

Brave Mobots use Representation: Emergence of Representation in Fight-or-Flight Learning

Chris Thornton
Cognitive and Computing Sciences
University of Sussex
Brighton
BN1 9QH
UK

Email: Christopher.Thornton@firenet.uk.com
Tel: (44)1273 678856

May 21, 2003

Keywords: Representation, Learning, Emergence

Abstract

The paper uses ideas from Machine Learning, Artificial Intelligence and Genetic Algorithms to provide a model of the development of a ‘fight-or-flight’ response in a simulated agent. The modelled development process involves (simulated) processes of evolution, learning and representation development. The main value of the model is that it provides an illustration of how simple learning processes may lead to the formation of structures which can be given a *representational* interpretation. It also shows how these may form the infrastructure for closely-coupled agent/environment interaction.

1 Introduction

In recent years there has been a rapid development in our understanding of the detailed mechanisms underlying the emergence of intelligent behaviour. Research on Machine Learning (ML) has introduced powerful new models of lifetime learning [1]. Research on Genetic Algorithms (GAs) and Artificial Life (ALife) has advanced our understanding of the role played by evolution [2]. And research in Artificial Intelligence (AI) and psychology has provided powerful illustrations of the way in which the emergence of intelligent behaviour

depends on the development of representational structure [Muggleton Inductive Logic Programming 1992; 3]. Unfortunately, despite this progress, the ‘big picture’ is still elusive. Ideas from one community do not interface easily with ideas from another.

If we look at the relationship between learning theories and representation theories, for example, we find that the learning processes investigated by the Machine Learning community tend *not* to produce ‘representational structure’ of the type commonly studied by the AI community. In fact, very often, learning processes are not constructive in any sense. Where they do involve the generation of structure, it tends *not* to have any obvious representational interpretation. The interface between models of learning and models of evolution is equally problematic. Many of the more powerful models of learning tend to be supervised in nature. This means they assume the existence of an all-knowing ‘teacher’ — an assumption which is hard to accommodate within standard evolutionary scenarios.

The present paper launches an all-out assault on this problem. It introduces a novel learning method (‘explicitation’) and shows how its unsupervised, constructive nature allows it to be used to link learning with both evolutionary and representation-development scenarios. The main body of the paper provides a demonstration of this. It presents a computational model of the development of a ‘fight-or-flight’ behaviour in a simulated agent (an ‘animat’). It shows how evolution, learning, agent/environment interaction and representation-development effectively ‘cooperate’ so as to achieve a particular behavioural competence.

2 The fight-or-flight task

In the model presented by the paper, a simulated agent (an ‘animat’) develops the ability to produce a fight-or-flight response to a certain type of attack. The agent which forms the main focus of the model is known as the ‘prey’. This agent has the sensory-motor properties of a simulated mobot [cf. 4] similar to the widely used ‘Khepera’ mobot [5]. The prey has two wheels situated at either end of a transverse axle, and a castor at the rear to provide a stable, tripod wheelbase. The wheels can be driven forwards or backwards at differing speeds thus achieving any blend of forwards, backwards and rotational movement.

Within the model, the prey agent interacts with another, mobot-like agent, known as the ‘predator’. The interactions take place in a simple artificial world which is a flat, continuous, rectangular arena. The arena always contains one prey agent and one predator agent. In Figure 1 (a), these two agents are shown as circular shapes. The larger circle represents the predator. The smaller circle represents the prey. Each agent has an arrow which shows its current direction of motion.

The prey has a single range-finding sensor and a one-cycle memory (i.e., it is able to remember its previous sensor input). In each time cycle of the simulation, the sensor produces a real value between 0 and 1 which varies monotonically with the *proximity* of the nearest obstacle along a ray pointing directly ahead, see

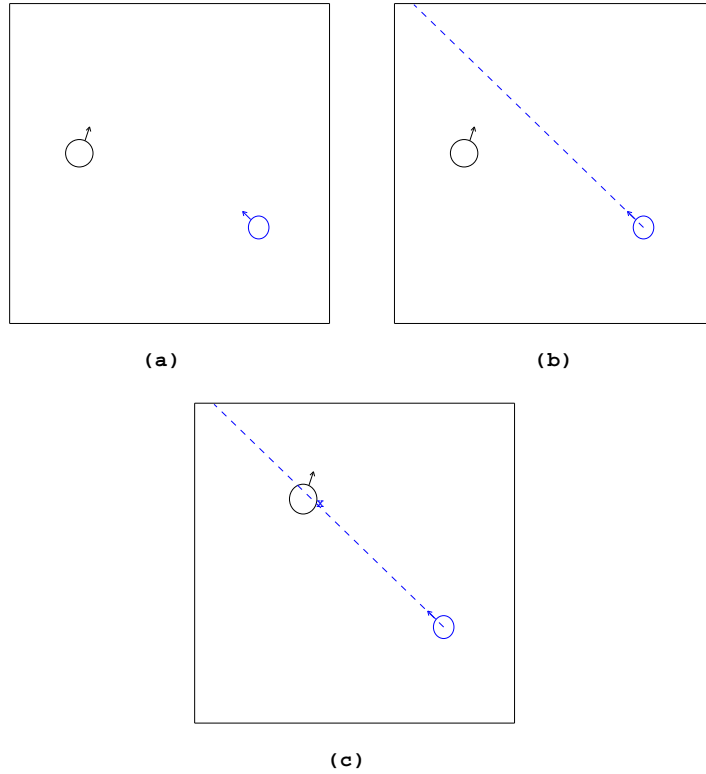


Figure 1: Basic agent scenarios.

Figure 1(b). The border of the arena does not constitute an obstacle and thus never affects the sensed proximity value. However, it may reflect the proximity of the other agent if the position of that agent falls somewhere along the sensor ray, see Figure 1(c). If no obstacle intersects the sensor ray, the sensor returns a zero value. The larger of the two agents, known as the **predator**, has a variety of sensors. However, since the model is concerned with the development of behaviour in the prey, these are not relevant and will not be discussed.

Within the simulation, the predator's goal is to destroy the prey as quickly as possible. The prey's goal is to survive while, at the same time, minimising movement. The predator attempts to achieve its goal by implementing a 'search-and-destroy' strategy. If it finds itself facing away from the prey, it turns towards it; see Figure 2(a) and (b). (Note that in these displays, the lightly shaded agent shapes mark out previous positions, and thus constitute a visible trail.) If it finds itself facing directly towards the prey, it moves forward: Figure 2(c). It may approach the prey at either a steady or accelerating pace. If the former, then it will swerve away at the last moment, provided the prey is facing directly into the attack, see Figure 2(d). If it approaches at an accelerating pace, it

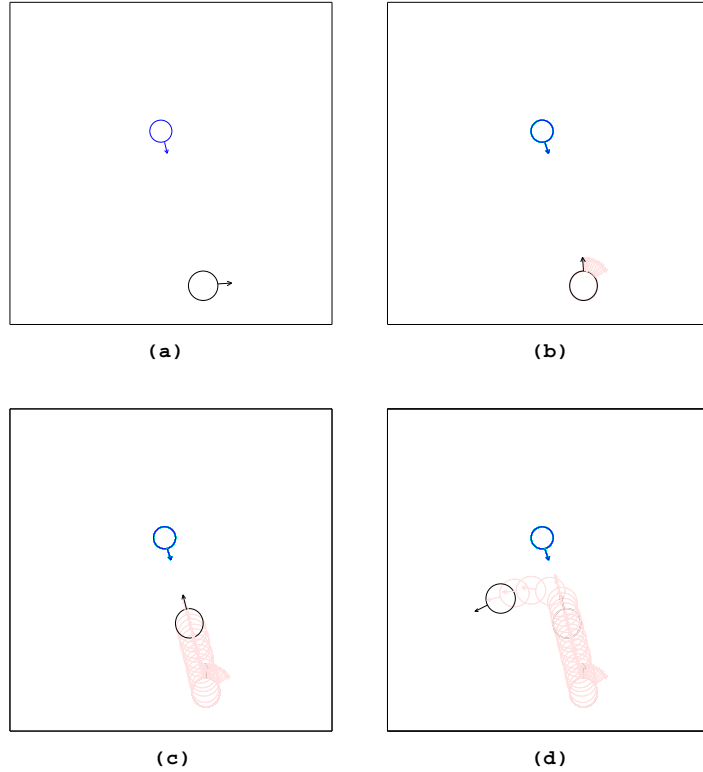


Figure 2: The basic predator/prey scenario.

will continue moving forwards regardless of where the prey is facing, eventually colliding with, and destroying it (not shown). Whenever a predator approaches, the prey should thus turn towards the approach and remain static. If the rate of approach turns out to be increasing, the prey should flee. The prey is thus required to implement a rudimentary ‘fight-or-flight’ strategy.

3 Previous work on fight-or-flight

What developmental or evolutionary story might account for the emergence of this sort of fight-or-flight response? To date, work on fight-or-flight has largely been confined to the ethology community. This research has focussed on the *implementation* of the behaviour rather than its development [cf. 6; 7]. There is also work of relevance in the animal learning literature but this is concerned with the phenomenon of behavioural ‘freezing’ rather than fight-or-flight itself [cf. 8; 9]. Where the issue of the development of behavioural responses is confronted explicitly [cf. 10], there are few operational models and none as yet that deal

explicitly with fight-or-flight.

Generic work on the evolution of intelligent behaviour has typically focussed on the genetic algorithm model of Holland [11; 12] and the classifier system of Wilson [13]. Work on behaviour learning has concentrated on algorithmic models such as recursive decision-tree generation [14] or neural-network models such as backpropagation [15] or competitive learning [16]. However, there has been no previous attempt to produce a computational model which shows the role played by evolutionary, learning and representation-construction processes in the development of this behaviour.

4 Evaluating standard learners

Can the development of the fight-or-flight behaviour in agents be modelled *solely* in terms of lifetime learning? To investigate this, various learning methods were tested for their ability to acquire the fight-or-flight response. Experiments were conducted using C4.5¹ and standard backpropagation. Both of these are supervised methods and thus require explicit training examples. In the experiments carried out a training set of 1000 input/output pairs was used and these were derived direct from a running simulation.

The input in each pair was a vector of numbers corresponding to the prey's current sensory input and its corresponding input in the previous cycle. (The second input constitutes the learner's 'memory' of the first.) The output in each case specified the wheel motions for an appropriate execution of the fight-or-flight behaviour. Thus, in those cases where the input vector was derived in the context of an oncoming, accelerating predator, the target outputs specified the wheel motions for a flight response. In the testing of backpropagation standard parameters for learning rate and momentum (0.01 for learning rate and 0.9 for momentum) were used in an ordinary feed-forward architecture of three, fully interconnected layers. Architectures involving from 3 to 30 hidden units in the middle layer were tested, but it was found that the number of hidden units made little difference to the final performance.

The two methods were both tested for their ability to generalise. The quality of generalisation with respect to unseen cases was examined, as was the degree to which the learning methods were able to reproduce the fight-or-flight response in a replication of the original simulation. In both cases, both methods performed poorly. The average results obtained are summarised in Table 1. The error shown here is RMS error on a testing set of 1000 cases. The 'Deaths' column shows the number of deaths sustained by a trained agent in a simulation of 2000 time steps, in which the predator agent produced 12 aggressive (accelerating) approaches.

¹C4.5 is an improved version of the well-known ID3 algorithm [14].

	Error	Deaths
C4.5	0.209	12
Backpropagation	0.670	8

Table 1: Performance of C4.5 and backpropagation on fight-or-flight.

5 A hand-crafted implementation

The failure of the standard learning methods to learn the fight-or-flight task might be due to the fact that this task is just ‘too hard’ to learn. But this seems unlikely since the behaviour is actually fairly easy to hand-craft. The essence of the fight-or-flight response is the discrimination between two *rates* of approach. For an agent with proximity sensors, an approach is essentially a pattern of increasing proximity values. An *accelerating* approach is a pattern in which the difference between the values is increasing through time.² So any implementation for this behaviour must discriminate these two types of effect.

Let us consider how we might develop the implementation as a neural network. First, we must implement the agent’s one-cycle memory and an easy way to do this is to arrange for the input stream to have a built-in ‘echo’. We arrange for the input data to be a set of two-element vectors, such that each vector is made up of the current and previous sensor input. Given this input stream, our network must obviously have two input neurons. But how to arrange the rest of the network? One possibility is illustrated in Figure 3. In this network, the right wheel motor has a forwards bias and the left motor has a negative bias. There is also an internal neural network in which activation flows upwards (or ‘backwards’ through the mobot). The single proximity sensor is connected by an excitatory connection to the left input neuron of the network (the right input neuron contains the previous proximity input) and there are inhibitory connections from the main input neuron to both motors.

The internal structure of the network functions to provide the detection of accelerating approaches. The neurons in the first layer of the network determine whether the difference between the current proximity (labelled $p@t$) and previous proximity ($p@t-1$) is above or below some given value. The two neurons actually shown in the figure have one positive and one negative weight each. The absolute difference between the weights is 0.1 but the signs are different in the two cases. The left-hand neuron will thus be activated just in case the difference between the two proximity values is less than 0.1 (the label is therefore $d<0.1$) and the right-hand neuron will be activated just in case the difference between the two proximity values exceeds 0.1 (the label is thus $d>0.1$). When the difference is exactly 0.1, the two neurons are thus *both* inactive. Since they both have inhibitory connections to the neuron shown in the third layer, and

²Formally, an approach is a **first-order** dynamic effect (i.e., a change in a variable over time) while the accelerating approach is a **second-order** effect (i.e., a change in a *change* in a variable over time).

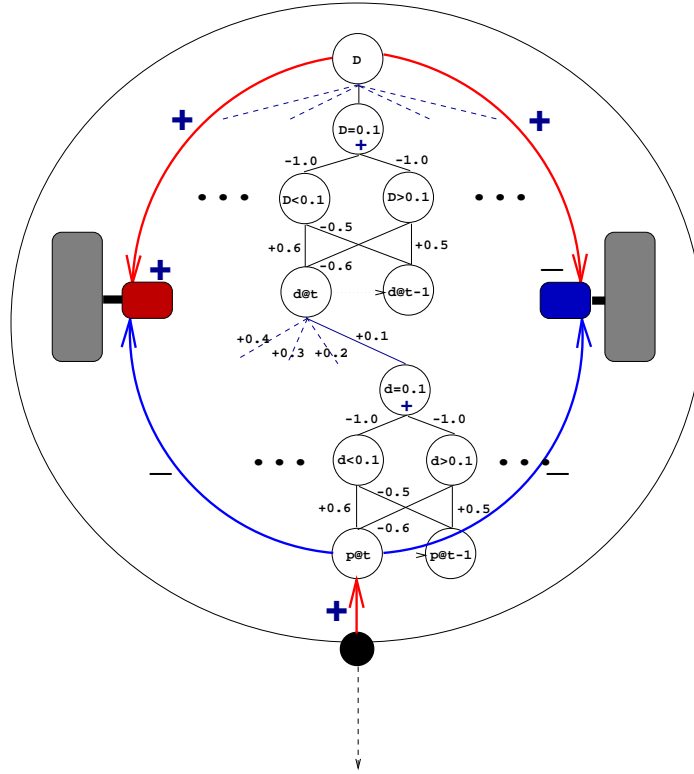


Figure 3: A hand-crafted fight-or-flight architecture.

since this neuron has a positive bias, it will always be activated except when the difference between the two proximity values is exactly 0.1. It thus forms a detector for this difference and has the label $d=0.1$.

This basic difference-measuring architecture is repeated (but not shown) for all other relevant difference values (0.2, 0.3, 0.4 etc.). Suitably weighted connections are then used to ensure that the activation of the left-hand neuron in the fourth layer precisely reflects the real difference between the two proximity values. The first four layers of the network thus function to measure the first-order dynamic effect we are concerned with; i.e., the current rate of approach as measured in terms of the difference between the current and previous proximity input.

Our ultimate aim is, of course, that the network should measure the *change* in the rate of approach, i.e., the acceleration. Fortunately, this can be accomplished in almost exactly the same way we measured the change in the proximity. The top four layers of the network repeat the strategy implemented by the bottom four layers — only now using the measured difference value (i.e., rate of approach) as the basic input. The node containing the measured rate of

approach is labelled \mathbf{dot} and the node containing the previous difference value is $\mathbf{dot-1}$. The top neuron of the network effectively measures the acceleration in the current rate of approach (\mathbf{D}). The establishment of excitatory connections from this neuron to both wheel motors ensures that the mobot will ‘jump’ whenever an approach is made with a suitable high rate of acceleration see Figure 4. The network thus implements a rudimentary ‘fight-or-flight’ response.




Figure 4: Flight response.

6 Learning by explicitation

As we saw in section 4, the experiments with standard learning methods produced poor generalisation performance. However, it turns out that we can obtain quite reasonable performance using a learning method called ‘explicitation’ [17]. By embedding this method within an ongoing evolutionary scenario and by allowing processes of agent/environment interaction to take place within each agent’s lifetime, we can obtain performance which is *reliably* good. The explicitation learning procedure is unusual in that it is both unsupervised *and* constructive. Moreover, the structures that it generates as a part of normal processing tend to have a representational interpretation.

The learning method incorporates a neural-network component and an algorithmic or ‘symbolic’ component. The neural network component has the task of exploiting statistical effects and the algorithmic component has the task of implementing the recursive exploitation of relational effects, as per the task

analysis of [18]. The latter uses a bias which effectively restricts its attention to relationships involving constant differences. A set of data were considered to exhibit a relational effect (i.e., to belong to a relationship) just in case they could be arranged into a linear order such that each variable would show a constant difference from datum to datum. Data satisfying this constraint were said to exhibit a **linear signature**. (As an efficiency measure, the relational exploitation was directed towards the statistical effects identified at any given layer, rather than to the relevant original data. Thus the learning always adopted a coarse-grained view of available data.)

The foundation for the neural network (statistical) component was the well-known unsupervised procedure of **competitive learning** [16]. This is a general and robust method. However, like many unsupervised processes, it is unable to directly exploit low-order³ statistical effects.⁴ The basic procedure of an unsupervised learning method involves grouping inputs together according to similarity. This process ‘discovers’ cases in which a set of input values typically co-occur and thus tends to expose n th-order associations between variable values (where n is the number of values making up a complete input). However, the process *ignores* possible associations of order less than n , i.e., ones which do not involve the complete set of input variables taken together.

Unsupervised methods can usually be adapted to overcome this deficiency. In the case of competitive learning [16] the adaptation is straightforward; it merely involves ‘double-weighting’ each connection in the network [19]. The original competitive learning process effectively searches for input clumps in the input space. The double-weighted method searches for clumps in *subspaces* of the input space. It not only finds associations between input values but also determines which values are associated. It is thus able to exploit m th-order statistical effects in an n -dimensional input space, where $m < n$.

6.1 Recursive exploitation of relational effects

To obtain the hybrid explicitation procedure, the adapted competitive learning regime was integrated with an algorithmic component capable of implementing relational searches. This component was configured so that it would become active at the point where all statistical effects in a given data source had been exploited. Its task was to test the discovered associations for the presence of linear signatures and, for each signature found, introduce a new variable whose values would provide an evaluation of the identified relationship. This was achieved by arranging for instantiations of the variable to reflect the projection of new data upon the relevant data line. These new variables, and their memory buffers, then form the input variables for the next layer of the network.

The method can be visualised as a network building operation which alternates between a statistical exploitation phase in which new competitive nodes are produced and a relational exploitation phase in which new variables are

³The order of a statistical effect is the number of variables involved.

⁴An ‘effect’ for present purposes is simply any observable probability whose value deviates markedly from chance.

produced. This is visualised in Figure 5. Initially the network contains just two input nodes or variables. The double-weighted competitive learning process is then used to produce a layer of competitive nodes (Figure 5 b). (Competitive learning usually operates with a fixed set of nodes. However, for present purposes it was more convenient to allow nodes to be added incrementally until such time as all statistical effects had been exploited.) The relational exploitation process is then used to derive higher-level variables (Figure 5 c) which gives rise to the creation of a higher level of competitive nodes (Figure 5 d).

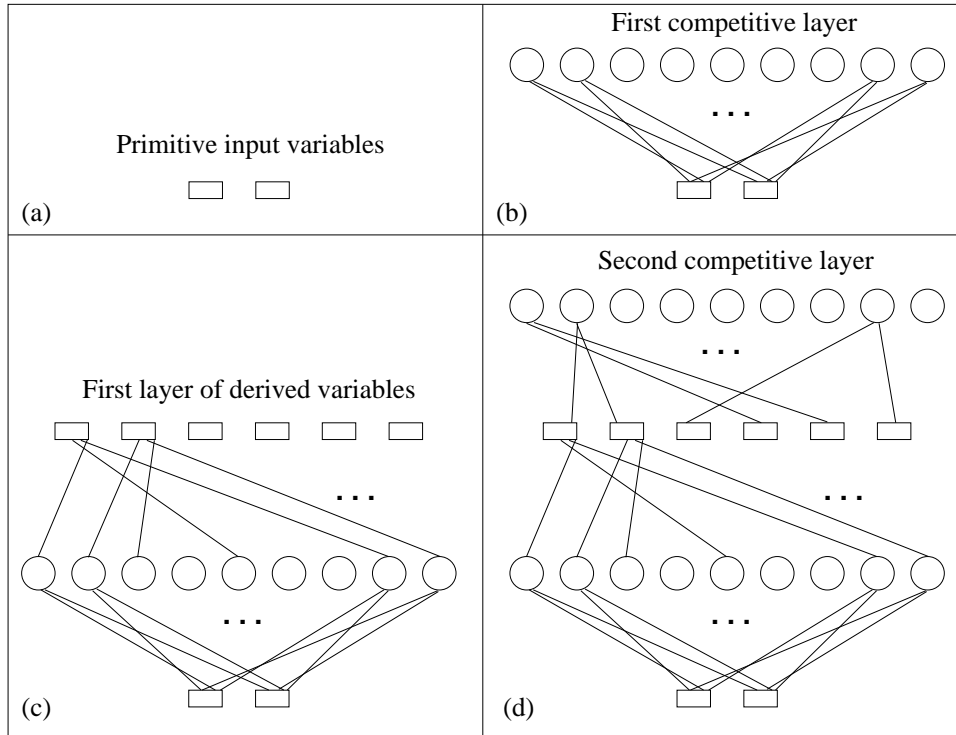


Figure 5: Explication using neural-network construction.

To summarise, the learning method used in the model is a variant of the competitive learning regime. It uses a recursive process of relational exploitation to incrementally extend the original neural network. The most important features of the method are as follows.

- Network connections are always double-weighted so as to enable exploitation of low-order statistical effects.
- The identification of a linear signatures at any layer always leads to the production of a new node at the next level of the network. The action of the network is configured so as to instantiate such variables with a

value reflecting the geometric projection of the current input upon the line carved out by the relevant signature. Such variables thus provided an approximate measure of the relevant relationship.

- The learning is fully incremental. Nodes and layers are added to the network up until the point at which all effects — statistical and relational — are fully exploited.
- The learner’s one-cycle memory is implemented by providing each main variable with a buffer which always hold its previous value.

7 The model

We now turn attention to the model itself. This takes the form of a simulation program which takes about 10 minutes to run on a single-user Sun SPARC 1+. As we will see, the phenomena it generates while doing so include simulations of evolutionary processes, learning processes and interactions between predator and prey agents.

The simulation divides up into a number of phases. In the initial phase, the world contains a single prey agent and a single predator agent. The predator agent exhibits the ‘non-aggressive’ behaviour pattern described above. If it finds itself facing away from the prey, it simply attempts to turn towards it; see Figure 3 (a) and (b). If the predator finds itself facing directly towards the prey, it moves forward at a ‘slow’ or a ‘fast’ pace: Figure 3 (c). If the prey is directly facing the predator when the predator is about to collide with the prey, the predator turns away: Figure 3 (d). If the prey is facing away, the predator continues moving forwards, colliding with the prey and destroying it (not shown).

To begin with, the prey agent has a very simple internal architecture. It has a single proximity sensor and this is connected by inhibitory or excitatory connections to either of the two wheel motors, both of which may have an independent bias. (Thus it is an example of Braitenberg’s ‘vehicle 2’ [20].) An example of this basic architecture is shown in Figure 6. This plan view shows the two wheels and drive motors (the dark, rectangular shapes close to the boundary of the agent), the sensor (a small dark circular shape on the lower part of the boundary) and the connections between them. There is a plus sign above the right motor which denotes the fact that the motor has a forwards (positive) bias. Conversely, the left motor has a minus sign which denotes a backwards (negative) bias. There are excitatory connections from the sensor to both wheel motors. With these biases and sensor-motor connections, the agent will tend to move towards whatever excites its sensor, tending to veer to the left at all times [20].

During the initial phase, the prey is subject to an evolutionary process involving asexual reproduction. The prey’s internal architecture is specified using a simple genetic encoding, i.e., a ‘genotype’. Each time a prey is destroyed through collision with the predator, a slightly mutated copy of its genotype is




Figure 6: Early prey architecture.

generated and this is used as the specification for an **offspring**. This offspring then replaces its parent in the arena. The predator agent remains unchanged throughout this process. We can think of this in terms of evolution having no effect on the predator. (The role of the prey’s genotype in this initial phase is very limited: it only has to specify the existence/non-existence of the sensor-motor connections and the biases for the motors. However, the genotype length is variable and later in the model it is required to specify a more complex internal configuration.)

To survive attacks by the predator, the prey must rotate continuously until such time as the predator is sensed, and then cease all motion. Because of the forwards orientation of the prey’s sensor, this strategy ensures that the prey stops at the point where it is facing directly towards the predator. This effectively ‘blocks’ any oncoming attack: Figure 7 (a). The only internal architecture that generates this behaviour is the one shown in Figure 7(b). Here, the right-track motor has a permanent forwards bias (i.e., tends to run forwards unless explicitly stopped) and the sensor is connected by inhibitory connections to both track motors. Thus the prey is caused to rotate when there is no sensory

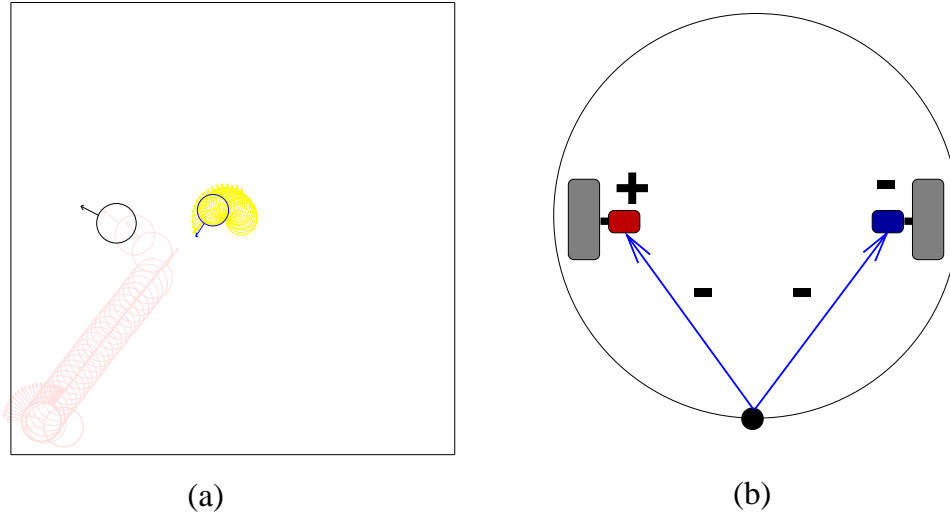


Figure 7: Simple survival strategy.

input and to halt when there is. This implements the required survival strategy. The initial evolutionary phase continues until such time as a prey with this configuration of connections and biases is generated. From this moment on, the attacks of the predator are always successfully blocked. As a result the ‘lifetime’ of the prey is sufficient to enable learning to take place.

7.1 Events in the initial learning phase

As we have seen, the prey has a single sensor which senses the proximity of the nearest obstacle (e.g., the predator) along a ray pointing directly ahead. Predator attacks, in the initial phase of the model, take the form of rapid advances towards the prey, which are then aborted if the prey turns out to be facing towards the oncoming predator. Thus the prey’s input environment is a sequence of proximity values, a sample of which is as follows.

```

0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59
0.59 0.59 0.59 0.59 0.60 0.60 0.61 0.62 0.63 0.64 0.66 0.67 0.68
0.69 0.70 0.71 0.72 0.73 0.74 0.76 0.77 0.78 0.79 0.80 0.81 0.82
0.84 0.85 0.86 0.87 0.88 0.91 0.00 0.00 0.00 0.00 0.00 0.00 0.00
0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00

```

Note the prevalence of zero (0.00) values, the runs of increasing values (signifying an oncoming predator) and the runs of identical values (signifying a predator moving into a head-on position).

Learning takes place in the prey using the hybrid explicitation method. As mentioned, all data variables in the learning process are buffered, which means that the learning actually has access to a data stream based on two variables: the sensor value itself and the previous sensor value. The statistical-exploitation aspect of the learning leads to the discovery that particular pairs (i.e., two-value sequences) of sensor values occur with high frequency. Examples include (0.63, 0.64) and (0.77, 0.78). The existence of these is a consequence of the fact that the predator tends to advance on the prey at a steady rate, thus smoothly increasing its measured proximity from cycle to cycle.

The exploitation of these effects results in the formation of a number of competitive nodes (neurons). By the time the learning method has discovered all frequently occurring pairs of values, the network contains 33 nodes (see the lower part of the network shown in Figure 8). The learning process now begins to ‘search’ for relationships among the patterns (centers) that the neurons respond to. It quickly discovers that the patterns exhibit two linear signatures. The predator tends to approach at either a fast or a slow pace and these two styles of approach give rise to two signatures, one of which is associated with small constant case-to-case variable differences and the other of which is associated with large, constant, case-to-case variable differences.

The learning responds to the discovery of these two signatures by creating two **derived variables**. These are connected up to the relevant nodes come to form what are in effect virtual sensors — one measuring the proximity of a slowly-approaching predator and the other measuring the proximity of a rapidly approaching predator. Again, the learning introduces buffers for both derived variables, see the middle layer of the network in Figure 8.

7.2 The second learning phase

At this stage in the model, the predator begins to vary its behaviour. When approaching a prey, it now does so either ‘aggressively’ (i.e., at an *increasing* pace) or ‘non-aggressively’ (at a *steady* pace). The prey’s sensory environment thus changes. A sample sequence of sensor inputs is as follows.

```
0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59 0.59
0.59 0.59 0.59 0.59 0.62 0.66 0.72 0.73 0.76 0.81 0.83 0.87
0.93 0.95 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
```

In addition to the primitive input environment, there is now also the derived (recoded) input environment for the prey to consider. This is set of instantiations produced in the derived variables (and their buffers) which is produced when sensor inputs are fed to the network. When the learning is applied to these new data a further layer of nodes is created. Through training, these come to respond to specific patterns of values in the derived variables.

Recall that the first layer of derived variables (and their buffers) effectively measure the proximity of a particular type of approach (i.e., a fast or slow




Figure 8: Final network structure.

one). The main patterns that they exhibit (and which are therefore discovered by the learning) thus include the pattern created by an accelerating but distant predator, the pattern created by a constant-paced, distant predator, the pattern created by a close, accelerating predator and so on. Again, due to the fact that sensed proximity always increases, buffer values are always slightly below current values. Thus, the nodes responding to accelerating approaches exhibit patterns of values which show a marked relationship (another linear signature).

The detection of this relationship leads directly to the construction of a higher-level variable. This effectively measures the proximity of an accelerating approach and thus comprises an ‘accelerating-approach’. Two other higher-level nodes are also generated at the same time. The final network is as shown in Figure 8.

7.3 The final evolutionary phase

The generation of these final nodes exhausts the exploitation of statistical and relational effects. The final evolutionary phase now begins. Instantiations of

the left, topmost variable reflect the proximity of an accelerating approacher. If no accelerating approacher is sensed, then this variable will have a zero instantiation. If one is sensed, this variable will have a non-zero instantiation. This node thus constitutes an accelerating-approach ‘detector’.




Figure 9: Final implementation.

The prey genotype is now considerably more complex. It may specify excitatory or inhibitory connections from any of the nodes in the network to either of the two motors. As evolution continues, different connection schemes are tried and many prey are destroyed. Eventually, a genotype is produced which specifies a combination of inhibitory connections from the main sensor variable, and excitatory connections from the accelerating-approach node to both of the two motors, as shown in Figure 9. With these connections and biases, the prey ‘jumps’ whenever it is confronted with an obstacle moving at an accelerating pace directly towards it.

The final behaviour of the prey is illustrated in Figure 10. Its response to a steady-paced approach is to remain stationary (to ‘face in’ to the attack, see Figure 10 (a)) while its response to an accelerating approach is to jump away (Figure 10 d). Thus the prey implements a rudimentary of fight-or-flight

response. With the successful emergence of this behaviour, the simulation terminates.

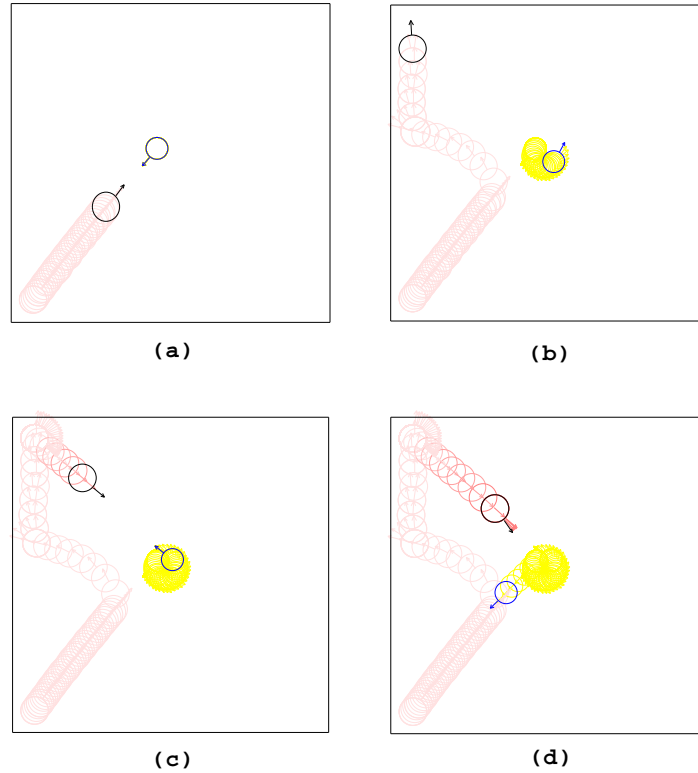


Figure 10: Production of flight response.

8 Representational interpretation

As we have seen, the learning and evolutionary processes simulated in the model result in the production of a large neural network structure within the prey. I now show that this structure can be given a representational interpretation. The internal structure of the network architecture developed in the prey agent(s) is illustrated schematically in Figure 11. Here, variables are shown as rectangular shapes and nodes as oval shapes in the usual manner. All of the nodes and variables shown in the figure function as relational detectors or measures of some sort. All are given a label which attempts to characterise the property they measure or detect. The notations \mathbf{t} and $\mathbf{t-1}$ are used to distinguish between current and previous properties. The character \mathbf{c} is used to denote an unknown constant value.

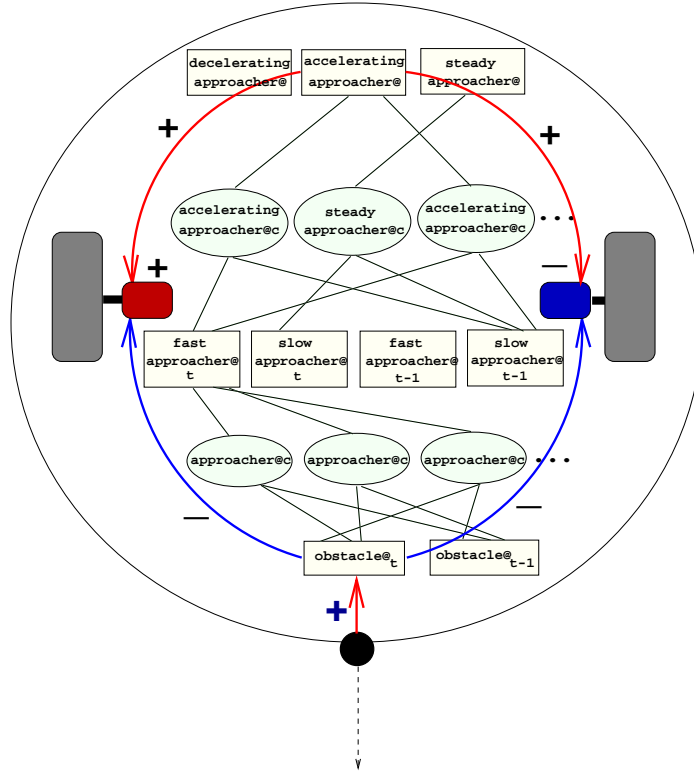


Figure 11: Representational structure of a successful prey architecture.

The bottom, left variable is simply the channel for the current sensor input. Its label is thus `obstacle@t` since the sensor is an obstacle proximity sensor and thus provides a measure of the proximity of an obstacle at the current time. Competitive nodes at the first level of the network respond to particular patterns of values across the main input variable and its buffer. Each one thus detects an approach taking place at a particular proximity; the label used is thus `approach@c`. (Note that all nodes at this level respond to events extending across time `t` and `t-1` and thus have no time subscript.)

The variables making up the third level of the network are derived from the signatures detected in the effects captured by the second-level nodes. They thus measure the proximity of an approach occurring at a particular rate (i.e., either fast or slow). The labels applied are thus `fast-approacher@t`, `slow-approacher@t` etc. The nodes at the fourth layer of the network respond to approaches which maintain a constant or a changing description over the third-layer variables, i.e., which either stay slow or fast, or change from one rate to the other. I therefore use the labels `accelerating-approacher@c` and `steady-approacher@c` here. Finally, the variables at the fifth and top layer of the network capture signatures

among the effects identified at the layer below. They thus provide measures of the proximity of an accelerating, decelerating or steady approach. The labels used here are thus `accelerating-approacher@` `decelerating-approacher@` and `steady-approacher@`.

Should we view this structure of detectors as a ‘representation’? If we address the question in a narrow-minded way, the answer has to be no. The structure created by the simulated processes is not a representational structure in the normal sense. It is not a frame structure. It is not an ISA hierarchy. It is not a script. It is not a Prolog database. And it is not a primal sketch, or an object-centered description.

On the other hand, the structure clearly *does* deal in terms of symbolic values. The current value of one of the `accelerating-approach@c` detectors, for example, is not an explicit (or even statistical) property of the raw input. It is a derived value which is used by other parts of the network — and the next layer of variables in particular — as a stand-in for an abstract property of the environment. The value is thus quite literally treated by the rest of the network as a symbol or sign of that property.

The same goes for the nodes and variables in layers two, three and five. Each node or variable in these layers effectively instantiates a symbol for an abstract feature of the environment. It does so by processing its inputs, which are themselves either raw sensor data or symbols for less abstract features of the environment. The network as a whole then is a kind of processing device. But the data it deals in are symbols for more-or-less abstract and in most cases *dynamic* features of the current environment. The dynamic state of the network is thus an active ‘re-representation’ — a multi-levelled recoding — of events and processes in the environment.

Rich and detailed though the final representational structure is, it nevertheless does not in any sense constitute a ‘fully-detailed world model’. Nor is it imposed by any external agency. Rather, it emerges as a natural result of the interactions between processes and agencies operating both externally in the environment and internally within the prey. The main driving force behind the representational development in the prey is in fact the interplay between the assimilatory and accommodatory forces active within the lifetime learning [cf. 21]. Presented with sensor data, the network effectively *assimilates* that data to its current, internal representation. Over time there is also *accommodation*. This comes in two forms: structural accommodation, involving the creation of new nodes, variables and layers; and non-structural accommodation involving the tuning of network weights. As a result, there is the construction of a sequence of *redescriptions* or recordings of the prey’s current sensory environment, much in the manner of Karmiloff-Smith’s ‘representational redescription’ process [3]. The general conclusion then is that the structure *does* have a representational interpretation and that it is precisely the one shown in Figure 11.

9 Final comments

The paper has presented a model of the development of a fight-or-flight behaviour in a simulated agent. Because the model makes use of the explicitation learning method, which is both unsupervised and constructive, it is able to present a developmental picture in which evolutionary processes effectively cooperate with learning processes in the formation of representational structure. Because the model uses processes which are fully incremental, it is able to show how an agent/environment interaction can modulate and guide an underlying evolutionary process. And because the learning is implemented in the form of a neural-network process, the learning process is at least reconcilable with current models of brain mechanism.

The model thus provides a ‘big-picture’ story about the way in which the development of complex intelligent behaviours might involve evolutionary processes, learning processes, agent/environment interaction and representation development. The learning method forms a bridge between the GA paradigm on the one hand and the representationalist paradigm on the other.

Unfortunately, in its present manifestation the model has many shortcomings. As a computational implementation, it is not as robust as one would like. It also incorporates as an ad hoc feature the device of one-cycle memory (i.e., data buffering). The selection of a linear-signature bias during relational learning is also weakly motivated. But the most serious deficiency is probably the fact that the model makes use of what is in effect, a pre-scripted sequence of events. It requires that the predator implement certain changes in behaviour at certain points in the process. If these behavioural alterations do not take place on cue, the learning process is derailed.

However, it is hoped that all these deficiencies will be remedied in the ongoing development of this work. Future versions of this model will aim to show how a fight-or-flight response emerges in a more natural, unscripted scenario, involving multiple predator and prey agents in addition to other forms of process and contingency. Work is currently in progress towards this end.

10 Acknowledgements

Thanks to Dave Cliff and Jim Stone for helpful input on an earlier draft.

References

- [1] Shavlik, J. and Dietterich, T. (Eds.) (1990). *Readings in Machine Learning*. San Mateo, California: Morgan Kaufmann.
- [2] Cliff, D., Husbands, P. and Harvey, I. (1993). Evolving visually guided robots. In J. Meyer, H. Roitblat and S. Wilson (Eds.), *From Animals to Animats: Proceedings of the Second International Conference on Simulation of Adaptive Behaviour (SAB92)*. MIT/Bradford Books.

- [3] Karmiloff-Smith, A. (1992). *Beyond modularity: a developmental perspective on Cognitive Science*. Cambridge, Ma.: MIT Press/Bradford books.
- [4] Nolfi, S., Floreano, D., Miglino, O. and Mondada, F. (1994). How to evolve autonomous robots: different approaches in evolutionary robotics. In R.A. Brooks and P. Maes (Eds.), *Proceedings of Artificial Life IV* (pp. 190-197).
- [5] K-Team, (1993). *Khepera users manual*. Lausanne: EPFL.
- [6] Feder, M. and Lauder, G. (Eds.) (1986). *Predator-Prey Relationships. Perspectives and Approaches from the Study of Lower Vertebrates*. Chicago and London: University of Chicago Press.
- [7] Gerlai, R. (1993). Can paradise fish (macropodus opercularis anabantidae) recognize a natural predator? an ethological analysis. *Ethology*, 94 (pp. 127-136).
- [8] Bolles, R. (1979). *Learning Theory* (2nd edition). New York: Holt.
- [9] Flaherty, C. (1985). *Animal Learning and Cognition*. New York: Knopf.
- [10] Barnett, S. (1973). *Ethology and Development*. Heinemann Medical.
- [11] Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
- [12] Goldberg, D. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley.
- [13] Wilson, S. (1991). The animat path to AI. In J. Meyer and S.W. Wilson (Eds.), *Proceedings of the First International Conference on the Simulation of Adaptive Behaviour (From Animals to Animats)* (p. 16). Cambridge: MIT Press.
- [14] Quinlan, J. (1986). Induction of decision trees. *Machine Learning*, 1 (pp. 81-106).
- [15] Rumelhart, D., Hinton, G. and Williams, R. (1986). Learning representations by back-propagating errors. *Nature*, 323 (pp. 533-6).
- [16] Rumelhart, D. and Zipser, D. (1986). Feature discovery by competitive learning. In D. Rumelhart, J. McClelland and the PDP Research Group (Eds.), *Parallel Distributed Processing: Explorations in the Microstructures of Cognition. Vol I* (pp. 151-193). Cambridge, Mass.: MIT Press.
- [17] Thornton, C. (1995). Compression, dilation and the redescriptive role of explicitation. In J. Hoc (Ed.), *Proceedings of the European Conference on Cognitive Science* (pp. 19-30).
- [18] Thornton, C. (1995). Measuring the difficulty of specific learning problems. *Connection Science*, 7, No. 1 (pp. 81-92).

- [19] Thornton, C. (1994). Unsupervised learning with the soft-means algorithm. *Proceedings of the World Congress on Neural Networks*. Vol. 4 (pp. 200-205). San Diego.
- [20] Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. London: The MIT Press.
- [21] Boden, M. (1979). *Piaget*. Fontana Modern Masters, Fontana Press.