Distributed Action Selection by a Brainstem Neural Substrate: An Embodied Evaluation*

Mark Humphries and Tony Prescott

Adaptive Behaviour Research Group, Department of Psychology,
University of Sheffield, UK
{m.d.humphries, t.j.prescott}@shef.ac.uk

Abstract. Theoretical approaches to the problem of action selection in autonomous agents often contrast centralised and distributed selection schemes. Here we describe a neural substrate for distributed action selection in the vertebrate brain-stem, the medial reticular formation (mRF), which may form a evolutionary precursor to centralised schemes found in the higher brain. We evaluate its competence as a selection device for robot control in a simulated resource co-ordination task, and use a genetic algorithm to evolve the mRF's inputs and internal structure. Some configurations of the mRF could sufficiently co-ordinate actions to maximise the robot's energy, but this is critically dependent on a high rate of energy acquisition, which leaves an animal (or agent) susceptible to food shortages. Thus, the inflexibility of the mRF as a distributed selection mechanism may have provided impetus for the evolution of more complex, centralised, selection mechanisms in the brain.

1 Introduction

A generally effective strategy for designing controllers of autonomous agents is to reverse-engineer biological systems that have evolved as solutions to the control problems. One such problem is action selection: a mortal agent must continuously choose and co-ordinate behaviors appropriate to both its context and its current internal state if it is to survive. Animals necessarily embody successful solutions to the action selection problem. Thus, it is natural to look at what parts of the central nervous system — the neural substrate — have evolved to carry out the action selection process.

Recent proposals for the neural substrate of the vertebrate action selection system have focussed on the basal ganglia - a set of fore- and mid-brain nuclei whose input, output, and inter-connections seem to be consistent with a central (as opposed to distributed) resource switching device [1,2]. Decerebrate animals, altricial (helpless at birth) neonates, and lateral hypothalamic rats do not have fully intact or functioning basal ganglia, but are capable of expressing spontaneous behaviors and co-ordinated and appropriate responses to stimuli. During

* This research was supported by the EPSRC (GR/R95722/01), a Wellcome Trust VIP Award, and the European Union Framework 6 ICEA project.

S. Nolfi et al. (Eds.): SAB 2006, LNAI 4095, pp. 199–210, 2006.

decerebration the entire brain anterior to the superior colliculus is removed leaving only the hindbrain intact. Yet, the chronic decerebrate rat can, for example, spontaneously locomote, orient correctly to sounds, groom, perform co-ordinated feeding actions, and discriminate food types [3,4]. Such animals clearly have some form of intact system for simple action selection. We have recently argued that, of the potential candidate structures left intact in the brainstem of decerebrate animals, the medial reticular formation (mRF) is the most likely substrate of a generalised (if limited) action selection mechanism [5,6,7].

We first evaluated the single existing computational model of RF function — a landmark model proposed by Warren McCulloch and colleagues [8] — in both simulation and embodied form (on a version of the same task used here) [5]. However, inevitably, given its age, several aspects of the model were incorrect or implausible, or omitted features known from more modern studies of the mRF. We thus turned to reviewing the modern literature on the mRF, and found that the organisation of the mRF's inputs and outputs, and the functional properties of its cells, are all consistent with the action selection proposal [7]. We began addressing the question of how the mRF represents and resolves the competition between actions by synthesising the neurobiological data to determine the mRF's internal structure [9] — see Fig. 1. The mRF is made up of stacked cell clusters, each cluster a mix of projection and inter-neurons. The projection neurons are excitatory, project a long axon to the motor centres in the brainstem and spine, and contact cells in other clusters via collaterals from that long axon. The inter-neurons are inhibitory, and project only within their own cluster. We outlined a novel quantitative anatomical model that generated networks with this structure and we found that the networks had small-world properties [9].

Potential configurations of the mRF as an action selection system were explored by simulation of a new computational model whose connectivity was based on the anatomical model [6,7]. We found that a *sub-action* configuration most effectively supports selection: the projection neurons of each cluster represents a component of an action, and a coherent behavioural response is created by clusters recruiting other clusters which represent compatible components (in addition, incompatible components are suppressed by inhibition of their representing clusters, which occurs via activation of that cluster's inter-neurons).

The neurobiological and simulation data indicate that the mRF is a distributed selection mechanism, from which the selection of actions is an emergent phenomenon. This can be contrasted with the basal ganglia, which are a centralised control structure, selecting actions on the basis of inputs from multiple command systems. It thus appears that evolution has seen fit to produce both forms of selection structure that are often counter-posed in theoretical discussions [10]. In this paper we extend the assessment of the mRF's capabilities as an action selection mechanism by testing our new computational model of the mRF in an embodied form. The aim of this work was to determine whether or not the mRF was capable of carrying out action selection independently from other neural systems that may be involved in the intact animal, and to shed some light on the complexity of task that the mRF could cope with.

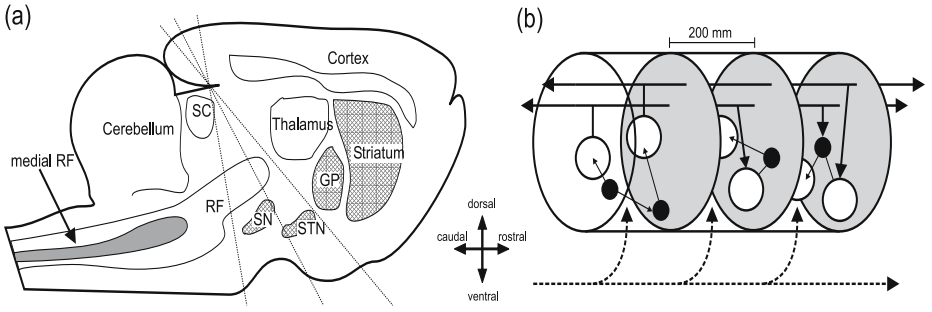

Fig. 1. Anatomy of the vertebrate medial reticular formation (mRF). Directional arrows apply to both panels. (a) The relative locations of major nuclei and structures including the basal ganglia (hashed) and the medial reticular formation (RF) shown on a cartoon sagittal section of rat brain. The dashed lines show the location of the common decerebration lines — all the brain rostral to the line is removed, leaving hindbrain and spine intact. GP: globus pallidus. SN: substantia nigra. STN: subthalamic nucleus. SC: superior colliculus. (b) The proposed mRF organisation: it comprises stacked clusters (3 shown) containing projection neurons (open circles) and inter-neurons (filled circles); cluster limits (grey ovals) are defined by the initial collaterals from the projection neuron axons. The projection neurons receive input from both other clusters (solid black lines) and passing fibre systems (dashed black line). The inter-neurons project within their parent cluster. Reproduced from [7].

2 Methods

We begin by describing the computational model of the mRF, which forms the basis for the robot controller, then describe the task on which the robot is evaluated, the input and output of the controller, and the form of the genetic algorithm used to evolve the mRF.

2.1 A Population-Level Model of the mRF

We do not have sufficient space to fully describe the anatomical model which provides the connectivity data for the computational model — see [9]. It is sufficient to note that the model has six parameters: the number of clusters N_c; the number of neurons per cluster n; a proportion ρ of those are projection neurons, the rest are inter-neurons; the probability of a projection neuron sending a connection to another cluster $P(c)$; the probability of that connection then contacting any neuron in that cluster $P(p)$; and the probability of an inter-neuron contacting any other neuron in its own cluster $P(l)$. Each of these parameters were limited to a range of values sourced from the neurobiological data: the specific values used in this paper are thought to be the most realistic. There are six sub-actions within the task to represent, hence $N_c = 6$; the other parameters were set to: $n = 50$, $\rho = 0.7$, $P(c) = 0.25$, $P(p) = 0.1$, and $P(l) = 0.1$ (see [9] for further details). The result of creating a particular instance of this model is a network of linked nodes.

We use a population-level computational model of the mRF. In this approach, populations of neurons are treated as a statistical ensemble, assuming functionally meaningful sub-groups of neurons cannot be further distinguished. Thus, the model is a set of simplified ordinary differential equations describing the change in the normalised mean firing rate of each population over time. Given the proposed cluster structure, and the hypothesis of projection neurons encoding the action representation, the most natural division of the mRF is into separate populations of projection and inter-neurons for each cluster. The computational model thus has two vectors encapsulating its behaviour: the projection neuron activity $\boldsymbol{c}$ and the inter-neuron activity $\boldsymbol{i}$. Each vector element is a population: c_k is the normalised mean firing rate of the kth cluster's projection neuron population, and i_k is the normalised mean firing rate of the kth cluster's inter-neuron population. These activities evolve according to the differential equations given in [6,7]. Input to the model is described by vector $\boldsymbol{u}$, where each element u_k is the scalar summation of all external input to cluster k from sensory and internal monitoring systems, and which thus represents the salience of that cluster's represented sub-action.

The connections between the populations are defined by the underlying network generated by the anatomical model: variables A_{jk} and C_{jk} are the mean number of contacts from cluster j to, respectively, the projection and inter-neurons of cluster k; b_k and d_k are the mean number of contacts from inter-neurons in the current cluster k to, respectively, the projection and inter-neurons in that same cluster. Figure 2 shows a schematic of the mRF population-level model, which further explains its structure, and details of the parameters which are optimised by a genetic algorithm (see below).

2.2 The Energy Task

We have previously evaluated bio-mimetic computational models on an energy-based task [11,5]. In those evaluations, fixed action patterns were selected by the models: for example, avoiding an obstacle was a complete pattern in which the robot reversed, turned, then moved ahead in a different direction. However, as our simulations have shown that the mRF is most likely to represent components of actions [6], we here decompose the action patterns into their constituent parts.

The form of the task is as described in [5]: a mobile robot explores an arena with a grey coloured floor (representing neutral) upon which are laid two white and two black tiles. The robot controller has six state variables: the states of the left and right bumpers, B_{L} and B_{R}; the values of the bright and dark infra-red floor sensors L_{B} and L_{D}; potential energy P_{E} (which is recharged on black tiles); and energy E (which is recharged on white tiles by consuming potential energy). Both the internal variables P_{E} and E were limited to the range [0,1].

The change δP_{E} in potential energy when recharging on a black tile for T_{eat} seconds is

$$\delta P_E = E_{\mathrm{rate}}\, T_{\mathrm{eat}}\, L_{\mathrm{D}} \ . \tag{1}$$

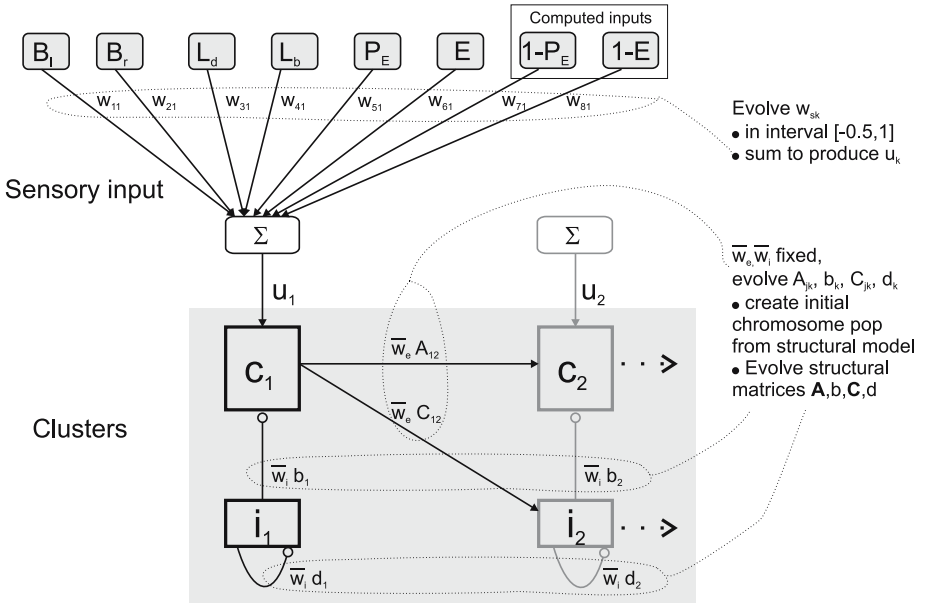

Fig. 2. Schematic of the population-level mRF model. Two clusters are shown; only the first cluster's (c_1) inputs and outputs are depicted. The state variables from the robot are weighted and summed to produce a scalar input signal that represents the urgency of request for the sub-action represented by the cluster. (In the GA, the weights are evolved to obtain the optimal balance of sensory input signals required for that sub-action.) The internal structural parameters $\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{C}, \boldsymbol{d}$ are initially derived from the anatomical model — forming the initial population — and are then evolved as well. The mean excitatory weight $\bar{w}_{\mathrm{e}}$ is constant, and the mean inhibitory weight $\bar{w}_{\mathrm{i}}$ is computed for each model (see text). These values represent the average synaptic efficacy between neurons. Thus, the emphasis is on evolving the *structural* properties of the mRF.

The increase in energy δE and decrease in potential energy $\delta P_{\mathrm{E}}^{-}$ when recharging from stored potential energy on a white tile for T_{digest} seconds is

$$\delta E = E_{\mathrm{rate}}\, T_{\mathrm{digest}}\, L_{\mathrm{B}}, \quad \delta P_{\mathrm{E}}^{-} = -E_{\mathrm{rate}}\, T_{\mathrm{digest}}\, L_{\mathrm{B}} \ . \tag{2}$$

The initial experiments set the acquisition (and conversion) rate $E_{\mathrm{rate}} = 0.027$, following our prior work [11,5].

In the original version of the task, the robot had four actions available to it: *Wander*: a random walk in the environment, formed by forward movement at a fixed speed followed by a turn of a randomly selected angle; *Avoid Obstacles*: reverse movement, followed by a turn away from the object; *Reload On Dark*: stop on a black tile and charge potential energy; *Reload On Light*: stop on a white tile and charge energy by consuming potential energy.

We decomposed these into the following six sub-actions: move forward, move backward, turn left, turn right, recharge potential energy, and recharge energy.

Hence the mRF model used had six clusters, one per sub-action. This decomposition is based on studies of the mRF's control of movement: for example, in the lamprey there are separate mRF neuron sub-populations whose activity drives moving forward, moving backward, and turning [12]. The recharging sub-actions are distinct from the original fixed actions of reloading, as they do not include commands to stop movement: the model must co-ordinate the stopping of movement with the selection of recharging at the appropriate time.

All robot simulations were performed in Webots 4 (Cyberbotics). One robot simulation time-step is one second. At the beginning of each run the robot was initialised with $E = 1$ and $P_{\mathrm{E}} = 0.5$ and placed at a particular location in the arena. Regardless of the action(s) selected, energy E is depleted at a constant rate of $E_{\mathrm{met}} = 0.002$ unit/s, corresponding to a fixed metabolic rate. Therefore, if no recharging of energy occurred then the minimum survival time was 500 seconds.

At each time-step, the instantiated mRF computational model receives inputs computed as described below, and is run until it reaches equilibrium, or until $t = 0.5$s, whichever occurs first. The model was solved in discrete-time, using a zero-order hold approximation, with a time-step of 0.001s. Each run is initialised with the final state of the previous run, so that the model's dynamics are effectively continuous. Its output is then converted into the activity of the corresponding spinal motor centres. The behavioural vector of the robot is then computed by aggregating these signals. We describe each of these processes in turn: a schematic of the controller is given in Fig. 3b.

2.3 mRF Input

Our previous robot work based on the energy task used complex salience equations: computed levels of urgency of each action, based on functions of each sensory variable. However, mRF inputs are mostly directly from sensory and internal monitoring systems, so there is little scope for complex neural processing of those signals. We thus assume that the kth cluster's input u_k is given by a summed weighted input of the robot's sensory variables, energy variables, and the inverse of the latter (information on falling energy levels, or at least the volume of used gut capacity, is available to the mRF [4]). These inputs are shown in Fig. 2. The resulting total is re-scaled so that $u_k \in [-0.5, 1]$, following neurobiological data on the input to the mRF: see [6,7] for more detail.

2.4 Interpreting mRF Output

The mRF has direct control over vertebrate central pattern generators (CPGs), the circuits which drive limb and jaw movements [13]. Increasing activation of reticulo-spinal neurons causes the onset of and then increasingly rapid locomotion [12], which corresponds to the onset of oscillations and their increasing frequency within the locomotor CPG. Simulations have explained the seemingly paradoxical result that some increases in activation can terminate locomotion: given sufficient input, the oscillations terminate and a stable state is resumed

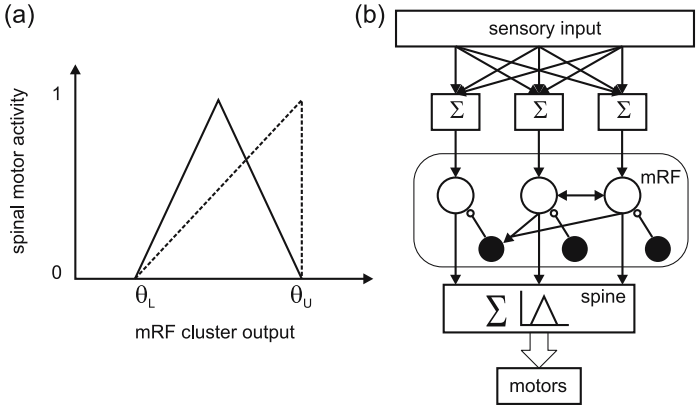

Fig. 3. (a) Activity in the spinal motor plants is a non-linear function of mRF drive. Two transfer functions, based on neurobiological data, are evaluated in this paper: "dual" (dashed line) and "triangle" (solid line), which give non-zero output between two thresholds (θ_{L} and θ_{U}). (b) Cartoon schematic of the robot controller. Sensory variables are weighted and summed, the total then input to the mRF computational model (projection neuron populations, white circles; inter-neuron populations, black circles). The output of each cluster is filtered through the spinal transfer function, then aggregated to produce a motor vector. This then drives the wheels and sets the rate of change for both energy and potential energy.

[14]. There thus appears to be separate thresholds for the onset and termination of activity in the CPG. It is not clear whether the increasing reticulo-spinal activity causes a continuous increase in CPG activity until a sudden, discontinuous, stop, or whether it causes an increase to some maximum followed by a decrease (see Fig 3a) — both will be examined.

The output vector $\boldsymbol{c}$ of the mRF computational model is thus converted to a spinal-command vector by $m_k = M(c_k)$, where M is one of the output transfer functions, "dual" or "triangle", shown in Fig. 3a. Lower and upper thresholds for M were set at $\theta_{\mathrm{L}} = 0.1$ and $\theta_{\mathrm{U}} = 0.8$ for the results reported below. The spinal-command vector then gates the contribution $\boldsymbol{s}^k$ of each sub-action to the final behaviour vector $\boldsymbol{b}$: each sub-action contribution being either the requested drive of the robot's wheels, or the requested quantity of energy and/or potential energy to change, as detailed in Table 1. The final behaviour vector thus has four elements: the first and second elements are the motor speed sent to the left and right wheels respectively, the third element is the quantity to update E, and the fourth element is the quantity to update P_{E}.

2.5 Form of the Genetic Algorithm

We detail the features of the genetic algorithm (GA) used to search the space of mRF models. Following our previous work, we define our fitness function as the mean E over a fixed time window of 2500 seconds *after* the minimum survival period had elapsed [5]. (We do not use survival time as this is unbounded —

Table 1. Behavioural contributions of each sub-action: the first two elements are the requested motor commands sent to the left and right wheels of the robot; the second two are the requested changes in energy and potential energy, whose values are given by (1)-(2). The motor speeds for the turn sub-actions are the values necessary to turn the robot 180^o in one time-step.

Sub-action	Vector $\boldsymbol{s}$
Move forward	[80 80 0 0]
Move backward	[-40 -40 0 0]
Turn left	[10.5 -10.5 0 0]
Turn right	[-10.5 10.5 0 0]
Recharge potential energy	[0 0 0 δP_{E}]
Recharge energy	[0 0 δE $\delta P_{\mathrm{E}}^{-}$]

the robot may never expire). Our resulting fitness function naturally falls in the range [0,1], with 1 indicating maximum fitness.

An initial population of 200 chromosomes was created: 200 mRF anatomical models were instantiated with the parameters given above, and the connection matrices derived from them; the sensory input weights were randomly chosen from the interval [-0.5,1]. Every subsequent chromosome population had 50 members. Each chromosome of a population was converted into the set of connection matrices, and the resulting mRF model evaluated on the energy task. The population was then ranked by fitness level, and the best 20 chromosomes retained. From this remaining population, 30 pairs of chromosomes were randomly chosen for mating: from each pair, a new chromosome is created by conjoining the two chromosomes at a randomly chosen split point. Thus, a new population of 50 chromosomes results (20 parents, 30 offspring). The new population is subjected to mutation, where each element is changed to another value within its preset interval (defined below) with a probability of 0.05. The top chromosome of the parent population is never mutated, so that the most fit parent is always retained intact (elitism). Once all pairings and mutations have been carried out, the resulting population is again evaluated on the energy task. This process was iterated until the termination condition was reached, that the top chromosome was unchanged for 20 consecutive generations.

The models were encoded as a real-valued chromosome of 120 elements, broken down into: 48 input weights (8 sensory inputs $\times$ 6 clusters); 60 inter-cluster connections (total number of non-zero elements in $\boldsymbol{A}$ and $\boldsymbol{C}$); and 12 intra-cluster elements (total number of elements in $\boldsymbol{b}$ and $\boldsymbol{d}$). The intervals over which each element could be mutated were limited as follows. We used an interval of [-0.5,1] for the sensory weights. Each element of $\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{C}, \boldsymbol{d}$ was limited in the interval $[0{,}3 \times E(x)]$ where $E(x)$ is the expected value of each element, given the connection probabilities used in the underlying anatomical model.

The mean connection weights $\bar{w}_e$ and $\bar{w}_i$ were not optimised, to reduce the number of free parameters. We set excitatory weight $\bar{w}_e = 0.2$, as before [7]. Previously, we set inhibitory weight by $\bar{w}_{\mathrm{i}} = -\bar{w}_{\mathrm{e}} \times N_{\mathrm{e}}/N_{\mathrm{i}}$, where N_{e} and N_{i} are the total number of excitatory and inhibitory connections in the network

(see [7] for explanation): here we sum together the elements of $\boldsymbol{A}$ and $\boldsymbol{C}$ to compute N_{e}, and sum together the elements of $\boldsymbol{b}$ and $\boldsymbol{d}$ to compute N_{i}.

3 Results

3.1 Assessment of the Fitness Landscape

To assess the fitness landscape to be explored by the GA, we performed a series of Monte Carlo simulations: 1500 mRF models were instantiated, using the anatomical model parameter values given above, and each was assessed for its fitness as a controller for the robot. We thus hoped to gain some insight into the distribution of fitness over all possible models.

Three random searches were conducted: for the first we used the "dual" transfer function; for the second we increased the rate of energy acquisition by an order of magnitude so that $E_{\mathrm{rate}} = 0.27$ (we ran this search because, when observing robot behaviour on individual model trials, we noted that the rates of recharging energy and potential energy were, consistently, roughly 10% of those used in the previous work [11,5], due to the gating of the sub-actions by the mRF output); for the third we retained the same increased E_{rate}, but used the "triangle" transfer function. For each search we computed the empirical cumulative distribution function (EDF) [1] of model fitness, shown in Fig. 4 as a function of fitness error (i.e. 1-fitness).

The first search found a maximum fitness of 0.165 and its EDF shows that most models had a fitness close to zero, and were therefore not surviving long beyond the minimum survival period. For the second search, the resulting fitness distribution had a markedly increased maximum of 0.518, but again the EDF shows that most models had low fitness. For the third random search, the fitness distribution had a further increased maximum of 0.897; however, the EDF again shows that the overwhelming dominance of low-fitness models remained. From these random search results we concluded that, though infrequent, there were forms of the model that a GA-based search could potentially find and optimise.

3.2 Increasing Energy Acquisition Rates Increases Fitness

We confirmed the first random-search results by conducting a series of GA-based searches with the original acquisition rate ($E_{\mathrm{rate}} = 0.027$) and "dual" transfer function: altering numerous parameters of the GA search (the initial population, the retained population, the number of offspring), or features of the underlying model (using un-scaled salience inputs, using hand-coded sensory-input weights, changing the transfer function thresholds), never resulted in a maximum fitness that markedly exceeded that of the random-search by the time the search terminated.

[1] An empirical cumulative distribution function is an estimate of the underlying cumulative distribution function, each probability estimated by $P(x)$ = (number of observations $\leq x$) / (total number of observations).

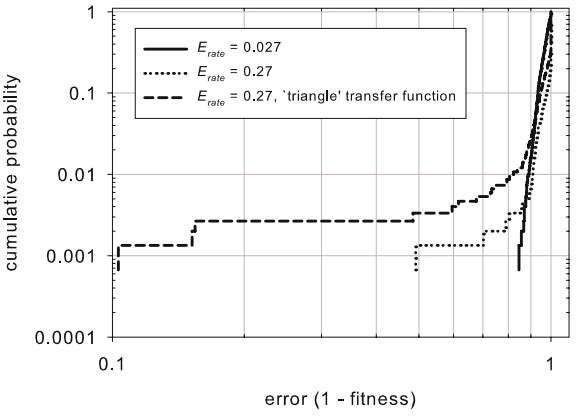

Fig. 4. Distributions of model fitness over a large random sample of all possible models in three different set-ups of the robot task. The distributions are shown as EDFs (see text), plotted on log-log axes as a function of fitness error (1 - fitness), to allow plotting of zero fitness on the log-scale. Increasing E_{rate} increased maximum fitness, but the distribution was always dominated by low fitness models.

Subsequent GA-based searches using the increased energy acquisition rate of $E_{\text{rate}} = 0.27$ confirmed the results of the other random-searches: for both "dual" and "triangle" transfer functions the maximum fitness was considerably increased following the increase in E_{rate}. Indeed, the maximum fitness achieved during the searches either exceeded ("dual", fitness = 0.714) or equalled ("triangle", fitness = 0.883) that found by the corresponding random search. The mean fitness of each population did not increase over the course of the generations (Fig. 5a), consistent with the dominance of low-fitness models in the random searches.

3.3 Dependence of Fitness on the Rate of Energy Acquisition

We then assessed the dependency of the performance of the GA-based search on the value of E_{rate}. A search was run using the GA set-up described in Sect. 2.5, and models with the "dual" transfer function, for each step increase of 0.027 from the initial value of $E_{\text{rate}} = 0.027$ up to $E_{\text{rate}} = 0.27$ — there were thus 10 searches. A linear regression showed a significant increase in maximum fitness ($r = 0.8426$, $p < 0.01, n = 9$; one outlier), and a significant increase in final generation mean fitness ($r = 0.725$, $p < 0.05, n = 10$) as a function of increasing E_{rate} (Fig. 5b). However, the increase in mean fitness was not of the same order of magnitude as the increase in maximum fitness, indicating that the majority of the population remained at low fitness regardless of E_{rate}.

4 Discussion

The complexity of the task seems to be a difficult one for the mRF models to solve, given the low fitness of the vast majority, but we have shown that mRF structures

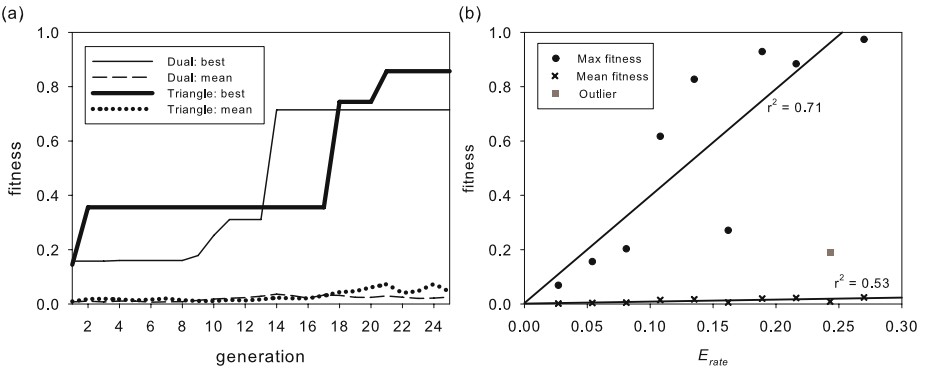

Fig. 5. (a) A small but high-fitness set of mRF models can evolve over time given a sufficiently high energy acquisition rate ($E_{\text{rate}} = 0.27$), using either the 'dual' or 'triangle' output functions. However, there is little increase in fitness of the majority of the population. (b) The effect of increasing E_{rate} on the maximum and mean fitness of the GA-based model search. Increasing E_{rate} is significantly correlated with an increase in both the maximum fitness and the final generation's mean fitness. Mean fitness did not show the same order of magnitude increase as maximum fitness.

can select action components sufficiently well to co-ordinate energy gathering and acquisition. We should not be surprised that mRF models with high evolutionary fitness were difficult to find: our previous similar evaluations of bio-mimetic models were based on fixed action patterns, whose salience was a complex function of sensory variables; the necessity of creating emergent actions from components, based only on direct sensory input, makes the task far more difficult for the mRF model. Indeed, if most structural configurations of the mRF could support efficient action selection, then why would more complex neural systems have evolved to deal with the same problem, rather than co-opting the existing solution? Nevertheless, it is testament to the potential computational power of even the most "basic" of brain structures that the mRF model was successful at all.

The success of the mRF as an action selection system is dependent on the energy acquisition rate (it may equivalently be dependent on the metabolic rate E_{met}, which will be explored in future work). Interestingly, an increase in E_{rate} did not result in a correspondingly large general increase in model fitness: it appears that it only promoted the models which had the potential to successfully co-ordinate the robot's behaviour (Fig. 5b).

The use of a fixed acquisition rate seems valid: data from studies of decerebrate rats suggests that they are unable to adaptively alter their rate of food acquisition during periods of food deprivation [4]. This, in turn, suggests that a brain-stem dominant animal may be susceptible to fluctuations in food supply, and indeed may be inflexible in its response to other environmental changes: future work will test this idea more rigorously.

The results of this work neatly parallel the evolution of the vertebrate brain: some ancient species, such as the lamprey, have their motor behaviour domi-

nated by the reticulo-spinal system. Thus, the mRF seems to be a sufficient control system in some ecological niches. Yet most modern vertebrates have more complex neural systems that combine to control their behaviours. Studying the integration of these more complex, centralised, systems with the lower-level, distributed, mRF system may provide further insight into potential designs for control architectures of autonomous agents.

References

1. Redgrave, P., Prescott, T.J., Gurney, K.: The basal ganglia: A vertebrate solution to the selection problem? Neuroscience **89** (1999) 1009–1023
2. Prescott, T.J., Redgrave, P., Gurney, K.: Layered control architectures in robots and vertebrates. Adapt. Behav. **7** (1999) 99–127
3. Berntson, G.G., Micco, D.J.: Organization of brainstem behavioral systems. Brain Res. Bull. **1** (1976) 471–483
4. Grill, H.J., Kaplan, J.M.: The neuroanatomical axis for control of energy balance. Front. Neuroendocrin. **23** (2002) 2–40
5. Humphries, M.D., Gurney, K., Prescott, T.J.: Is there an integrative center in the vertebrate brainstem? A robotic evaluation of a model of the reticular formation viewed as an action selection device. Adapt. Behav. **13** (2005) 97–113
6. Humphries, M., Gurney, K., Prescott, T.: Action selection in a macroscopic model of the brainstem reticular formation. In Bryson, J.J., Prescott, T.J., Seth, A.K., eds.: Modelling Natural Action Selection. AISB Press, Brighton, UK (2005) 61–68
7. Humphries, M.D., Gurney, K., Prescott, T.J.: Is there a brainstem substrate for action selection? Phil. Trans. Roy. Soc. B. (2006) in press.
8. Kilmer, W.L., McCulloch, W.S., Blum, J.: A model of the vertebrate central command system. Int. J. Man-Mach. Stud. **1** (1969) 279–309
9. Humphries, M.D., Gurney, K., Prescott, T.J.: The brainstem reticular formation is a small-world, not scale-free, network. Proc. Roy. Soc. B. **273** (2006) 503–511
10. Maes, P.: Modeling adaptive autonomous agents. In Langton, C.G., ed.: Artificial Life, An Overview. MIT Press, Cambridge, MA (1995) 135–162
11. Girard, B., Cuzin, V., Guillot, A., Gurney, K.N., Prescott, T.J.: A basal ganglia inspired model of action selection evaluated in a robotic survival task. J. Integr. Neurosci. **2** (2003) 179–200
12. Deliagina, T.G., Zelenin, P.V., Orlovsky, G.N.: Encoding and decoding of reticulospinal commands. Brain Res. Brain Res. Rev. **40** (2002) 166–177
13. Noga, B.R., Kriellaars, D.J., Brownstone, R.M., Jordan, L.M.: Mechanism for activation of locomotor centers in the spinal cord by stimulation of the mesencephalic locomotor region. J. Neurophysiol. (2003)
14. Jung, R., Kiemel, T., Cohen, A.H.: Dynamic behavior of a neural network model of locomotor control in the lamprey. J. Neurophysiol. **75** (1996) 1074–1086