*, **4**, 215-232.