in Proceedings of 5th International Symposium on Artficial Life and Robotics. AROB.

SlugBot: A Robotic Predator in the Natural World.

Ian Kelly, Owen Holland and Chris Melhuish.

Intelligent Autonomous Systems Engineering Laboratory, University of the West of England,
Faculty of Engineering, Bristol, BS16 1QY, England.
Ian2.Kelly@uwe.ac.uk

Abstract

A key aspect of the autonomy of living things is their ability to find and use sources of energy in the natural environment. Clearly, any comprehensive attempt at producing artificial life should demonstrate an equivalent capability; equally clearly, so should any truly autonomous robot. To date, both Alife agent simulations and robotic implementations have used environments and energy sources much too simple or structured to allow such equivalence to be claimed. This paper describes recent progress on an attempt to break free of these limitations by developing the world's first artificial predator – a robot which lives free on agricultural land, hunting and catching slugs, and fermenting the corpses to produce the biogas which is its sole source of energy.

Introduction

One of the most impressive facts about living creatures is that for much or all of their lives they are truly autonomous: they are able to survive and operate successfully in an unstructured environment without requiring any assistance whatsoever. Autonomy, whether biological or artificial, can be thought of as consisting of two major aspects: computational autonomy, and energetic autonomy. Computational autonomy refers to the ability to determine and carry out actions independently, whilst some of these actions may be related to the acquisition of energy, most will be concerned with other aspects of system operation. In the context of a biological system, the other aspects are associated with survival and reproduction, such as avoiding predators, finding shelter, grooming, finding mates, and so on. Energetic autonomy refers to the independent ability to maintain the internal availability of energy above the lethal minimum for sufficiently long periods to enable the system to achieve its mission, which in the biological context corresponds to securing the effective propagation of its genes. This ability does not merely involve making correct decisions in order to secure the raw energy; it also includes the conversion of the raw energy source into a usable form.

In the context of artificial life, a typical investigation of so-called 'autonomous agents' might involve simulated agents attempting to survive and reproduce in a world containing spatially localised elements corresponding to 'food' and 'predators', and so on. Although such abstract studies have been genuinely useful in exploring the dynamics of such situations, all fall very far short of the complexities faced by an actual animal in the real world. Perhaps the closest approach was made by Tyrrell [1], in a study within the area of adaptive behaviour rather than artificial life. Tyrrell reviewed and summarised the challenges faced by animals in surviving, and devised a simulation environment which included representatives of each class of problem. He then compared the various methods of action selection which had been proposed, in the contexts of biology and adaptive behaviour, as being able to support appropriate decision making within this environment. Three things are clear from his study:

- Existing artificial life simulations do not contain all of these problem types
- Even his environment is a gross simplification compared to any real environment
- None of the action selection systems examined comes close to the performance of any real creature

Over the last two decades, the design and control of autonomous robots has formed a major area of academic and industrial research. There are now many examples of robots or automated mobile systems (such as missiles, smart torpedoes, and some spacecraft) which achieve an apparently high degree of energetic and computational autonomy. Such systems carry enough fuel for their mission or can use radiant energy from their environment, and can control themselves 'intelligently' without human intervention. Some automated cleaning and materials handling AGVs use opportunity battery charging to achieve a degree of apparent autonomy. Several academic research groups have constructed robot environments which feature a 'powered floor', giving the possibility of indefinitely extended operation. However, on reflection, it is clear that most of these so-called 'autonomous' robots still require some explicit or implicit intervention from humans in order to carry out their tasks. Forms of human intervention include supplying information and energy, physically assisting the robot, and modifying the environment to suit the robot. The issue of autonomy has been finessed, rather than having been confronted and overcome.

This is by no means the first identification of the lack of autonomy in artificial agents – similar observations in a slightly different context were perceptively articulated by Steels [2] several years ago. However, the project