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 described here represents the first serious attempt to design and construct a robot system with energetic and computational autonomy comparable to an animal system. We have sought to guarantee the comparability by constraining a robotic system to obtain its energy in the same way as most animals - by finding and 'digesting' organic material in an unstructured environment.

This decision immediately brings a host of problems in its train. For instance, natural resources of this type are found only in certain types of places, and may only appear transiently, depending on time of day, season, and weather conditions; most importantly, the resources are destroyed by being used. The control of the foraging process must therefore take into account issues such as where and when to look for food, when to revisit an area where food was previously found, when to abandon a site which is not producing food, and so on. Once found, the organic resources, or food, need to be converted to an appropriate form of energy for storage and use. We propose to convert the organic material to electricity by using a mixture of biological and advanced engineering technologies: an initial and quite standard anaerobic fermentation process will be used to obtain biogas, and the biogas will then be passed through a specially developed tubular solid oxide fuel cell which generates electricity directly. Modern fuel cells of this type can produce electricity from biogas containing methane at concentrations of 20% or less [3].

The nature of the organic material to be used imposes further constraints. One option would be to use vegetable matter, which means that the robot would correspond to a grazing animal. While grazing is not a trivial problem, especially when resources are scarce, we did not think that such a choice would sufficiently challenge the requirement of computational autonomy. We therefore took the second option open to us: the source should be a mobile animal species, since the technical problems of predation are certainly severe enough from the computational viewpoint.

The choice of a prey species was guided by several considerations. Most obviously, the prey species should be reasonably plentiful and not require rapid pursuit, which would be difficult to achieve over soft ground at a reasonable energy cost. It is also likely that there would be some minimum energy cost associated with finding, catching, and consuming any individual creature, however large or small, and so the prey species should be large enough to give a significant margin over the minimum energy expended. Finally, in order to conform to ethical considerations, it should be an invertebrate pest species already subject to lethal control measures, so that the system would be doing something of actual use that would have to be done anyway. All of the above criteria are met by the slugs found on agricultural land in the UK, especially *Deroceras reticulatum* [4]. They are slow moving, abundant, relatively large, and extremely destructive - UK farmers spend over £20m per annum on buying and spreading molluscicides [5]. Slugs are also potentially more suitable for fermentation than some other possible target species: they do not have a hard external shell or exoskeleton and have a high moisture content.

Slugs are quite general pests, but perhaps do most damage to crops requiring a well cultivated seedbed, such as winter wheat or potatoes [6]. Following cultivation and planting, such ground is very soft; this means that a heavy fermentation vessel could not be moved over it without consuming large amounts of energy. The fermentation vessel/fuel cell system will not therefore be part of the robot, but will be in a fixed location, and the robot will bring slugs to it, and collect power from it. (This strategy is reminiscent of certain social insect colonies, where the insects bring raw organic material – which they cannot digest - to a ‘fungus farm’ in the nest, and feed on the fungus, which they can digest.) Inevitably, the fermenter will require a certain amount of energy to cover operating overheads, and our initial calculations indicated that a single robot would be unlikely to be able to gather sufficient energy to service both its own and the fermenter's energetic requirements; a multiple robot system must therefore be envisaged. Rather than being a handicap, this confers many potential advantages: for example, some search tasks might be performed more efficiently by several communicating robots than by a single robot. More importantly, the multiple approach gives a potential for achieving system reliability through redundancy. Interestingly, a robot that has run out of power away from the recharging station could be rescued by receiving power from another robot.

We have sketched out the basic concepts, and the logic behind the overall system design. The remainder of this paper describes the progress made to date in the design and construction of the robot hardware. Work on the fermenter and fuel cell system is to begin shortly, and initial field trials are scheduled for this year.

The Robots

To enable the robots to operate in a real unmodified environment and obtain their energy from the very limited resource represented by slugs, the robots will have to be energy efficient, reliable, and able to operate outside on soft and perhaps irregular ground. (They will also need to be protected from the weather, slug slime, and mud.) Energy efficiency can be achieved in several ways:

- by constructing the robots using light but strong materials like carbon fibre and aluminium
- by using decentralised modern low power controllers and electronics where possible - instead of a single high speed central processor - thus allowing currently unused devices to be shut down
- by the use of physical designs and control strategies designed to optimise efficiency

For example, the energy requirements for hunting and catching slugs have been minimised in the following way. Any system designed for finding and capturing slugs will require some means of searching the surface of the ground to detect slugs, and some means of picking slugs up from the ground. In our design, the sensor used

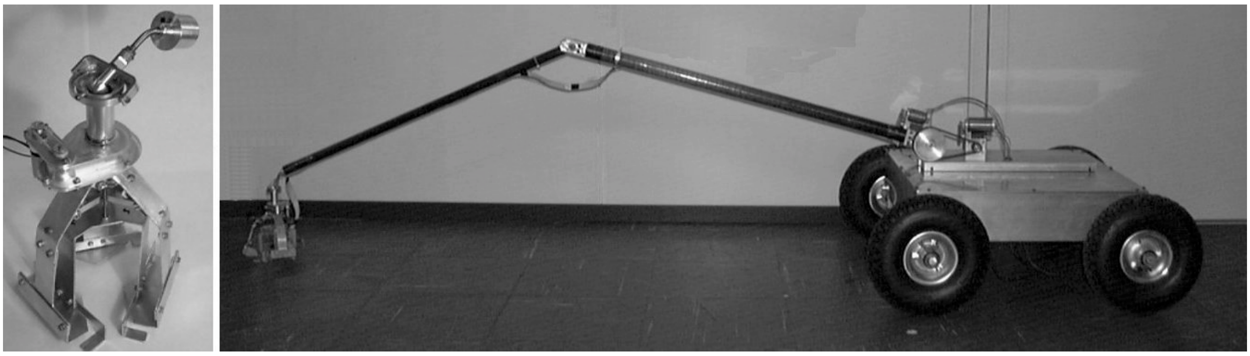

Fig. 1. Prototype three fingered gripper with wiper blades and compliance gimbal (left), and the arm and gripper system mounted on a turn table (right)

for detecting slugs, and the gripper used for catching them, are both located at the end of a long articulated arm. The energy required to move the end of this arm for a certain distance across the surface is much less than the energy required to move the whole robot through the same distance over a soft and muddy field. The robots will hunt for slugs by firstly moving to the centre of the region to be searched. The ground around the robot will then be scanned by rotating the arm around the body of the robot and gradually extending it, guiding the sensor, and hence the gripper, in a continuous spiral trajectory starting from the robot and moving gradually outwards. During this phase, the sensor and gripper will be actively maintained at a constant height above the ground. This spiral strategy covers the maximum amount of ground possible, and also reduces the power consumption of the scanning operation by conserving angular momentum. When a slug is found, it will be picked up by the gripper (which will already be in the correct location) and transferred to an on-board storage container. The arm will then be returned to its former location so that scanning can continue. When all the slugs at that location have been collected, the robot will move to a new location and repeat the procedure. (When the onboard slug container is full, or when the robot needs power, or when it appears that more energy will be used in hunting than is expected to be gained from it, the robot will return to the fermenter, where it will deposit any slugs and perhaps recharge itself.)

The optimal length of the arm is a function of the spatial density and distribution of slugs, and also of the power required to operate arms of different lengths and weights in relation to the power required to move the robot. Obviously, the longer the arm, the more ground it can cover without having to move the robot, but the more power will be required both for scanning and for picking up slugs. Our calculations showed that with between 1 and 10 detectable slugs per square metre, an arm with a nominal base length of around 1.5m would be the most efficient, but where there are many more slugs a shorter arm would be more efficient. On the basis of these results and some initial slug counts, we have opted for the 1.5m long arm. (At full extension, it measures 1.86m from the centre of the robot to the centre of the gripper.)

Other desirable attributes of the arm are that it should be light (and hence energy efficient), stiff, and easily controlled, and that it should be possible to manoeuvre the tip accurately. If possible, it should also have a simple construction, increasing its reliability, making it easier to manufacture and reducing its cost. (The total budget for the project, including labour and materials, is around $100,000.) The final version, designed and fabricated by Martin Scull of the Faculty of Engineering, consists of two 0.75m tubular sections, with a hinged joint between them (see Fig. 1, right). In order to allow the arm to rotate around the whole robot, it is mounted on a powered turntable located in the centre of the robot's chassis. The chassis is large enough to maintain stability in all directions when the arm is fully extended and loaded; the wheelbase dimensions were minimised by mounting the robot's batteries on the turntable on the opposite side from the arm, thus acting as a counterbalance. The arm is constructed from aircraft grade carbon fibre tube to meet the requirements of lightness and stiffness. To keep the arm structure light and the inertia low, the motor and gearbox required to provide movement at the elbow joint are mounted on the turntable; the drive to the elbow joint is transmitted by a lightweight toothed belt inside the arm. Since the numbers of slugs available on the surface peak in the early evening and just before dawn, the rate of gathering them during these periods must be as high as possible. To this end the arm motors and gearboxes were selected so that the arm can move from fully retracted to fully extended, or vice-versa, in under 1.5 seconds. Self locking worm gearboxes provide the required motor speed reduction, and allow the arm to be held in position without consuming any energy. All motors are fitted with optical encoders to provide position and velocity feedback.

The arm's end-effector is a robust lightweight gripper capable of picking up and releasing both wet and dry slugs, regardless of their size, sliminess, and orientation, and in the presence of irregularities in the substrate. Several design iterations were necessary; the final version consists of three fingers at 120 degrees spacing, operated by a single miniature motor (see Fig. 1, left). As the fingers close, they meet underneath the slug so that it can be lifted; when the gripper is opened, wiper blades ensure the slug's release, however much slime is present.

Slugs are detected and targeted visually, and the camera carrying out this function is mounted in the centre of the gripper, away from slugs and mud, and ideally positioned to provide accurate feedback of the position of the gripper with respect to a target slug. To ensure that the view from the camera is always perpendicular to the ground, regardless of the arm's extension, the whole gripper assembly normally hangs freely on a gimbal. It is also possible to lock the gripper assembly in a fixed position, thus stopping any swinging, by fully opening the gripper. Underneath each of the three wiper blades is a plate which produces passive alignment with the contours of the ground when the gripper descends over a slug, ensuring that all three blades move under the slug when the gripper closes. The gripper mechanism will be protected from the weather, mud, and slime by a flexible cover.

Since the fermentation and recharging station will be in a fixed location, it is imperative that the robots are always able to locate and return to it. The nature of the terrain means that wheel slip and the irregularity of the substrate will preclude any extended use of odometry; we will take advantage of non-biological technology here by using a combination of the Differential Global Positioning Satellite (DGPS) system, and an active infrared localisation system [7], [8]. DGPS will also be used for mapping the locations of grazed areas, so that good patches can be found again, and over-grazing can be avoided. (In fact, this last point may not be a problem: a study undertaken for this project [9] found that the removal of all surface slugs from a field location every few days does not, in the medium term, appear to reduce the number of available surface slugs. It appears that a large reservoir of underground slugs is always present; as slugs are removed from the surface more slugs migrate upwards to replace them.

Although the fields in which trials will be undertaken are quite large and empty, obstacles may be present from time to time, and the robot will carry the normal complement of obstacle avoidance sensors required by any mobile robot. Obstacle detection will be achieved using a combination of ultrasonic sonar and, as a last resort, bump sensors. In addition, two sets of miniature ultrasonic sonar transceivers will be placed in the gripper: one set will point downwards so that the sensor for slug detection can be kept at the optimal distance from the ground regardless of any irregularities in the ground, and a second set will face outwards to detect any obstacles in the path of the gripper.

Detecting Slugs

We intend to operate the robots on crops of winter wheat, for about three months from the preparation of the seed bed. There are several reasons for this: it is the period when slugs are most numerous, when they are largest, and when they do most damage; it also offers us reasonably level ground with a relatively sparse crop cover. However, even under these relatively benign conditions, the detection of slugs is a difficult practical problem, mainly because of the presence of non-slug objects such as living and dying vegetation, stones, and lumps of soil. Many different types of sensors could be used for this task; we have opted for a vision based system since it offers the best combination of size, weight, cost, and effectiveness. VLSI Vision Ltd. produce a monochrome CMOS image sensor that is lightweight, relatively low power (<175mW), of adequate resolution (164 by 124 pixels), and sensitive (down to 0.1 Lux). It has a digital interface, and the maximum frame rate of 60 Hz enables reasonably high scan speeds for the arm. This image sensor also has adjustable automatic exposure control, and can calculate the average image intensity of the last frame, and perform pixel level thresholding using an adjustable threshold.

To reduce the complexity of the problem of the visual detection of slugs, we have developed a simple method of filtering out image components deriving from soil, living and dead vegetation, and stones. Since slugs are active mainly from dusk to dawn, some form of illumination of the imaged area will be required. This opens up the possibility of using a combination of coloured light and optical filtering to increase the relative visibility of slugs whilst decreasing the visibility of vegetation and earth. We have found that this can be achieved by using red light from extreme brightness LEDs, and placing a matching filter in front of the image sensor. Under these lighting conditions both vegetation and soil appear dark, whilst the slug *Deroceras reticulatum* reflects red light and thus appears very bright in the received images. Fig. 2 shows a 32mm long *Deroceras*, together with some grass, under white light - note the relative brightness of the slug and grass. Fig. 3 is the same image (except for some movement of the slug) taken using the red light and filter combination - the grass now appears dark whilst the slug is bright. Fig. 4 shows how the image of the slug can be picked out from the background by applying a simple pixel-based threshold function:

$$\text{Threshold} = (c[\text{Average Image Intensity}] - k)$$

where c and k are constants, to the red illuminated image of Fig. 3. There is an added benefit, in that the threshold scheme does not detect slugs under 15mm in length, and so filters out small slugs which in fact cost more in energy to retrieve than they can possibly yield. The final and relatively simple stage is the identification of bright patches which are of the correct size and shape.

As regards hardware, we use a low-power PAL to perform the high-speed image collection and compression; the data are then passed to a Scenix low-power 8-bit PIC-like microcontroller for shape recognition. The average image intensity and pixel-wise

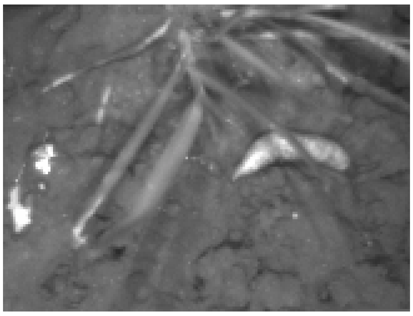

Fig. 2. Under white light

Fig. 3. Under red light

Fig. 4. After thresholding

thresholding are both performed directly by the image sensor, with the Scenix microcontroller calculating the pixel threshold using the function given above. Each image is buffered into an 8-bit wide static RAM by the PAL, with eight thresholded pixels being stored in each location of the SRAM.

The Robot Control Architecture

The robot system is quite complex: there are several motors (gripper, elbow, shoulder, turntable, drive, steering) which must be precisely controlled, and several sensors (shaft encoders, imaging, obstacle avoidance, scan level, battery level, limit switches etc.) which must be monitored. Instead of centralising all processing, we have instead opted for a fairly distributed system with the emphasis on local processing, allowing us to use a low-speed three-wire serial (I^2C) bus to link all of the major subsystems. However, the high-level control of the robot is still handled by a single processor. By using a relatively decentralised approach, the overall system complexity is decreased, and the reliability should be increased. For example, all motion controllers share the same design – a PIC microcontroller handles communication and basic processing, and a Hewlett-Packard HCTL-1100 intelligent motion controller manages the shaft encoders and motor drives; each motor controller has a unique address on the I^2C bus.

Since the arm is capable of moving from fully retracted to fully extended (or vice-versa) in under 1.5 seconds, it is potentially hazardous, and we have had to confront some safety issues. Each closed loop motion controller incorporates a simple fail-safe protection system in the form of an independent watchdog. A separate retriggerable monostable controls the enable input of each motor driver. If a monostable does not receive a logic high to low transition every 20ms from its PIC microcontroller (due to a crash, or loss of communication) the associated motor will be shut down. In addition, both the PIC and the HCTL-1100 will stop the motor if it is driven into one of its limit switches. As final layers of defence, each motion controller board features motor current monitoring, and also has an emergency stop input capable of being activated by the high-level microprocessor as well as a stop command from a remote radio transmitter.

Control Strategy

Although the detailed engineering of the robot system presents enormous challenges, the deepest problems are expected to be at the level of high level control and decision making: when should the robot do what? This is the action selection problem examined by Tyrrell in [1]; it is well known within biology and adaptive behaviour, but has not been confronted by robot designers because the situation of a typical animal is radically different from that of a typical robot. In the industrial and military context, a robot or similar automated system usually has a single goal, and is rarely faced with a real choice of what to do next. In contrast, animals and our robots, which are in the same free-living open-ended situation, always have several simultaneous goals.

From the robot's point of view this includes such goals as: gathering slugs; recharging batteries; not becoming lost or stuck; always having enough charge to be able to return to the refuelling point; maintaining the functionality of its sensors and effectors by grooming.

Unfortunately, the action selection problem has not yet been solved: it is not known in general terms how to design a system which always selects the action which will maximise the expected survival time. Our only option, as designers and engineers, is to find and implement a computationally feasible solution which gives at least adequate performance. What makes the situation even more challenging is that, even with unlimited supplies of slugs, our system will at best be on the borderline of survivability, because our scheme for energy recovery is likely to be at least an order of magnitude worse than any biological system, and so the performance requirements are likely to be even more severe than those which an animal living entirely on slugs would face.

The chosen solution has been to adopt an existing model of motivation and action selection in animals and robots - the d,r,k model ([10], [11], [12]). This has been developed over several years, and is uniquely suitable for application on robots because it has been developed by considering animals as if they were robots, and vice-versa. It is computationally extremely simple – all that is necessary is the calculation at frequent intervals of a numerical value for every possible action in the current situation. Although there are very many possible actions, the speed of a modern microprocessor makes it possible to perform such calculations very quickly. The factors involved in the calculation are also intrinsically simple: they involve the resources which the robot requires (slugs, cleanliness of sensors, knowledge of its location in relation to the fermenter, power in its batteries), the actions it can perform in the current situation to acquire or maintain each of these resources, and the rate at which the performance of each such action can be expected to produce the relevant resource. The main implementation problem centres on the last factor: an animal can generate accurate expectations through instinctive knowledge, or through learning. For example, an animal which lived on slugs would have very accurate expectations of how many slugs might be available in a certain type of location at a certain time of year under particular weather conditions. All we can give the robot at present is some very inaccurate expectations, and the hope that it will be able to learn to improve them.

Conclusions

Although significant engineering challenges remain, we have shown that it is possible to conceive of a credible plan for constructing a robotic predator, and that the execution of the initial parts of the plan has proved feasible. The project has already served to focus our thoughts on the real problems associated with true autonomy and self sufficiency in robots.

References

[1] Tyrrell T (1993), Computational Mechanisms for Action Selection. PhD Thesis, University of Edinburgh, UK

[2] Steels L (1995), When are robots intelligent autonomous agents? Journal of Robotics and Autonomous Systems 15, pp. 3-9

[3] Staniforth J, Kendall K (1998), Biogas powering a small tubular solid oxide fuel cell. Journal of Power Sources 71, pp 275-277

[4] South A (1992), Terrestrial Slugs: Biology, Ecology, and Control. Chapman and Hall

[5] Glen DM (1994), Ecology of Slugs in Cereals in Relation to Crop Damage and Control. In: Leather SR, Walters KFA, Mills NJ et al (eds), Populations, and Patterns in Ecology, Intercept, pp. 163-171

[6] Glen DM, Milsom NF, Wiltshire CW (1998), Effects of Seed-bed Conditions on Slug Numbers and Damage to Winter Wheat in a Clay Soil. Annals of Applied Biology, pp. 115, 177-190

[7] Kelly ID, Keating DA (1996), Flocking by the Fusion of Sonar and Active Infrared Sensors on Physical Autonomous Mobile Robots. Conf. on Mechatronics, pp. 1/1-1/4

[8] Kelly ID (1997), The Development of Shared Experience Learning in a Group of Mobile Robots. Ph.D. Thesis, Department of Cybernetics, University of Reading, UK, pp 2/4-2/12

[9] Glover C (1998), Reinvasion of Cleared Areas by the Slug Deroceras reticulatum. Internal Report. Long Ashton Research Station, Bristol

[10] McFarland DJ, Boesser T (1993), Intelligent Behaviour in Animals and Robots. MIT Press, Cambridge, MA

[11] McFarland DJ, Spier E (1997), Basic Cycles, Utility, and Opportunism in Self-Sufficient Mobile Robots. Robotics and Autonomous Systems 20, pp. 179-190

[12] Spier E, McFarland D (1998), Possibly Optimal Decision Making Under Self-Sufficiency and Autonomy. Journal of Theoretical Biology