Reafferent or redundant:
How should a robot cricket use an optomotor reflex?

Running title: A robot cricket with an optomotor reflex

Barbara Webb & Richard Reeve

Centre for Cognitive and Computational Neuroscience,
University of Stirling,

Corresponding author:
Barbara Webb
Dept. of Psychology,
University of Stirling,
FK9 4LA, U.K.
Ph:+44 (0)1786 466841
Fax:+44 (0)1786 467641
Email: b.h.webb@stir.ac.uk

ABSTRACT

A general problem in understanding the mechanisms underlying animal behaviour is the integration or interaction of different sensorimotor systems. In Webb & Harrison (2000a,2000b) we investigated the addition of an optomotor reflex to a sound-localising robot modelled on cricket behaviour. Böhm et al. (1991) proposed a simple additive mechanism to explain how the cricket combines the two behaviours. Problems implementing this on the robot led us to propose an alternative inhibition mechanism, which proved effective. Here we directly compare these two possibilities and several further alternatives. First, in a simulation of the open-loop paradigm used by Böhm et al we demonstrate that there are at least five algorithms (including 'efferent copy') that may adequately account for the data they present. We then consider possible neural implementations of several of these schemes, and test them in robot experiments. The results suggest that inhibition is both neurally plausible and effective as a means of combining these behaviours in real sensorimotor situations.

Keywords: sensorimotor integration, cricket, robot, optomotor, efferent copy

1 Introduction

For an animal, or a mobile robot, rotation of the entire visual field is usually the result of self-rotation. Consequently, unintended self-rotation can be corrected by turning in response to visual rotation signals. This 'optomotor reflex' is a well-studied behaviour, especially in insects (Gotz, 1975). A particular area of interest has been the interaction of the optomotor response with other behaviours, for example, visual tracking in flies (Srinivasan & Bernard, 1977; Wagner, 1986; Egelhaaf, 1987; Heisenberg & Wolf, 1988; Kirschfeld, 1994) and hoverflies (Collett, 1980), auditory escape in locusts (Robert & Rowell, 1992) and chemical tracking in moths (Willis & Carde, 1990). The experimental inspiration for the robot implementation presented here was reported by Böhm, Schildberger and Huber in 1991. They investigated the interaction between cricket phonotaxis (sound-localising) behaviour and the cricket's response to visual stimuli, including a rotating visual field.

The range of systems studied and the diversity of experimental methods have led to rather divergent conclusions about how two, potentially conflicting, orientation responses should be combined. While it seems plausible that different animals might solve this problem in different ways, not all studies have considered all alternatives. It is useful, for this reason, to clarify what schemes might be employed, and what behavioural or physiological evidence might be used to distinguish them. To some extent, modelling has already been used to address these issues, but usually at a very abstract level. We started from a different basis, that is with models incorporating specific details of the individual sensory motor systems to be combined (auditory and optomotor tracking) that were implemented in hardware and used to control a robot required to react to realistic environmental signals. This methodology gives complementary insights to other modelling approaches and the results of experimental investigations on the animal itself (Webb, 2000).

The results of the experiments on the cricket by Böhm et al (1991) led them to conclude that the "turning tendency [of the cricket to both stimuli] can be explained as the weighted sum of the turning tendencies evoked by the two individual stimuli". In previous work, Webb & Scutt (2000) had developed a robot system that could reproduce cricket phonotaxis and Harrison & Koch (1998, 1999) had built a robot that could reproduce optomotor behaviour using an aVLSI chip to perform efficient visual processing. It seemed straightforward to try combining the two sensory systems on one robot platform and using a weighted sum of their outputs to control its behaviour. However, we found (Webb & Harrison, 2000a) that the optomotor response tended to interfere with the phonotaxis behaviour. Essentially the problem was that each turn towards the sound would produce a clear optomotor stimulus, which would cause the robot to 'correct' itself and turn away from the sound again. This unsatisfactory result was an empirical demonstration of the problem theoretically formulated by von Holst and Mittelstadt in 1950: how can an animal with an optomotor reflex make intentional turns without automatically correcting (and thus negating) them? This problem was not encountered by the crickets in the Böhm et al study because their behaviour was measured under open-loop conditions, which would not produce the normal visual feedback.

One obvious and easily implemented solution to this problem is to have the turning response to sound inhibit the optomotor response, by setting its weighting in the sum temporarily to zero. This kind of 'switching' behaviour has been shown in several animal systems e.g. in response to 'escape' signals in the locust (Robert & Rowell, 1992) and during pursuit turns in the housefly (Srinivasan & Bernard, 1977). Using this solution, we carried out a series of trials to see whether phonotaxis behaviour was improved (described in detail in Webb & Harrison 2000b). A summary of these results is given in figure 1. It can be seen that the optomotor reflex always improves tracking (increasing the directness of the path to the speaker) but that this improvement is only slight in the case of normal motor control, either from the center or the sides. When we added a constant bias to the robot's movement, to emulate motor asymmetries seen in real crickets, tracking was significantly worse when performed by phonotaxis only. Adding the optomotor response enabled the robot to compensate completely for the motor bias and track normally.

This mechanism was simple to implement and has some biological plausibility. Nevertheless there are alternative schemes, which also have some biological support. Collett (1980) describes several. In 'efferent copy', the expected signal optomotor signal resulting from a turn is subtracted from the actual signal. In 'follow-on', the intended turn is actually controlled via the optomotor response by injecting the inverse of the expected signal, so that the optomotor system in correcting for the apparent signal executes the desired turn. Or, within an 'additive' scheme, the size of intended turns can be increased to compensate for the expected optomotor feedback.

As Collett (1980) shows, these three schemes algorithmically all reduce to addition of the two signals with appropriate gains. This implies that our original (additive) system might have worked if we had simply scaled the weighting of phonotaxis to compensate for the optomotor feedback. However as Collett also shows, the schemes are not equivalent when considered at the more detailed level of the temporal dynamics of the different reflexes. This was demonstrated by the fact that we could not find suitable additive gain parameters for the robot. For example, turning twice as fast in response to sound simply led to a stronger optomotor signal, which because of the inherent delays in the optomotor low-pass filter, tended to reach its maximum just as the phonotaxis turn ended. In fact, even in the implemented inhibition system, this 'residual' activity in the optomotor signal tended to make the robot turn back a little after each turn. This is because the visual signal was still being integrated during the turn even though the response was inhibited. It would seem a better idea to suppress the signal before the integration stage. We did not think of doing this on the robot because the integration was occuring in the hardware of the optomotor sensor.

If we add these two possibilities for inhibition - 'pre-integration' and 'post-integration' - to the three alternatives described by Collett (efferent copy, follow-on, scaled addition) we have five possible algorithms for combining the optomotor and phonotaxis responses. Our first study looks at whether any of these explanations can be ruled out for the cricket by a simulation of the original open-loop experiment by Böhm et al. In study 2 we report the results of trying to use the robot in an analogous task. In study 3, we consider the issue of how these different algorithms might map onto to motor control circuits using a new and more realistic neural simulation on the robot, and evaluate their performance in the normal closed-loop situation, i.e. tracking sound sources in the lab environment.

2 Study 1

2.1 Methods

Using MatLab™ we simulated a 'treadmill' experimental paradigm for cricket phonotaxis. The cricket was assumed to be fixed above an air-suspended ball, so that any turning movements it made would be recorded as angular velocity of the ball. Forward movement does not bring it any closer to the sound source so is disregarded. The simulation can be run either in closed loop (i.e. the angular velocity of the cricket will change its subsequent direction relative to the sound source) or open loop (i.e. its direction remains fixed). The starting direction relative to the sound source, the number of cycles to run, and any imposed optomotor stimulation can be specified. At each point in time, the position of the sound source and the optomotor stimulation are used to calculate the angular velocity for the next time step, as detailed below. The mode of integration of the sound and optomotor signals can be chosen from the following six possibilities:

1. 'no opto' - the optomotor input (imposed or self-generated) is ignored.

2. 'additive' - the two inputs are simply added with appropriate gains. Compared to no opto, the ear gain was doubled to try to compensate for the expected optomotor signal.

3. 'Pre-integration inhibition' - when turning in response to a sound signal, the optomotor signal is suppressed at the *input* to a low-pass filter.

4. 'Post-integration inhibition' - when turning in response to a sound signal, the optomotor signal was suppressed at the *output* of a low-pass filter. This corresponds to the integration algorithm previously tested on the robot.

5. 'Efferent copy' - when turning in reponse to a sound signal, a corresponding opposite signal was added to the optomotor signal to cancel out the expected input that would be generated by a turn.

6. 'Follow on' - turning to sound signals was implemented indirectly by adding double the opposite signal to the optomotor signal so that the required turn would occur as part of the optomotor response.

The specific algorithm for determining the (intended) angular velocity on each cycle was as follows:

$\dot{\theta}_t = g_p \times f_p(ear_t) + g_o \times f_o(opto_t)$ if mode = efferent or pre-integration

$= 2 \times g_p \times f_p(ear_t) + g_o \times f_o(opto_t)$ if mode = additive

$= g_o \times f_o(opto_t)$ if mode = follow-on, or post-integration and $ear_t = 0$

$= 2 \times g_p \times f_o(ear_t)$ if mode = post-integration and $ear_t \neq 0$

where f_p, f_o are low-pass filters with the phonotaxis-filter time-constant faster than the optomotor-filter time-constant, g_p, g_o are gain terms (in tests here both set = 1), and

$ear_t = 1 \quad if \quad \theta_{t-1} > 10$ degrees

$= -1 \quad if \quad \theta_{t-1} < -10$ degrees

$= 0 \quad if \quad -10 < \theta_{t-1} < 10$ degrees

(i.e. the signal is +1 for a right turn to sound on the right, -1 for a left turn to sound on the left, 0 if the sound is in the center and no turn occurs) and

$opto_t = \dot{\theta}_{t-1}$ if mode=additive, post-integration, or pre-integration and $ear_t = 0$

(i.e. the optomotor signal is determined by the angular velocity of the previous time step)

$= \dot{\theta}_{t-1} + g_p \times f(ear_t)$ if mode=efferent

$= \dot{\theta}_{t-1} + 2 \times g_p \times f(ear_t)$ if mode=follow on

$= 0$ if mode=pre-integration and $ear_t \neq 0$

A constant imposed optomotor stimulation could also be added to the opto-input. Also, it was assumed that the phonotaxis response would be intermittent because the normal sound signal is intermittent (i.e. there are gaps between syllables & chirps (see below) and the robot doesn't always detect the sound). This was simulated by having the ear input set = 0 (whatever the direction) for 10 cycles out of every 20 (note 10 cycles in the simulator were roughly equivalent to 100 cycles in the robot, thus one simulated cycle represents about 10ms of real time).

2.2 Results

Figure 2 illustrates the time course of behaviour in the simulation in a closed-loop situation, starting at 60 degrees from the sound. In each case the 'cricket', as would be expected, makes a series of turns towards the sound till the angle falls below 10 degrees. The plots illustrate how the latency and stability of the response is affected by the different schemes. The most noticable effects are the smoothed phonotaxis response in the follow-up scheme and the 'bounce-back' in the additive and post-integration suppression schemes, when each turn to sound tends to stimulate correction by the optomotor system. However the differences are slight and it seems unlikely that in any experiment on the animal it would be possible to detect such subtle differences amongst the general noise of the cricket's motor response. In particular the behaviour with pre-integration switching and efferent copy look almost identical.

To replicate experiment of Böhm et al (1991) we ran the simulation in open-loop conditions. Each integration mode was tested under nine possible combinations of three sound directions (+60, 0 or -60 degrees) and three imposed optomotor stimulus values (+1, 0 and -1). Figure 3 shows the resulting average angular velocity over 200 cycles. Several things are obvious from these plots. First is that the pattern of response under the additive, efferent copy, and follow-on modes is identical; in each case the visual rotation shifts the turning rates in the expected direction. Second, the pattern of the response under the two inhibition modes is identical but not the same as the other schemes. Rather than a uniform increase or decrease in rotation velocity in response to the imposed optomotor signal, the effect 'levels off'. The mean rotation velocity is the same for 0 or +60 degree sound when there is a positive optomotor stimulus, and for the 0 and -60 degree sound when there is a negative optomotor stimulus. This is because the response is continuous turning at the same rate, either because there is only optomotor driven turning (at 0 degrees) or the phonotactic turns are alternating with the optomotor turns (at 60 degrees). Finally (and not surprisingly) the optomotor rotation has no effect in the 'no-opto' situation.

It is interesting to compare these results to those reported by Böhm et al (1991) in the equivalent experiment on the cricket (figure 4). Although they describe the curves as a simple additive shift (which would be consistent with the results for additive, efferent or follow on) in fact there seems to be some evidence that a similar 'flattening' of the curves for the same direction of optomotor and phonotactic response occurs as we see for the 'inhibition' schemes. Bearing in mind that there is substantial noise in the cricket data, and that the shapes of the simulation curves depend to some extent on simplifications and assumptions in the model (e.g. the choice of rotation velocity in response to sound, and the linear relationship between rotation velocity and optomotor response) we can only safely conclude that the results for the cricket are not clearly inconsistent with any of the given schemes (apart from 'no opto'). For example, the 'flattening' could be explained under the additive scheme as the cricket reaching an upper limit on its possible turning velocity. A finer analysis of the cricket behaviour than taking the average rotation velocity would clarify whether there is any evidence of 'switching', e.g. if the distribution of directions when the two signals conflict has two peaks or wider variance than when they coincide.

3 Study 2

Our aim in this study was to follow-up the results of the simulation using the real auditory and optomotor inputs of the robot, but preserving the 'open-loop' paradigm. Rather than build a treadmill for the robot, we maintained the the fixed orientation of the sound source by attaching the speaker to the robot, and used an imposed rotation of the robot to produce the optomotor stimulus. The 'loop' was opened by recording the intended motor output of the robot in response to the sound and optomotor stimulus without actually sending these commands to the motors. We then tested the robot using the same six integration algorithms as the simulation.

3.1 Methods

The hardware and software used in these experiments was largely the same as that used by Webb & Harrison (2000b). For convenience we will summarise the main features here.

3.1.1 Hardware

The robot base used in the following experiments was a Koala™. The chassis is approximately 30x30x20cm, and has three tyred wheels on each side driven by a pair of DC motors; it is intended to be capable of movement on outdoor terrain. The processor is a Motorola 68331@22MHz; programmable in C. The auditory and visual processing was carried out by custom built analog circuits (described in more detail below) designed to mimic the animal's sensors, interfaced to the processor through A-D ports. The sensors were mounted, pointing forward, side by side on top of the robot, giving them a height above the ground of approximately 30cm. See figure 5.

The auditory circuit is shown schematically in figure 6 (see (Lund et al., 1997; Lund et al., 1998) for further details.). It is based on the pressure-difference receiver instantiated by the cricket's tympani, spiracles and trachea (Michelsen et al., 1994). Two microphones separated by 18mm (1/4 wavelength of the carrier frequency – 4.7kHz - of cricket song) receive and amplify the sound. The signal from the left microphone is delayed by 53 microseconds (1/4 the phase of 4.7kHz) and then subtracted from the right; and vice-versa (the delay and relative weighting of the two signals is programmable). The amplitude of the resulting waves is direction dependent. It is measured using a RMS circuit, resulting in two analog signals representing the amplitude of vibration of the posterior tympani in the cricket.

The optomotor circuit is shown schematically in figure 7 (see Harrison & Koch (1998); (Harrison & Koch, 1999) for further details). It is based on the Hassenstein-Reichardt (1956) motion detector proposed for the fly. The chip used in our experiments contains a 24 x 6 array of photoreceptors. A local measure of motion is computed between adjacent pairs of photoreceptors in each of six rows across the chip. First, these signals are bandpass filtered to remove the DC illumination levels. In the Reichardt model, photoreceptor signals are delayed, then correlated with non-delayed signals from neighboring photoreceptors. The chip uses the phase lag inherent in a lowpass filter as the delay, and multiplier circuits as the correlators. This motion detection is performed in opponency, and the results across the chip are summed. The results of the chip are lowpass filtered (tau = 100 ms) to remove residual pattern dependencies from the response. All of these operations are performed on a single analog VLSI chip that dissipates less than 1 mW of power. The output is a single analog signal that increases for rightwards motion and decreases for leftwards motion. The chip was fitted with a lens (focal length 2.6mm) resulting in a total visual angle of about 30 degrees.

3.1.2 Software

The software used is a modification of that described in Webb & Scutt (2000) which simulates the cricket's auditory neurons using a simple state-based integrate-and-fire neuron model. The left and right ear signals produced by the auditory circuit are summed in respective auditory interneurons. These fire with a latency and firing rate correlated to the amplitude of the signal. The auditory interneuron that fires first excites an ipsilateral motor neuron, and inhibits the effects of the opposite auditory-motor connection. Synaptic depression between the auditory and motor neuron means that several successive sound onsets are required to excite the motor neuron above threshold, making it sensitive to the temporal pattern in the sound source. Spikes in the motor neurons are signals to turn to the right or left. Thus, a left spike sets an *ears_signal* variable to -1, which is reset to 0 after a fixed time (around 400 ms). If another left spike occurs in that time, the time is extended. Right spikes set *ears_signal* to +1. If a right spike occurs during a left turn (or vice versa), the current turn is stopped and the opposite turn started.

The output of the optomotor circuit was differenced from a reference value measured before the robot was moved at the beginning of each trial. Subsequently, a positive *opto_signal* indicated rightwards visual motion and a negative *opto_signal* leftwards visual motion (usual range from -100 to +100) and the turning rate of the robot could be modulated proportionally to compensate for the rotation. Note this function was programmed directly, not implemented as part of the neural simulation.

The final motor output (the speed of left ($\dot{l}$) and right ($\dot{r}$) motors) was calculated as:

$$\dot{l}_t = b + g_o \times opto\_signal + g_p \times ears\_signal$$

$$\dot{r}_t = b - g_o \times opto\_signal - g_p \times ears\_signal$$

Where b is a constant base-speed, in these experiments set to zero (i.e. no forward motion), the ears gain g_p was set at 20 (which would normally cause a rotation at approximately 12 degrees per second) and the optomotor gain g_o at 0.14, which would normally produce good compensation for any visual rotation. Note, as described above, that this motor output was not actually used in the open-loop situation to drive the motors but instead was recorded as the intended angular velocity of the robot.

To implement the different integration modes, the *ears_signal* and *opto_signal* were modified as follows:

No-opto: the *opto_signal* was set to 0 at all times.

Additive: the *ears gain* g_p was doubled to try to compensate for the expected optomotor feedback.

Pre-integration: if the *ears_signal* was non-zero, the input to the opto-integrator was set to zero; and the *opto_signal* was the current output of the integrator. Note that this required us to modify the original hardware to access the opto-motor signal from the chip before it passed through low-pass integration. We did this by altering the time-constant of filtering on the chip to 1/100th of its previous value, and then adding a software integration process to the incoming signal, scaled to resemble the original filter.

Post-integration: if the *ears_signal* was non-zero, the *opto_signal* (i.e. the output of the integrator) was set to 0, but any input from the sensor was still integrated.

Efferent copy: the *ears_signal*, taken from the output of the auditory processing, was also fed as efferent copy (i.e. inverted and scaled) to the input of the opto-integrator, the output of which was the *opto_signal*. The value for the optomotor efferent copy was estimated from tests on the robot i.e. we measured the average size of the optomotor signal produced from normal turns in response to a phonotaxis signal.

Follow_on: the *ears_signal* was doubled and fed to the input of the opto-integrator, and the *ears gain* was set to zero so that output of the optomotor system drove the response.

3.1.3 Experimental methods

The robot behavior was tested in the normal lab environment. The robot was tested with each of the six integration modes under the nine possible condition combinations as follows. The speaker was attached to the robot using a plastic ruler, and thus sat about 30cm from the ears, and could be positioned at approximately 70 degrees, 0 degrees, or -70 degrees. The auditory stimulus was a simulated cricket song. A single ‘syllable’ of the male cricket song is a 20ms burst of almost pure 4.7kHz sine wave. The syllables occur in groups of 4, with intersyllable gaps of 20ms; and the groups (‘chirps’) repeat at approximately 2.5 Hz. This sound was simulated using a customized program under Linux, writing directly to the sound card. The imposed visual rotation was either zero or +/- 12 degrees per second, implemented by setting the left and right motors of the robot to + or -20. Each trial lasted 20 seconds, which for rotation trials meant the robot made a 360 degree turn. The intended motor outputs (i.e. the left and right motor speeds calculated from the *ear_signal* and *opto_signal*, not the imposed rotation speeds) were recorded every 10 cycles (approx 100Hz), and the average intended rotation speed calculated by subtracting the mean(right_speed) from the mean(left_speed).

3.2 Results

In figure 8 the average angular velocity of the robot has been plotted for the same sets of conditions as those shown for the simulation in figure 3. The first difference to note is that, when the sound direction was zero, the robot (unlike the simulation) often responded by turning, with a slight bias to turning left. Hence the points in all the graphs for the zero sound direction are rather varied. Nevertheless it can still be observed that the additive and efferent copy graphs are very similar, with the curves shifted in the expected directions. The follow-on graph shows higher velocities, but this is just a scaling effect from feeding the phonotaxis response through the optomotor system. At +/-70 degrees sound direction, the same pattern of change, in response to the optomotor stimulus, as for additive and efferent copy can be seen.

The other striking difference is in the 'inhibition' paradigms, where the added optomotor stimuli appears to have little effect. This is because the phonotaxis algorithm on the robot used a turning duration (=400ms) long enough to last through the gaps between chirps (=400ms). With the nearby sound source being detected very reliably, the result was that the phonotaxis was continuously suppressing the optomotor response, hence altering the optomotor stimulus had no effect. We had not anticipated this as it had been evident, in closed-loop tracking, that the robot did not normally respond so continuously to sound as to preclude any activity of the optomotor system. It is obvious that *if* the cricket's response to sound is continuous during open-loop tracking of a nearby sound, then the results of Böhm et al are not consistent with the possibility that the cricket completely inhibits its response to optomotor stimuli when turning to sound. However, there is no clear basis for assuming that this is, in fact, true for the cricket, although it turned out to be true for our robot model.

4 Study 3

The results of the previous simulation and robot experiments have indicated that, at the algorithmic level, it is difficult to determine which mechanism is used by the cricket, at least in open-loop testing. In the previous experiments in closed-loop on the robot (Webb & Harrison, 2000b) we concluded that the additive scheme seemed unlikely because gain adjustments would not be sufficient to compensate for the interaction between the behaviours. Inhibition seemed to work reasonably well, but under the conditions in study 2 resulted in phonotaxis completely dominating the behavioural response. Efferent copy would seem preferable because it does not simply ignore all optomotor input during phonotaxis, but only the expected input.

Another way to address this issue is to consider what neurophysiological constraints might help distiguish the different options. Unfortunately, direct evidence of the neural mechanisms underlying the integration of these behaviours is not available for crickets. However, we can at least investigate the relative plausibility of the mechanisms, i.e. how straightforward is it to design a neural circuit that carries out the different functions? Are the five options we have described equally valid if we consider the problem from the bottom up rather than top down?

4.1 Methods

4.1.1 Robot hardware and software

The same robot base and sensory systems were used as in Experiment 1, although the sensors were positioned differently with the camera mounted centrally and the ears placed on the front of the robot, about 5cm above ground level. The software controller was a new neural simulation that has been developed to improve the accuracy of the cricket model. Details of this model are given in Reeve & Webb (2002). The main alterations are:

i) the use of a spiking neural model more similar to those described by described by Koch (1999), in which the neuron is considered to be an RC circuit with a fixed membrane capacitance and membrane conductance and a base potential across the membrane to which it will decay exponentially in the absence of external input. If the membrane potential rises above a threshold, the neuron will 'fire', sending a spike to any output synapses. The synapses are modelled as a variable conductance with a battery potential which corresponds to the reversal potential of the ion channel opened by the synaptic neurotransmitter. Synapses are characterised by a reversal potential, a delay (corresponding to the sum of possible axonal, neurotransmitter, and dendritic delays), a variable time course for the exponential decay of conductance, as well as a more standard 'weight' determining the standard conductance change in the affected ion channel. They also allow short-term adaptation of the weight.

ii) a neural circuit for phonotaxis that includes more detail from cricket neurophysiology. The input to the left and right auditory interneurons consists of eight 'parallel fibres' from the ear sensors which encode the sound in poisson distributed spike trains proportional to the amplitude. A second pair of auditory interneurons receive the same input but provide cross-inhibition (based on identified neurons in the cricket, Horsemann & Huber, 1994) which serves to sharpen the difference and also acts as a gain control mechanism. The output from the auditory neurons passes through two levels of 'brain' neurons (based on those described by Schildberger, 1984) which filter for the temporal pattern. The output of these neurons is a spike, approximately once per chirp, on the side corresponding to the loudest and clearest sound pattern, indicating a turn in that direction is needed to approach the sound.

iii) the output from the phonotaxis circuit is fed in to a neural circuit to control the movement of the robot, designed to reproduce some of the specific characteristics of the cricket's behaviour. It is based on the motor controller used by Chapman (2001) in modelling escape behaviour and draws on evidence of motor control mechanisms in other insects. Forward movement is produced by a mutually excitatory neural pair that act as a burst generator (BG) when initiated by an incoming spike. The length of the burst is limited by the eventual activation of a STOP neuron which inhibits the BG neurons. One trigger for movement is the ambient light level. The other is the firing of the left or right brain neuron corresponding to a phonotaxis signal. These also act via a right or left turn neuron (RT &

LT) to modulate the forward velocity by appropriate excitatory and inhibitory connections to the 'left forward' and 'right forward' output.

The optomotor system was implemented neurally as follows (see figure 9). The sensor is used to produce four spike trains, two for each direction of motion. These are summed in two optomotor interneurons, which act as the integration stage previously carried out in hardware. 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 e.g. leftward visual rotation will result in leftward rotation of the robot. 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.

This circuit as described so far corresponds to an additive scheme of integration, with the output from the motor neurons being a 'sum' of the inputs from the auditory and the optomotor processing circuits. The obvious way to convert this to an 'inhibitory' scheme is to add inhibitory synaptic connections between the phonotaxis interneuron and the optomotor interneuron corresponding to the expected turn direction (dotted lines in fig 9).

This, however, is not precisely the same as the previous inhibition scheme. 'Pre-integration inhibition' would require a synapse-on-synapse blocking of the inputs to the optomotor interneuron; and 'post-integration inhibition' would require similar blocking of the connections from this interneuron to the motor neurons. Instead, we are using 'shunting inhibition'; thus it counteracts any optomotor excitation up to the strength of the inhibition. It is also direction specific (as hypothesised for the interaction of tracking and optomotor behaviour in the fly by Heisenberg & Wolf (1988)). This means the robot might still repond to optomotor signals during a phonotactic turn if the signal is in the opposite direction to that expected, or is much larger than expected. Consequently this scheme also has some of the character of efferent copy.

It is hard to determine how *precise* efferent copy could be implemented in a simple neural circuit. In theory, to produce the correct cancellation, it would be necessary for the system to be able to predict the exact size and time-course of the expected optomotor signal. As this signal is dependent on the exact turning behaviour, the spatial frequency and contrast of the scene, and the properties of the sensor, these would all have to be predicted, implying the existance of a complete internal 'forward model' of the motor system, environment and sensory system. Obviously it is not viable to implement such a model with a few simple neural connections. Hence we have not implemented a separate 'efferent copy' mechanism in the following tests. However we did consider "follow-on" interaction, in which the direct connection from the phonotaxis circuit to the motor steering output is disabled and instead the steering response to sound is controlled via the optomotor system, by exciting the appropriate optomotor integration neuron.

4.1.2 Experimental methods

The robot was tested for its phonotaxis performance in the normal lab environment using

1. phonotaxis only

2. phonotaxis and optomotor combined additively

3. phonotaxis and optomotor with inhibition of the latter by the former

4. phonotaxis and optomotor, with phonotaxis controlled by a 'follow-on' input to the optomotor system.

For each condition, we ran a series of 30 trials, from 3 different starting positions relative to the speaker: directly ahead at a distance of 2 metres, facing the speaker; and to the left or to the right at a distance of 1.5 metres, facing an angle of approximately 30 degrees to the speaker. The sound source was as described before. The speaker was placed on the floor of the lab. No special soundproofing or other controls for noise or echoes were used. The visual stimulus was simply the lab furniture. The robot path was recorded using an overhead camera tracking system and this data was stored in synchrony with all the internal sensing and neural data.

In a second set of experiments we compared the behaviour of the robot under the different schemes when it had a superimposed 'random turn' behaviour. This was generated by using a Poisson-distributed input, via an interneuron, to produce random spikes in the left or right 'turn' neurons that are usually activated by the phonotaxis circuit.

To look for a significant difference in the directness of tracks we adapted the measure used by Schul (1998) for crickets tracking on a treadmill. This was calculated as follows:

- From each track, the x,y coordinates corresponding to each phonotactic or random turn made by the robot were extracted.
- Each successive pair of coordinates was used to define a vector with

$$distance_i = \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2} \qquad heading_i = \arctan(\frac{y_i}{x_i}) - \arctan(\frac{y_i - y_{i+1}}{x_i - x_{i+1}})$$

- And the normalized mean vector for the track was then calculated as:

$$magnitude = \bar{x}^2 + \bar{y}^2 \quad angle = \arctan(\frac{\bar{y}}{\bar{x}})$$

$$\text{where } \bar{x} = \frac{\sum distance_i * \cos(heading_i)}{length} \qquad \bar{y} = \frac{\sum - distance_i * \sin(heading_i)}{length}$$

$$length = \sum_i distance_i$$

The angle of the normalized mean vector indicates the average heading of the robot relative to the speaker during the track, and the magnitude is a measure of the amount of variance around that direction, such that a mean vector of angle = 0 and magnitude = 1 would indicate a completely direct path to the speaker from the starting position.

- Overall directness is then scored as:

$$D = magnitude \times \cos(angle) \times tracktime$$

This combines the cosine of the angle of the mean vector (which varies from 1 to 0 as the robot deviates from heading towards the speaker), the length of the mean vector (which varies from 1 to 0 as the robot deviates more around the mean angle) and a measure corresponding to the time taken to do the trial: $tracktime = \frac{\text{minimum time to do trial}}{\text{actual time to do trial}}$ (which varies from 1 to 0 as the robot takes longer to finish the trial). We took the 'minimum time' to be the time it should take the robot to complete the trial if it moved on a straight line from the starting position to the speaker at a constant 10cm/second.

4.2 Results

As described in the methods, we compared the tracks of the robot from different starting positions, first under ordinary motor conditions and then with the motor output randomly disturbed, using the four different forms of optomotor integration: phonotaxis only, additive, inhibition and follow-on.

In figure 10a it can be seen that the robot is able to track reliably to the sound source using phonotaxis only. This confirmed that the new neural circuitry has the expected capability - the robot makes repeated corrections to produce a 'zig-zag' path to the sound that is comparable to the behaviour of crickets. The average time to reach the speaker was 39 seconds (s.d. = 6s).

The character of the tracks is altered by the various different modes for integrating the optomotor response. In the additive case (figure 10b) the centre tracks look straighter, but the robot seems less able to correct for the deviations in heading occuring towards the end. The tracks from the side look more curved than zig-zag. From observing the robot it appears that it is less able to adjust its heading in response to sound because each turn is resisted by the optomotor stimulation it produces. This can be seen in a plot of the membrane potential of the motor, phonotactic-turn and optomotor neurons taken from one of the tracks where the sound was initially on the right (fig 11). The burst generator produces burst of spikes that are fed to the left and right motor neurons. The sound is detected on the right and this causes the right motor to be inhibited and the left to be further excited, so the robot starts turning right. However the turn causes excitation of the clockwise optomotor neuron, which inhibits the left motor and excites the right causing a turn back in the original direction. The robot is still able to reach the sound eventually through a series of small corrections, but it is noticable that from the sides there is an average of 18.75 phonotaxis-driven turns in one direction (to the side on which the sound is at the start) vs. only 1.2 turns in the other direction. In the phonotaxis-only mode the turns were more evenly distributed in each direction (14.95 vs 11.55).

In the case of inhibition (10c) the tracks from all starting positions appear straighter than the original phonotaxis-only plots. The robot can turn at the outset of the track to face the sound - unlike the additive case, these turns are not counteracted by the optomotor system. As can be seen in the plots of neural activity (fig 12), the optomotor response is held below the level of firing while the robot responds to sound, and only when an optomotor stimulus occurs between responses to sound is a reverse turn activated. It appears the robot is able to use the optomotor response to keep it on a straighter course toward the speaker, with the corrective turns fairly evenly distributed but fewer of them (10.8 vs. 7.1).

The most obvious feature of the results from the follow-on scheme (10d) is that the paths consist of longer straight segments. However the robot is making fewer turns in this case because where normally each spike of the output neurons from the phonotaxis circuit is sufficient to trigger a turn directly, two or more successive spikes are needed to trigger a turn when passed through the optomotor neuron. This can be seen in the neural plots (fig 13): the right and left motor outputs remain the same, as each spike in the sound detector contributes to an increase in membrane potential in the optomotor sensor, eventually resulting in a turn. During the turn, the sound and optomotor stimuli balance out, so the problems caused using the additive scheme do not occur. Though this overall results in rather direct looking paths, they have less resemblance to the cricket. The paths are also not really comparable to the others because the time course of turning in response to sound has been altered by driving it through the optomotor filter, so they have not been included in the following analysis.

Using the directness measure described in the methods, we statistically compared the tracks produced under 'no-opto', 'add' and 'inhibit' conditions. The graph in fig 14 shows that the mean directness was increased by both forms of optomotor integration, for paths from each starting position, and this was a significant main effect in a two-way ANOVA (F=25.7, p<.001). The inhibit condition appears to increase the directness more than the additive condition; however post-hoc Scheffé tests showed no significant difference between these two conditions (p =0.199), with both significantly different from the no opto (P<.001). Tracks from the right were significantly more direct than from the left or centre (as is evident in the plots of tracks, and confirmed statistically F= 10.34, p<.001). This presumably reflects some difference in the amount of auditory distortion from this direction, rather than a better visual background for the optomotor response, as the difference in start position was consistent across the three optomotor modes (there was no interaction between start-position and optomotor-mode, F=0.93 p =.45).

As discussed in the introduction, the optomotor response is not necessarily all that useful to a robot or animal capable of maintaining a straight course, although here it has been shown to help the robot stay 'locked on' to the speaker direction. A better test of the contribution of the optomotor response is when the course is disturbed. In a real environment this might be due to uneven terrain, wheel slip etc. In the next set of experiments we simulated such disturbances by including random inputs to the left and right motors as described in the methods. In figure 15a the effect of these random turns on the robot doing phonotaxis without optomotor control is evident: the tracks become substantially less direct and sometimes fail to reach the speaker.

Using the additive optomotor integration we see that the randomness of the tracks is reduced, but the robot still misses the sound on some occasions. Using inhibitory integration improves the tracks substantially, so that they closely resemble those without random disturbance. Performing the equivalent statistical analysis on the directness measure, we find that the latter tracks are significantly more direct than either additive or no-opto, which do not differ (main effect F=15.0, p<.001; Scheffé comparisons no-opto vs. inhibitory, p.<.001, inhibitory vs additive p =.003, additive vs no-opto p=.167). As before, tracks from the right are more direct (F=27.4, p<.001) and there is no interaction (F=2.3, p =0.07).

Overall these results suggest that the inhibitory implementation is more effective than the additive implementation, and produces more 'cricket-like' tracks than the follow-on scheme. It compensates well for random disturbance of the robot's tracks while not interfering with the intended phonotactic turns.

5 Conclusion

This paper has taken three different approaches to addressing the issue of how a robot or cricket, performing phonotaxis, should integrate an optomotor response. We first examined, in simulation, whether the available behavioural data from the cricket could allow us to chose between a number of alternative algorithms. We then repeated these experiments with real inputs, using hardware and software that mimicks the sensory processing of the biological system. Finally we considered whether or how these algorithms might be mapped onto a plausible neural model of motor control and compared their operation under real tracking conditions including random motor disturbance.

Five integration schemes were discussed. The first is simple weighted summation of the outputs, which potentially suffers from the problem that the behaviours in question might interfere with one another. The second and third schemes anticipate this problem by having one behaviour (phonotaxis) inhibit the other (optomotor). They differ in the point at which the inhibibition is imposed: one inhibits the input to the optomotor controller, and the other inhibits the ouput. The fourth scheme is the 'classic' solution of using efferent copy to try to explicitly counterbalance the expected interaction of the two behaviours. The fifth scheme, 'follow-on' effectively does the same, but this time by having one behaviour (phonotaxis) controlled by converting it to a 'pseudo' optomotor signal that is then combined with any actual signal to determine the final motor output.

The simulation demonstrated that the open-loop paradigm used in cricket experiments does not distinguish between the additive, efferent copy and follow-on algorithms (figure 3). The cricket data (figure 4) does not, on close inspection, appear to be a simple addition of a constant, optomotor-induced, shift to the response curve induced by sound as originally claimed by Böhm et al. The effect looks more similar to that seen in the results of the simulated inhibition schemes. However, it could also be accounted for by any of the other schemes with the addition of an assumption of maximum turning velocity. The resemblance to the inhibition schemes occurs because we assume the animal in normal phonotaxis is regularly switching between sound-following behaviour and optomotor behaviour. In implementing open-loop behaviour on the robot, this latter assumption was not met. The nearby sound source (attached to the robot) was detected very reliably, produced a directional response even when positioned near the mid-line, and each response lasted longer than the natural gap between chirps in the stimulus. Consequently the imposition of optomotor stimulus had little or no effect on the response curves for the inhibition schemes on the robot (figure 8). The follow-on scheme was distinguished in these robot experiments by producing higher velocities but this could be modified simply by changing the scaling of the phonotaxis input to the optomotor integrator. The additive and efferent copy algorithms could not be distinguished.

These results reveal some of the pitfalls in trying to reason from behavioural results to underlying mechanisms of control. Using time averaging to summarise data, in this case, made it impossible to distinguish whether the insect is switching between behaviours (which are averaged in the results) or actually performing an averaged behaviour in response to two stimuli. Fairly straightforward analysis of the behaviour (such as looking at directional distributions) should make it possible to distinguish between these possibilities, but may not easily reveal the differences between averaging, efferent copy and follow-on control. The latter might be distinguished by careful considerations of behavioural dynamics - i.e. if the animal can respond much faster to sound than optomotor stimuli, then it seems unlikely the sound response is being controlled via modulation of the optomotor response. It should also be possible to test for efferent copy by investigating the insect's response to incorrect feedback.

However such experiments might be complicated by indications that insects can quickly detect when they are in open vs. closed loop and adjust their behaviour accordingly (Heisenberg & Wolf, 1988).

In the final robot experiments the implementation of phonotaxis was more strongly constrained, by copying elements of known neurophysiology in the cricket. Consequently there were fewer arbitrary assumptions reducing concerns about possible artifactuality of the results. The optomotor reflex and the integration schemes were also implemented as neural circuits rather than as algorithms. Additive integration was straightforward as the outputs from phonotactic and optomotor interneurons could both modulate the excitation of the same motor neurons to influence turning behaviour. Inhibitory integration could have been implemented in several different ways. The most straightforward seemed to be to have the phonotactic output make an additional inhibitory synaptic connection onto the appropriate optomotor interneuron. However, this results in a scheme that does not fall neatly under the previous classification. Inhibition is represented in our model by a biologically realistic 'shunting' mechanism, which will counteract excitation up to a certain level. Thus, this form of inhibition corresponds to 'ignoring' optomotor excitation up to the amount expected during a turn, in the direction expected during a turn. It thus begins to resemble the efferent copy scheme, although it lacks the subtlety of attempting to exactly predict the size and time-course of optomotor excitation, which a proper efferent copy mechanism would need to do. Finally, the follow-on scheme could be implemented in a fairly direct way by having the phonotactic output directly excite the optomotor interneuron that would lead to the appropriate turn.

This latter scheme proved problematic when tested on the robot simply because the output signals from the phonotactic circuit were not appropriately scaled for the optomotor system, so it responded more rarely to sound. While this could have been adjusted in our model, it points out the inherent unlikeliness of this solution as a practical hypothesis for the animal: to have its responses to one modality filtered by the parameters and time-constants of another modality does not seem an effective mechanism. It seems particularly unlikely in this case, where crickets are very capable of performing the sound localisation behaviour in the dark. Of the remaining alternatives, the additive scheme revealed the same problem as in previous experiments i.e. the optomotor response tended to fight against the intended turns of the robot, resulting in curved tracks from the sides, and worse behaviour when random turns were introduced. By contrast, the inhibitory scheme was very effective in increasing the directness of the robot's tracks under both normal and randomly-disturbed conditions.

We would therefore suggest, on the basis of all our experiments, that some variant of inhibitory interaction is currently the most plausible hypothesis for how the cricket integrates phonotaxis and optomotor behaviour. This leaves us with a number of intriguing questions. Is it possible that insect brains actually carry out 'forward-modelling' to produce efferent-copy cancellation of expected input? Many current theories of primate motor control use such models, but as yet it is not clear if the same arguments really apply to insect behaviour. What range of possibilities exist between 'full' forward models and very simple inhibitory interactions? What neural circuitry might support such functions? These issues will form the focus of our future work.

References

Böhm, H., Schildberger, K., & Huber, F. (1991). Visual and acoustic course control in the cricket *Gryllus-bimaculatus*. Journal of Experimental Biology, 159, 235-248.

Chapman, T. (2001). Morphological and Neural Modelling of the Orthopteran Escape Response. Unpublished doctoral dissertation, University of Stirling, U.K..

Collett, T. (1980). Angular tracking and the optomotor response: an analysis of visual reflex interaction in a hoverfly. Journal of Comparative Physiology, 140, 145-158.

Egelhaaf, M. (1987). Dynamic properties of two control systems underlying visually guided turning in house-flies. Journal of Comapartive Physiology, 161, 777-783.

Gotz, K. G. (1975). The optomotor equilibrium of the Drosphilia navigation system. Journal of Comparative Physiology, 99, 187-210.

Harrison, R. R., & Koch, C. (1998). An analog VLSI model of the fly elementary motion detector . in M. I. Jordan, M. J. Kearns, & S. A. Solla (editors), Advances in Neural Information Processing Systems 10 (pp. 880-886). Cambridge, MA: MIT Press.

Harrison, R. R., & Koch, C. (1999). A robust analog VLSI motion sensor based on the visual system of the fly. Autonomous Robotics, 7(3), 211-224.

Hassenstein, B., & Reichardt, W. (1956). Systemtheoretische Analyse der Zeit-, Reihenfolgen-, und Vorzeichenauswertung bei der Bewungsperzeption des Rüsselkäfers. Chlorophanus. Z. Naturforschung, 11b, 513-524.

Heisenberg, M., & Wolf, R. (1988). Reafferent control of optomotor yaw torque in *Drosophila-melanogaster*. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology, 163(3), 373-388.

Horseman, G., & Huber, F. (1994). Sound localisation in crickets. I. Contralateral inhibition of an ascending auditory interneuron. Journal of Comparative Physiology A, 175, 389-398.

Kirschfeld, K. (1994). Tracking of small objects in front of a textured background by insects and vertebrates - phenomena and neuronal basis. Biological Cybernetics, 70(5), 407-415.

Koch, C. (1999). Biophysics of Computation. Oxford: Oxford University Press.

Lund, H. H., Webb, B., & Hallam, J. (1997). A robot attracted to the cricket species *Gryllus Bimaculatus*. P. Husbands, & I. Harvey (editors), Fourth European Conference on Artificial Life (pp. 246-255). Cambridge MA: MIT Press.

Lund, H. H., Webb, B., & Hallam, J. (1998). Physical and temporal scaling considerations in a robot model of cricket calling song preference. Artificial Life, 4(1), 95-107.

Michelsen, A., Popov, A. V., & Lewis, B. (1994). Physics of directional hearing in the cricket *Gryllus bimaculatus*. Journal of Comparative Physiology A, 175, 153-164.

Reeve, R. & Webb, B. (2002) New neural circuits for robot phonotaxis., EPSRC/BBSRC International Workshop Biologically-Inspired Robotics: The Legacy of W. Grey Walter Hewlett-Packard.

Robert, D., & Rowell, C. H. F. (1992). Locust flight steering .2. Acoustic avoidance maneuvers and associated head movements, compared with correctional steering. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology, 171(1), 53-62.

Schildberger, K. (1984). Temporal selectivity of identified auditory interneurons in the cricket brain. Journal of Comparative Physiology , 155, 171-185.

Schul, J. (1998). Song recognition by temporal cues in a group of closely related bushcricket species (genus Tettigonia)). *Journal of Comparative Physiology A* , 183, 401-410.

Srinivasan, M. S., & Bernard, G. D. (1977). The pursuit response of the housefly and its interaction with the optomotor response. Journal of Comparative Physiology, 115, 101-117.

von Holst, E., & Mittelstaedt, H. (1950). Das Reafferenzprinzip. Wechselwirkungen zwischen Zentralnervensystem und Peripherie. Naturwissenschaften, 37, 464-476.

Wagner, H. (1986). Flight performance and visual control of the free-flying housefly (*Musca domestica*). III Interactions between angular movement induced by wide- and smallfield stimuli. Philosophical Transactions of the Royal Society of London B, 312, 581-595.

Webb, B. (2000). What does robotics offer animal behaviour? Animal Behaviour, 60(5), 545-558.

Webb, B., & Harrison, R. R. (2000a). Eyes and ears: combining sensory motor systems modelled on insect physiology . IEEE International Conference on Robotics and Automation (pp. 3913-3918). San Francisco April 24-28.
Reprint: 2000

Webb, B., & Harrison, R. R. (2000b). Integrating sensorimotor systems in a robot model of cricket behavior . Sensor Fusion and Decentralised Control in Robotic Systems III . Boston Nov 6-8: SPIE.
Reprint: 2000

Webb, B., & Scutt, T. (2000). A simple latency dependent spiking neuron model of cricket phonotaxis. Biological Cybernetics, 82(3), 247-269.

Willis, M. A., & Carde, R. T. (1990). Pheromone-modulated optomotor response in male gypsy moths, lymantria-dispar l - upwind flight in a pheromone plume in different wind velocities. Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology, 167(5), 699-706.

Figure captions

Figure 1: Summary of results from Webb & Harrison (2000b). Mean and standard deviation of the directness of 10 tracks by the robot to a sound source, with and without optomotor correction. Solid lines represent a central starting position, dotted line a side starting position, and dashed line a motor biased condition. Only the latter reveals a signicant improvement with the optomotor response.

Figure 2: Simulated closed loop behaviour under different integration schemes. Regular turns are made towards the sound. These are smoothed in the follow-on condition and cause optomotor corrections in the additive and post-integration conditions.

Figure 3: Simulated open loop behaviour under different integration schemes, with different sound directions and different directions of visual rotation. Additive, efferent copy and follow-on schemes produce identical results.

Figure 4: Open loop behaviour of the cricket. Redrawn from Böhm et al (1991). The curves resemble the results for the integration conditions in figure 3.

Figure 5: The Koala robot used in experiments 2 and 3. The optomotor chip is mounted behind the lens, and the microphones are at the front.

Figure 6: Schematic of the analog electronic circuit for processing sound.

Figure 7: Schematic of the analog VLSI circuit for the optomotor response.

Figure 8: Open loop behaviour of the robot.

Figure 9: Neural circuit to integrate the optomotor and phonotaxis behaviours. OA and OC integrate the optomotor signal and modulate right forward (RF) and left forward (LF) neurons, which are also activated by a burst generating circuit (BG, GO and STOP) and modulated by the brain neuron (BN) output from sound processing (dashed lines). The inhibitory scheme adds inhibitory connections between BN and OC/OA (dotted lines). The follow-on scheme removes the direct connection of BN to RF/LF and instead has each BN excite the relevant OC or OA neuron (not shown) to indirectly modulate the behaviour.

Figure 10: Tracks of the robot during phonotaxis a) with no optomotor input b) with optomotor response simply added to the phonotaxis response c) with phonotaxis inhibiting the optomotor response d) with phonotaxis controlled by follow-on input to the optomotor response.

Figure 11: Neural response for additive integration

Figure 12: Neural response for inhibitory integration

Figure 13: Neural response for follow-on integration

Figure 14: Comparison of the directness of tracks under different integration conditions. Both additive and inhibition schemes improve directness of the tracks.

Figure 15: Tracks of the robot during phonotaxis, with random disturbance a) with no optomotor input b) with optomotor response simply added to the phonotaxis response c) with phonotaxis inhibiting the optomotor response d) with phonotaxis controlled by follow-on input to the optomotor response.

Figure 16: Comparison of the directness of tracks, with random disturbance, under different integration conditions. Only the inhibition scheme significantly improves directness.

Figure 1

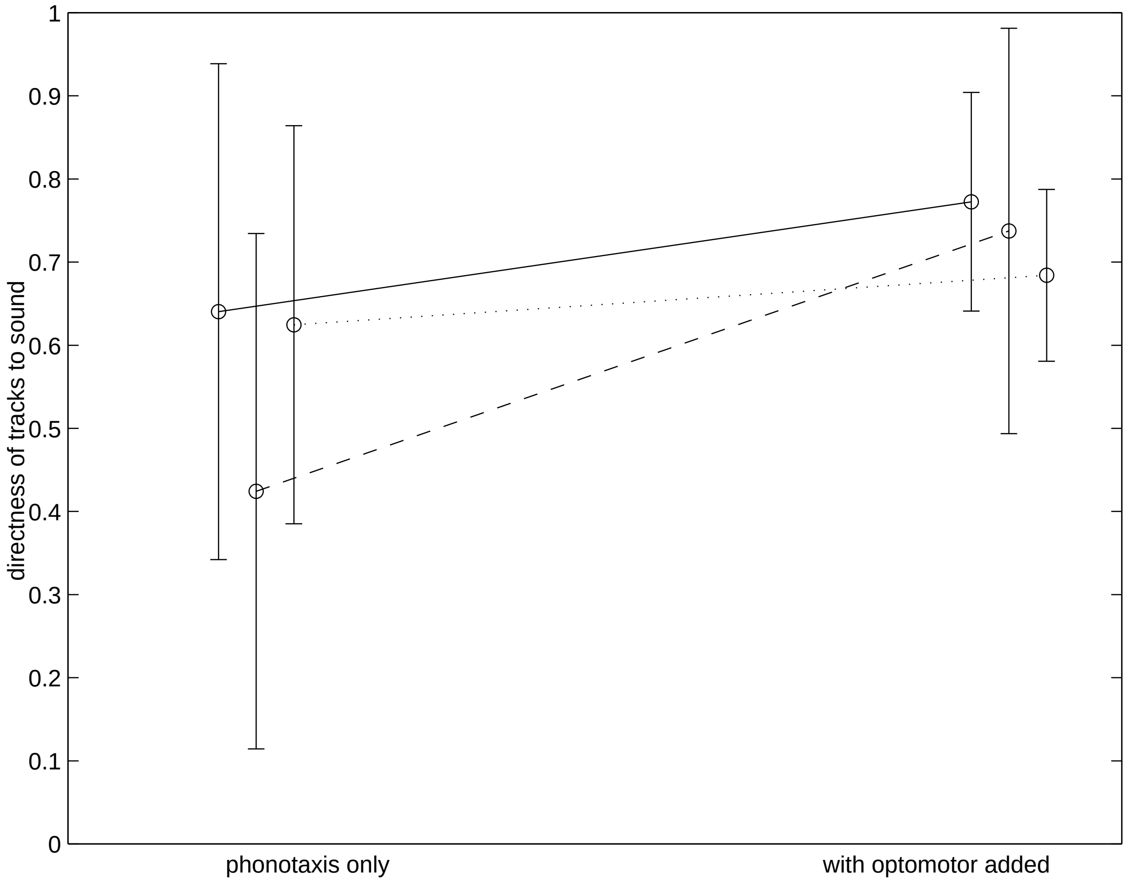

Figure 2

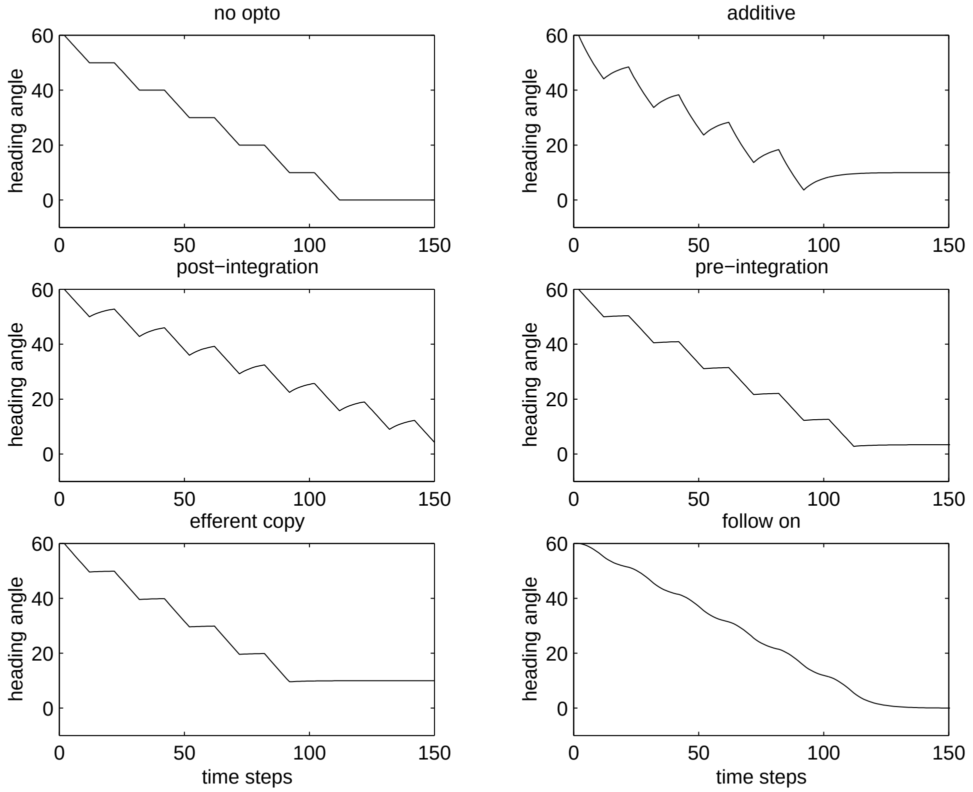

Figure 3

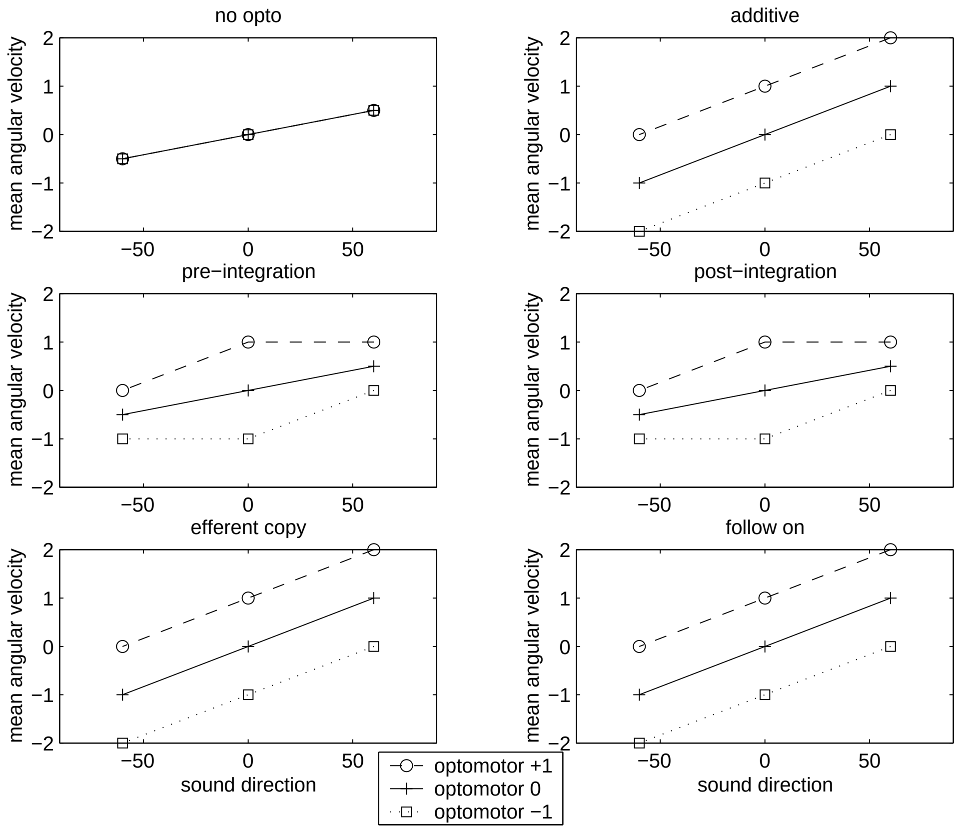

Figure 4

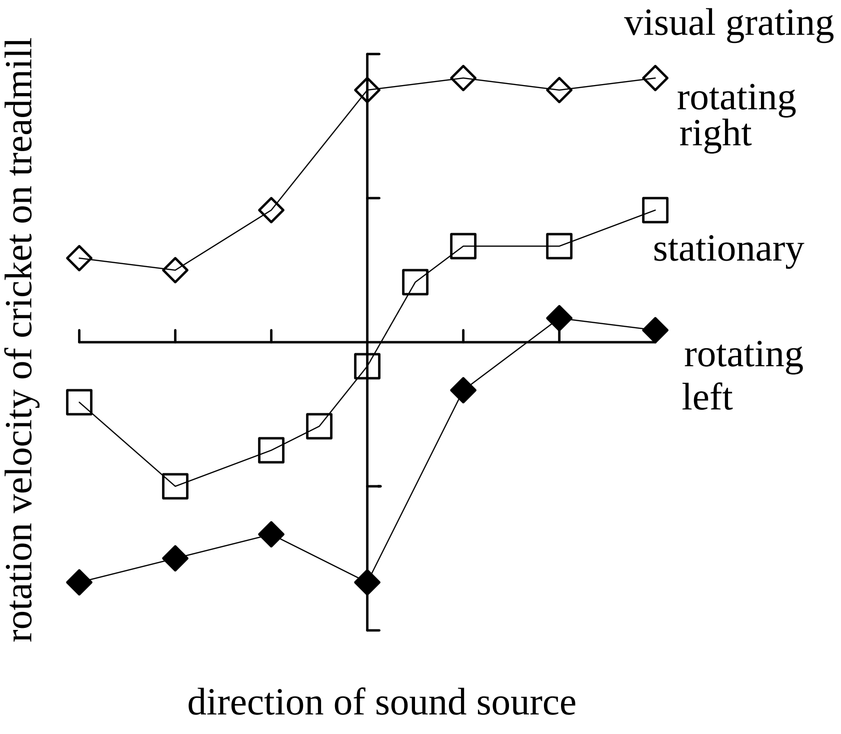

Figure 5

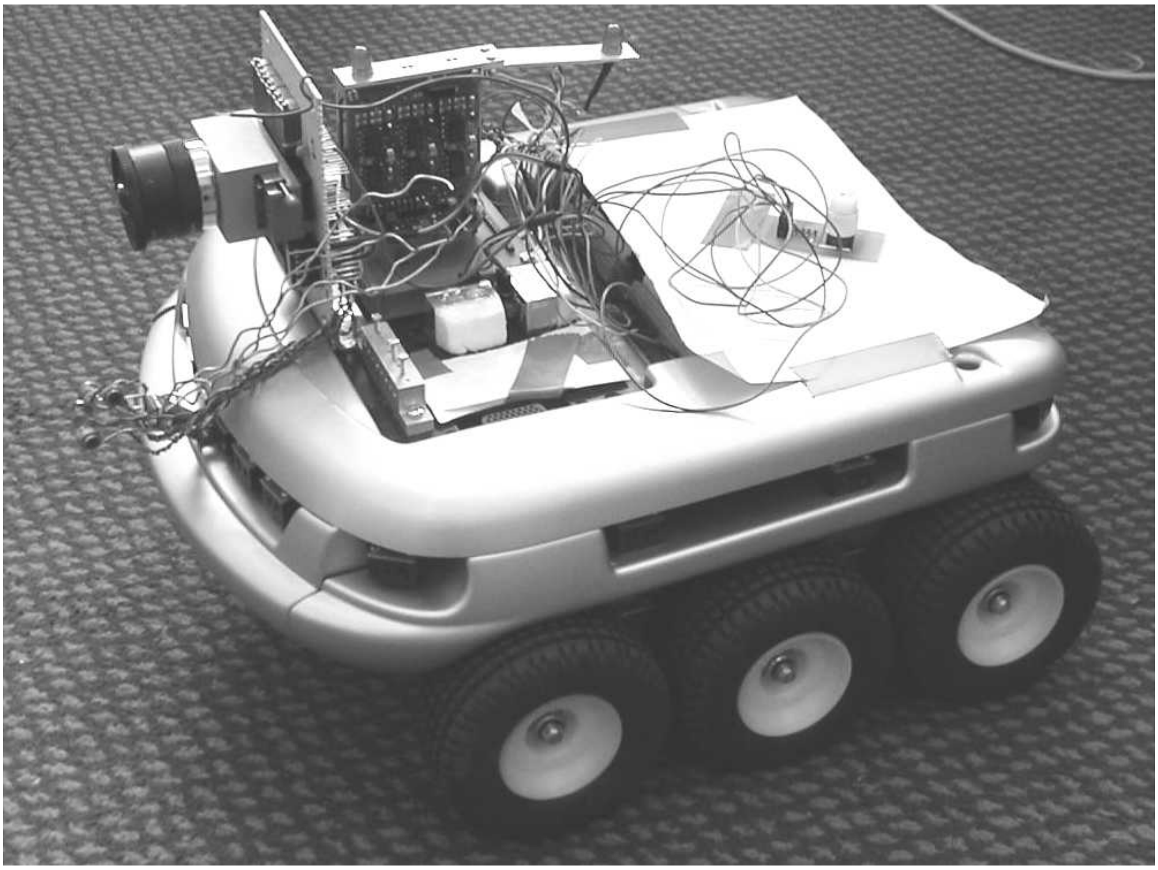

Figure 6

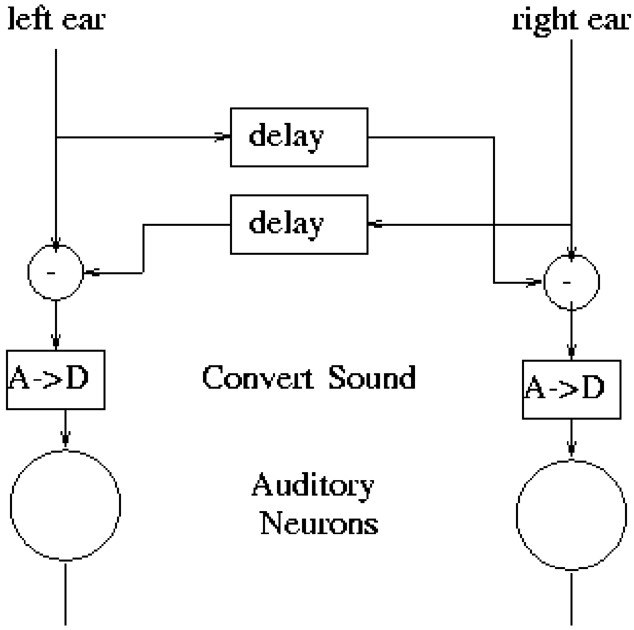

Figure 7

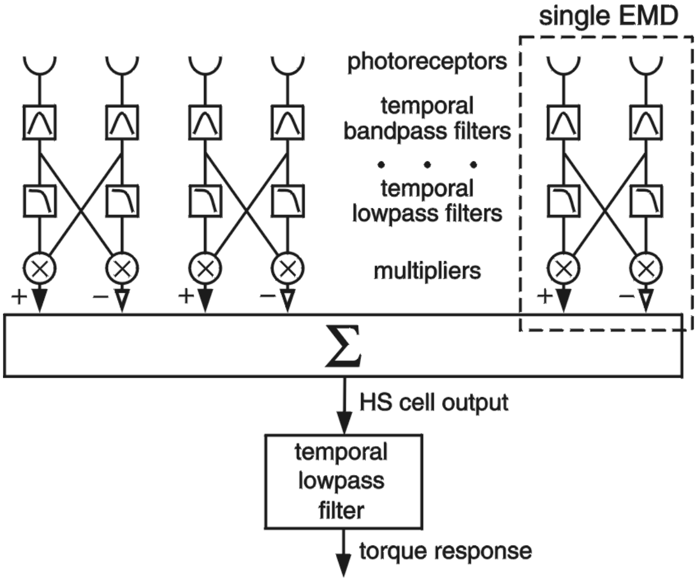

Figure 8

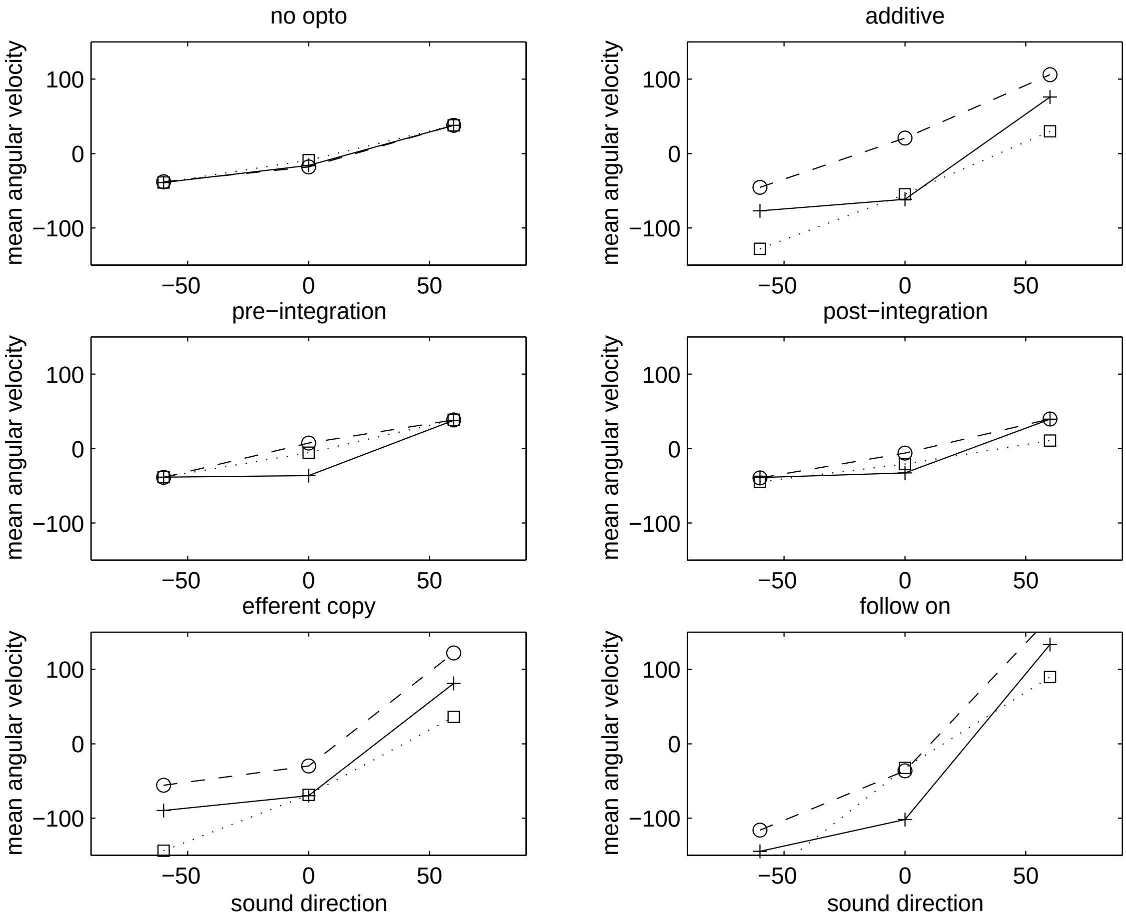

Figure 9

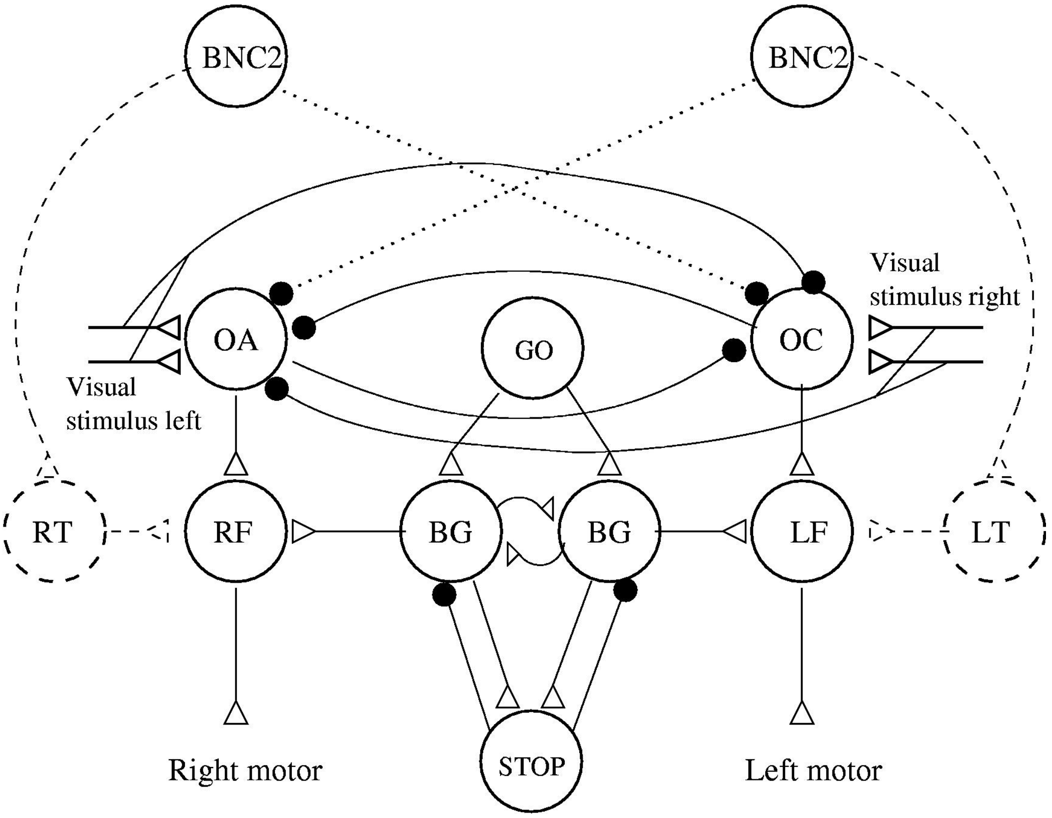
BNC2
BNC2
Visual
stimulus right
OA
GO
OC
Visual
stimulus left
RT
RF
BG
BG
LF
LT
Right motor
STOP
Left motor

Figure 10 a) b)

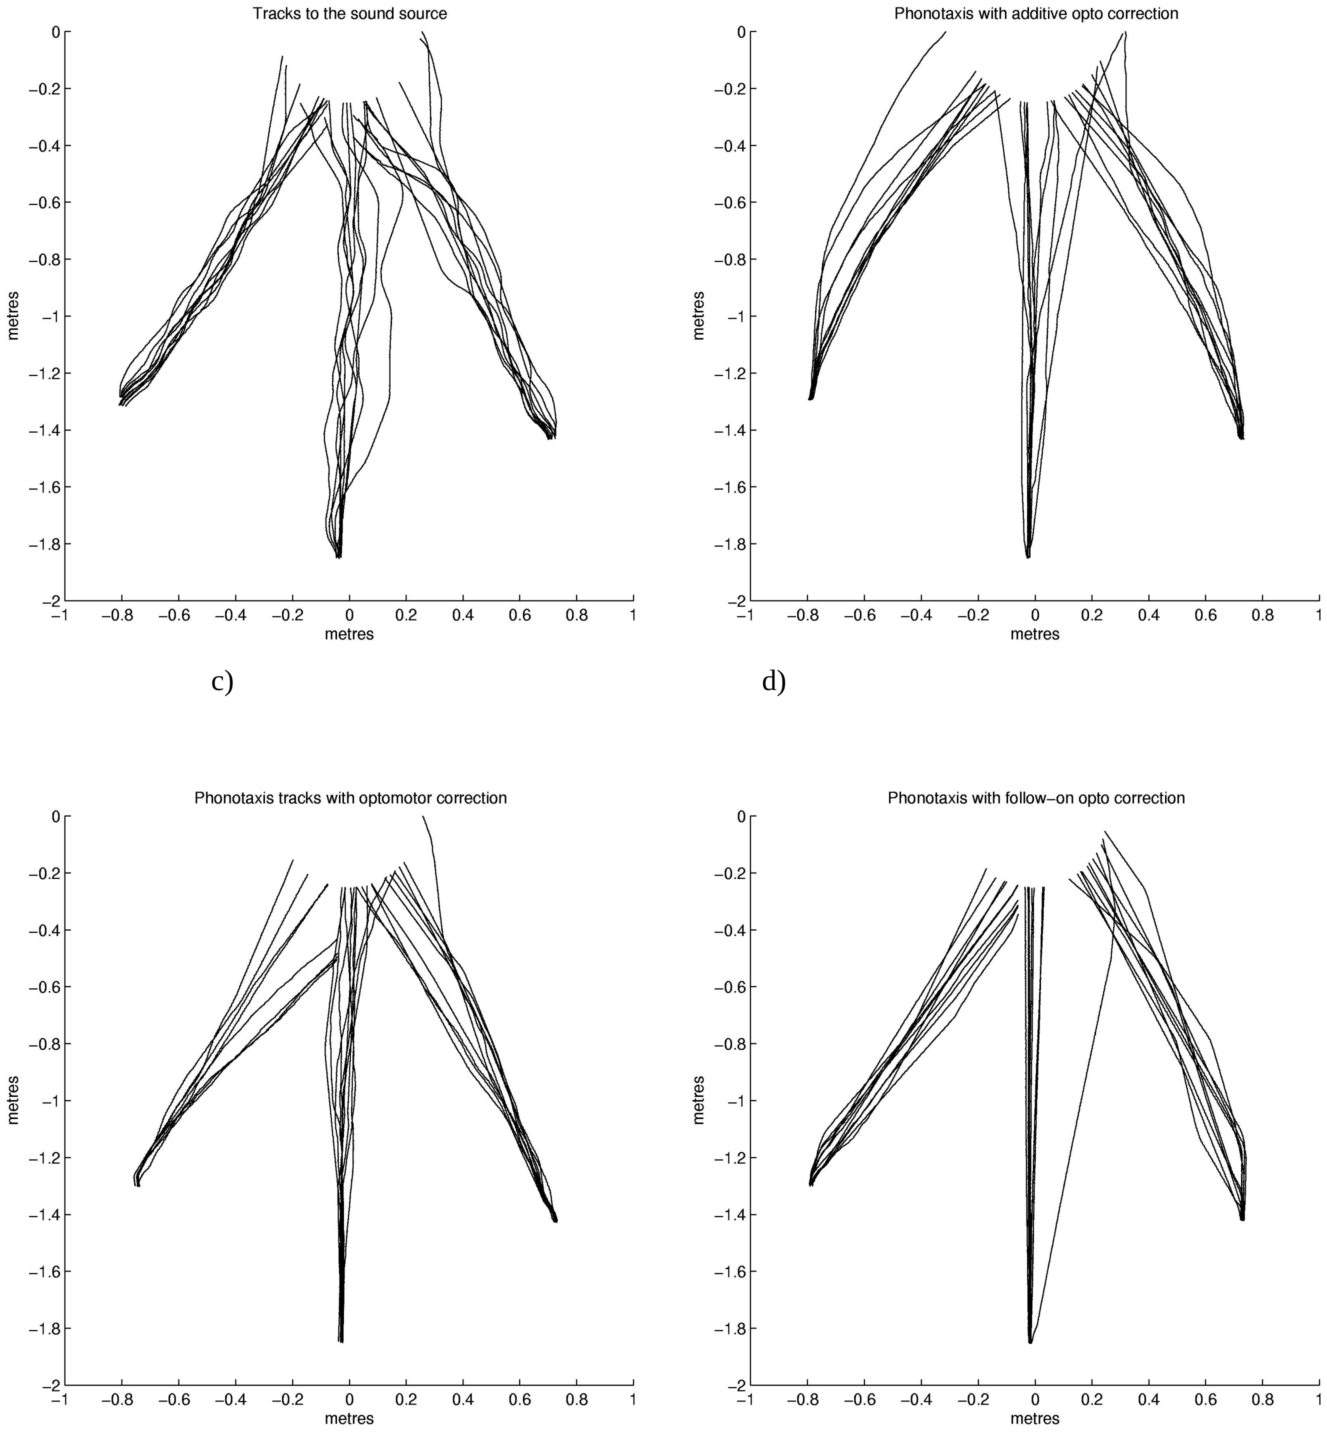

c) d)

Tracks to the sound source

Phonotaxis with additive opto correction

Phonotaxis tracks with optomotor correction

Phonotaxis with follow−on opto correction

Figure 11

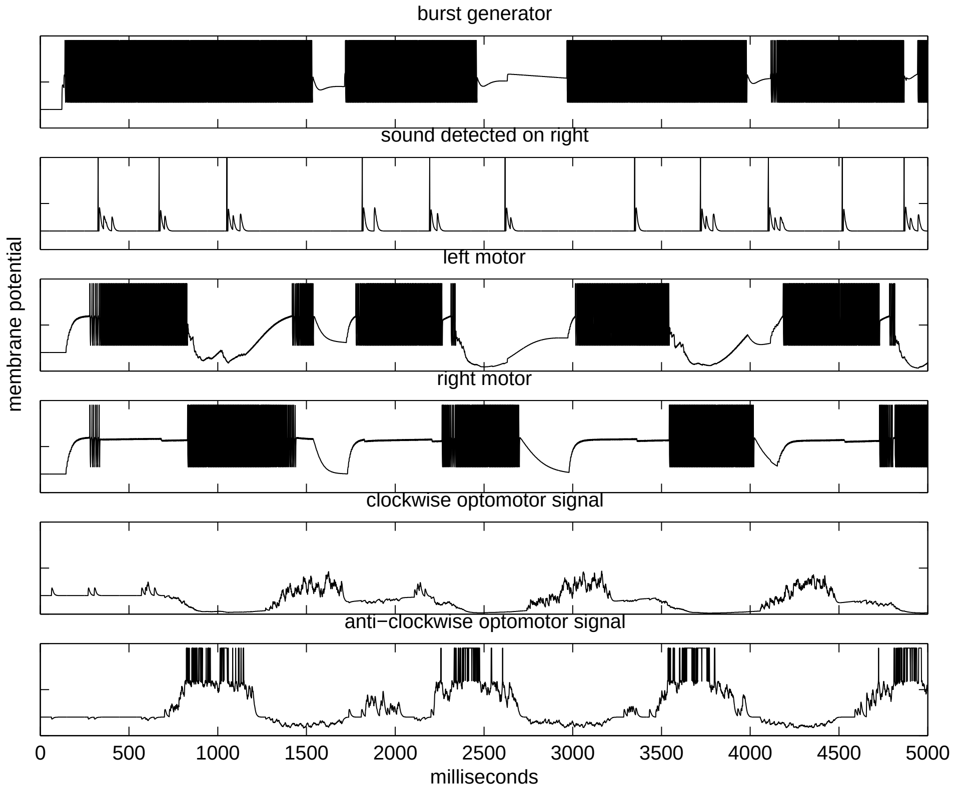

Figure 12

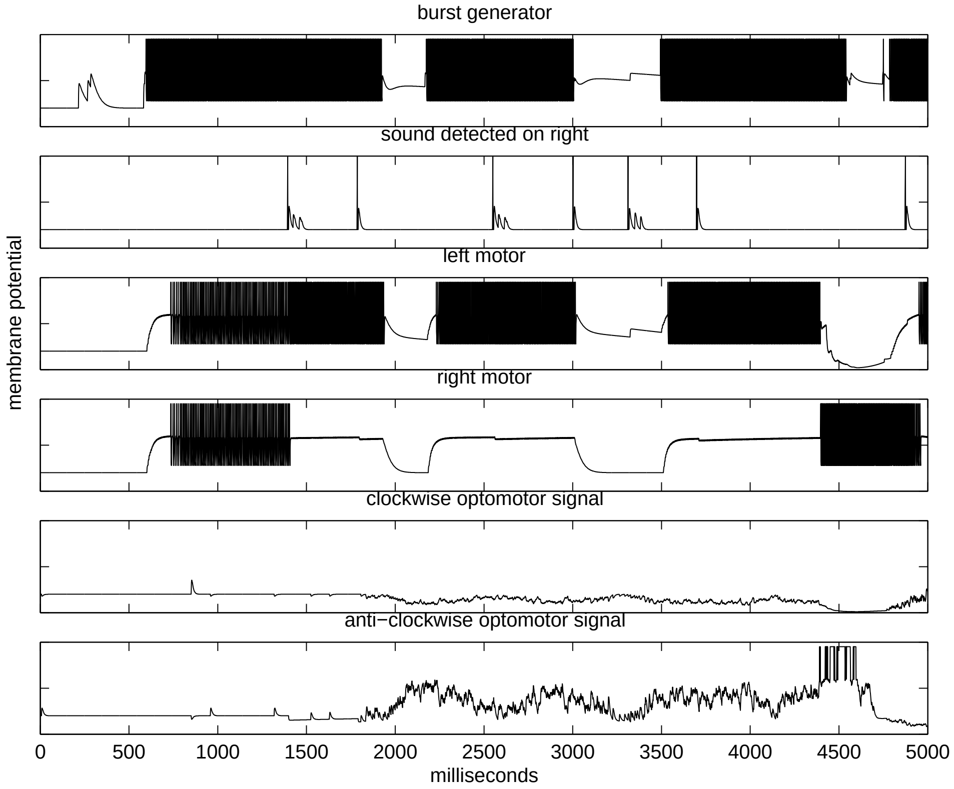

Figure 13

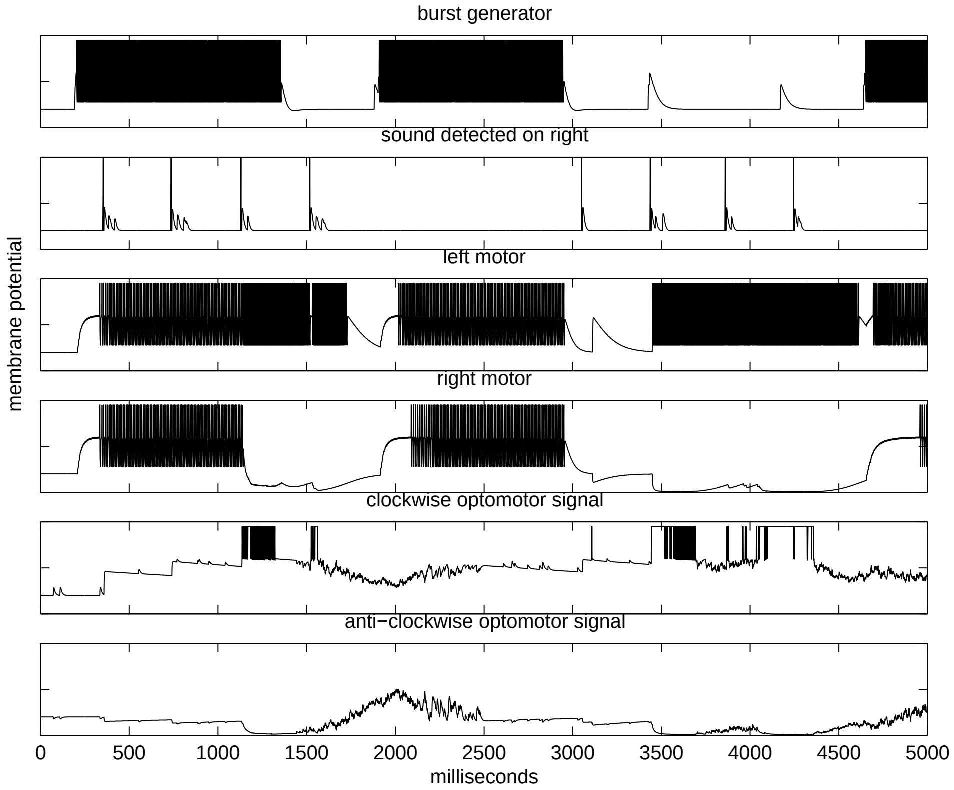

Figure 14

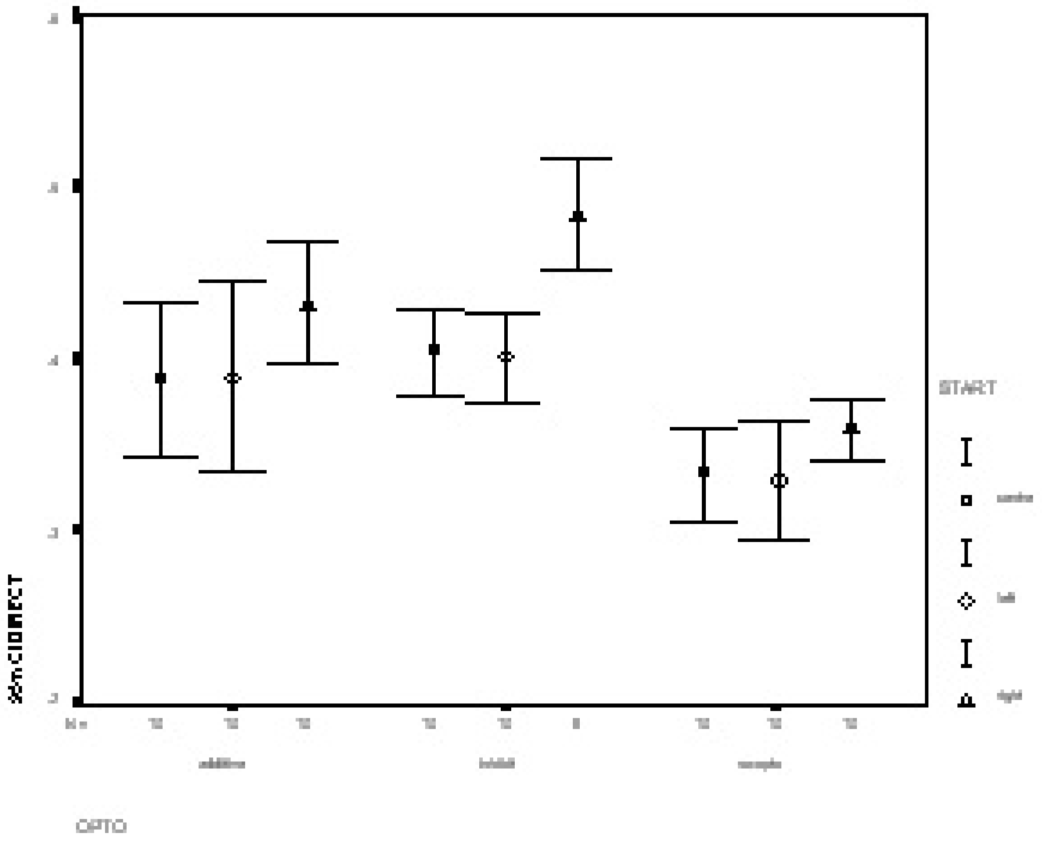

Figure 15

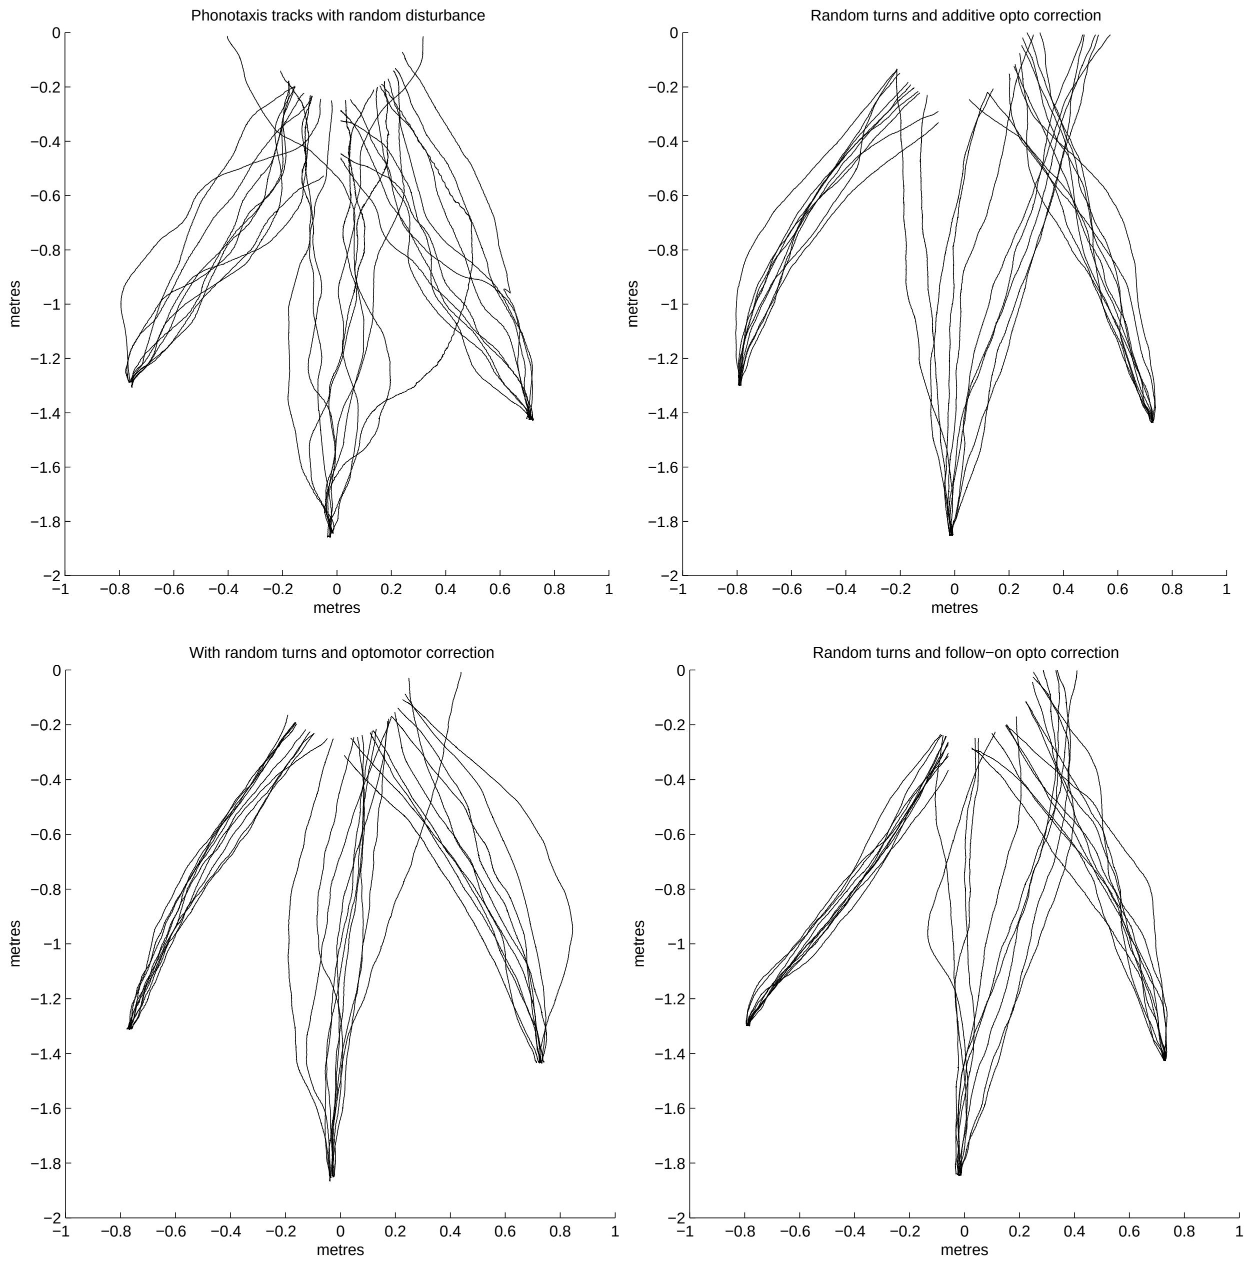
Phonotaxis tracks with random disturbance
Random turns and additive opto correction
With random turns and optomotor correction
Random turns and follow−on opto correction
metres
metres
0
−0.2
−0.4
−0.6
−0.8
−1
−1.2
−1.4
−1.6
−1.8
−2
−1
−0.8
−0.6
−0.4
−0.2
0
0.2
0.4
0.6
0.8
1

Figure 16

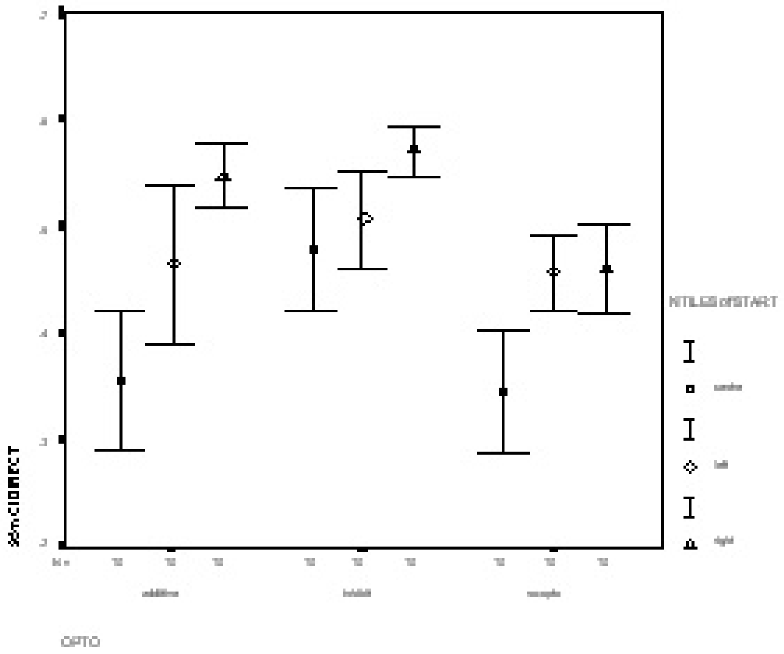

.7
.6
.5
.4
.3
.2
95% CI CORRECT
NTILES of START
centre
left
right
additive
inhibit
OPTO