Defining the Action Selection Problem

Toby Tyrrell

email: lrtt@uk.ac.ed.cns
Centre for Cognitive Science
University of Edinburgh
Edinburgh EH8 9LW
Great Britain

Abstract

There has been a lack of progress in the field of action selection due to an incomplete understanding of the problem being faced. The differing nature of constituent parts of the action selection/'time-allocation' problem has not been properly appreciated. Some common sub-problems, such as obtaining food and avoiding predators, are described in terms of the demands they make on an animal's time. The significant differences between these sub-problems are highlighted and a *classificatory scheme* is proposed, with which sub-problems can be categorised. The need to take into account the full range of different sub-problems is demonstrated with a few examples. A particular shortcoming shared by all of the more well-known action selection mechanisms, from both robotics and animal behaviour, is described.

Introduction

Action selection is the choosing of the most appropriate action out of a set of possible candidates. The term *action* here refers to one of a set of mutually exclusive entities at the level of the behavioural final common path. That is, their demands on the effectors of the animal are such that only one of them may be executed at any one time. But how should the term 'most appropriate' be defined? The most appropriate action is that which maximises the number of copies of the animal's genes in future generations (assuming that an animal is just a carrier/propagator of its genes, as argued in (Dawkins, 89)).

There are two ways in which an animal can try to maximise the number of copies of its genes in future generations: (1) it can try to reproduce as often as possible (and so create new individuals with many of its genes), and (2) it can try to bring about the reproduction of other individuals which share many of its genes (e.g. offspring, siblings, parents).

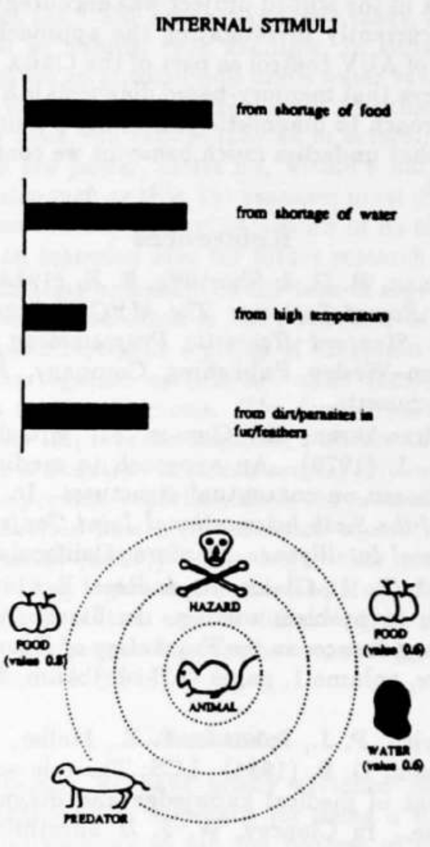

Figure 1: The action selection problem – what is the most appropriate action?

Obviously (1) engenders a further goal of surviving so as to be able to reproduce in the future, and (2) engenders a further goal of helping close genetic relatives to survive, so that they too can reproduce in the future. Therefore the 'purpose' of the mechanism selecting an animal's actions (and indeed of the rest of the animal as well) is to allow it to survive and to cause it to mate successfully as often as possible, and to help closely-related conspecifics to do the same.

This fairly abstract, high-level problem can be split into several *sub-problems*. For instance, the problem of survival leads to sub-problems such as obtaining enough food, avoiding predators, regulating body temperature, etc.

The following list contains some of the action selection sub-problems that commonly occur for different animals: obtaining food, obtaining water, regulating body temperature, avoiding predators, avoiding hazards in the environment, cleaning/preening, sleeping somewhere safe at night, and mating.

The problem of action selection can be thought of as that of allocating the animal's available time to different sub-problems. The nature of the time-allocation problem as a whole depends on the form of the individual demands on the animal's time made by each sub-problem. In this paper a description is given of the nature of some of these sub-problems. The different sub-problems are compared and a set of terms with which to describe them is presented. The usefulness of this tentative classificatory scheme is then demonstrated by showing that some proposed mechanisms for selecting actions are not able to deal successfully with some different types of sub-problem, or with some combinations of the different types.

In the rest of this paper the abbreviation AS will be used for action selection.

Classifying Sub-Problems of the Action Selection Problem

Several common sub-problems of the AS problem have been implemented in a complex simulated environment (Tyrrell & Mayhew, 91), and different AS mechanisms have been tested to see if they can cope with the different sub-problems present. Experience with this simulated environment has given the author some familiarity with the different types of sub-problem and with the problems of designing a mechanism to deal with them.

Some well-known mechanisms are (1) Maes (Maes, 91), a spreading activation network with two 'waves' of input from perception and motivations respectively and with a node for each behaviour; (2) Tinbergen (Tinbergen, 50 & Tinbergen, 51), a hierarchical model in which decisions are made at progressively lower levels until an action has been selected; (3) Lorenz (Lorenz, 50), the hydraulic model which uses an analogy with a reservoir of water to reproduce many of the phenomena of AS; and (4) the drive model (see (Hull, 43) or the description in (McFarland, 85)) which calculates a 'drive' for each sub-problem and then selects an action to satisfy the sub-problem with the highest drive. More explanation of these or other mechanisms (e.g. (Brooks, 86), (Ludlow, 80), (Halperin, 91)) cannot be given here.

Explicit planning systems are not considered here because (1) they presuppose a rather high degree of intelligence on the part of the animal, and (2) they are computationally and intellectually intractable in complex environments about which they only receive incomplete and unreliable information. It should also be noted that many of the mechanisms presented here are able to perform *implicit planning*. They can produce sequences of appetitive and consummatory actions in order to satisfy needs, but can interrupt these if urgent alternative actions are required.

Figure 2 shows some graphs of activations over time for the six example sub-problems now considered:

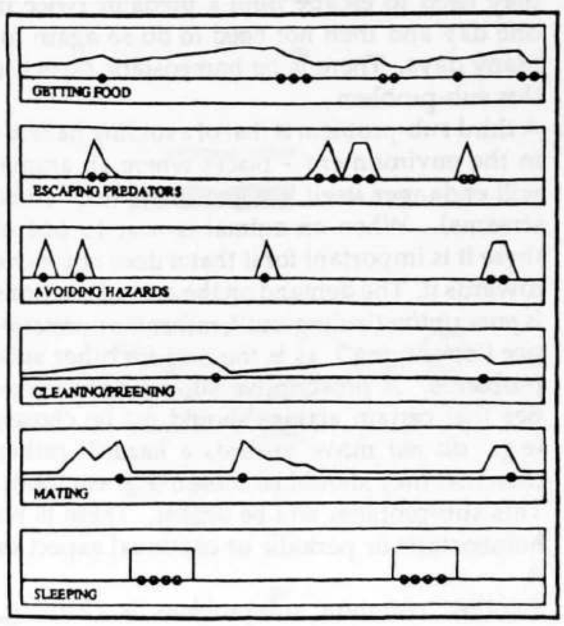

Figure 2: 'Longitudinal Profiles' for the six example sub-problems (after (Maes, 91)). The graphs show the 'urgency' of the sub-problem over time. The circles denote instances when the sub-problem in question determines the animal's action.

1. Trying to obtain enough food is the most commonly covered sub-problem in the ethological literature. The likelihood of the animal trying to obtain food should depend on both external stimuli (whether it sees food in the vicinity) and on internal stimuli (whether it is short of food). It is also a recurring, continual problem. An animal will need to eat a certain average amount each day. The need to obtain food will in general place a low-urgency demand on the animal's time, unless the animal is particularly short of food, since the animal is not likely to die suddenly or become injured if it does something else. Finally, food intake is a homeostatic problem.
2. A second common sub-problem is the need to escape predators. This is usually urgent and overriding, in that if the animal does not attend to it immediately then the consequences could be fatal. It is also highly dependent on external stimuli but not at all on internal stimuli. The priority which the animal should assign to trying to escape from a predator is highly dependent on whether the animal senses any predators, and if so how close they are. This sub-problem is a non-periodic, non-continual sort in that there is no pattern to how often the animal will need to attend to it. An animal may need to escape from a predator twice in one day and then not need to do so again for many days. There is no homeostatic aspect to this sub-problem.
3. A third sub-problem is that of avoiding hazards in the environment - places where an animal will endanger itself if it goes there (e.g. cliffs, streams). When an animal is near to one of these it is important for it that it does not move towards it. The demand on the animal's actions is *proscriptive* ('ruling-out'), rather than *prescriptive* ('specifying'), as is the case for other sub-problems. A proscriptive sub-problem specifies that certain actions should *not* be chosen (e.g. do *not* move towards a hazard) rather than that they should be chosen (e.g. eat food). This sub-problem will be urgent. There is no homeostatic or periodic or continual aspect to it.
4. Another common sub-problem is cleaning, preening or grooming. Most animals need to spend some time every so often to remove dirt/parasites from their fur, clean and oil their feathers, or whatever. This will not be an urgent activity, since it will not be crucial to the animal to pay attention to it at any particular moment in time. It will tend to occur most frequently at moments when no other activity is urgently required. It is continual in that the need for it will recur frequently. It is dependent on internal stimuli but not on external ones.
5. A fifth sub-problem is that of mating. External stimuli are important n that the animal should attach more priority to this sub-problem when a potential mate is perceived (assuming an animal which makes occasional matings with different mates and which forms no long-term partnerships). It is sometimes periodic and related to internal stimuli (e.g. menstrual cycles) and sometimes not. This sub-problem will be prescriptive, non-continual and non-homeostatic. It will probably be fairly urgent but the level of urgency in relation to the other sub-problems will depend on factors such as how often opportunities for mating arise and how much longer the animal can expect to live.
6. A final sub-problem is that of the animal needing to return to its den and sleep there at night. This is periodic (since it will occur every 24 hours). It will have increasing urgency as nightfall approaches. It is non-homeostatic and prescriptive.

Six different common sub-problems that compete for a 'share' of the animal's time have now been described. It does not matter so much that some of this set may not be completely general or that the description of them may not be completely accurate. The important point is that the competing demands on an animal's time vary in their nature. Some of the ways in which they vary have been highlighted by the preceding discussion. Past discussions of AS, and past proposals for AS mechanisms, have not taken fully into account the variety of different types of sub-problems. The study of AS/time-allocation has been held back by a lack of understanding of the different possible sub-problems and a vocabulary to describe them. A tentative vocabulary is now proposed here.

- **Homeostatic *v.* non-homeostatic** - a homeostatic sub-problem contains an internal variable which has a desired 'set-point' (optimal value), or at least a desired range of values. The behaviour of the animal will always act so as to return the value of the variable towards the set-point or range of values (Toates, 80).
- **External stimulus dependent *v.* external stimulus independent** - the urgency with which certain sub-problems should 'demand' the animal's attention is dependent on the appearance of certain external cues (e.g. getting water on

the stimulus of a water source, escaping predators on the appearance of a predator). Other sub-problems are independent of external cues (e.g. cleaning).

- **Internal stimulus dependent *v.* internal stimulus independent** - as for above, except that the important factor is an internal cue (e.g. body heat is too high, not enough food in the animal's stomach).
- **Periodic *v.* non-periodic** - some sub-problems such as sleeping at night are highly periodic, with the desirability of paying attention to them rising and falling with a regular rhythm.
- **Continual *v.* occasional** - some sub-problems need to be attended to frequently and the need for them keeps recurring (e.g. cleaning, getting food/water). They will need to be carried out several times each day. They are often internal stimulus dependent. Others only occur very occasionally and are usually external stimulus dependent (e.g. escaping predators, mating).
- **Degree of urgency** - some external stimulus dependent sub-problems arise only occasionally but are extremely urgent and over-ridingly important when they do occur (e.g. escaping predators, maybe mating). There will be significant consequences for the animal (in terms of future expected genetic fitness) if the sub-problem is not allowed to influence the action the animal selects. Some sub-problems generally have a fairly low urgency (e.g. cleaning), and tend to take over only when none of the more urgent sub-problems are relevant.
- **Prescriptive *v.* proscriptive** - most sub-problems require a certain set of actions to be carried out (e.g. find food, approach it then eat it), whereas others (e.g. avoid hazards) only require that certain actions should *not* be carried out.

This list is almost certainly incomplete in that there are other ways in which sub-problems can vary, but the most important differences are contained here.

Application of the New Classificatory Scheme

The previous section developed a list of terms that can be used to describe the sub-problems of an AS/time-allocation problem. Some well-known AS mechanisms will now be considered which do not produce optimal selections for some of the types of sub-problems outlined (although they perform satisfactorily in most respects). This is because they were designed without a full awareness of the differences that could occur in sub-problem parts of the AS problem.

Example 1 – Improper Combination of Proscriptive and Prescriptive Needs

Most of the proposed mechanisms do not take account of the fact that an AS sub-problem which only proscribes certain actions (e.g. the need to avoid a hazard prohibits movement towards that hazard) still leaves a wide range of actions which can be performed to the advantage of the animal. It can be seen in figure 3(a) that if the animal moves to the top left then it will both *not move towards* the hazard, and *move towards* the food.

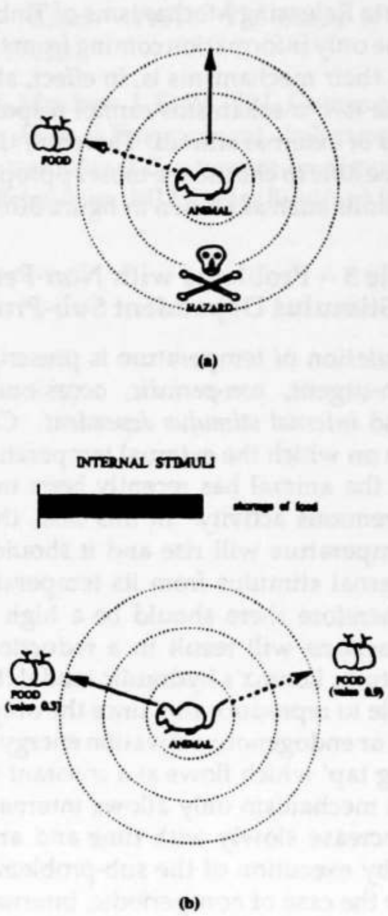

Figure 3: Example situations in which certain mechanisms may select the wrong action. In both (a) and (b) the dashed arrow indicates the optimal action and the solid arrow indicates the sub-optimal choice that might be made.

Most mechanisms cannot generate the optimal action because they select a particular sub-problem as the most urgent/appropriate, and *then* decide which action is most relevant for that sub-problem only. Mechanisms which have this fault are Maes, Tinbergen, Lorenz and the drive model.

Example 2 – Improper Selection between External Stimulus Dependent Sub-Problems

Consider the situation shown in figure 3(b). Here the animal is short of food, and there are two food sources with differing utility to the animal. Some mechanisms, e.g. Tinbergen and Maes, do not allow the *sizes* of external stimuli to affect the priority of a sub-problem. The information from the environment is expressed only in terms of logical conditions about the environment which can be either true or false (the environmental propositions of Maes, or the Innate Releasing Mechanisms of Tinbergen). Because the only information coming from the environment in their mechanisms is, in effect, along binary links, the two mechanisms cannot respond to varying sizes of external stimuli. Therefore they will not always be able to choose the most appropriate action in situations such as shown in figure 3(b).

Example 3 – Problems with Non-Periodic, Internal Stimulus Dependent Sub-Problems

The regulation of temperature is prescriptive, usually non-urgent, *non-periodic*, occasional, homeostatic and *internal stimulus dependent*. Consider an occasion on which the external temperature is fairly hot and the animal has recently been undertaking some strenuous activity. In this case, the animal's body temperature will rise and it should receive a high internal stimulus from its temperature receptors. Therefore there should be a high likelihood that its actions will result in a reduction of body temperature. Lorenz's hydraulic model (Lorenz, 50) is not able to reproduce this since the only source of internal, or endogenous activation energy is from his 'dripping tap' which flows at a constant rate. Since Lorenz's mechanism only allows internal variables which increase slowly with time and are then decreased by execution of the sub-problem, it cannot work for the case of non-periodic, internal stimulus dependent sub-problems.

Discussion – Making Decisions at the Wrong Level

One point to arise out of the previous section concerns the *level* at which decisions should be made. Lorenz's mechanism, Maes's mechanism, the drive model and Tinbergen's mechanism share a common attribute. Initially they all make a decision as to which sub-problem of the AS problem is the most appropriate (usually by calculating a value for each candidate sub-problem and picking the sub-problem with the highest value), and then later they make a decision as to what action is most appropriate for that sub-problem only. There is a two-stage process: (i) select a *sub-problem* (e.g. obtain food) to attend to, and (ii) select the most appropriate *action* for that sub-problem (e.g. move in a certain direction, eat food). While the mechanisms are fine in most respects, this common attribute would seem to be incorrect.

As described in the first example of the previous section, this approach leads to a shortcoming in that the needs of prescriptive and proscriptive sub-problems cannot be combined. The problem is more general than that though. No 'compromise' actions can be selected, whether the compromise is between prescriptive and proscriptive sub-problems, between two prescriptive sub-problems, or between two proscriptive sub-problems. This is because only one sub-problem is taken into account when selecting the best action. So, for instance, in the situation in figure 1 (ignoring the predator), the animal would not be able to choose the best action, that of moving to the right (medium-valued water, medium-valued food and not moving to the hazard). Instead, since it would only consider the most urgent need – to avoid the hazard – it would move directly away from that.

One AS mechanism which addresses the above problem and seems as if it would be able to cope with the whole range of AS sub-problems is (Rosenblatt & Payton, 89). This mechanism deserves further attention but cannot be considered here.

Conclusions

In the introduction to this paper the terms *action selection problem*, *sub-problem*, and *action selection mechanism* were defined. In §2 some of the most common sub-problems of an animal's action selection problem were described. This led on naturally to a set of descriptors for action selection sub-problems. The usefulness of this *classificatory scheme* was then shown in §3, in which several shortcomings with various action selection mechanisms were described using the new vocabulary.

The importance of the classificatory scheme presented in this paper is not just that it gives some convenient labels for describing action selection sub-problems. Rather, the importance lies in the whole way of thinking that it engenders. Action selection is the problem of dividing an animal's time amongst a number of sub-problems of differing nature. Past suggestions for action selection mechanisms have not come to terms with the wide range of differing sub-problems. Any valid action selection mechanism needs to be able to cope with all of them, and to interweave their demands on the animal's time/actions efficiently. Progress in the study of action selection has been hampered because this point has not been fully appreciated. It is to be hoped that the theory presented in this paper will enable more progress to be made.

Acknowledgements I would like to thank David Willshaw and John Hallam for their help and advice during the time I was carrying out this work. I would also like to thank Janet Halperin and Bridget Hallam for commenting on drafts of this paper. This project was supported by SERC (grant no. 89310818).

References

Brooks, Rodney A. 1986. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation* 2:14–23.

Dawkins, Richard. 1989. *The Selfish Gene* (2nd edition). Oxford University Press.

Halperin, J. R. P. 1991. Machine Motivation. In Proceedings of the First International Conference on Simulation of Adaptive Behaviour. MIT Press/Bradford Books.

Hull, Clark L. 1943. *Principles of Behaviour: An Introduction to Behaviour Theory.* D. Appleton–Century Company.

Lorenz, Konrad. 1950. The Comparative Method in Studying Innate Behaviour Patterns. *Symposia of the Society for Experimental Biology* 4:221–268.

Lorenz, Konrad. 1985. *Foundations of Ethology.* Heidelberg: Springer-Verlag.

Ludlow, A. R. 1980. The Evolution and Simulation of a Decision Maker. In *Analysis of Motivational Processes.* Toates, F. M. & Halliday, T. R. eds. Academic Press.

Maes, Pattie. 1990. How To Do The Right Thing. *Connection Science Journal* (special issue on hybrid systems). 1(3).

Maes, Pattie. 1991. Bottom-Up Mechanism for Behaviour Selection In An Artificial Creature. In Proceedings of the First International Conference on Simulation of Adaptive Behaviour. MIT Press/Bradford Books.

McFarland, David. 1974. Time-Sharing as a Behavioural Phenomena. In *Advances in the Study of Behaviour (4).* Lehrman D. S., Rosenblatt J. S., Hinde R. A. & Shaw E. eds. New York: Academic Press.

McFarland, David. 1985. *Animal Behaviour.* Pitman.

Rosenblatt, Kenneth J & Payton, David W. 1989. A Fine-Grained Alternative to the Subsumption Architecture for Mobile Robot Control. In Proceedings of the IEEE/INNS International Joint Conference on Neural Networks.

Tinbergen, Nikko. 1950. The Hierarchical Organisation of Nervous Mechanisms Underlying Instinctive Behaviour. *Sympos. Soc. Exper. Biol.* 4:305–312.

Tinbergen, Nikko. 1951. *The Study of Instinct.* Clarendon Press.

Tyrrell, T. & Mayhew, J. E. W. 1991. Computer Simulation of an Animal Environment. In Proceedings of the First International Conference on Simulation of Adaptive Behaviour. MIT Press/Bradford Books.