

New neural circuits for robot phonotaxis

BY RICHARD REEVE AND BARBARA WEBB

*Centre for Cognitive and Computational Neuroscience, Department of Psychology,
University of Stirling, Stirling, FK9 4LA, Scotland, U.K.*

W. Grey Walter built robotic systems to improve understanding of biological systems. In that tradition, this paper reports ongoing work on a robot model of cricket sound localisation. The main advances are the inclusion of a much larger range of neuroethological detail, and the investigation of multimodal influences on the behaviour. The former allows exploration of the functionality of identified neurons in the insect, including the possible roles of multiple sensory fibres, mutually inhibitory connections, and brain neurons with pattern-filtering properties. The latter focuses on the inclusion of an optomotor stabilisation response, and how this might improve tracking particularly under conditions of random disturbance.

Keywords: neural circuits, cricket phonotaxis, robot phonotaxis

1. Introduction

But if the performance of a model is to be demonstrably a fair imitation of cerebral activity, the conditions of stimulation and behaviour must be equally comparable with those of the brain.

[Walter, 1961]

W. Grey Walter is rightly credited as a forerunner in many areas of the study of brain and behaviour. In the field of biologically-inspired robots, it is important to remember that his ingenious devices were seriously intended as working models for understanding biology, a 'mirror for the brain' that could both generally enrich our understanding of principles of behaviour (such as the complex outcome of combining simple tropisms) and be used to test specific hypotheses (such as Hebbian learning). The research described here follows directly in this tradition, with the focus on building robot systems as a means of exploring biological questions.

We, like Walter, are interested in the complexity of behaviour that can result from the interaction of a small number of neural elements. However, Walter chose to explore this in the context of rather abstract, generic characterisations of animal behaviour such as 'exploring', 'approaching' and 'avoiding'; whereas we have chosen to focus on modelling the specific neural circuit underlying a particular animal behaviour, the sound localisation ability of the cricket. We consider the modelling of this 'simple' insect behaviour to be the focal point for investigating a range of interesting, interconnected general issues for biology. As illustrated in figure 1, these include: the functional significance of low-level neural properties; the importance of physical embodiment as a solution and constraint on behaviour; the possible connections between biological solutions and conventional engineering approaches

to sensorimotor control and integration; and how we can scale up from single behaviours to explain the flexible interaction of multiple adaptive systems within an organism:

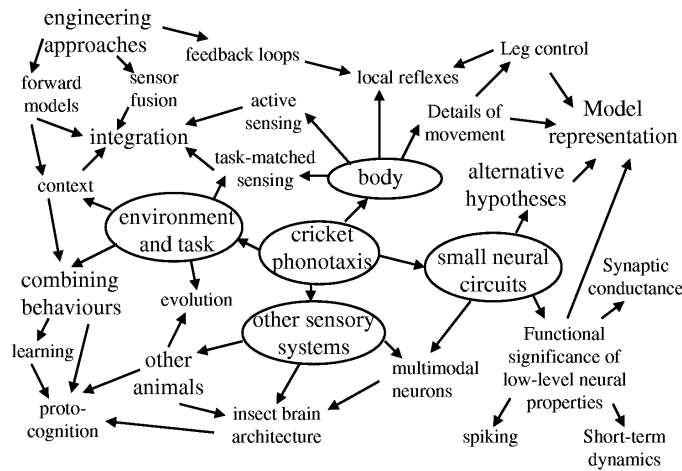


Figure 1. Cricket phonotaxis is the central focus for exploring a range of general themes in behavioural neuroscience.

Female crickets can locate conspecific males by moving towards the species-specific calling song the males produce by opening and closing their wings. Typical male songs consist of groups of short sound bursts (*e.g.* for *Gryllus bimaculatus*, four 20 ms bursts of 4.7 kHz sound) produced several times a second. Females appear to be particularly selective for the repetition rate of bursts (syllables) within each group (a chirp). The neuroethology of this system has been extensively studied [reviewed in Pollack, 1998] and we have built a series of robot models that have led to interesting reinterpretations of this data. In [Webb, 1995, Lund et al., 1997, 1998] it was shown that the recognition and localisation aspects of the task could be closely interlinked, allowing a surprisingly simple controller to produce the same kind of selective approach behaviour in a robot as was observed in the cricket. In Webb and Scutt [2000] the robot hardware was interfaced with a spiking neural network simulation, and it was shown that the algorithmic controller explored previously could be captured in a four neuron circuit (figure 2). Although the essence of this circuit is a Braitenberg-like connection between the input on each side and the motor outputs [Braitenberg, 1984], the actual function is more subtle, *e.g.* it exploits spike timing and dynamic synaptic properties to be selective for the temporal pattern of signals in a manner resembling the female cricket.

Effective biological modelling requires repeated cycles of implementation, testing and refinement. The previous robot's neural circuit could reproduce much of the cricket behaviour, but this is not sufficient to conclude that the cricket's neural circuit is the same. There may be alternative wirings that work as well or better, or are more consistent with the anatomical and physiological data. For example, the inhibitory cross-connections from prothoracic auditory neurons (AN) to opposite motor neurons (MN) in the model are not particularly plausible, and may not




Figure 2. The four neuron circuit used in previous work [Webb and Scutt, 2000]. The auditory neurons (AN) represent the first stage of processing in the cricket's prothoracic ganglion in which a small number of neural pairs receive direct input from the auditory nerve. The motor neurons (MN) produce an output signal to turn left or right depending on which AN neuron fires first. The AN-MN synapses exhibit depression, so the response is best for a signal with appropriate temporal patterning.

be necessary to create a winner-take-all response to the stronger auditory signal. Instead the AN or MN could be mutually inhibitory. There is neurophysiological evidence of mutual inhibition at the prothoracic AN level, mediated by so-called omega neurons (ON). The previous model did not include the adaptation that occurs in the auditory nerve and prothoracic neurons. It appears this property may serve to enhance the onsets of the syllables, which could help explain the apparent insensitivity of the female to variation in syllable length (Hedwig, pers. comm.), so it should be included in the new model. In addition, the cricket is unlikely to have such a direct connection between the auditory input and motor output. There is information about intermediate stages of processing in brain neurons from Schildberger [1984] that can be more directly addressed. Finally, the basic neural model used in the earlier work was itself somewhat arbitrary; so a representation of spiking dynamics that links more closely to current theoretical neuroscience would improve the biological plausibility of the circuits implemented on the robot. In section 2 (b) we describe such a model, and in section 3 (a) discuss the results of including the specific neural properties mentioned above.

An important motivation for using a robot model, rather than a simulation, is

to test the behaviour of a given neural circuit in the context of real environmental interaction [Webb, 2001]. An issue that then needs to be addressed is whether the movements of the robot replicate the speed and size of movements of the cricket; and how this relates to possible motor neuron circuitry. A first point to note is that crickets behave somewhat differently when performing this task in the dark vs. in the light. In the light, forward velocities are higher and the animal makes fewer/shorter stops [Weber et al., 1981]. This effect is probably both motivational (being in the light represents a predation risk) and a result of having visual cues available to aid navigation, as will be discussed below. In arena tests in the light, walking velocities are in the order of 5–20 cm per second, with stops typically occurring every few seconds, although some animals may not stop at all. In the dark, stopping becomes more regular. Turns in response to sound occur both after stopping and during walking, on average every two seconds. A turn made when at an angle of more than 30 degrees from the sound source will almost always be in the correct direction; but turns also occur, in a random direction, when the angular error is less than 10 degrees, or with increasing reliability to sounds between the midpoint and 30 degrees [Rheinländer and Blätgen, 1982]. It has been suggested the size of turn is fairly consistent at around 30–40 degrees [Oldfield, 1980]. In section 3 (b) we present a motor circuit that produces similar behaviour when implemented on the robot.

We are working towards the implementation of a robot system able to perform phonotaxis outdoors. As a step in this direction we are investigating the issue of how different sensory modalities involved in navigating through real environments should be integrated. It is generally agreed that the biological solution to the issue of integrating behaviours is unlikely to take the classic form of sensor fusion into a general world model over which a problem solving process produces an output. On the other hand, there are many alternative architectures that seem plausible: subsumption-type architectures that hierarchically suppress different reflexes; action selection schemes, either centralised or distributed; combining basis functions; or using some form of forward model to distinguish afferent and reafferent signals. We are exploring how simple neural interactions at the motor circuit level might be used to reproduce some of the observed behavioural effects of interacting modalities in the cricket, in particular, the effects of light level on the dynamics of approach behaviour (discussed above), and the use of an optomotor reflex to stabilise tracks, especially under circumstances of random disturbance.

The optomotor reflex is shown when an animal rotates if stimulated by global rotation of their visual field. This will produce compensation for accidental deviations in heading. Böhm et al. [1991] tested crickets in an open-loop treadmill experiment with competing sound and visual rotation signals. They concluded that the animal simply added the separate turning tendencies. However, in [Webb and Harrison, 2000] we found that such a scheme faces problems in the real, closed loop situation. If the responses to sound and vision are simply allowed to run in parallel, then whenever the robot makes an intended turn towards sound, a strong visual rotation stimulus is produced, and the optomotor reflex will cause the robot to make an unwanted correction. A simple inhibitory interaction, *i.e.* suppressing the optomotor response when responding to sound, was shown to be one way to solve this problem, but a number of alternative solutions can be suggested, such as the well known concept of 'efferent copy' signals to cancel out expected feedback.

In the final section of this paper we compare three different neurally implemented schemes for combining the behaviours.

2. Implementation

The robot base used in the following experiments is the KoalaTM; a six-wheeled base measuring approximately 40x30x20 cm with two drive motors. This choice of platform is due to the need to mount multiple sensory systems and to accommodate future aims to do phonotaxis out of doors, neither of which is practical on the small robot (KheperaTM) previously used. Although the actuation bears little direct resemblance to the cricket it is possible to use this robot to replicate appropriate speeds and turning rates. Another current project (Horchler, submitted) is to implement phonotaxis on one of the cockroach-inspired walking robots being developed by Quinn and Ritzmann [1998].

(a) Sensory systems

The main sensory systems on our robot are dedicated hardware designs for processing sound and visual signals, based closely on biology. The auditory system mimics the pressure difference receiver properties of the cricket ear [Michelsen et al., 1994] to produce a highly directional signal for a specific frequency range. Details are in [Lund et al., 1997]. The visual system is an analog VLSI implementation of the Hassenstein-Reichardt model of elementary motion detection, developed by Harrison and Koch [1998]. The directional response is summed across an array of detectors providing a signal for full-field motion, which resembles the response of certain tangential cells in the lobular plate of the fly brain, associated with the optomotor response. Cells with similar responses to visual rotation have been found in the cricket.

We are making use of the onboard Koala sensors to derive two further signals that resemble (with less accuracy) information available to the cricket. The ambient light level is being measured, which serves as an indicator of the amount of cover in the animal's current position – dark is safety. The infrared reflection detectors are being used as antennae, that is, to provide a signal giving approximate distance to obstacles in a particular direction, which is indicated by mechanical sensing of the bend point in insect antennae [Camhi and Johnson, 1999]. We have also constructed and interfaced a number of additional sensors, to detect tilt, body contact, wind direction, and compass heading. However, in this paper we only use the auditory, ambient light† and optomotor sensors.

(b) Neural model

The overall behaviour of the neurons is similar to single compartment 'leaky integrate and fire' models, but based more closely on the models described by Koch [1999]. Their state is a representation of the potential difference across the neural membrane (the membrane potential). They have a base potential, to which they will decay in the absence of external input. Synapses attach to them and raise and

† Due to problems in controlling the light level reliably during experiments and interactions with the optomotor, we simulated the inputs from this sensor.

lower their potential when their pre-synaptic neurons spike. If the potential rises above a specific level, the neuron will fire, sending a spike to any connected output synapses; then the potential will reset to some lower value, and the neuron will enter a refractory period when it will be unable to receive synaptic input.

The model synapses are conductance based, which is to say they model the conductance and battery potential of the ion channels which open when a synapse is activated, pulling the membrane potential towards the battery potential with a strength proportional to the conductance. This is more biologically realistic than the standard 'charge-dump' or current injection synapses used in most artificial neural networks, which directly raise or lower the membrane potential by depositing a small or large packet of charge into the neuron.

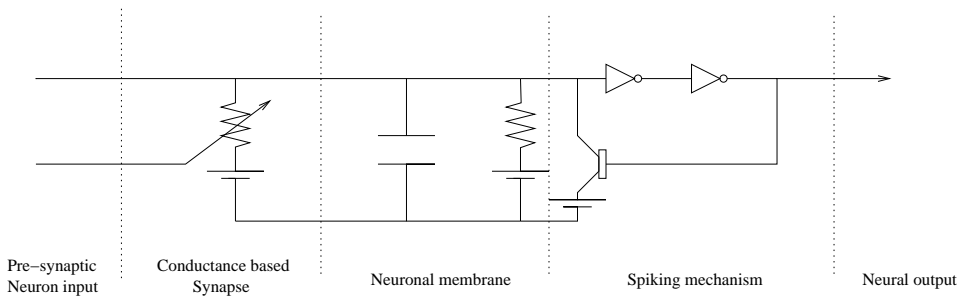


Figure 3. A simplified electrical model of a neuron with synapses

Koch (*ibid.*) explains neuron and synapse models with these features in terms of their electrical properties, thus describing neurons and synapses based on electrical circuits (this makes it very suitable for VLSI implementation). The features described above are given in the simplified circuit shown in Figure 3. The neuron is an RC circuit with a fixed (membrane) capacitance, and basal potential and conductivity, to which the neuron will decay exponentially at a rate determined by the capacitance and conductance in the absence of external input. The synapse model also has an electrical implementation: it is modelled as an additional (variable) conductance and battery, which will pull the post-synaptic membrane potential towards the potential of the battery with a strength determined by the conductance. This conductance decays exponentially to zero, but receives a boost when a spike arrives from the input neuron. The size of the boost to the conductance is effectively the 'weight' of the connections.

These 'weights' are further affected by short-term facilitation and depression mechanisms (not shown in the figure) of variable size and time constants. This allows the synapses to affect the post-synaptic neuron more or less strongly depending on whether they have recently received other spikes or analogue inputs. Indeed, it is possible to combine the effects so that, for instance, the synapses will respond increasingly strongly for the first few spikes, and then that response will die away to nothing until the synapse has had a period of time to recover. They also all have variable synaptic delays (the time from the firing of the presynaptic neuron to the reception of the excitation or inhibition by the post synaptic neuron). In

real neurons this delay involves axonal, synaptic and dendritic delays. These are all combined as a single variable in the model.

Despite this level of complexity, the neural model remains analytic under a very broad range of conditions, and the synaptic model in isolation is always closed form. Thus we are able to specify constraints and calculate exact values for parameters directly from the equations without an optimisation step, which enormously simplifies the parameter search task when compared to most sophisticated neural models.

For cases where the above neural model is too complex (for instance for sensory neurons attached to the ears), we use a simpler model of a poisson distributed spike train proportional to the synaptic input (in this case the amplitude of vibrations in our artificial tympanum) with optional depression. Again the closed form solutions to this model allow us to calculate directly the required depression parameters to match the data we have for parallel sensory fibre spike rates.

(c) *Experimental methods*

The neural simulation system described above is embedded as part of a C++ program for running robot experiments. This system deals with the serial transfer of sensory data and motor data to and from the robot, and is designed for easy transfer between different robot bases and the addition of arbitrary new sensor inputs or motor outputs. It can also be run without the robot, using previously recorded or artificially generated data files as input, which is particularly useful when tuning circuit parameters. It automatically incorporates data from an overhead camera tracking system (or alternative tracking mechanisms) into the data record produced from running the robot.

The results described below were produced in two ways: analytical and simulation results; and results from running the robot. The latter experiments occurred in the normal lab environment, within a 2x1.6 metre space determined by the overhead camera's field of view. The sound was a computer-generated song. The standard song consists of 20 ms bursts (syllables) of 4.7 kHz sound, with 20 ms gaps between them, grouped into a four-burst chirp followed by 340 ms silence. This is the characteristic calling song of male *Gryllus bimaculatus*. In some experiments the pattern was varied to produce different 'syllable repetition intervals': the length of the chirp was held constant at 250 ms, and the length of syllables within the chirp varied from 5 ms to 45 ms, with intersyllable gaps of the same length as the syllable.

3. Results

The process of designing the specific neural circuits to test on the robot was highly interactive – moving back and forth between issues of neural data and functional efficacy. We have therefore chosen to explain the various elements of the circuits with immediate reference to the relevant results. We first discuss the neural connections underlying the processing of the sound signal, then the motor control, and finally the combination of these circuits with optomotor behaviour.

(a) *The phonotaxis circuit*

The previous neural model tested on the robot focussed on one pair of identified ascending interneurons (AN1) in the cricket's prothoracic ganglion that appear to be critical for phonotaxis. AN1 respond best to sound at the calling song carrier frequency, and clearly encode the pattern of the song in their spiking response. Hyperpolarising one of the pair leads to a change in walking direction. In the current model we incorporate more of the neural data from the cricket, motivated by functional hypotheses about the various neural elements, aiming to reproduce specific response properties. The new model is illustrated in figure 4.




Figure 4. The new neural circuit for phonotaxis, based on cricket neuroanatomy, as described in text. 'AN', 'ON' and 'BN' refer to specific 'ascending', mutually connected 'omega', and 'brain' neurons that have been identified in the cricket nervous system.

First, we have represented the parallel sensory fibres that provide input from the tympanum (on the cricket's legs) to the prothoracic ganglion. In the cricket there are around 50–60 sensory neurons, with perhaps half of these tuned to the calling song frequency [Esch et al., 1980]. There is some evidence of range fractionation for different sound amplitudes. Recording from the leg nerve shows fairly fast time constants (in the order of a millisecond) and thus clear copying of time varying patterns, with some adaptation within the first few milliseconds of response to a sound burst, levelling off to a steady rate. We are currently using eight such sensory inputs per ear (as opposed to only one previously), and have found this has the advantage of increasing reliability by reducing the total variability of the input, and also enables the coding of amplitude differences over a larger range. The plots in figure 5 show the group response of the simulated auditory nerve, which has been tuned to resemble the cricket data.

Second, we have included a second pair of neurons that receive input from the auditory nerve, based on the omega neurons (ON) described in the cricket [Wohlers and Huber, 1981]. These are mutually inhibitory and also inhibit the opposite ascending interneurons. The most obvious function of these connections is to increase

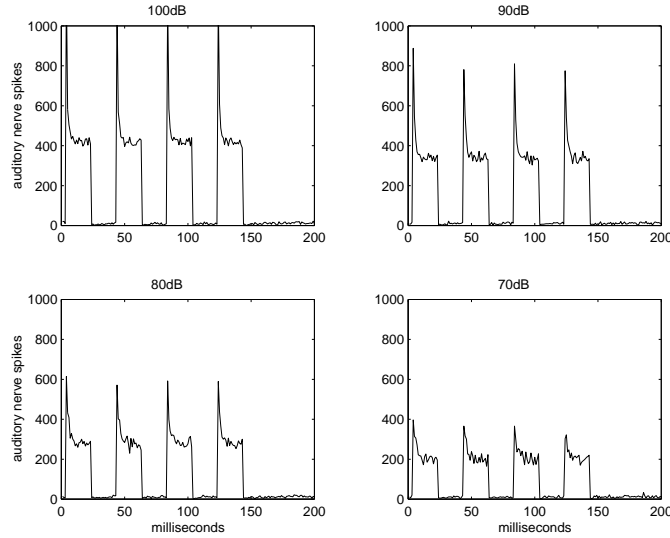


Figure 5. The summed spiking response of multiple units in the simulated auditory nerve to the normal cricket song pattern at different amplitudes. Note the peak at syllable onsets at the inter-syllable interval used here caused by recovery from adaptation during the gap between syllables.

the difference in activation between the two sides, to emphasise the directionality of the response. Elegant experiments from Horseman and Huber [1994] have shown that inhibition from the opposite side makes a difference of several spikes per chirp to the AN1 response. Figure 6 illustrates the replication of this effect in our model (compare to their figure 3). Here the decrease in firing rate to an ipsilateral sound level is proportional to the contralateral sound level. Although the decrease in firing rate is not large, its role may be in gain control, so that for varying amplitudes of sound (experienced as the cricket or robot approaches the sound source) the response adapts to encode the relevant difference between the ears within a similar range of firing rates, without saturation. Another effect of the inhibitory cross-connection is that, due to the delay introduced by the additional synapse, the onset of the syllable is less affected than the remainder of the syllable. This serves to enhance the onset encoding and effectively cleans up the signal so that the encoding of the sound pattern becomes clearer in the ascending neuron firing pattern for real sound stimuli.

Third, we have included a stage of neural processing in the 'brain' before signals are sent to the motor output. The most comprehensive study of the role of cricket brain neurons in phonotaxis is provided by Schildberger [1984], who suggests a possible filtering circuit for syllable rate recognition by the female. He identifies two main classes of auditory responsive cells: BNC1 which appears to get direct input from AN1; and BNC2 which get input via BNC1. These neurons vary in their response to the pattern of sound. BNC1d appears to require a minimum syllable duration near to the typical calling song before it reaches threshold, which makes it a lowpass filter for the syllable rate, assuming a constant duty cycle. BNC2b appears to spike around once per syllable, which makes it highpass *i.e.* as the




Figure 6. Effect of inhibitory cross connections on the firing rate of AN1. The ipsilateral firing rate decreases as the contralateral sound level increases.

syllable rate decreases the firing rate will also decrease. BNC2a shows a bandpass filtering effect, responding at somewhat less than a spike per syllable for normal rates but producing fewer spikes for slow rates or fast rates. Schildberger argues that the response of BNC2a reflects an 'AND'ing of the output of BNC2b and BNC1d, to produce a neural recognition signal for the appropriate sound pattern.

We initially attempted to copy these response patterns in our simulated neurons, but there were several difficulties. In particular, it seemed impossible to tune neuron or synapse parameters to distinguish very long syllables from very rapid syllables, as they both produce the same response – more or less continuous firing – in the AN1 neurons. After some systematic searching of the parameter space we found that a lowpass filter could be produced by having a moderately depressing synaptic connection from AN to BN1. This would fire several spikes at the onset of syllables. If the syllable rate was fast, the depression would not recover and fewer spikes would occur. The ideal syllable rate was the upper limit for which clear enough gaps occurred for the synapse to recover. At slower rates, the recovery would be greater and the chance of spiking at syllable onsets increased; however as the time between onsets increases the number of spikes per chirp starts to decrease. A second depressing synaptic connection from BN1 to BN2, with a relatively slow time constant, performs temporal summation of the BN1 output so that these slower responses were not usually sufficient to produce a spike, making BN2 a bandpass filter for the correct syllable rate. In figure 7 we show the response of our 'BN1' and 'BN2' neurons to song patterns identical to those used by Schildberger, *i.e.* using equal length chirps with syllable repetition intervals (SRI) ranging from 10 ms to 90 ms. It can be seen that BN1 has a moderately bandpass response which is sharpened by BN2; the response of BN2 is very similar to the BN2a neuron reported by Schildberger. In the top plots which show results with simulated sound inputs,

the apparent best tuning is to an SRI of 26 ms. However, using the same neural parameters with real sound input, as shown in the lower plots, syllables at this rate are not so clearly coded, and the best response in BN2 moves to an SRI of between 42 ms and 58 ms.

The need for a clear pattern makes the BN2 neurons much more likely to respond to clearer firing patterns in AN, which generally correspond to the louder side, particularly after the cross-inhibitory effects of ON. We found it was not necessary to include any explicit mechanism for comparing the firing rates or latencies at the BN level. We could simply take a spike in the left or right BN2 as indicating the need to turn in that direction to approach the sound.

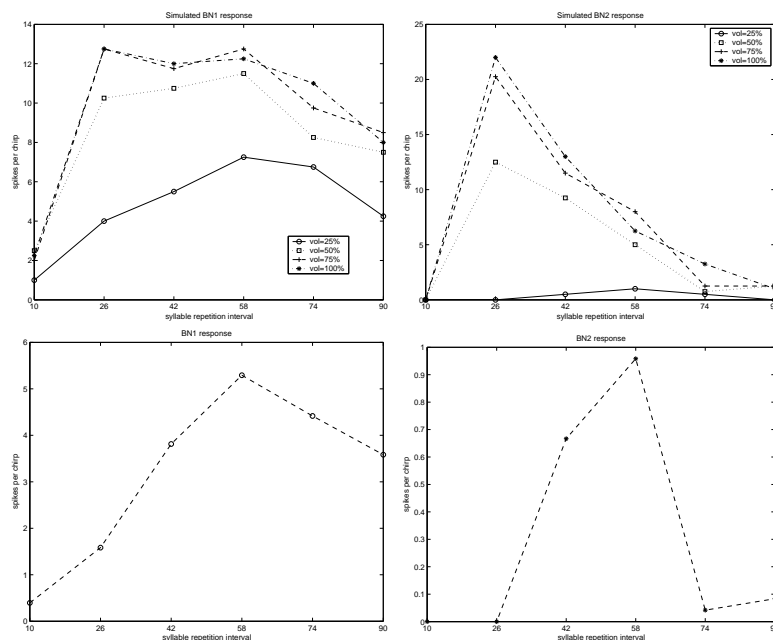


Figure 7. Spike rates of simulated brain neurons to songs with different syllable repetition intervals. Upper plots, simulated inputs; lower plots, real sound inputs. Left, the BN1 neuron shows a lowpass/bandpass response; right, the BN2 neuron shows a bandpass response.

(b) Motor circuit

Our aim in designing the motor circuit was to reproduce the specific kinds of moves, turns and stops observed in crickets performing phonotaxis, with as few neural connections as possible. There is unfortunately little direct data on the neural connectivity from the brain to the motor control circuits for sound localisation in the cricket. Staudacher and Schildberger [1998] describe some 200 descending neurons, some of which show a response to sound, with their response often 'gated' by whether or not the animal is walking. One of these neurons has a firing rate that correlates with the walking direction of the animal. Current models of motor

control for six-legged walking in other insects [*e.g.* Cruse et al., 1998] suggest several features to include in our circuit. Motor patterns tend to be self-sustaining through local feedback or Central Pattern Generator (CPG) activity, so that only a trigger signal is required to start the movement. Steering appears to be modulated by fairly simple turn signals from the brain interacting with this pattern generator so as to modify limb movements appropriately. The neural circuitry we have used is adapted from the scheme devised by Chapman [2001] in his robot model of cricket escape behaviour. It is illustrated in figure 8. The paired burst generators (BG) will, when initiated by an incoming spike, mutually activate each other to produce a continuous burst of spikes that go to right and left forward neurons (RF and LF) and drive the robot forward. The length of the burst is limited by the eventual activation of the STOP neuron which inhibits the BG neurons. One trigger for movement is the ambient light level. The other is a spike in the left or right BNC2. These act via a right or left turn neuron (RT or LT) to additionally modulate the forward velocity by appropriate excitatory and inhibitory connections to RF and LF.

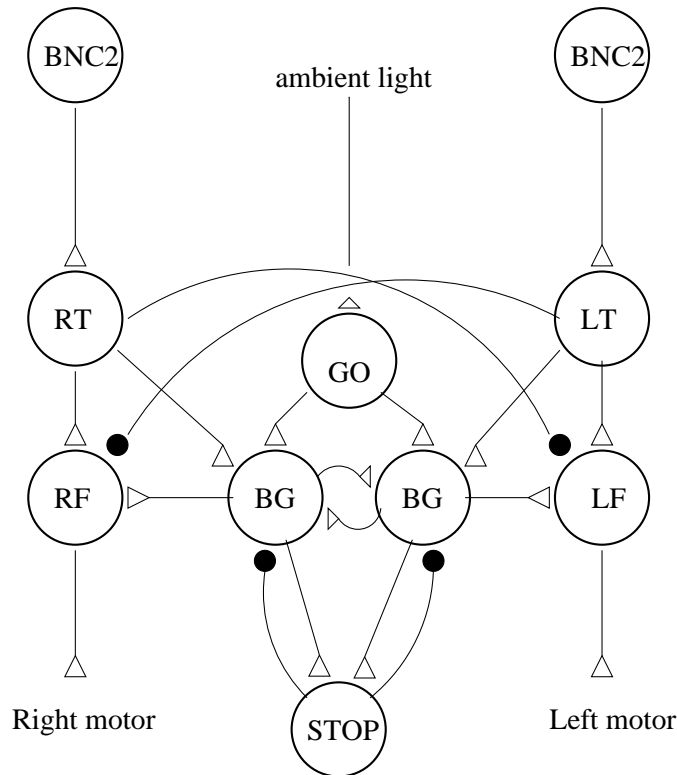


Figure 8. Neural circuit for motor control. The BG pair form a burst generator that excites right forward (RF) and left forward (LF) neurons when triggered by RT, LT or GO till deactivated by STOP. BNC2 is the output from the auditory circuit in figure 4 and RT and LT produce right and left turns respectively by excitatory and inhibitory modulation of RF and LF.

Figure 9 shows the performance of the robot using the combined phonotaxis and motor circuits to track cricket song. A total of 30 trials were run, from 3 different starting positions relative to the speaker. It can be seen that the robot is capable of successfully locating the sound source, producing a cricket-like path that zigzags towards the sound.




Figure 9. Tracks of the robot to a cricket song from three starting positions (thirty tracks in total). The sound source is at (0,0).

(c) Integrating different modalities

The same motor circuitry needs to subserve other sensorimotor behaviours that are to be integrated on the robot, including the default behaviour when sound is not present (taken for now to be random search for safe dark places), and the optomotor response of turning so as to stabilise visual slip. For example the average speed of the robot should increase with the ambient light as the GO neuron in the motor circuit receives more stimulation. The instantaneous forward speed of the robot under five different light levels, with or without sound, is shown in Figure 10. Sound causes no movement unless there is a light level of at least 10, and that level of light is too low to stimulate movement without sound. The pattern of movement in response to sound is changed by the light level – in brighter light the robot makes fewer stops.

In figure 11 we show how the optomotor sensor has been integrated into the motor control. The sensor is used to produce four spike trains, two for each direction of motion. These are summed in two optomotor interneurons, which effectively

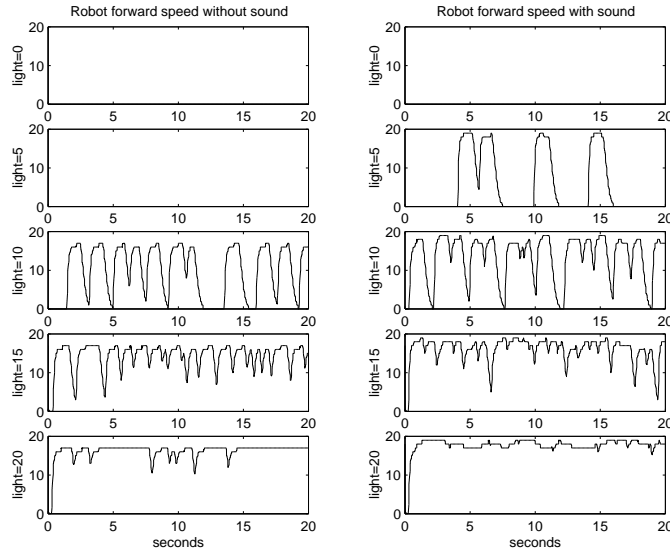


Figure 10. Forward speed of the robot's motors under different light levels: left without sound; right with sound.

smooths the signal. The interneurons mutually inhibit one another. The output from the interneurons steers the robot in the appropriate direction by excitatory inputs to the left or right 'forward' neurons. For example, if someone walks in front of the robot it tends to follow this movement, and if it veers off course in anyway (*e.g.* through wheel slip) it will compensate, and maintain a straight heading. Note that unless the robot is already moving (due to light or sound inputs) the optomotor input is not enough to move the robot on its own, *i.e.* it only modulates movement. In figure 12 we show how the optomotor response can correct for randomly imposed turns (as might be produced by environmental disturbance such as uneven ground surface, or for a flying insect, wind gusts) to produce relatively straight paths.

In addition, there is an inhibitory connection between the phonotaxis output (BNC2) and the optomotor interneuron corresponding to the expected turn direction. This will suppress the optomotor response that would otherwise 'correct' intended turns towards sound. Note that unlike the algorithm in [Webb and Harrison, 2000] the suppression is direction specific (as hypothesised for the interaction of tracking and optomotor behaviour in the fly by Heisenberg and Wolf [1988]). The suppression is the result of shunting inhibition, which will counteract any optomotor excitation up to the strength of the predicted input, but might still enable a response if the signal, even though in the expected direction, is much larger than expected. However this is not equivalent to efferent copy, which would require prediction of the exact size and time-course of the expected optomotor signal, so as to cancel it out. This would seem to require a more complex neural mechanism to implement it. Indeed, it is theoretically impossible to make a precise prediction because this depends on the actual spatiotemporal pattern of contrast encountered during the turn, which is unknown in a non-uniform and unfamiliar environment. In figure 13 we show the results of running the robot under the same conditions




Figure 11. Left or right visual movement excites opto-clockwise (OC) or opto-anticlockwise (OA) respectively to produce an appropriate turn. BNC2, as before, activate the motor circuit via RT and LT (dashed lines, not all connections from 8 shown) but also inhibit OC or OA (dotted lines). In the additive scheme, these inhibitory connections are left out.




Figure 12. Left, paths of the robot with randomly imposed turns (no sound source) and right, the same conditions with the random turning substantially corrected by the opto-motor reflex.

as figure 9, but with the optomotor response also active. It is evident that the robot produces more direct paths to the sound source with fewer corrections. The advantage of including an optomotor response is even more evident by looking at the robot's behaviour when it has random turns imposed, while trying to track sound (figure 14). The paths with phonotaxis alone diverge much further from the correct heading direction, and on several occasions it misses the sound source. With the optomotor response added, the tracks are much more direct and the random disturbance is hardly evident.




Figure 13. Paths of the robot integrating phonotaxis and optomotor responses to make a more direct approach to the sound.

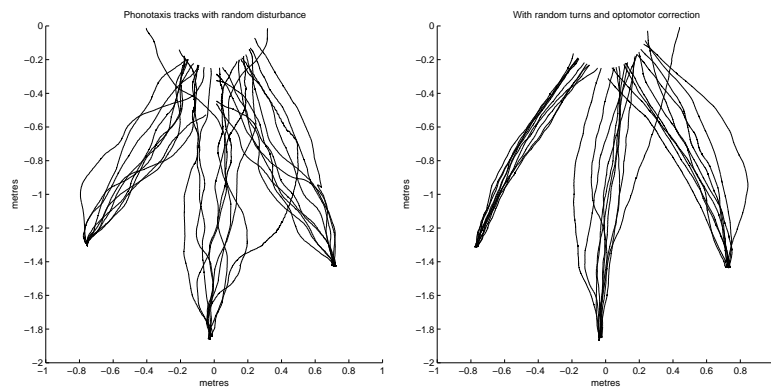


Figure 14. Comparison of paths when random turns are imposed: left, phonotaxis on its own produces very indirect paths; right, with the optomotor reflex, the random turns have almost no effect on the approach to sound.

While this appears to be a fairly successful way to combine the behaviours, a number of alternatives are possible. Efferent copy has already been mentioned, but would require more complex neural processing. Two possibilities can be easily tested by small modifications to the circuit shown in figure 11. The first is 'additive' interaction in which the connections that suppress the optomotor response during

turns to sound are disabled, so the LF/RF activity simply reflects the summed input of optomotor and phonotactic corrections. Note that this was the mechanism originally suggested by Böhm et al. [1991] on the basis of their open loop tests on the cricket. The second is 'follow-on' interaction, shown in figure 15, in which the direct connection from LT/RT to LF/RF is disabled and instead the steering response to sound is controlled via the optomotor system, by exciting the appropriate optomotor integration neuron. This idea was first proposed (in the context of fly tracking behaviour) by Mittelstaedt (1951 cited in Collett [1980]).




Figure 15. As in figure 11, left or right visual movement excites opto-clockwise (OC) or opto-anticlockwise (OA) respectively to produce an appropriate turn. BNC2 now activate the motor circuit only via OC or OA (dotted lines).

Figure 16 shows the behaviour of the robot under these two alternative schemes. While the 'follow-on' tracks appear quite direct, this is slightly misleading because the robot is actually making fewer corrections. Instead of turning to each spike in RT or LT, at least two consecutive spikes on the same side are needed to cause a turn when the signal is filtered though OA or OC. Consequently, if the robot is headed in roughly the right direction, it is unlikely to adjust its course. This might seem fairly effective behaviour for the robot, but it is not such a good resemblance to the behaviour of the cricket as the previous mechanism. The 'additive' tracks are also moderately successful, but there is a discernable tendency to approach the sound in an arc rather than an oscillating track, again showing less resemblance to the cricket. In observing the robot it is evident that this result is due to the robot repeatedly turning towards the sound, then turning away a little in response to the optomotor signal produced by its turn to the sound. The average time taken to arrive at the sound is increased.

Figure 17 shows the performance of the additive and follow-on schemes when the robot also has to deal with random turns. Both work fairly well to remove the effects of the disturbance (as might be expected: the optomotor behaviour in response to a random turn is the same in all three schemes). However both produce some tracks that miss the sound source, suggesting the interaction with phonotaxis




Figure 16. paths of the robot using (left) the 'follow-on' scheme and (right) the 'additive' scheme to integrate the optomotor reflex with phonotaxis.




Figure 17. as for figure 16 with random turns imposed on the robot during phonotaxis.

is less successful. To examine this, figures 18 and 19 show the average number of attempted turns made in response to sound (produced by counting the spikes in RT and LT) under the various conditions. Figure 18 corresponds to paths starting in the centre position. All three schemes significantly reduce the number of corrective turns needed significantly in the non-random condition. In the random condition, the number of corrective turns needed is significantly greater for additive mode than for the original suppressive interaction. Figure 19 shows the paths starting from the left or right, and separately indicates the number of turns needed in each direction. The suppression and follow-on schemes show similar results, but the additive scheme and the no-opto require significantly more turns. The additive scheme also has a significantly higher imbalance in the ratio of turns in each direction, compared to all the other schemes. This is a reflection of the curved paths produced by the additive scheme (discussed above) compared to the oscillating paths produced by phonotaxis on its own.

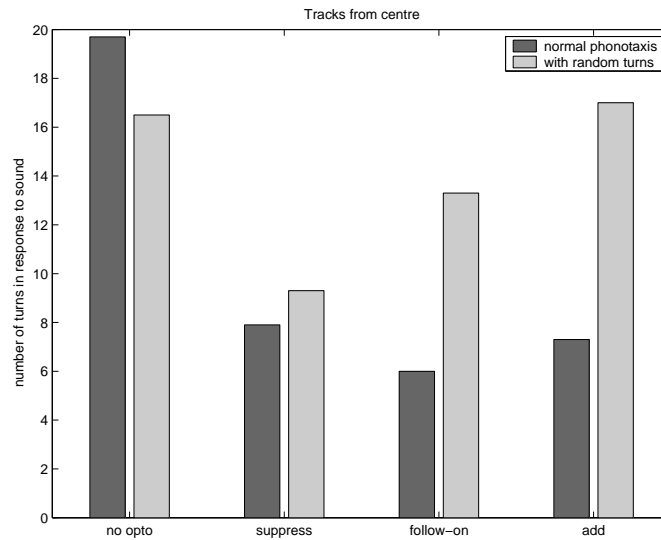


Figure 18. The number of attempted turns to sound during paths from the central starting position. Dark bars, normal paths; light bars, paths with random turns. All the optomotor integration schemes reduce the number of turns needed under normal conditions, but under random conditions, progressively more turns are required by the follow-on and additive schemes.

4. Conclusions

In this paper we have used small networks of neurons to control a robot so that it tracks a sensory source. Is this any advance on W. Grey Walter’s ‘tortoises’, demonstrated 50 years ago? Despite the (humbling) similarity, we believe there are a number of significant differences. The sensor systems for processing sound and visual motion are far more sophisticated, and allow us to replicate the real signals that animal brains have to deal with. As yet, the motor system is not much more complex; this is one of the issues we will be addressing in future work. The neural model, as well as containing more detail, is far easier to use. We can copy and tune arbitrary circuits, and the behaviour of all the elements can be recorded during behaviour for full analysis. The circuits we are building are directly based on neurophysiological data, so we can test biologically relevant hypotheses, with results that can translate directly back to experiments in biology. Nevertheless we are indebted to Walter’s pioneering insight that understanding the brain is best done within the context of a body.

References

- H. Böhm, K. Schildberger, and F. Huber. Visual and acoustic course control in the cricket *gryllus bimaculatus*. *Journal of Experimental Biology*, 159:235–248, 1991.
- V. Braitenberg. *Vehicles: experiments in synthetic psychology*. MIT Press, Cambridge, MA, 1984.

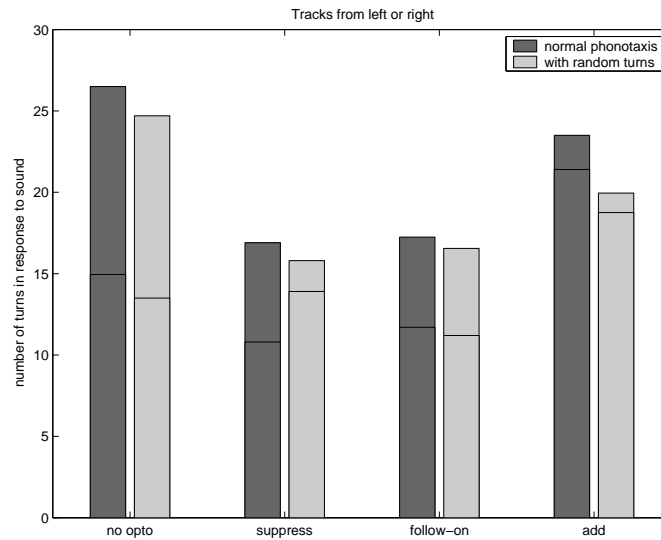


Figure 19. The number of attempted turns to sound during paths from the left or right. Bars are split to show the number of turns in each direction, relative to the starting position i.e. the lower half of each bar represents the sum of left turns when started on the left and right turns when started on the right, the upper half all turns in the opposite direction. The additive scheme makes almost as many turns as phonotaxis on its own, and they are not equally often in each direction, indicating repeated corrections in one direction rather than alternating corrections.

- J. Camhi and E. Johnson. High-frequency maneuvers mediated by tactile cues: antennal wall-following in the cockroach. *Journal of Experimental Biology*, 202: 631–643, 1999.
- T. Chapman. *Morphological and Neural Modelling of the Orthopteran Escape Response*. PhD thesis, University of Stirling, Stirling, U.K., 2001.
- T. Collett. Angular tracking and the optomotor response: an analysis of visual interaction in a hoverfly. *Journal of Comparative Physiology A*, 140:145–158, 1980.
- H. Cruse, T. Kindermann, M. Schumm, J. Dean, and J. Schmitz. Walknet – a biologically inspired network to control six-legged walking. *Neural Networks*, 11 (7-8):1435–1447, 1998.
- H. Esch, F. Huber, and D. W. Wohlens. Primary auditory interneurons in crickets: physiology and central projections. *Journal of Comparative Physiology*, 121:27–38, 1980.
- R. R. Harrison and C. Koch. *Advances in Neural Information Processing Systems*, volume 10, chapter An analog VLSI model of the fly elementary motion detector, pages 880–886. MIT Press, Cambridge, MA, 1998.

- M. Heisenberg and R. Wolf. Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *Journal of Comparative Physiology a – Sensory Neural and Behavioral Physiology*, 163(3):373–388, 1988.
- G. Horseman and F. Huber. Sound localisation in crickets. I. Contralateral inhibition of an ascending auditory interneuron. *Journal of Comparative Physiology A*, 175:389–398, 1994.
- Christof Koch. *Biophysics of Computation*. Oxford University Press, 1999.
- H. H. Lund, B. Webb, and J. Hallam. A robot attracted to the cricket species *gryllus bimaculatus*. In P. Husbands and I. Harvey, editors, *Proceedings of 4th European Conference on Artificial Life*, pages 246–255. MIT Press/Bradford Books, MA., 1997.
- H. H. Lund, B. Webb, and J. Hallam. Physical and temporal scaling considerations in a robot model of cricket calling song preference. *Artificial Life*, 4(1):95–107, 1998.
- A. Michelsen, A. V. Popov, and B. Lewis. Physics of directional hearing in the cricket *gryllus bimaculatus*. *Journal of Comparative Physiology A*, 175:153–164, 1994.
- B. P. Oldfield. Accuracy of orientation in female crickets, *teleogryllus oceanicus* (Gryllidae): dependence on song spectrum. *Journal of Comparative Physiology*, 141:93–99, 1980.
- G. S. Pollack. *Comparative Hearing: Insects*, chapter Neural Processing of Acoustic Signals, pages 139–196. Springer, Berlin, 1998.
- R. D. Quinn and R. E. Ritzmann. Construction of a hexapod robot with cockroach kinematics benefits both robotics and biology. *Connection Science*, 10:239–254, 1998.
- J. Rheinländer and G. Blätgen. The precision of auditory lateralization in the cricket, *gryllus bimaculatus*. *Physiological Entomology*, 7:209–218, 1982.
- K. Schildberger. Temporal selectivity of identified auditory interneurons in the cricket brain. *Journal of Comparative Physiology*, 155:171–185, 1984.
- E. Staudacher and K. Schildberger. Gating of sensory responses of descending brain neurones during walking in crickets. *Journal of Experimental Biology*, 201:559–572, 1998.
- W. G. Walter. *The Living Brain*. Penguin Books, Harmondsworth, Middlesex, 1961.
- B. Webb. Using robots to model animals: a cricket test. *Robotics and Autonomous Systems*, 16(2-4):117–134, 1995.
- B. Webb. Can robots make good models of biological behaviour? *Behavioural and Brain Sciences*, 24(6):1033–1094, 2001.

- B. Webb and R. R. Harrison. Integrating sensorimotor systems in a robot model of cricket behavior. In *Sensor Fusion and Decentralised Control in Robotic Systems III*. SPIE, Boston, 2000.
- B. Webb and T. Scutt. A simple latency dependent spiking neuron model of cricket phonotaxis. *Biological Cybernetics*, 82(3):247–269, 2000.
- T. Weber, J. Thorson, and F. Huber. Auditory behaviour of the cricket I. Dynamics of compensated walking and discrimination paradigms on the Kramer treadmill. *Journal of Comparative Physiology A*, 141:215–232, 1981.
- D. W. Wohlers and F. Huber. Processing of sound signals by six types of neurons in the prothoracic ganglion of the cricket *gryllus campestris l.* *Journal of Comparative Physiology*, 146:161–173, 1981.