Article:

Humphries, M.D., Gurney, K. and Prescott, T.J. (2005) Is there an integrative center in the vertebrate brain-stem? A robotic evaluation of a model of the reticular formation viewed as an action selection device. Adaptive Behavior, 13 (2). pp. 97-113. ISSN 1059-7123

https://doi.org/10.1177/105971230501300203

Humphries, M.D. and Gurney, K. and Prescott, T.J. (2005) *Is there an integrative center in the vertebrate brain-stem? A robotic evaluation of a model of the reticular formation viewed as an action selection device.* Adaptive Behavior, 13 (2). pp. 97-113. ISSN 1059-7123

Is there an integrative center in the vertebrate brain-stem? A robotic evaluation of a model of the reticular formation viewed as an action selection device

Mark D. Humphries, Kevin Gurney, and Tony J. Prescott
{m.d.humphries, k.gurney, a.j.prescott}@sheffield.ac.uk

Adaptive Behavior Research Group, Department of Psychology, University of Sheffield, Sheffield, S10 2TP, UK.

Corresponding author: Tony Prescott. Phone: +44 (0)114 222 6547. Fax: +44 (0)114 276 6515

Neurobehavioral data from intact, decerebrate, and neonatal rats, suggests that the reticular formation provides a brainstem substrate for action selection in the vertebrate central nervous system. In this article, Kilmer, McCulloch and Blum's (1969, 1997) landmark reticular formation model is described and re-evaluated, both in simulation and, for the first time, as a mobile robot controller. Particular model configurations are found to provide effective action selection mechanisms in a robot survival task using either simulated or physical robots. The model's competence is dependent on the organization of afferents from model sensory systems, and a genetic algorithm search identified a class of afferent configurations which have long survival times. The results support our proposal that the reticular formation evolved to provide effective arbitration between innate behaviors and, with the forebrain basal ganglia, may constitute the integrative, 'centrencephalic' core of vertebrate brain architecture. Additionally, the results demonstrate that the Kilmer et al. model provides an alternative form of robot controller to those usually considered in the adaptive behavior literature.

Keywords: Action selection, reticular formation, genetic algorithms, robot

1 Introduction

A functioning, mortal, autonomous agent within a stimulating environment must choose and co-ordinate behaviors appropriate to both the received stimulus data (via its external sensory systems) and its current internal state if it is to survive. For an animal, inappropriate selection and co-ordination may lead to death via, for example, starvation or attack from a predator. For a mobile robot, the same errors may lead to damage or loss of power. The problem may be framed as one of *action selection*: given all available pertinent information, and a repertoire of potential actions, what mechanism should the agent employ to select the most appropriate action(s)?

If we wish to build agents (particularly animats) which adaptively and robustly function in a complex environment, a generally effective strategy is to reverse-engineer biological control systems that have already solved the problem. The process of evolution has produced animals that embody a set of competent solutions to the action selection problem; their competency is demonstrated by the individual animal's continued survival in the short-term, and the perseveration of the species in the long-term.

Ethological models of action selection are reverse-engineered from observations of animal behavior in response to varying contexts and stimuli (for example Baerends, 1970). They postulate abstract control systems coordinating elementary behaviors and thus have a strict hierarchy of control. Robot control implementations of this form of ethological model are able to perform moderately well in a simulated environment (Tyrell, 1993), but are typically restricted by their inflexibility. Control architectures based on ethological principles and using less-rigid hierarchical structures can overcome this problem (Tyrell, 1993; Blumberg, 1994).

The level of hierarchical structure necessary for a robot control system is a continuing debate in the adaptive behavior literature (Bryson, 2000; Maes, 1995; Blumberg, 1994; Tyrell, 1993). There is a corresponding debate on the necessity for a specialized selection device as opposed to selection emerging from the architecture, what we may call centralized versus distributed selection. Fully heterarchical (flat) control systems necessarily contain emergent selection. Thus, adaptive behavior researchers proposing robot control systems are compelled to address both these issues in the design of their systems.

Our approach has been to create robot control architectures by reverse-engineering the animal's central nervous system: key components of the controller are accurate models of neural systems known to be involved in action selection. By examining the functional architecture of neural systems involved in action selection, we can determine how a particular successful solution to the action selection problem works, and state with some confidence how the mechanism is organized. The degree of hierarchical or heterarchical structure and of centralized or distributed selection within our robot controllers is specified by the biology rather than by the search for some optimally-operating controller (in turn, the performance of our robots will contribute data on the appropriateness of these design choices for control architectures). Thus, the neural-based approach gives us a starting point for candidate robot control systems. Moreover, implementing our neural models of action selection systems as control architectures is also a strong test of those underlying models: if the robot performs poorly, then our model is likely to be incorrect. This paper reports work which continues our investigation of the functional architecture of action selection in the vertebrate brain.

2 Layered architecture, centralized control

A review of the vertebrate central nervous system's global organization and evolution led us to propose that there is strong evidence for a neural substrate of a layered control architecture (for details see Prescott, Redgrave, & Gurney, 1999). The rat's defense system illustrates just such a layered architecture in that the reactions to increasingly complex classes of stimuli are determined by brain structures which are higher up the neuraxis. This implies a distributed control system in which multiple stimuli from more than one class can be processed in parallel. A constraint is that any given animal has a final common motor pathway: the connections of the spinal cord and the number of muscle groups limit the set of actions that can be expressed simultaneously. Therefore, some mechanism is required to reduce the actions represented by the outputs of the layers to just those actions capable of simultaneous expression. In terms of the centralized versus distributed selection question, there are three alternative mechanisms: the higher layers may suppress the responses of the lower layers given appropriate stimuli, there may be competition between layers, or there may be a central selection device.

2.1 The vertebrate basal ganglia as an action selection device

We have proposed that, given the constraints of biological tissue, a central rather than distributed selection system would be the preferred solution implemented by a neural substrate (Prescott et al., 1999). Briefly, our argument runs as follows. Typically, distributed selection mechanisms, formed by reciprocally inhibitory links between n behaviour-representing nodes, contain $n(n-1)$ links and grow as $2n$ for every additional node. By contrast, a central selection device which is reciprocally connected with all n nodes (thus allowing control over the expression of each node's represented behavior) requires just $2n$ links, and grows by 2 for each additional node. Thus, a central selection device is more economical in both the number of connections required and the cost of adding nodes. Such economy of wiring appears to be a priority for the central nervous system (Cherniak, 1994).

There exists a group of structures in the vertebrate brain, the *basal ganglia*, that have the necessary inputs, outputs, and internal connectivity to function as just such a central switching system and which are intimately involved in behavioral control (Redgrave, Prescott, & Gurney, 1999). Computational modeling of the intrinsic basal ganglia circuitry has demonstrated that it is capable of resolving competition between action-representing signals (saliences) such that the basal ganglia output expresses the selection of the most appropriate action(s) and suppresses the others (Gurney, Prescott, & Redgrave, 2001). We refer to this as the GPR model hereafter. Using the GPR model as a control architecture for mobile robots has demonstrated that the ability to resolve individual selection competitions results in coherent sequencing of behavior in both a foraging task (Montes-Gonzalez, Prescott, Gurney, Humphries, & Redgrave, 2001) and a survival task (Girard, Cuzin, Guillot, Gurney, & Prescott, 2003).

2.2 Central control does not extend to lower levels of the neural substrate

We know that the basal ganglia cannot be the only action selection system operating in the vertebrate brain. Altricial[1] neonates and decerebrate rats and cats have a limited behavioral repertoire that can be expressed in the absence of basal ganglia (in altricial neonates it is not connected; in decerebrates it has been lesioned). For example, neonatal rats

---

[1] helpless at birth

have a complete set of ingestive behaviors, such as lapping, by postnatal day three (Hall, 1979) and can spontaneously groom by two weeks postnatal (Berridge, 1994); decerebrate rats can also spontaneously groom, locomote, feed in a coordinated manner, and have intact fear, escape, and defensive responses (Berntson & Micco, 1976). Thus, some neural structures within the intact brainstem must also be capable of functioning as a limited action selection system.

Neuroscientists have long suspected this to be true. The eminent neurobiologist and cybernetics pioneer Warren McCulloch proposed the mode selection[2] hypothesis of brainstem function: he identified 25 or so incompatible general modes of behavior common to all vertebrate animals, such as sleeping, fighting, grooming, and fleeing. An animal could be considered in a particular mode if its central nervous system was primarily focused on executing components of that mode. McCulloch proposed that the core of the reticular formation (RF), the neural structure at the center of the brainstem, was the substrate of the mode selector. As we detail below, his justifications for this proposal remain valid and, therefore, we agree that the RF is the potential neural structure for a selection mechanism in the brainstem.

In a landmark paper, Kilmer, McCulloch, and Blum (1969) presented a computational model of RF function which demonstrated mode selection in simulation. This was the first computational model explicitly constrained by the known anatomy and physiology of a neural structure, and its importance is reflected in the continued citation of the model to the present day (Leibetseder & Kamolz, 2004; Delgado, Mira, & Moreno-Diaz, 1989; Barto, 1985). Remarkably, the general structure of the model is still consistent with more modern data on RF organization and neuron morphology. Moreover, there have been no alternative quantitative models of the RF published in the interim. Indeed, the model has been recently revised by Kilmer himself (Kilmer, 1997). Thus, this model is of more than historical interest: it remains a valid model of the RF, and a valid selection mechanism. Given the existence of this computational model (which we shall refer to as the Kilmer-McCulloch model) of RF function, we propose to replicate and test it to assess its suitability to form the action selection mechanism for the lower levels of the layered-architecture.

To recap, our primary hypothesis is that the RF and basal ganglia form separate action selection mechanisms, which must interact in the intact vertebrate brain to generate coherent sequences of behaviors. Indeed, these structures together form what has been termed the brain's "centrencephalic core", a network of centralized brain structures that co-ordinate and integrate the activity of neural centers throughout the brain (Penfield, 1958; Thompson, 1993). The exact form of the interaction is open to investigation, given the paucity of data on the relationship between behavior and basal ganglia–RF connectivity. However, before implementing a complete neural model which contains both basal ganglia and RF components, we must first assess the hypothesis of RF as action selector.

2.3 Objectives

The aims of this study were: (1) to assess the functional capabilities of the Kilmer-McCulloch model in simulation as a guide to subsequent robot experiments; (2) to implement the Kilmer-McCulloch model as a robot control architecture, and compare its performance to that of alternative controllers. We would thus be able to assess both the hypothesis of action selection by the RF, and the general suitability of the model as a robot controller; (3) to determine if any versions of the Kilmer-McCulloch model could perform well as a robot controller independent of any neural-modeling constraints (we use here a genetic algorithm to search the space of Kilmer-McCulloch model variants).

On an historical note, McCulloch and his colleagues were keen to emphasize that the model of RF could be used as a robot controller, an idea that was echoed in Kilmer's recent paper. Thus, the testing of this model as a robot control architecture allows us to fulfill McCulloch's original hope for the model for the first time.

3 The Kilmer-McCulloch model

3.1 Anatomy of the RF

We briefly outline the anatomical features of the RF (summarized in Figure 1), predominantly based on the Scheibel's neuron-staining studies (Scheibel & Scheibel, 1967), which led McCulloch to propose this region as the substrate for a plausible selection mechanism. These anatomical findings have been repeatedly replicated using a variety of staining and microscopic techniques (Jones, 1995; Newman, 1985; Bowsher & Westman, 1970).

The predominant neuron type in the medial core of the RF has a giant body and undifferentiated

---

[2] Mode selection is synonymous with what adaptive behavior researchers now term action selection

radially symmetric dendrites which do not significantly extend in the anterior-posterior direction. Their axons bifurcate, extending posteriorly to the spinal cord and anteriorly to the midbrain, giving off numerous collaterals along their trajectories. The collaterals form synapses in the extensive dendritic trees of other giant cells. These giant neurons thus form discrete discs or *modules* of overlapping dendritic fields which are connected by far-reaching axons.

The giant neurons receive extensive primary and secondary afferent ascending sensory input via the spinothalamic tract, sensory trigeminal nuclei, vestibular nuclei, and other brainstem relay nuclei. Thus, the modules are in a position to sample from every sensory system available to an animal. This array of inputs has led to the medial RF being termed the 'integrative core' of the brainstem.

3.2 The computational model

We shall detail the recent revision (Kilmer, 1997) of the Kilmer-McCulloch model as it is more amenable to replication than the original version. The anatomical features reviewed above form the basis of the model design shown in Figure 2. Each of the U modules of the model corresponds to an anatomical module of the RF described above; the internal computations of a module are detailed below. Sensory input to the RF is represented by S sensory systems, where the kth output of each S_j represents that system's estimate of the probability of behavior k being selected (and is therefore in the range (0,1)). The model has M behaviors represented by M descending and M ascending connections from each module and M outputs from each S_j. Each module U_j receives M inputs, each input being from the corresponding mode output of a randomly selected (with equal probability) sensory system: for example, input $k = 1$ to module U_1 comes from output $k = 1$ of a randomly selected S_j.

The ascending/descending connections are analogous to the far-reaching bifurcating axons from the giant cells in each module. The probability of a module receiving an ascending or descending connection from a particular module was specified by a power law: for any given pair of modules (i, j), the probability P_{ij} of a connection from j to i was given by $P_{ij} = d_{ij}^{-r}$, where d_{ij} is the distance between the modules,

$$d_{ij} = (U + |i - j|) \bmod U, \qquad (1)$$

and the exponent r is some positive integer (we use $r = 2$ throughout). For each module connection, a

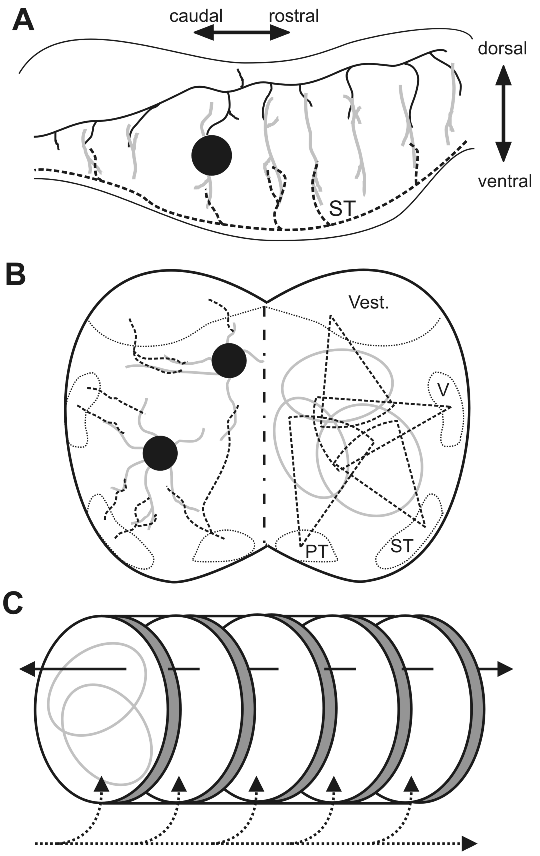

Figure 1: Schematic summary of the vertebrate reticular formation's anatomical organization. *A* Sagittal section through the brainstem; the dendritic trees (black lines) of the giant cells (single cell body shown) extend throughout the RF core along the dorso-ventral axis but extend little along the posterior-anterior axis. These dendritic trees contact axon collaterals of both ascending sensory systems (grey dashed line) and far-reaching axons of the giant cells (the axon of the depicted cell body is shown by the solid grey line). *B* Cross-section through the brainstem (dash-dot line indicates the midline). Left: dendrites extend radially about the giant cells' bodies, often preferentially directed to the axon collaterals extending from the passing sensory fiber tracts. Right: the giant cells' radial dendritic fields and passing sensory fiber tracts' axon collateralisation create overlapping fields of synaptic contact. *C* The Scheibel's summary of RF organization: the RF core is comprised of stacked disc-like modules containing giant cells, with limits defined by the dendritic extension from the cell bodies. The radial dendritic fields allow sampling of ascending and descending input from both other modules (solid grey line) and sensory systems (dashed grey line). Abbreviations: PT - pyramidal tract; Vest - vestibular complex; V - sensory trigeminal system; ST - spinothalamic tract.

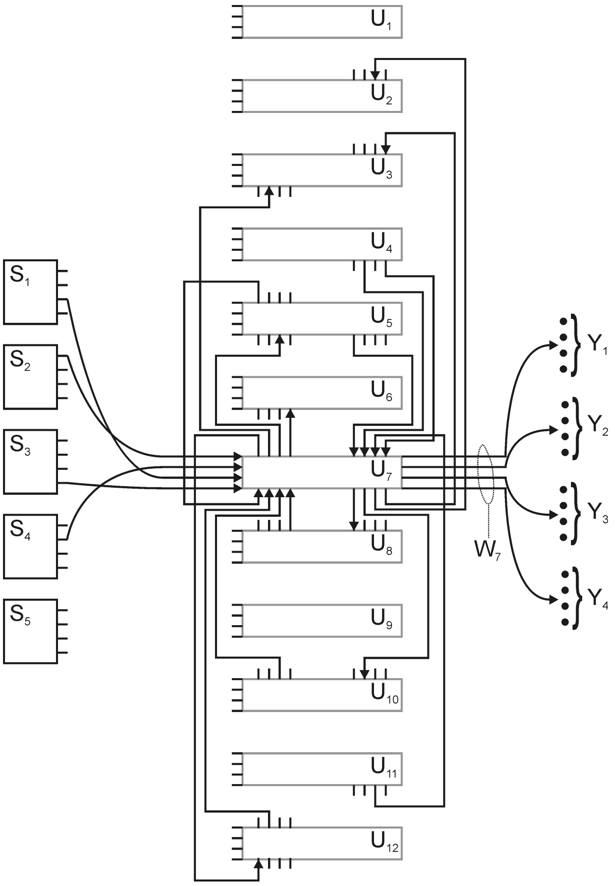

Figure 2: Schematic diagram of a version of the Kilmer-McCulloch model. This particular version has $S = 5$ sensory input systems, $U = 12$ modules, and parallel ascending/descending connections representing $M = 4$ modes. Only the connections to and from a single module are shown (from Kilmer, 1997).

source module was randomly selected and P_{ij} assessed; this was repeated until a connection was made.

The output W_i of each module U_i contributes its kth element to the vector Y_k (and thus each Y_k has U components); selection of behavior k (or *convergence* on that behavior) is signaled by the following conditions: at least $U(1-\delta)$ of Y_k's components have high values (H) and at least $U(1-\delta)$ of all other Y_j's components have low values (L), where $H \geq 1-\epsilon$ and $L \leq \epsilon$. In all that follows, we follow Kilmer and use values $\delta = 1/6$ and $\epsilon = 0.49$. With these specific values, convergence on behavior k would occur if more than 5/6 of the output vector Y_k's components were greater than 0.51 and if more than 5/6 of the components in each of the other Y vectors were less than 0.49.

3.2.1 Module design

We have abstracted a single artificial 'neuron' from the module description of the Kilmer-McCulloch model so that it may be more easily compared to artificial neurons in general use. There are M of these units in each module, one for each behavior, and their only point of interaction is at the normalization step (equation 3 below). Formally, the output p_k of the kth unit of module U_i can be expressed as

$$q_k^i = \frac{X_{sk}^2 + \Gamma(A_{dk}^2 + B_{bk}^2)}{1 + 2\Gamma} \qquad (2)$$

$$p_k^i = q_k^i / \sum_{j=1}^{M} q_j^i \qquad (3)$$

where Γ is the *coupling co-efficient*, X_{sk} is sensory input from the kth component of the sth sensory system, A_{dk} is the "descending" input from the kth component of the dth module, and, similarly, B_{ak} is the "ascending" input from the kth component of the ath module.

3.2.2 Operation of the Kilmer-McCulloch model

At $t = 0$ a new set of inputs are presented at the sensory systems S and a new set of S-U (sensory-system-to-module) and U-U (module-to-module) connections are created. The modules then compute their outputs for each behavior using the computations given by equations (2) and (3). Γ is set to zero at this first time-step. At each successive time-step the modules compute their outputs (the S values remain fixed) following two changes: first, Γ increases by 0.25 up to a limit of 2, where it remains thereafter; second, every U_i is randomly assigned a new module to receive each of its kth descending and ascending inputs from. When the convergence criteria are met the time-step T is recorded: the elapsed number of steps from $t = 0$ to $t = T$ is termed an *epoch*.

3.3 Simulation results

Our initial task was to run computer simulations of the Kilmer-McCulloch model and investigate Kilmer's (1997) claims for its dynamical properties. Specifically, he stated that the model would always converge (as defined above) within 30 time-steps: a value he claimed to be sufficiently rapid to demonstrate that the RF could support mode selection. We replicated the model version which he investigated, with $S = 5$, $U = 12$, and $M = 4$, and simulated it for 10000 epochs. Inputs from the five sensory systems were sampled from a uniform random distribution within the range (0,1).

The four modes were each selected in roughly equal proportion ($\sim$ 2100 times each), as would be expected when assigning the input values using a uniform random distribution. However, no convergence on any mode within the 30 time-steps occurred for 1501 epochs, which is roughly 15% of the total set of simulations. Therefore, Kilmer's claim that the model always converges is not entirely correct: there are some inputs for which the model does not quickly converge (which is not to say that it never converges).

3.3.1 Using fixed module-to-module connections

We have investigated many aspects of the Kilmer-McCulloch model in simulation (Humphries, 2003), but have space only to report the most pertinent results for the use of the model in the robot (see section 5). In particular, we wished to know if the random re-assignment of module-to-module connections at every time-step was necessary because (a) this operation was difficult to reconcile with plausible biological operations and (b) it made analysis of the model's dynamics impossible. Thus, the model was simulated for a further 10000 epochs, with the module-to-module connections randomly specified at the beginning of each epoch and then not changed. We found that the lack of module connection reassignment did not change the convergence proportions.

4 Embodying the Kilmer-McCulloch model

Given the failure of the Kilmer-McCulloch model to converge for all input sets, we may ask why we should continue to investigate it in an embodied form. The simple answer is that we cannot truly probe a model's capabilities using random noise: the inputs it receives when embodied in a real-world environment are not just noise, and thus may never stray into the regions of input-space which result in the non-convergence of the model. Moreover, we do not know *a priori* whether the non-convergence of the model is sufficiently great to prevent it from successfully coordinating actions in the long-term. Thus, we proceeded to test the Kilmer-McCulloch model as a robot controller.

Our GPR model of basal ganglia function (Gurney et al., 2001) performed well in a robot survival task in which the robot's goal was to survive by continually storing and recharging energy from specific locations indicated by colored squares (Girard et al., 2003). It is this task which we chose to assess the Kilmer-McCulloch model on, because it provides a set of quantitative measures, such as survival time, that can be used to assess the relative merits of different models. A simple winner-takes-all (WTA) selector and a random controller are also assessed on this task as control conditions against which the performance of the Kilmer-McCulloch model implementation may be compared.

4.1 The task

The form of the task is as described in Girard et al. (2003): a mobile robot explores an arena with a grey colored floor (representing neutral) upon which are laid two white and two black tiles. The robot continually consumes energy, and may recharge it from a separate energy store while stopped on a white tile; it may recharge the energy store when stopped on a black tile. When all energy has been used, the robot expires. The aim of the task is to maximize the lifetime of the robot.

We used a Hemisson robot (K-Team, Switzerland) for the real-world experiments and a Webots (Cyberbotics, Switzerland) simulation of the same robot and arena combination to both speed up data collection and allow us to test in an entirely noise-free environment, thus ensuring optimal performance from the WTA selector. From the robot's array of sensors we have used the two downward firing infrared sensors for determining floor color and the front-left and front-right infrared sensors to represent bumpers for compatibility with Girard et al's sensory variables.

4.2 The robot's state variables and action repertoire

The robot controller has six state variables available to it, four external and two internal: B_L and B_R represent the binary state of the left and right bumpers (a value of 1 represents contact); L_B and L_D, the *Brightness* and *Darkness* values (derived from the infrared floor sensors) are also binary and represent the floor color ($L_B = 1$ on white, $L_D = 1$ on black; $L_B = L_D = 0$ on neutral); P_E represents the potential energy (which is recharged on black tiles); and E represents the robot's energy (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).

Girard et al. (2003) specified the following equations for E and P_E changes. The change δP_E in potential energy when recharging on a black tile for

T_{eat} seconds is

$$\delta P_E = 0.027\, T_{eat}\, L_D.$$

The change δE in energy when recharging from stored potential energy on a white tile for T_{digest} seconds is

$$\delta E = 0.027\, T_{digest}\, L_B$$

and the corresponding decrease in P_E is

$$\delta P_E = -0.027\, T_{digest}\, L_B.$$

The robot has four selectable actions available to it, each of which take a fixed number of time-steps (one time-step is one second):

- *Wander*: a random walk in the environment, formed by forward movement at a fixed speed followed by a turn of a randomly selected angle (2 time-steps).
- *Avoid Obstacles*: a maneuver to re-enter open space; the robot moves backwards followed by either, if both bumpers activated, a turn of 180^o or, if one bumper activated, a turn of 45^o in the opposite direction to the activated bumper (2 time-steps).
- *Reload On Dark*: stop on a black tile and charge potential energy (1 time-step).
- *Reload On Light*: stop on a white tile and charge energy by consuming potential energy (1 time-step).

Regardless of the action selected, energy E is consumed at a constant rate of 0.002 unit/s. At the completion of the currently selected action, the controller uses the current sensory data to select a new action (we call this the *behavioral update*). If the controller is unable to resolve the selection competition to one of the above actions, then *Rest* is selected for 1 time-step, during which the robot is stationary but consumes energy at the same rate.

4.3 Implementing the controllers

In the original implementation of the task (Girard et al., 2003), the selection of the four actions was based on their saliences: values which indicate the level of urgency or motivation to perform that action. To calculate the saliences Girard et al. handcrafted the following equations

$$\begin{aligned}
S_W &= -B_L - B_R + 0.8(1-P_E) + 0.9(1-E)\\
S_A &= 3B_L + 3B_R\\
S_D &= -2L_B - B_L - B_R + 3L_D(1-P_E)\\
S_L &= -2L_D - B_L - B_R +\\
&\quad 3L_B(1-E)[1-(1-P_E)^2]^{1/2}
\end{aligned}$$

which are, respectively, the salience calculations for *Wander* (S_W), *Avoid Obstacle* (S_A), *Reload On Dark* (S_D), and *Reload On Light* (S_L).

These salience values are calculated at each behavioral update. The WTA controller thus simply selects the action with the highest salience value as the winner, and the robot executes that action. The random controller simply randomly selects one of the five possible actions with equal probability at each behavioral update.

For direct comparison with the WTA controller, and for ease of comparison with the previous work using the GPR basal ganglia model, we used the salience equations with the Kilmer-McCulloch model. Given that there are 4 actions, we required $M = 4$ mode lines to represent them. Thus, we used $S = 4$ sensory systems for which the corresponding output carried the salience value. That is, for S_1 output 1 had value S_W, for S_2 output 2 had value S_A, and so on. We used $U = 12$ modules to enable direct comparison with the simulated version used in section 3.

The squaring operation (equation 2) performed by the modules on their sensory inputs means that negative salience values would be incorrectly used. Thus, we threshold the salience equations using the Heaviside step function, so that the salience values are either positive or zero. In addition, we found in simulation that convergence often failed if there were zero-valued inputs; thus we added a small amount of noise to the sensory systems' outputs at each behavioral update, sampled from a Gaussian distribution with variance of 0.001.

At each behavioral update the salience values are calculated from the state variables and presented at the appropriate outputs of the sensory systems, as described above. The model is then run for 30 time-steps or to convergence, whichever is sooner. If convergence is not reached, then the *Rest* behavior is selected. Otherwise, the behavior on which the model converged is executed, where behavior 1 is *Wander*, behavior 2 is *Avoid Obstacle*, and so on.

4.4 Results

At the beginning of each run (whether simulation or real-world) the robot was initialized with $E = 1$ and $P_E = 0.5$ and placed at a random location in the arena. Therefore, if no recharging of energy occurred then the minimum survival time was 500 seconds.

We tested two different forms of the Kilmer-McCulloch model. The first was the original model: a new set of connections was created at each behavioral update in an identical manner to the simulated model described in section 3. The second was the

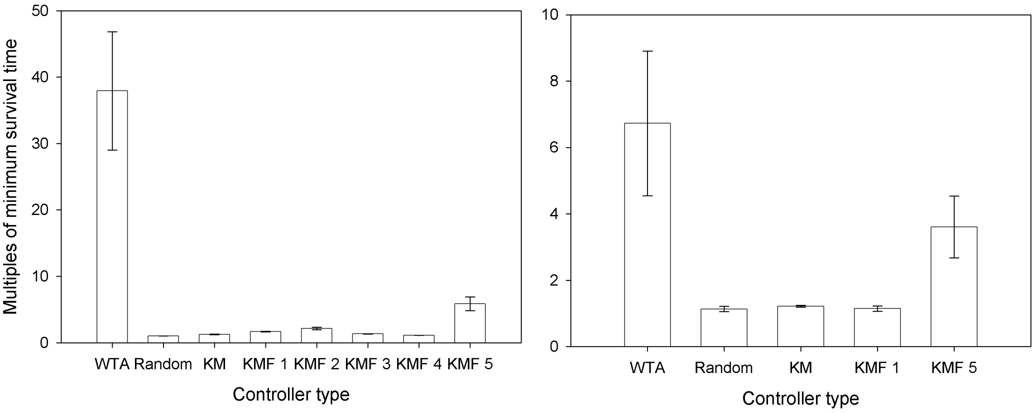

Figure 3: Mean survival times on the energy task for the robot controllers, expressed as multiples of the minimum possible time. Left: simulated task; robots using the winner-takes-all (WTA) controller survived for substantially longer than all other robot-and-controller combinations. Right: using the real robot; the relative performances of the controllers is the same as that for the simulated version. Error bars are ±1 S.E.

"fixed" model: a new set of connections was randomly created at the start of each run of the robot and then *not* altered at the start of each behavioral update so that the sensory-system-to-module and module-to-module connections were the same throughout. Thus, the "fixed" models created a set of random samples from the Kilmer-McCulloch model configuration space.

The majority of robot experiments were conducted in the Webots simulation environment, with confirming tests run on the Hemisson robot. We present first the major simulation results.

4.4.1 The Webots simulation of the energy task

We tested the random, WTA, original Kilmer-McCulloch model, and 5 different realizations of the "fixed" Kilmer-McCulloch model based controllers 20 times each. The robot started from a different position in the arena on each of the 20 tests, the sequence of positions being initially randomly chosen and then repeated for each controller. In the following, we denote the original Kilmer-McCulloch model as KM, and the five "fixed" models as KMF 1, KMF 2, and so on.

The robot consistently survived longest using the WTA controller, with a mean time of 18972 seconds; using the random and KM controller, the robot survived little longer than the minimum possible time (Figure 3). Using most variants of the "fixed" Kilmer-McCulloch model as the controller also resulted in comparatively low survival times. However, for one version (KMF 5) the robot survived considerably longer (mean 2937 seconds; maximum 7953 seconds) than all controllers other then WTA, and than the minimum survival time. Thus, some fixed configuration of connections for the Kilmer-McCulloch model can result in improved robot performance compared to the original Kilmer-McCulloch model.

We compared the robot's behavior patterns from the tests using the WTA, KM, and KMF 5 controllers, to determine what differentiated their performances. Specifically, we measured the mean duration (period of consecutive selection) and the mean frequency (per 100 seconds – thus providing a basis for comparing frequencies between simulations which lasted different periods) of each action's selection, averaged over all 20 simulations. Figure 4 summarizes these measures. The WTA controller selected the reloading actions (*Reload On Dark* and *Reload On Light*) for longer than the other controllers, with a correspondingly lower frequency of selection. By contrast, the KMF 5 controller selected the reloading actions for considerably less time, but far more frequently. The KM controller did not select the reloading actions for long periods nor did it select them frequently. These characterizations of behavior patterns are borne out by the sequences of behavior shown by the robot (Figure 5). Thus, the failure of the KM controller may be attributed to it not selecting the reloading actions frequently enough to compensate for the short duration of their selections. We note from the behavior analysis (Figure 4) that *Rest* is never selected by either the original or "fixed" Kilmer-McCulloch model versions. Therefore, convergence of the model always occurs given real-world sensory

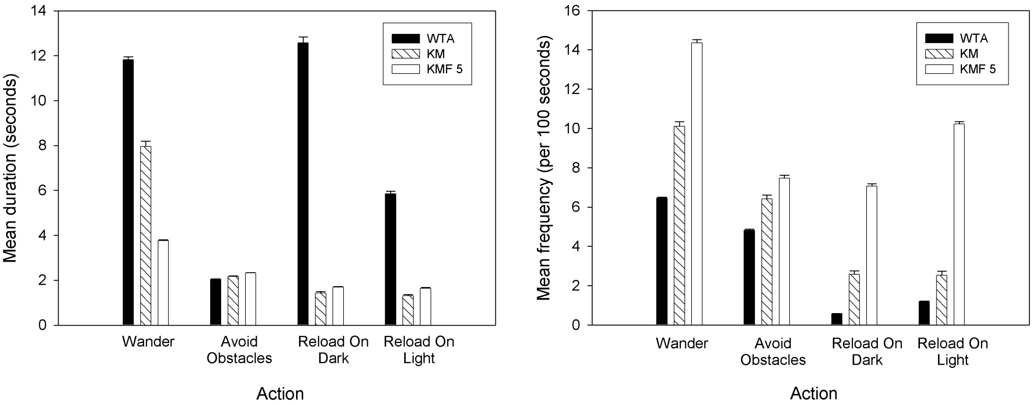

Figure 4: Mean frequency per 100 seconds (left) and mean duration (right) of each behavior for three controller types: the winner-takes-all (WTA) controller, the original Kilmer-McCulloch (KM) model based controller, and the "fixed" KM model controller KMF 5 (the version which survived the longest). Error bars are ±1 S.E.

data, rather than inputs of pure noise. This result emphasizes the importance of testing neural models in embodied systems as well as in simulation, for the results of section 3.3 may have led a researcher to immediately abandon the model.

4.4.2 Real-world experiments

We used the real robot and arena to repeat a subset of the experiments described above, using the random, WTA, KM, KMF 1, and KMF 5 controllers. Each controller was tested three times, the robot starting from a random location in the arena on every test. To keep the experiment times reasonable, we removed the black and white tiles after 5000 seconds and allowed the robot to expire; this only occurred for the WTA controller. Nevertheless, the relative performance of the controllers is the same as for the simulated robot (Figure 3).

4.5 Summary

Our results were consistent across the simulated and real robot tests; we base our comments on the simulated versions as they were repeated more often. We used a random controller to demonstrate that the design of the arena is a sufficient test of the other controllers. If a random selection of actions was sufficient for the robot to survive for considerably longer than the minimum time, then we would not be able to compare the other controllers on the basis of survival time, as their performance would be indistinguishable from chance. Therefore, because the robot using the random controller survived no longer then the minimum time, we are able to compare the other controllers' performances.

The implementation of the original Kilmer-McCulloch model as a robot controller failed as an action selection mechanism on the simple behavioral task we used. This is evidence that the original model is not an adequate model of the action selection capabilities of the RF, assuming that the RF is the action selection mechanism of the brainstem. Moreover, as it performed no better than the random controller yet, unlike that controller, had information-carrying inputs (saliences), we conclude that the computations of the original model actually degraded the input information.

However, given that we found "fixed" versions (KMF 5) of the model which resulted in considerably greater survival times than the minimum time, it is an open question whether or not there are other configurations of the model which could perform as well as the WTA algorithm when used as a robot controller. If these configurations exist, we would like to know what structural features make them successful (and how these map onto the behavior patterns described above), and how easy it is to find them.

5 Optimising the Kilmer-McCulloch model

5.1 Possible modifications

We know that some aspect of the Kilmer-McCulloch model's structure can be fixed so that a robot controlled by such a model performs competently during the energy task (section 4.4.1). In addition, we know that the randomization of module-to-module

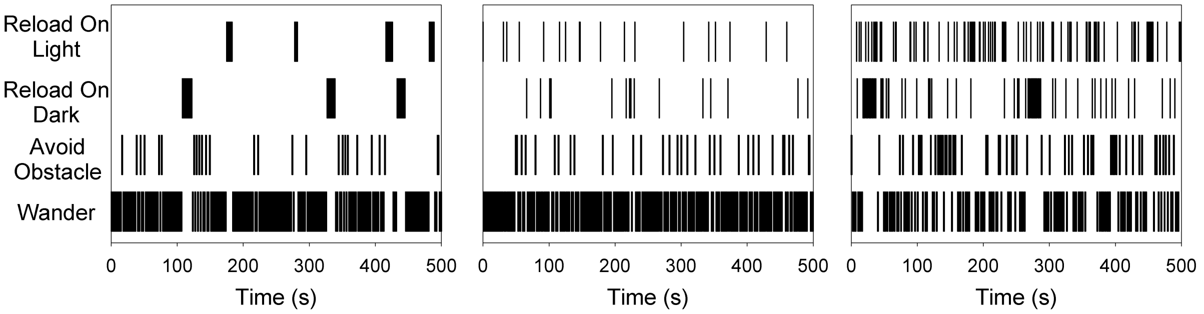

Figure 5: Example behavioral sequences from the simulated Hemisson, showing the selected actions and their durations for the first 500 seconds of the robot test. Left: the winner-takes-all (WTA) controller. The robot has infrequent but long selections of the reloading actions; Middle: the original Kilmer-McCulloch model; Right: the best performing "fixed" Kilmer-McCulloch model (KMF 5). The robot has frequent short selections of the reloading actions. The robot started at the same position and orientation in the arena for the tests which generated these sequences.

connections made no difference to the ability of the model to converge (section 3.3). Therefore, we hypothesize that it is only the particular configuration of sensory-system-to-module connections which crucially determine the selection capabilities of the Kilmer-McCulloch model.

The following describes how we searched the configuration space using the Webots simulation of the robot task (carrying out such a search using the Hemisson would have been too time-intensive). For all the robot tests, the module-to-module connections were randomly but evenly distributed, such that every mode output from a module contacts exactly two other modules (one an ascending, and one a descending connection). Thus, the outputs of each module were evenly sampled by the other modules. This connection set was maintained across all the robot experiments discussed below.

5.2 Evolving structure using a genetic algorithm

The model used for the robot has $S = 4, U = 12, M = 4$. There are $MU = 48$ sensory-system-to-module connections to specify, where each connection takes an integer value specifying its originating sensory system. Therefore, there are $4^{48} = 7.9928 \times 10^{28}$ possible combinations of connections. We cannot reasonably explore such a large parameter space through random search, so our strategy is to evolve the connections using a genetic algorithm (GA) and examine the structure of the models which have the greatest genetic fitness. The following describes our design choices for the chromosome, measurement of fitness, and form of algorithm.

The chromosome for the GA is straightforward: it has 48 elements, one for each sensory-system-to-module connection, which can take integer values in the range (1,4) specifying the sensory system from which the connection is made. The first four elements specify the connections to module 1, the next four elements specify the connections to module 2, and so on.

The key design choice for a GA is the measurement of genetic fitness which is used to rank the chromosomes in order of performance and to decide which produce offspring and which are removed. For this study, we are applying the GA to a problem which has a direct biological parallel, that of evolving an action selection mechanism. Therefore, the obvious choice for measure of genetic fitness would be survival time. However, as we have seen in section 4.4.1, it is possible for a robot to survive in excess of 20000 seconds and so such a GA would in practice take many days, even in simulation; moreover, given that survival time is unbounded, it is possible that a robot will never expire and, therefore, will not be assigned a fitness value.

As an alternative to survival, it seems reasonable to suppose that a biological controller could be attempting to maximize the agent's (or animal's) energy in the short-term thus ensuring that the agent is able to reproduce. In other words, the selection pressure is exerted on the evolution of controllers which maximize energy rather than survival time (as the latter subsumes the former: a controller which minimizes energy will inevitably lead to short survival times). Therefore, we measure mean E over a fixed time window of 3000 seconds; we demonstrate below that this fitness measurement is viable. (The time window was chosen to be considerably greater than the best survival time of the original Kilmer-McCulloch model, thus heavily penalizing any controllers which allowed the robot to expire). Our resulting fitness function naturally falls in the range (0,1), with 1 indicating maximum fitness.

The algorithm was specified as follows. An initial population of 20 chromosomes was created, each element chosen randomly from the possible integer range (1,4). For every chromosome population, each chromosome in turn was converted into a set of sensory-system-to-module connections, and the resulting model evaluated on the energy task.

The population was then ranked by fitness level, and the best 10 chromosomes retained. From this remaining population, 10 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 20 chromosomes results (10 parents, 10 offspring).

The new population is subjected to mutation, where each element is changed to one of the other possible integer values 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 10 consecutive generations (iterations of evaluation-ranking-selecting-mating-mutating).

5.3 Results

5.3.1 Fitness of the other controllers

As a basis for comparison, it was necessary to determine representative fitness measurements for the WTA, original Kilmer-McCulloch model, and random controllers. These were computed by averaging the fitness measurement (defined above) over 20 runs of the robot test used in section 4, each run again starting from a randomly selected position in the arena. The resulting mean fitness values were: WTA, 0.6669; original Kilmer-McCulloch model, 0.1006; random, 0.0852.

5.3.2 Using normal inputs

For the first GA test we used the inputs defined in section 4.3, with salience values on the appropriate mode output from the sensory system and low-valued noise on the others. The GA terminated after 24 generations with the best chromosome having a fitness of 0.9203 (found on generation 14), which was considerably more than all other controllers. Moreover, from the total population, over all generations, 52 (out of 480) chromosomes produced Kilmer-McCulloch controllers with higher fitness then the WTA controller. Repeated runs of the GA produced similar results. Thus, the GA found numerous versions of the Kilmer-McCulloch model which had greater fitness than the WTA controller.

5.3.3 Using probabilistic inputs

We wished to see if the model could be evolved to handle a harder version of the task, using noisier inputs to the model than those tried previously. If a successful model could be evolved, this would further demonstrate its capabilities as a general architecture for robot action selection.

The GA proceeded as described above. To create noisier inputs to the model in a consistent manner, the output of a sensory system was interpreted as a probability vector: every output value indicated the probability of that action being selected. The salience values were calculated as before (section 4.3), then normalized with respect to their maximum value. The other 3 outputs for each sensory system were then randomly assigned values that would make the total output for each sensory system equal one.

Using this input scheme, the GA terminated after 25 generations, with the best chromosome having a fitness of 0.7099 (found on generation 15). This is better than (random, original Kilmer-McCulloch), or roughly equal to (WTA), the fitness that the other controllers were able to achieve on the simpler normal-input task. Again, repeated runs of the GA produced similar results.

5.3.4 Energy-based fitness translates to survival time

To justify the comparison between the fitness of the previously tested controllers (in section 4) and the evolved versions, we must demonstrate that the mean-energy based fitness measurement is a suitable alternative to measuring survival time directly. We do this by assessing the survival times of the robot controlled by Kilmer-McCulloch models decoded from the most-fit chromosomes.

The most-fit chromosome evolved for the normal- and probabilistic-input GAs were tested on the robot experiments described in section 4: the Kilmer-McCulloch model structure was decoded, the robot started from 20 random locations in the arena, and run each time until expiration. The mean survival times (±S.E.) for the three controllers were: WTA 18972 ± 4461 s; from the most-fit normal-input chromosome 105023 ± 20465 s; from the most-fit probabilistic-input chromosome 3023 ± 382 s. Thus, the greater fitness of the normal-input evolved Kilmer-McCulloch model is reflected in its survival time. Its maximum survival time was 307841 seconds, which is roughly 85 hours (or five and a half days), demonstrating the necessity of using a fitness measurement other than survival time

To verify that the GA was finding a generally-useful model, the above test was repeated once on the real Hemisson robot, using the most-fit chromosome from the normal-input GA as the basis for the model. The robot survived in excess of 6000 seconds, thus demonstrating that the evolved model-based controller was at least equivalent to the WTA controller when used on the real robot. (We could not, of course, verify that the robot was able to survive to the same extent that it did in simulation). Therefore, using a simulation-based GA produced a controller that could be successfully used by the real robot.

5.3.5 Structure of the model from the best chromosomes

The pertinent structural feature for the normal inputs (and, as it turns out, for the probabilistic inputs) is the mapping of the salience-carrying outputs of the sensory systems to the appropriate inputs on the modules. To illustrate, consider a chromosome for which the first four elements are [3 1 2 4]. Figure 6 demonstrates the decoding of this part of the chromosome, with module U_1's inputs mapped thus: the first output of S_3 connects to input 1; the second output of S_1 connects to input 2; the third output of S_2 connects to input 3; and the fourth output of S_4 connects to input 4. It is only this last connection which carries a salience value (for *Reload On Light*). We may thus characterize a chromosome-encoded structure by the number of such salience connections from each sensory system N_s, where s is the number of that sensory system.

For the most-fit chromosomes from both the normal- and probabilistic-input GAs, we find that $N_1 < N_2 < N_3 < N_4$. That is, the sampling of *Reload On Light* salience is greater than the sampling of *Reload On Dark*, which in turn is greater than that for *Avoid Obstacle*, which in turn is greater than that for *Wander*. Thus, it appears that the GA has produced a mapping of connections which favors the reloading behaviors.

Yet, the behavioral patterns of the robot when assessing the most-fit chromosomes (Figure 7) show that, for both input types, this mapping of input connections results in a behavior pattern roughly similar to that resulting from the WTA controller. The main difference, compared to WTA, is a much shorter duration of *Reload On Dark.*

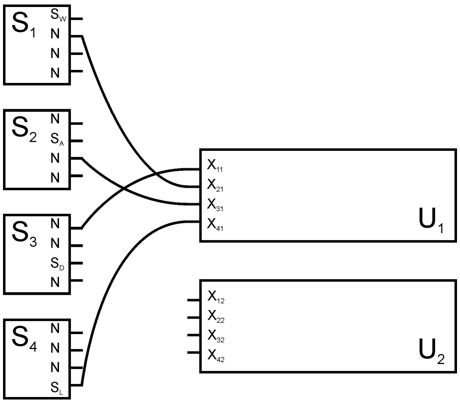

Figure 6: Decoding a chromosome into a Kilmer-McCulloch model structure. The first four elements of this chromosome are [3 1 2 4], which code the sensory system inputs to the first module U_1. Only the fourth element specifies a connection which carries a salience value, the others specify connections with noise (N).

The behavioral analysis also shows that both of the most-fit chromosomes encoded models which did not always converge. Robots using controllers based on these models enacted the *Rest* behavior at least once during each of the 20 runs. For the probabilistic-input evolved model-based controller, the frequency of occurrence was roughly the same as that of the reloading behaviors. These results are evidence that fixed Kilmer-McCulloch models which do not always converge may actually be optimal for real-world tasks.

5.3.6 Summary

For both input forms, the comparatively few generations (14 for normal, 15 for probabilistic) that were required to find the best chromosome, and the number of chromosomes better than WTA indicates that, although the configuration space is massive, many configurations work as robot controllers. Repeated robot tests using the best evolved chromosomes demonstrated the flexibility of the Kilmer-McCulloch model, as the analysis showed that it could cause a WTA-like behavior pattern as well as the frequent-selection behavior pattern seen in the previous robot test (section 4).

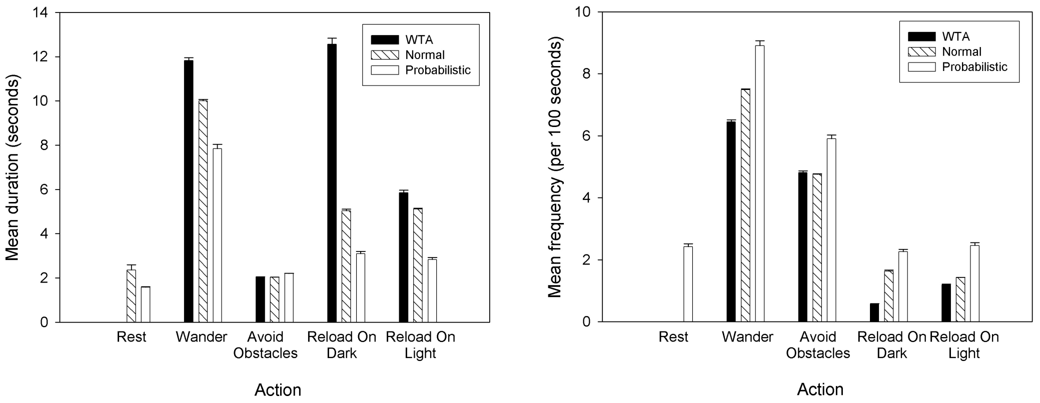

Figure 7: Behavioral statistics of the most-fit Kilmer-McCulloch models. Left: mean duration of behaviors. Right: mean frequency of behaviors. These statistics show that the most-fit Kilmer-McCulloch models on both input tests caused a behavior pattern roughly similar to that caused by the winner-takes-all (WTA) controller, though with considerably shorter durations of *Reload On Dark*. Error bars are ±1 S.E.

6 Discussion

The preceding results have demonstrated that: (1) in simulation, the original Kilmer-McCulloch model often, but not always, rapidly converges on a particular selection, and convergence does not depend on the continual random re-connection of module-to-module connections; (2) in a robotic evaluation of action selection capability, the original Kilmer-McCulloch model is inferior to a simple winner-takes-all algorithm. Random sampling of the space of possible fixed Kilmer-McCulloch architectures provides an existence proof that fixed versions can perform better than the original version; (3) using a genetic algorithm, the space of possible Kilmer-McCulloch architectures can be rapidly searched to find versions which out-perform the winner-takes-all algorithm; (4) even when using a noisier input representation, the genetic algorithm can still find Kilmer-McCulloch model architectures which perform competently; (5) robots using the best evolved architectures as controllers have very different behavioral patterns to those using the best randomly-generated architectures as controllers, evidence of the flexibility of the general Kilmer-McCulloch architecture.

6.1 The reticular formation as action selector

By our own criteria, the performance of the original Kilmer-McCulloch model as a robot controller is evidence that either it is not an accurate model of the RF or that the RF is not the action selection mechanism of the vertebrate brainstem. We believe the former to be true: a simple modification of the model – stopping all random re-assignment of connections – was able to produce competent robot controllers. Yet this modification increases the biological plausibility of the model: it is unlikely that such a neural structure, with indirect control over such basic physiological processes as respiration and heartbeat (Yates & Stocker, 1998), has a place for large degrees of randomness in its structure. Thus, we argue that it is the *operational* features of the Kilmer-McCulloch model that are not accurate, whereas the *architectural* features are accurate to the extent that we have tested them. The tentative conclusion therefore must be that the reticular formation, as conceptualized by the Kilmer-McCulloch model architecture, is a candidate for the action selection mechanism of the vertebrate brainstem.

6.2 Implications for the model of vertebrate action selection

How then to proceed with the development of the general model of the vertebrate action selection system? If we accept that the RF is the candidate action selection mechanism of the brainstem, we must address the issue of how the basal ganglia and RF interact to produce coherent sequences of behavior. From the available neuroanatomical data, it appears that the basal ganglia outputs do directly contact regions of the medial RF (Schneider, Manetto, & Lidsky, 1985), but are principally relayed via the pedunculopontine nucleus, a structure which is not represented in the Kilmer-McCulloch model (Delwaide, Pepin, De Pasqua, & Noordhout, 2000). Thus, we wish to extend the model to incorporate the necessary neural structures to test the form of interaction.

In doing so, we may incorporate more recent data on the medial reticular formation that has added detail to the Scheibel's studies which formed the basis for the Kilmer-McCulloch model. The existence of small- and medium-size neurons (Newman, 1985), and of neuromodulator (noradrenaline and serotonin) receptors in the giant cells' dendritic fields (Stevens, McCarley, & Greene, 1994; Kobayashi, Matsuyama, & Mori, 1994), has been demonstrated and these may add to the computational capabilities of an RF model. In addition, it must be admitted that the operations performed by the Kilmer-McCulloch model are somewhat removed from the now-traditional forms of neural modeling, particularly in the model "neuron" as we have abstracted it. A replication of the selection results using more realistic model neurons would provide further evidence for the RF-as-action-selector. Thus, we propose to construct a more directly biologically-constrained model of medial reticular formation function, with the hypothesis that such a model would show the same basic selection operations as the Kilmer-McCulloch model.

6.3 General action selection mechanisms

Our robot's performance demonstrates that the Kilmer-McCulloch model can form an action selection mechanism for artificial agents. Finding the best models using the GA required only a quick search of the space of all possible model configurations; and, even then, the GA found many configurations which performed better than the alternative controller types. We conjecture that such rapid optimization makes the fixed Kilmer-McCulloch model suitable for a wide range of robot tasks. It remains to be seen if the model can adapt to a changing environment using a continuous learning method rather than a GA.

We are *not* claiming that the Kilmer-McCulloch model will always perform better than the WTA algorithm. To fully demonstrate such a claim, we would have to optimize the salience equations to determine what their maximal possible fitness would be: for example, we could replace the salience equations' constants with variables that can be optimized using a GA, evaluating the fitness with a robot using the WTA controller. However, this does not alter our finding that the Kilmer-McCulloch model has potential as general action selection mechanism.

6.3.1 Potential advantages of Kilmer-McCulloch type architectures

It is worth noting how the Kilmer-McCulloch model fits into the adaptive behavior debates on hierarchy and centralized control. The model is a distributed selection architecture, but does not require inhibition: for the purposes of simulation, the mode decision was made by summing over output vectors. For a model with U modules and M modes, there are $3UM$ links. Adding a new mode requires an additional $3U$ links. However, this is a constant rate for a given model, whereas global reciprocal inhibition nets grow at an increasing rate with each additional node.

Most, if not all, current models of action selection in the adaptive behavior literature have some form of modular decomposition of the agent's behavioral repertoire (Bryson, 2000). That is, they have discrete functional modules which each represent an action or group of actions. The Kilmer-McCulloch model offers an alternative to modular representations: the behaviors are distributedly represented by connections rather than functional units within the selection mechanism. This offers a great advantage over a modular representation to hardware or biological implementations as damage to part of the system does not result in the loss of ability to represent an action in the selection mechanism.

6.4 Conclusions

What general lessons may we take from this study? Certainly there is a reaffirmation of the importance of embodying neural models. The simulated version of the Kilmer-McCulloch model did not always converge on a selection. Yet the initial robot tests demonstrated that real-world inputs are sufficiently limited (or structured) to ensure convergence for some versions of the model. And, as it turned out, the optimally-performing versions – determined by a genetic algorithm – did not require the model to always converge. The occasional inability to make a decision seems to have been a worthwhile trade-off for ensuring consistently high levels of energy.

Unintuitive results such as these demonstrate the usefulness of reverse-engineering robot controller from biological substrates. The biological action selection mechanisms may have built-in features that solve problems we are unable to anticipate, or may demonstrate the efficiency or utility of a design methodology that we had not considered. The Kilmer-McCulloch model's take on the reticular formation anatomy has provided both of these, by not always requiring a decision, and by demonstrating a modular architecture without modular representation.

We have yet to determine whether the evolution of the modern vertebrate brain has deemed the brainstem selection mechanism sufficiently useful to build upon, or if it has found a better solution within the structures of the basal ganglia. However, given that functions supported by other lower neural structures, such as the superior colliculus's role in vision, are maintained in the modern vertebrate brain it is likely that the reticular formation continues to form a crucial part of the vertebrate action selection mechanism.

Acknowledgments

We thank Jonathan Chambers for writing the initial genetic algorithm code, and Nathan Boddy for running pilot tests of the robot. This work was funded by the EPSRC under grant GR/R95722/01.

References

Baerends, G. (1970). A model of the functional organization of the incubation behaviour. *Behaviour Supplement*, *17*, 263–312.

Barto, A. (1985). Learning by statistical cooperation of self-interested neuron-like computing elements. *Human Neurobiology*, *4*, 229–256.

Berntson, G., & Micco, D. (1976). Organization of brainstem behavioral systems. *Brain Research Bulletin*, *1*, 471–483.

Berridge, K. (1994). The development of action patterns. In J. Hogan & J. Bolhuis (Eds.), *Causal mechanisms of behavioural development* (pp. 147–180). Cambridge: Cambridge University Press.

Blumberg, B. (1994). Action selection in Hamsterdam: Lessons from ethology. In D. Cliff, P. Husbands, J. Meyer, & S. Wilson (Eds.), *From animals to animats 3: Proceedings of the third international conference on simulation of adaptive behavior* (pp. 22–29). Cambridge, MA: MIT Press.

Bowsher, D., & Westman, J. (1970). The gigantocellular reticular region and its spinal afferents: a light and electron microscope study in the cat. *Journal of Anatomy*, *106*(1), 23–36.

Bryson, J. (2000). Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, *12*, 165–189.

Cherniak, C. (1994). Component placement optimization in the brain. *Journal of Neuroscience*, *14*(4), 2418–2427.

Delgado, A., Mira, J., & Moreno-Diaz, R. (1989). A neurocybernetic model of modal co-operative decisions in the Kilmer -McCulloch space. *Kybernetes*, *18*(3), 48–57.

Delwaide, P., Pepin, J., De Pasqua, V., & Noordhout, A. de. (2000). Projections from basal ganglia to tegmentum: a subcortical route for explaining the pathophysiology of Parkinson's disease signs? *Journal of Neurology*, *247 Suppl 2*, 75–81.

Girard, B., Cuzin, V., Guillot, A., Gurney, K. N., & Prescott, T. J. (2003). A basal ganglia inspired model of action selection evaluated in a robotic survival task. *Journal of Integrative Neuroscience*, *2*(2), 179–200.

Gurney, K., Prescott, T., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. *Biological Cybernetics*, *85*, 401–423.

Hall, W. (1979). Feeding and behavioral activation in infant rats. *Science*, *205*(4402), 206–9.

Humphries, M. (2003). *A critique of the Kilmer-McCulloch model of reticular formation function. ABRG 4.* (Tech. Rep.). Dept Psychology, University of Sheffield, UK.

Jones, B. (1995). Reticular formation: Cytoarchitecture, transmitters, and projections. In G. Paxinos (Ed.), *The rat nervous system, 2nd edition* (pp. 155–171). New York: Academic Press.

Kilmer, W. (1997). A command computer for complex autonomous systems. *Neurocomputing*, *17*, 47–59.

Kilmer, W., McCulloch, W., & Blum, J. (1969). A model of the vertebrate central command system. *International Journal of Man-Machine Studies*, *1*, 279–309.

Kobayashi, Y., Matsuyama, K., & Mori, S. (1994). Distribution of serotonin cells projecting to the pontomedullary reticular formation in the cat. *Neuroscience Research*, *20*(1), 43–55.

Leibetseder, M., & Kamolz, T. (2004). Are depressive persons capable of describing changes in their reactions without being able to explain them? A proof of a cybernetic hypothesis of depression. *Psychopathology*, *37*(2), 86–91.

Maes, P. (1995). Modeling adaptive autonomous agents. In C. Langton (Ed.), *Artificial life* (pp. 135–162). Cambridge, MA: MIT Press.

Montes-Gonzalez, F., Prescott, T., Gurney, K., Humphries, M., & Redgrave, P. (2001). An embodied model of action selection mechanisms in the vertebrate brain. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat, & S. Wilson (Eds.), *From animals to animats 6: Proceedings of the sixth international conference on simulation of adaptive behaviour* (pp. 157–166). Cambridge, MA: MIT Press.

Newman, D. (1985). Distinguishing rat brainstem reticulospinal nuclei by their neuronal morphology. I. Medullary nuclei. *Journal fur Hirnforschung*, *26*(2), 187–226.

Penfield, W. (1958). Centrencephalic integrating system. *Brain*, *81*, 231–234.

Prescott, T., Redgrave, P., & Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, *7*, 99–127.

Redgrave, P., Prescott, T., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*(4), 1009–1023.

Scheibel, M., & Scheibel, A. (1967). Anatomical basis of attention mechanisms in vertebrate brains. In G. Quarton, T. Melnechuk, & F. Schmitt (Eds.), *The neurosciences: A study program* (pp. 577–602). New York: The Rockefeller University Press.

Schneider, J., Manetto, C., & Lidsky, T. (1985). Substantia nigra projection to medullary reticular formation: relevance to oculomotor and related motor functions in the cat. *Neuroscience Letters*, *62*(1), 1–6.

Stevens, D., McCarley, R., & Greene, R. (1994). The mechanism of noradrenergic alpha 1 excitatory modulation of pontine reticular formation neurons. *Journal of Neuroscience*, *14*(11), 6481–6487.

Thompson, R. (1993). Centrencephalic theory, the general learning system, and subcortical dementia. *Annals of the New York Academy of Sciences*, *702*, 197–223.

Tyrell, T. (1993). *Computational mechanisms for action selection.* PhD, University of Edinburgh.

Yates, B., & Stocker, S. (1998). Integration of somatic and visceral inputs by the brainstem: functional considerations. *Experimental Brain Research*, *119*(3), 269–275.