Jakobi, N., Husbands, P. and Smith, T.M.C. **Robot Space Exploration by Trial and Error**. In Koza *et al.* (editors). *Proceedings of the Third Annual Conference on Genetic Programming: GP98.* pp807-815. Morgan Kaufmann. 1998.

Robot Space Exploration by Trial and Error

Nick Jakobi
AnimatLab
Ecole Normale Superieure
Paris, France
nick.jakobi@ens.fr

Phil Husbands
COGS and CCNR
University of Sussex
Brighton, UK
philh@cogs.susx.ac.uk

Tom Smith
CCNR
University of Sussex
Brighton, UK
toms@cogs.susx.ac.uk

ABSTRACT

This paper argues that evolutionary robotics (ER) techniques can act as useful and potentially wide ranging tools in the scientific investigation of adaptive behaviour. After discussing the kinds of investigations ER can play a central role in, a concrete example is presented. We conclude that these kinds of studies are not only scientifically useful, but are necessary for the field to develop as an engineering methodology for autonomous robotics.

1 Introduction

Evolutionary robotics (ER) involves evaluating, over a number of generations, whole populations of autonomous robot control systems specified by artificial genotypes. These are interbred using a Darwinian scheme in which the fittest individuals are most likely to produce offspring. Fitness is measured in terms of how good a robot's behaviour is according to some task-based evaluation criterion. This particular flavour of new-wave autonomous robotics was originally proposed as an automatic alternative to hand design of control systems e.g. [12]. The field has largely proceeded by throwing up many individual instances of evolved controllers for fairly simple behaviours. These examples have usually been contrived to fit in with the sensorimotor constraints of a particular robot and have been based on individual researchers' favourite style of control system.

We argue that the field has now built up enough experience, tools and methods to engage in more principled exploratory work. This paper describes a number of types of investigations that can be undertaken and gives a concrete example that we are currently engaged in. ER can be used to address, for instance, questions about necessary and sufficient mechanism underlying the generation of adaptive behaviour. We believe that such investigations can make real contributions to the science of adaptive behaviour, and that they are also crucially important to the development of ER as an engineering methodology. Without them the field will die.

2 What Types Of Questions?

Since Artificial Evolution (AE) is most commonly thought of as an optimization technique, it is difficult at first glance to appreciate how it may be used to ask questions about search spaces beyond the simple 'what is the optimal solution?'. In this section we look at three different questions AE can be profitably used to explore. In the next section we look at several classes of spaces underlying adaptive behaviour that we would like to ask such questions about.

The first type of question that AE can be used to explore concerns large unconstrained search spaces which we would like to get some general feel for. In this case AE can be used as a principled sampling technique that can provide information as to the nature of the solutions that may be *easily* found in the space. If we hypothesise that these will be solutions of a certain type, then this is confirmed if at least some of a series of runs of a GA find solutions of this type, and disconfirmed if none of them do. Note that this is not the same as saying that AE can be used to test whether solutions of a particular type exist *at all* within a space; the fallible nature of stochastic search in large unconstrained fitness landscapes means that a negative result implies no more than that solutions of a particular type are hard to find (possibly because they don't exist).

The second type of question that AE can be used to explore concerns more constrained search spaces in which we can be fairly confident that AE will always converge on solutions that are near-optimal. In this case AE can be used as a principled sampling technique that can provide information about what the 'better' solutions within a particular space might look like. If we hypothesise that these solutions will be of a certain type, then this is confirmed if at least some of a series of runs of a GA find solutions of this type, and disconfirmed if none of them do. Note that this is not the same as trying to use AE to find *the* optimal solution within a particular search space, which it may or may not do during this sort of exploration. The optimization properties of AE, in these circumstances, are just being used to get a general idea of the sort of solutions that a particular space is capable of.

The third type of questions that AE can be used to explore concern the shape and nature of the space itself. As such, this sort of question can be asked about both large unconstrained search spaces and smaller search spaces we have a greater knowledge of. For example, if we hypothesise that a search space contains many good solutions, all of similar fitness, then this is confirmed if each of a series of runs of a GA quickly converges on a different good solution, all of similar fitness. If all of the solutions that the GA converges on are of different fitnesses, or if it always converges on the same solution, then the hypothesis is false.

3 What Types of Spaces?

Having examined the sort of questions that AE can be used to investigate, this section looks at some of the different types of search space underlying systems capable of adaptive behaviour that we would like to ask such questions about. In each of the subsections below, a different search space is examined, and examples of how each of the three types of questions listed above can be applied to that search space are discussed. As we shall see in section 5, some of these examples refer to pieces of work that have already been performed or are in the process of being actively researched. The others should be seen as directions for future research.

3.1 Nervous-system Mechanisms

One area that Artificial Evolution can profitably be used to explore is the space of what we shall here refer to as 'nervous-system' mechanisms. Loosely speaking, these are the mechanisms that are responsible for generating output signals to the actuators in response to input signals from the sensors. In mobile robots many different paradigms, from artificial neural networks to hand-designed rule-based controllers have been employed. Below we look at examples of how the three questions of section 2 may be asked about these sorts of mechanisms.

Recent results in neuroscience suggest that the simple picture of a neuron as an electrical processing unit may be manifestly insufficient [21]. This leads to the hypothesis that there may be important dynamics present in real neural networks for the generation of adaptive behaviour which the simple neuron models that are currently being used for the control of mobile robots are incapable of instantiating. One way of approaching the question of what these dynamics might look like is to examine the sort of networks that can be produced by neural models with very different functionalities to those based on the integrate-and-fire paradigm. Artificial Evolution provides a principled way of doing this. Different neuron functionalities may be tested by allowing AE to evolve a series of networks of such units for a given task and analysing the resulting range of fitnesses and behaviours. This allows us to test hypotheses about the sort of networks that can *easily* be found within the space, and thus to get a feel for the ways in which the dynamics of the neurons may be combined to solve the problem. As such, evolution is being used as an exploratory tool to assist in answering questions of the first type talked about in section 2. An example of this sort of research is the work of Husbands et al on the possible dynamics of networks that use diffusing gases along with electrical impulses to effect interactions between neurons. This is discussed more fully in section 5.

A form of the second type of trial and error exploration mentioned in section 2 is performed all the time by a wide range of researchers in adaptive robotics. It occurs when a new architecture has been designed and the parameters are then tuned by trial and error until an optimal or near optimal parameter set is reached. Only then can the designer's hypotheses about the full functionality of the new system be confirmed or disconfirmed. In this case, applying AE involves no more than a principled method of automating this trial and error process. Consider an example in which we want to know what happens when three layers of neurons, each of a different type, are connected together. The three layers might be a winner-takes-all layer connected to a Hebbian Layer connected to a Kohonen map layer for example. Now although we may have a good idea of what such an architecture is capable of, and the sort of tasks it will be able to allow a robot to perform, there will be so many parameters involved that setting them manually may be too time-consuming a task to be practical. By applying Artificial Evolution according to the methodology discussed below in section 4, however, we may automatically *evolve* the parameters to achieve settings that are at least as good, and probably much better, than those that could have been obtained by hand. We may then proceed to examine whether our original hypotheses about the tasks that such a three-layered architecture is capable of were in fact well-founded. Examples of this kind of work can be found in [4, 18].

The third type of exploration discussed in section 2 concerning the shape and nature of the fitness landscape, can have particular relevance to the search of a space of nervous-system mechanisms. Whether the landscape is rugged or smooth, for instance, can have a profound effect on the efficacy of other adaptive processes apart from evolution, such as life-time learning.

3.2 Sensori-Motor Morphologies

As well as the dynamics of the nervous-system, the physical dynamics of the embodied agent within its environment also play a major role in determining the overall behaviour of the system [23, 1]. For instance, the type and number of the sensors, and the way in which their geometry resonates with that of their environment, has a direct bearing on the modalities by which features of the environment can affect nervous-system input. On the output side, the physical dynamics of the actuators determine the ways in which the agent may interact with the environment to generate behaviour. As has been shown by several experimenters, the dynamics of such interactions may often be complicated and subtle enough to provide much of the 'processing' underlying adaptive behaviour that is all too commonly assumed to be generatable only by nervous-systems [24, 25]. There is, as yet, no principled technique for predicting the outcome of such interactions, or choosing the sensor or body morphology most suited to a given robot task. Artificial Evolution, however, provides a method of exploring spaces of sensori-motor morphologies that may provide insights other wise unobtainable. Below, we examine how the three types of exploration discussed in section 2 may be applied to sensori-motor morphology space.

An example of the first type of exploration is provided by the continuing work on evolving visual morphologies for the Sussex gantry robot first described in [11]. In this work, neural network controllers and the visual morphologies of their inputs are evolved simultaneously to make a visually guided robot perform a simple shape discrimination task. Artificial Evolution is here being used to explore how the geometries of evolving visual morphologies may causally interact with the geometries present in the environment to perform the task. Many novel mechanisms have been discovered through an extensive series of runs, and it is now becoming possible to say some general things about the possible solutions that exist in this rich space. Further details of the experiments are provided in section 5.

The second type of exploration discussed in section 2 may be applied to the space of sensori-motor morphologies in much the same way as for nervous-system mechanisms. For example, we may hypothesise that the emergent properties of superimposing arrays of Reichardt (movement) detectors [5] on top of each other, each with different spatial and temporal characteristics, may provide a mobile robot with the necessary optic flow information to perform a certain task. Again, the number of parameters involved may make setting them by hand impractical. Artificial evolution can thus be used to search parameter space for settings that are at least as good as those that could be derived by hand, and probably a whole lot better. We may then test the resultant visual morphology to see if our hypotheses were correct.

The third type of exploration discussed in section 2 concerns the shape and nature of the fitness landscape. Whether the landscape associated with a space of sensori-motor morphologies is rugged or smooth will again have consequences for the efficacy of adaptive processes other than evolution. If we want to use self-organising processes to fine tune evolved sensor morphologies, for instance, then we may need to know whether the areas around optima in the fitness landscape are smooth or rugged as far as the self organising operators are concerned. This can be explored through a series of runs of an evolutionary algorithm by sampling around the good solutions found using operators relevant to the self organising mechanisms.

3.3 Biological Models

As well as exploring mechanisms that may be used to underly adaptive behaviour in robots, Artificial Evolution may also be used to explore and test questions about adaptive behaviour in real animals. The sorts of questions that can be asked will be for the most part about *specific* behaviours and mechanisms, i.e. those of the real animal. Due to the fallible nature of AE, therefore, this sort of exploration will not always provide answers. If we are lucky, however, then we may be provided with insights into biological issues that would not have arisen otherwise. Below we look at examples of the three types of questions of section 2 as applied to spaces formed from biological models.

One of the reasons why we might want to do exploration of the first type discussed in section 2 is to see whether a behaviour that a particular space of mechanisms is *not* supposed to be able to generate does, in fact, lie within the space of behaviours generated by those mechanisms. This may then highlight the shaky foundations of some commonly held belief or dogma about how animals perform such behaviours. In [3], for example, the authors argue that the lack of internal representations within simple evolved visual machines implies that simple animals do not necessarily use internal representations to perform similar behaviours, as is commonly assumed. Further examples of using AE in this way to search spaces for counterexamples of commonly held beliefs, include showing how evolved robot controllers that select between actions do not necessarily employ the mechanisms that are commonly assumed to be present [22], and much of the work of Beer and colleagues [26, 2].

Exploration of the second type introduced in section 2 is more obviously applied to spaces generated by biological models. For example, a complex theoretical model might be hypothesised to explain a particular animal behaviour. If a simulation corresponding to the model could be made to exhibit appropriate behaviour then this would add considerable weight to the hypothesis. A simulation of this type, however, may again involve so many parameters that it is impractical to try setting them by hand. In this scenario, Artificial Evolution may be used to explore the parameter space and see whether the simulation can, in fact, be made to exhibit the behaviour in question. An example of this sort of hypothesis testing is provided in [14]. General hypotheses about the principles of biological nervous systems can also be explored in this way.

The shape and nature of the fitness landscape associated with a biological model, and how this relates to the model's 'evolvability', can be of obvious interest to the biologist since evolution is itself a biological phenomenon. Indeed, certain classes of biological model, such as those used to model secondary structure formation in RNA, have been invented precisely to study the performance of evolution within certain realistic fitness landscapes. Apart from analytic results pertaining to these models, one of the most important and obvious tools that can be used to study their dynamics is AE.

4 How ?

In this section we outline practical techniques and methods for exploring robot space using Artificial Evolution. There have been several recent innovations in this area which make the whole enterprise more feasible and these are discussed below. There is not room in this paper to explain any of these techniques in detail, however, and the reader should therefore regard this section as a guide to further reading rather than an instruction manual for applying AE to the exploration of robot space.

In general, the spaces we want to explore assume the role of phenotype space for the purposes of AE. In order to search these spaces therefore, we need a way of evaluating phenotypes, an encoding scheme by which phenotypes can be decoded from genotypes, and suitable genetic machinery (genetic algorithm, genetic operators and so on). Each of these three major constituents are examined in turn below.

4.1 Evaluation

The artificial evolution of control architectures for simple behaviours typically involves thousands of fitness evaluations and this can be a very time-consuming process. If these evaluations are performed on robots in the real world then they must be done in real time. If they are performed in simulation, then evolved controllers may not transfer into reality unless the simulation is so complex that all speed advantages are lost. Recently, Jakobi has proposed new ways of thinking about and building fast-running easy-to-design minimal simulations for the evaluation of robot controllers. This methodology is described in detail elsewhere [16], but below we offer a brief sketch here:

1. A small *base set* of robot-environment interactions that are sufficient to underly the behaviour we want to evolve must be identified and made explicit. A simulation should then be constructed that includes a model of these interactions. Since the base set will not contain all of the robot-environment interactions that can affect evolving controllers, some features of the simulation will have a basis in reality (the *base set aspects*), and some features will derive from the simulations implementation (the *implementation aspects*).

2. Every implementation aspect of the simulation must be randomly varied from trial to trial so that controllers are unable to rely on them to perform the behaviour. In particular, *enough* variation must be included so that the only practicable evolutionary strategy is to actively ignore each implementation aspect entirely.

3. Every base set aspect of the simulation must be randomly varied from trial to trial. The extent and character of this random variation must be sufficient to ensure that reliably fit controllers are able to cope with the inevitable differences between the robot-environment interaction model and reality, but not so large that they fail to evolve at all.

The power behind these ideas derives from the fact that we only have to model a sufficient number of real-world features, and these do not even have to be modelled particularly accurately. This means that such simulations can be easily constructed and made to run extremely fast. As long as the right amount of variation is included according to the methodology outlined above, controllers that evolve to be reliably fit will almost certainly transfer into reality.

4.2 Encoding-schemes

The choice of encoding scheme has a fundamental effect on the way in which AE searches phenotype space. If the wrong choice is made, AE can quickly be reduced to blind search. Below we adopt an engineering approach in our recommendations of which scheme to use to explore which types of space: the aim is to use a scheme that will find interesting phenotypes as quickly as possible.

The type of encoding scheme most suited to a particular problem depends fundamentally on whether the dimensionality of the search space to be searched is fixed or variable. If it is fixed, then a so-called direct encoding scheme is usually sufficient where each dimension of the search space corresponds to a specific location on the genotype. For phenotypes with a topology, such as neural networks, fixed amounts of phenotypic symmetry or repeated structure may be easily imposed by encoding more than one dimension of the search space at the same location on the genotype.

If we want to explore search spaces whose dimensionality is not fixed then life can become more difficult. If a direct encoding scheme is used then the nature of the phenotype may be dependent on the ordering of elements of the genotype. If so, small changes in genotype length may result in massive changes in phenotypes, rendering the fitness landscape so rugged that the efficacy of AE is vastly reduced. The choice of encoding scheme we recommend in this situation depends on the nature of the space that is being explored, but the key point is that the encoding scheme must be robust to the operators used by the genetic algorithm.

For phenotypes with a topology, we would recommend one of two types of encoding scheme. The choice of which of the two depends on whether the amount of phenotypic symmetry or repeated structure within the search space is fixed, or whether the amount of symmetry and repeated structure is one of the phenotypic aspects we want to use AE to explore. In the case of the former we would recommend a version of the simple spatially distributed encoding scheme described in [17]. Along with several previous encoding schemes [19, 13], this scheme uses the idea of a developmental space to make the shape of the phenotype invariant to the ordering of 'genes' on the genotype. It is, however, much simpler than previous schemes of this type, designed purely with engineering efficiency in mind. If the amount of phenotypic symmetry and repeated structure is one of the features of the space that we want to explore, on the other hand, then we would recommend using a version of Gruau's cellular encoding [8].

4.3 Genetics

For many who use genetic algorithms on more traditional optimization problems, the process of artificial evolution occurs as an initially random population converges upon a solution, gradually decreasing the genetic diversity until an equilibrium is reached. At this point, it is assumed that no significant change will occur and the process for all practical purposes is finished. If the search space to be explored is of fixed dimensionality then the same can be assumed: at some point the population will converge, the rate at which fitness increases will level off, and the run is over. In this sort of scenario, robot space exploration may be treated as a noisy optimization problem, and mutation, crossover and other operators may be applied using conventional genetic algorithm techniques [7].

As with encoding schemes (see above), the situation is made more complex if the space to be searched is not of a fixed dimensionality. This is because there is always the possibility that significant change can occur *after* the rate of genetic convergence has stabilized. Harvey has suggested that it is only after this initial convergence phase that the real business of open-ended artificial evolution may begin [10]. In the natural world, after all, evolution occurs as a dynamic equilibrium that adapts to environmental pressures in an open-ended way rather than a limited and finite search process with a start and a finish. If we want to do open-ended evolution of arbitrary complexity with variable length genotypes, Harvey suggests [9], then we should allow the evolutionary process to continue running long after the rate of genotypic convergence has stabilized.

There are a variety of ways in which genotypes can be allowed to change in length under evolutionary control. Probably the simplest is to employ operators that just add or delete genetic material with a small chance at each offspring event. Another method might be to allow crossover to occur at different points on each parent genotypes, thus producing two offspring genotypes of different lengths. Whichever method is used, however, no more than a slight change in length should be allowed to occur at each offspring event. Although it is important that there is sufficient generation and evaluation of new genetic material after the rate of convergence has stabilized, too much will overpower the selection pressure and reduce the adaptive abilities of the evolutionary processes to those of random search.

5 A Concrete Example

This section describes a series of experiments in which many of the kinds of questions introduced previously have been explored. All the experiments involve the evolution of a particular visual shape discrimination behaviour for the same robot. In each a control network and visual morphology (layout of visual sensors) have been concurrently evolved. Four different styles of control network have been used in the experiments. We have asked such questions as: are there commonalities between the evolved behaviours for the same task using different types of network? Are there commonalities in the underlying behaviour generating mechanisms? Are controllers for the task easier to evolve with some styles of network than with others? The robot and task will now be described before detailing the other aspects of the investigation and the results found.

5.1 The Robot

This series of investigations made use of the Sussex Gantry Robot. In each case controllers were evolved using a minimal simulation. As explained earlier, such radical simulations run much faster than real time and have played a crucial role in allowing us to repeat the evolutionary experiments a sufficient number of times to be able to start answering some of the questions outlined in the previous paragraph. Controllers evolved in minimal simulation work perfectly on the real robot. For details see [15, 16].

The gantry-robot is shown in figure 1. The robot body is cylindrical, some 150mm in diameter. It is suspended from the gantry-frame with stepper motors that allow translational movement in the X and Y directions, relative to a co-ordinate frame fixed to the gantry. Such movements, together with appropriate rotation of the sensory apparatus, correspond to those which would be produced by left and right wheels. The visual sensory apparatus consists of a CCD camera pointing down at a mirror inclined at 45^o to the vertical (see figure 2).

The mirror can be rotated about a vertical axis so that its orientation always corresponds to the direction the 'robot' is facing. For full details see [11]. The gantry is a very useful apparatus for controlled experiments in the evolution of visually guided behaviours, but is probably best thought of as if it were a two wheeled mobile robot with a fixed video camera mounted on top.

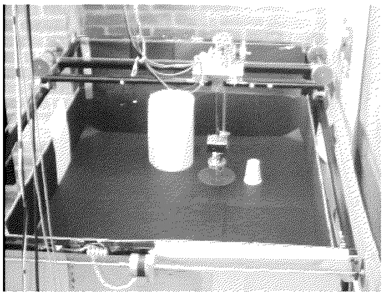

Figure 1 ***The Gantry viewed from above. The horizontal girder moves along the side rails, and the robot is suspended from a platform which moves along this girder.***

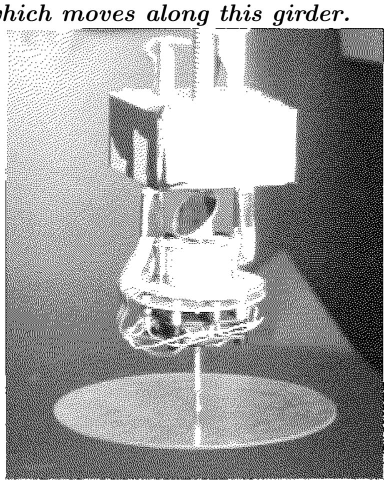

Figure 2 ***The gantry-robot. The camera inside the top box points down at the inclined mirror, which can be turned by the stepper-motor beneath. The lower plastic disk is suspended from a joystick, to detect collisions with obstacles.***

5.2 The Task

Control networks and visual morphologies were evolved for a target discrimination task. Two white paper targets were fixed to one of the gantry walls; a rectangle and an isosceles triangle with the same base width and height as the rectangle. Starting from a random position and orientation, the robot was required to move to the triangle while ignoring the rectangle. This was to be achieved under extremely variable and noisy lighting conditions in which the illumination intensity at any point in the gantry arena can vary by up to 100%. This was achieved by fixing a rig of spotlights above the gantry — the lights were randomly turned on and off at widely varying frequencies.

5.3 The Networks Investigated

The network size and topology, as well as various other properties detailed below, were under unconstrained evolutionary control in every experiment (i.e. arbitrarily recurrent networks were possible). So was the robot visual morphology, i.e. the way in which the camera image was sampled. This was achieved by genetically specifying the number and position of *single* pixels from the camera image to use as visual inputs. The grey scale intensity value of these pixels (normalised into range [0.0,1.0]) were fed into the network, one for each genetically specified visual input node in the net. This is illustrated in figure 3. Note this means that the evolved control systems were operating with extremely minimal vision systems, just a few single pixel values. Given the very noisy lighting conditions and the minimal visual input, this was a non-trivial task.

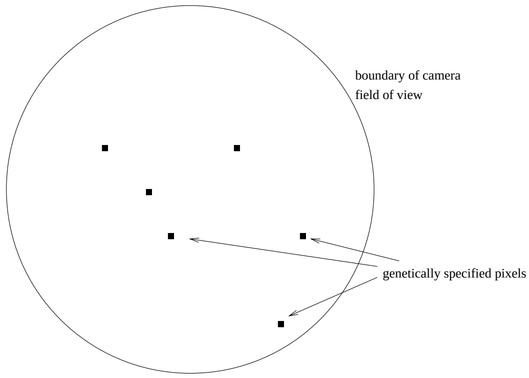

Figure 3 ***Evolved visual morphology. Visual input is taken only from the genetically specified single pixels. The rest of the camera image is thrown away.***

Veto Nets. These nets have been used at Sussex for various ER experiments over the past few years, including the original gantry work [11]. The nodes in this style of network use separate channels for excitation and inhibition. Real values in the range [0,1] propagate along excitatory links. The inhibitory (or veto) channel mechanism works as follows. If the sum of excitatory inputs exceeds a threshold, $T_v = 0.75$, the value 1.0 is propagated along any inhibitory output links the unit may have, otherwise a value of 0.0 is propagated. Any unit that receives a non zero inhibitory input has its excitatory output reduced to zero (i.e. is vetoed). In the absence of inhibitory input, excitatory outputs are produced by summing all excitatory inputs, adding a quantity of noise, and passing the resulting sum through a simple linear threshold function, $F(x)$, given below. Noise was added to provide further potentially interesting and useful dynamics. The noise was uniformly distributed in the real range [-0.1,0.1]. Each network had four motor neurons (left/right forward/backward).

$$F(x) = \begin{cases} 0, & \text{if } x \leq T_1 \\ \frac{x-T_1}{T_2-T_1}, & \text{if } T_1 < x < T_2 \\ 1, & \text{if } x \geq T_2. \end{cases} \quad (1)$$

Where T_1=0.0 and T_2=2.0.

Binary Nets. This style of network consists of nodes connected together by weighted links. Each unit used the transfer function given in equation 2. O_j is the output of the jth node and T_j is its threshold. The size and topology of the network was under evolutionary control, as were the connection weights, node thresholds and visaul morphology. As with the veto nets, four motor units were used. Thresholds were real numbers in the range [0.0,1.0], the weights, w_{ij}, were real numbers in the range [-2.0,2.0].

$$O_j = \begin{cases} 0, & \sum_i O_j w_{ij} < T_j \\ 1, & \sum_i O_j w_{ij} \geq T_j. \end{cases} \quad (2)$$

GasNets. The third class of networks used is completely new and inspired by the recent discovery that freely diffusing gases (in particular nitric oxide – NO) have important modulatory affects in nervous systems. Many nerve cells emit NO, under certain conditions, which then diffuses over relatively long distances. The presence of NO can trigger various cascades of chemical reactions that can have a wide range of affects on neuronal networks: changing intrinsic properties of nerve cells, changing synaptic transfer functions, turning hebbian style associative learning on and off, and many more [6]. These changes can occur at many different timescales. The Gas Nets described here are an abstraction of some of these interacting dynamical processes that give rise to forms of plasticity very different from those normally considered in ANNs. Two forms of Gas Net were used in this study. They are briefly described below. For fuller details see [20].

The 4 Gas Model. This style of networks consist of units connected together by excitatory links, with a weight of +1, and inhibitory links, with a weight of -1. The output , O_j, of a node j is a function of the normalised sum of its inputs, S_j, as described by equation 3. In addition to this underlying network in which positive and negative 'signals' flow between units, an abstract process loosely analogous to the diffusion of gaseous modulators is at play. Some units can emit 'gases' which diffuse and are capable of modulating the behaviour of other units by changing their transfer functions in ways described below. This form of modulation allows a kind of plasticity in the network in which the intrinsic properties of units are changing as the network operates. The networks function in a 2D plane; their geometric layout is a crucial element in the way in which the 'gases' diffuse and affect the properties of network nodes.

$$O_j = f(\overline{S_j}) \quad (3)$$

Where,

$$\overline{S_j} = \frac{(\sum_{p \in P_j} O_p - \sum_{n \in N_j} O_n + \sum_{k \in SEN_j} I_k)}{(np_j + nn_j + ns_j)} + R \quad (4)$$

In equation 4, P_j is the set of network elements with excitatory connections to element j. Likewise, N_j is the set of elements with inhibitory link to j, and SEN_j is the set of sensors connected to j. np_j, nn_j and ns_j are , respectively, the number of positive, negative and sensor connections to element j. R is the default activation of a node ($= 0.05$). Normalizing by dividing by the number of inputs keeps the summed input in the range [-1,1]. The transfer function, f, is defined in equation 5, its output range is [-1,1] given the restriction on the input range.

$$f(x) = \begin{cases} 0, & \text{if } x < 0 \text{ and } (a < 0 \text{ or } b < 0) \\ (x^a + x^b)/2, & \text{otherwise} \end{cases} \quad (5)$$

Where,

$$a, b \in PP = \{0.1, 0.2, 0.3 \ldots 0.8, 1, 2, 3 \ldots 9, 10\} \quad (6)$$

A wide range of output responses to a given input are possible, depending on the values of the parameters a and b. Default values of a and b for each node are set genetically, but are changed by diffusing gases as the network runs. It is genetically determined whether or not a node will emit one of four gases, and under what circumstances emission will occur (either when the 'electrical' activation of the node exceeds a threshold, or the concentration of one of the gases, genetically determined, in the vicinity of the node exceeds a threshold). For an emitting node, the concentration of gas at distance d from the node is given by equation 7. Here r is the genetically determined radius of influence of the node, so that concentration falls to zero for $d > r$. $TC(t)$ is a linear function that models build up and decay of concentration after the node has started/stopped emitting. The slope of this function is individually genetically determined for each emitting node, see [20] for full details. C_0 is a global constant.

$$C(d,t) = C_0 \times e^{\frac{-2d}{r}} \times TC(t) \quad (7)$$

The gas concentrations modulate the intrinsic properties of nodes in the network by changing the values of a and b. At every time step these values are updated as follows: at each node the value of a is linearly increased from its genetically set default by an amount proportional to the concentration of chemical 1 at the node, a is similarly decreased according to the concentration of chemical 2, b is changed in the same way by chemicals 3 and 4. For full details see [20]. The number of nodes, topology of connections and geometric layout of the nets are evolved, as are the gas emitting properties of nodes and the default values of a and b. Four motor nodes are used as with the previously described nets, and, of course, the visual morphology was concurrently evolved along with networks.

The 2 Gas Model. This style of network is very similar to the 4 gas model. This time there are only two gases and the transfer function at each node is of the form shown in equation 8[1]. Again the weights are restricted to be either +1 or -1. The value of b is genetically set at each node as is the default value of k. The two gases raise and lower the value of k in a similar manner to the way a and b are changed in the 4 gas model.

$$O_j = tanh([k \times \sum_i O_i w_{ij}] + b) \tag{8}$$

5.4 Results

As intimated earlier, the explorations described in this section are ongoing. We do not yet claim to have enough data to make strong and rigorous statements. However, we have now done between 5 and 15 runs for each kind of network and a number of interesting and suggestive things have already emerged. Figures 4–8 show typical evolved networks and visual morphologies for the different styles of network investigated. The genetic encodings used for the binary nets and the GasNets were very similar (see [15] and [20] respectively for full details). However, a simpler encoding, that was not suitable for the other styles of net, was used for the veto nets [11]. Distributed GAs with local selection and populations of size 100 were used for the GasNets. A slightly different version of the GA was used, with the same population size, for the binary networks. Our observations will now be briefly outlined.

Probably the most striking thing to emerge from our study is the fact that *all* of the successful[2] evolved controllers (more than 50 to date) employ one or both of *only two* (closely related) behavioural strategies. This even though a wide range of networks was used and the encodings and other aspects of the evolutionary machinery varied. The two strategies are illustrated in figure 9. The first strategy, illustrated to the left of the figure, involves moving until one of two strategically positioned visual inputs gives a high signal while the other gives a lower signal. The geometric layout of the sensors is such that this will only be reliably achieved when the robot is facing towards the triangle. The controllers illustrated in figure 4, 5, 6 and 8 use this strategy. The other strategy involves two vertically aligned visual sensors and is illustrated to the right of the figure. As the robot swings round towards a target, the bottom sensor will go high significantly earlier than the top sensor in the case of a triangle, but not for a rectangle. The controllers shown in figure 7 use this strategy which is based on the order in which sensors are excited. There is not enough space to describe the workings of the controllers in detail. However, it should be noted that many of them used a small number of additional visual inputs and various subtle internal dynamics to generate highly robust behaviours capable of coping with the extreme lighting conditions. In each case the visual morphology played a vital role. In each successful controller there was a perfect balance between the sensor geometry and robot motion resulting in active visual strategies. A traditional cognitive science perspective would think of the sensori capabilities as being passive and the sensor morphology as almost incidental; it is the internal processing where the real work is done. This is very clearly not the case in *any* of our evolved robots. The number and position of the visual inputs was under evolutionary control; it has clearly been demonstrated that very simple extremely low bandwidth sensors, when appropriately coupled to a dynamic controller, are sufficient for this kind of task.

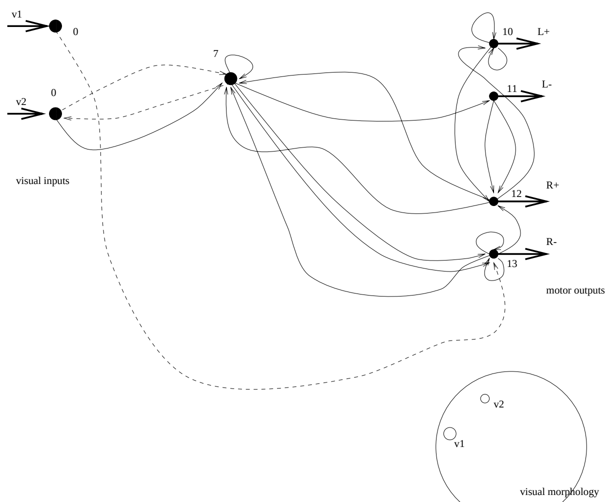

Figure 4 ***Evolved network and visual morphology for veto networks. Active part of network is shown. Solid lines are excitatory, dashed are veto.***

Another very striking observation we have made from the results available so far is to do with the speed of evolution of successful controllers. Evolving in a minimal simulation is in some sense harder than in reality (because of the extreme use of noise). Hence it would be expected that a greater number of evaluations would be needed to evolve a behaviour in such a simulation than in reality, although of course this is heavily offset by the speed at which the simulations run. The successful binary net and veto net controllers all took 6,000 or more generations to evolve. The successful 4 gas GasNet

[1] Thanks to Andy Philippides for suggesting this form.

[2] To count as successful a controller must move to the triangle on many (at least 30) *successive* trials on the real robot under full noisy lighting and with random relative positioning of the two targets on the gantry wall.

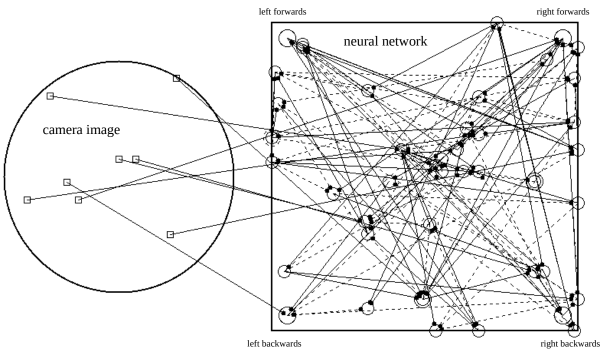

Figure 5 ***Evolved network and visual morphology for binary networks. Solid lines are excitatory, dashed are inhibitory.***

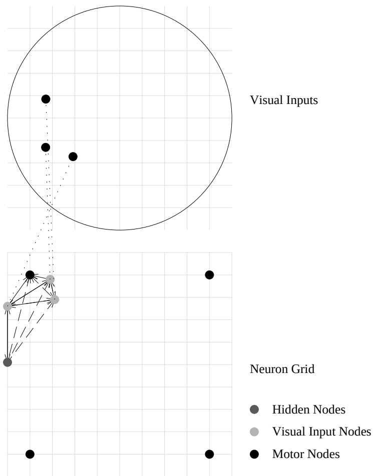

Figure 6 ***Evolved network and visual morphology for 4 gas GasNet. Solid lines are excitatory, dashed are inhibitory.***

based controllers all took between 1,000 and 3,000 generations. However, many of the successful 2 gas GasNet based controllers took less than 500 generations (with a few requiring up to 3,000 generations). We have data from more than 15 runs for the 2 gas GasNets, but less for the other styles. So what we can say at the moment is that we have strong suggestive evidence that we can evolve behaviours quicker (usually far quicker), in terms of numbers of evaluations needed, with GasNets than with the other styles tried. This is backed up by studies on different behaviours with another robot [20]. It seems that the 'behaviour space' that this class of nets

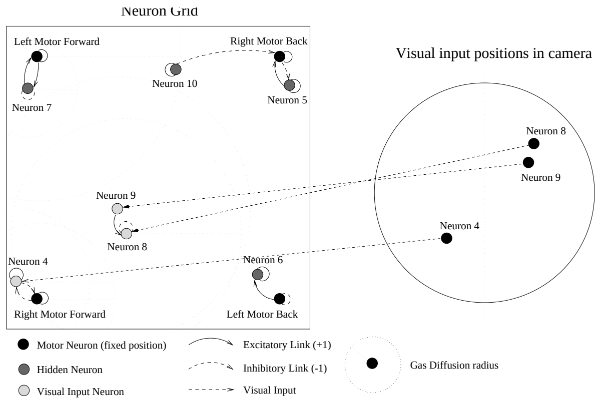

Figure 7 **Closed-loop 'tracking' two-gas model triangle finding network, see text for details. NB gas radii are shown only where used.**

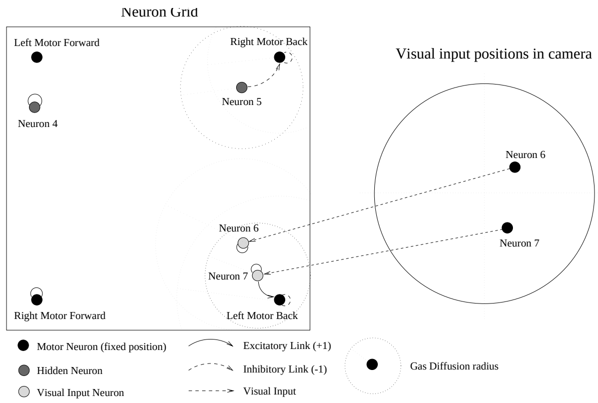

Figure 8 **Closed-loop 'tracking' two-gas model triangle finding network, see text for details. NB gas radii are shown only where used.**

generates is more dense. It is easier to find a route to a successful controller. This suggests that networks involving interacting, yet distinct, processes are a powerful alternative to more conventional connectionist thinking.

Something that quickly becomes obvious after glancing at figures 4– 8 is the difference in structural complexity of the evolved networks. The binary nets were the most structurally complex, while *all* successful evolved GasNet controllers were amazingly minimal. The veto net controllers were of intermediate complexity. Although the GasNets were very minimal, the modulatory interaction of the spreading and decaying gases and the sparse 'electrical' networks gave rise to the most sophisticated internal dynamics of any of the network classes.

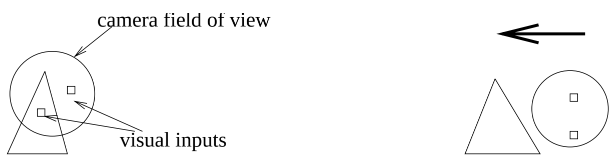

Figure 9 ***The only two classes of successful behavioural strategy that we have observed to date.***

Clearly it has been demonstrated that for each of the network types it is possible to find a successful controller. Importantly, we have learnt that a very unconventional class of network, namely the GasNets, are capable of generating adaptive behaviour and have many interesting properties. They look like a very good avenue to explore further.

There are many other aspects of these experiments that need further investigation, such as the role of the genetic encoding and the details of the GA. However, we feel that we have already gained valuable knowledge and insights that are feeding into current work in which we are attempting to take evolutionary robotics onto a higher plane of behavioural sophistication.

6 Conclusions

We have claimed that we now have enough tools and methods to employ evolutionary robotics techniques in a wider realm of exploration than has been practiced to date. We feel this is necessary if we are to make significant engineering advances in the field, but will also contribute to the science of adaptive behaviour. We argue that ER can act as a useful and potentially wide ranging tool in the scientific investigation of adaptive behaviour generating systems.

Acknowledgements: Thanks to Marco Dorigo, Emmet Spier, Jon Bird, Mike Wheeler, John Anderson and various members of CCNR for useful comments.

References

[1] R.D. Beer. A dynamical systems perspective on agent environment interaction. *Artificial Intelligence*, 72(1-2):173–215, 1995.

[2] R.D. Beer. A dynamical systems perspective on autonomous agents. *Artificial Intelligence*, 72:173–215, 1995.

[3] D. Cliff and J. Noble. Knowledge-based vision and simple visual machines. *Philosophical Transactions of the Royal Society, series B*, 352:1165–1175, 1992.

[4] Floreano D. and Mondada F. Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man and Cybernetics–Part B: Cybernetics*, 26(3):396–407, 1996.

[5] N. Franceschini, J-M. Pichon, and C. Blanes. Real time visuomotor control: from flies to robots. In *Proceedings of: International Conference on Advanced Robotics, Pisa*, 1991.

[6] J. Garthwaite. Glutamate, nitric oxide and cell-cell signalling in the nervous system. *Trends in Neuroscience*, 14:60–67, 1991.

[7] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, Massachusetts, 1989.

[8] F. Gruau. *Neural Network Synthesis using Cellular Encdoing and the Genetic Algorithm*. PhD thesis, Ecole Normale Superieure de Lyon, 1994.

[9] I. Harvey. Species adaptation genetic algorithms: the basis for a continuing saga. In F. J. Varela and P. Bourgine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 346–354, Cambridge, Massachusetts, 1992. M.I.T. Press / Bradford Books.

[10] I. Harvey. *The Artificial Evolution of Adaptive Behaviour*. PhD thesis, School of Cognitive and Computing Sciences, University of Sussex, 1993.

[11] I. Harvey, P. Husbands, and D. Cliff. Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J.A. Meyer, and S. Wilson, editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, volume 3. MIT Press/Bradford Books, 1994.

[12] P. Husbands and I. Harvey. Evolution versus design: Controlling autonomous robots. In *Integrating Perception, Planning and Action, Proceedings of 3rd Annual Conference on Artificial Intelligence, Simulation and Planning*, pages 139–146. IEEE Press, 1992.

[13] P. Husbands, I. Harvey, D. Cliff, and G. Miller. Artificial evolution: A new path for ai? *Brain and Cognition*, (34):130–159, 1997.

[14] A.J. Ijspeert, J.Hallam, and D. Wilshaw. Artificial lampreys: Comparing naturally and artificially evolved swimming controllers. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 256–265, Cambridge, Mass, 1997. MIT Press.

[15] N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6:(in press), 1997.

[16] N. Jakobi. *Minimal Simulations for Evolutionary Robotics*. PhD thesis, COGS, 1998.

[17] N. Jakobi and M. Quin. Some problems (and a few solutions) for open-ended evoltuionary robotics. In P. Husbands and J-A Meyer, editors, *Proceedings of Evorobot*. Springer Verlag, 1998, forthecoming.

[18] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 190–197. MIT Press/Bradford Books, 1994.

[19] S. Nolfi D. Parisi. Evolving artificial neural networks that develop in time. In F. Moran, A. Moreno, and J.J. Merelo, editors, *Advances in Artificial Life: Proceedings of the third European conferrence on Artificial Life.*, Berlin, 1995. Springer-Verlag.

[20] P.Husbands, T. Smith, N. Jakobi, M. O-Shea, A. Philippides, and J. Anderson. Brains, gases and robots. In *Proc. ICANN'98*, page (in press). Springer-Verlag, 1998.

[21] D. Purves and George Augustine. *Neuroscience*. Sinauer, 1997.

[22] A. Seth. Evolving action selection and selective attention without actions, attention, or selection. In R. Pfeiffer, editor, *From Animals to Animats 5, Proc. of 5th Intl. Conf. on Simulation of Adaptive Behavior, SAB'98*, page (in press). MIT Press/Bradford Books, 1998.

[23] T. Smithers. What the dynamics of adaptive behaviour and cognition might look like in agent-environment interaction systems. In T. Smithers and A. Moreno, editors, *3rd International Workshop on Artificial Life and Artificial Intelligence, The Role of Dynamics and Representation in Adaptive Behaviour and Cognition*, San Sebastian, Spain, 1994.

[24] B. Webb. Modeling biological behaviour or 'dumb animals and stupid robots'. In *Proceedings of the Second European Conference on Artificial Life*, pages 1090–103, 1993.

[25] R. Wehner. Matched-filters - neural models of the external world. *Journal of Comparative Physiology*, 161(4):551–531, 1987.

[26] Brian M. Yamauchi and Randall D. Beer. Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, 2(3):219–246, 1994.