

Integrating sensorimotor systems in a robot model of cricket behavior

Barbara Webb^{*a} and Reid R. Harrison^b

^aCentre for Cognitive and Computational Neuroscience, Dept. of Psychology, University of Stirling, FK9 4LA, U.K.

^bElectrical Engineering Department, University of Utah, Salt Lake City, Utah 84112, USA

ABSTRACT

The mechanisms by which animals manage sensorimotor integration and coordination of different behaviors can be investigated in robot models. In previous work the first author has built a robot that localizes sound based on close modeling of the auditory and neural system in the cricket. It is known that the cricket combines its response to sound with other sensorimotor activities such as an optomotor reflex and reactions to mechanical stimulation of the antennae and cerci. Behavioral evidence suggests some ways these behaviors may be integrated. We have tested the addition of an optomotor response, using an analog VLSI circuit developed by the second author, to the sound localizing behavior and have shown that it can, as in the cricket, improve the directness of the robot's path to sound. In particular it substantially improves behavior when the robot is subject to a motor disturbance. Our aim is to better understand how the insect brain functions in controlling complex combinations of behavior, with the hope that this will also suggest novel mechanisms for sensory integration on robots.

Keywords: sound localization, insect vision, biorobotics, sensorimotor integration

1. INTRODUCTION

Biorobotics is a field that not only takes ideas and inspiration from biology for robot building, but can also contribute to biology by building robots that are true models of animal systems. This means that the robot is used to implement and test hypotheses about the neural control and sensor/actuator interface to the environment that underlie the animal's behavioral capabilities. It could be said that "investigating animal behavior" is the "application" for which these robots are built¹⁻⁶. At the same time it is apparent that, as animals still substantially exceed in capacity any robot yet implemented, understanding their behavior is likely to contribute to robot technology in the future. One notable feature of most animals, compared to most robots, is the number and diversity of sensory systems they utilize. How they are capable of integrating these different systems in the control of behavior is the issue the current paper addresses.

To do so, we aimed to build on the previous work we have done in investigating individual sensorimotor systems, specifically the phonotaxis behavior of crickets and the optomotor response of flies. For both these systems it was found that incorporating appropriate peripheral sensory processing in the model resulted in a signal that was easily translated into the required motor action. That is, the control problem was solved largely by having sensors that act as 'matched filters'⁷ with respect to the environmental and task constraints. But do these mechanisms scale up to more complex behaviors or does the combination of different sensorimotor processes require more sophisticated central processing based on representations? Taking a bottom-up approach to this problem we looked at the issues raised in integrating both the phonotactic and optomotor responses on a single robot. In this paper we expand on preliminary work⁸, report experimental results on an opto-phono-taxis system, and suggest issues to be explored in further investigations.

* Correspondence: Email: b.h.webb@stir.ac.uk; Telephone:+44 1786 466841; Fax: +44 1786 467641

1.1. Cricket and robot phonotaxis

The phonotaxis behavior of crickets illustrates the principal of using a matched filter to simplify sensorimotor processing. Female crickets are able to locate males by walking towards the calling song they produce, and they do so by utilizing a unique auditory mechanism. The cricket's eardrums, located on the forelegs, are connected by a tracheal tube through which sound waves can propagate. Consequently the vibration of the eardrum reflects the combination of the direct sound and the delayed sound traveling from the other ear (see figure 1). This delay is tuned to the carrier frequency of cricket song (around 4-5kHz) such that the relative phase of the waves reflects the direction of the sound source, and thus the summed vibration has an amplitude that reflects the direction of the sound source⁹.

By building a robot model that incorporated this localization device (see more detail below) we could investigate the neural connectivity that allows the cricket to actually go towards the sound, and to discriminate the temporal pattern of the calling song. A common assumption is that these two aspects of the behavior are processed separately, that is, that the cricket has specific neural circuitry that filters for the temporal pattern, and uses the result to 'gate' the localization behaviour¹⁰. However, we proposed that the utilization of temporal (as opposed to spike rate) coding in the neural system could allow one mechanism to serve both functions. Briefly, a pair of auditory neurons sum the input and the one that fires first provides input to one of a pair of motor neurons which sum these inputs to determine a turn. By simple variation in neural time-constants, the motor response can be made selective to a bandpass of repetition rates in the signal, which is a characteristic of cricket behaviour^{11,12}. Full details of this model can be found in Webb & Scutt¹³.

This simple four-neuron mechanism when implemented on the robot was able to reproduce a surprising amount of the cricket's behavior. This included, not only a clear preference for a particular repetition rate (i.e. not responding if the sound did not have this pattern), but also the capability to choose and track one sound when similar or differing sounds were presented. When a song was presented from above (so that it provides no directional information) the robot still reacts like a cricket in showing distinctive meandering behavior with no preferred direction. If a song from above is augmented by playing a continuous tone from one side, crickets turn and track *away* from the continuous sound i.e. turning towards the side with the lower total sound amplitude, but less obscured song pattern¹⁴. Testing the robot in the same situation showed that it responded in just the same way¹³.

The cricket has around 100,000 times more neurons than this robot model - apparently few of these are needed to explain the range of phonotaxis behaviors we have described. However there are many ways in which the real cricket's behavior exceeds that replicated in the robot. For a start, the robot can only operate on a smooth, obstacle free surface, whereas a cricket in its natural environment must negotiate complex terrain (especially relative to its size), controlling a motor system with many degrees of freedom. Also crickets are obviously capable of many other behaviors, such as mating and reproducing, finding and ingesting food, and avoiding predators. Even when engaged in phonotaxis, other sensory stimuli have been shown to affect their behavior. As a starting point for investigating some of this behavioral complexity we considered first the experimentally demonstrated effects of visual stimuli on the crickets response to sound.

1.2. Visual influences on cricket phonotaxis

Crickets show a number of visually guided behaviors. It appears they can use visual information in maintaining a steady course direction: through an optomotor response that produces corrective movements to global rotation in the visual field; through tracking of a visual target; or even by maintaining a constant angle relative to skylight polarization. Weber et. al.¹⁵ compared the behavior of crickets tracking sound in the light and in the dark. Without visual information the cricket produced a zigzag path to the sound, meandering 30°-60° from the sound direction. In the light, it produced relatively straight runs in the speaker direction, with deviations of only around 6°. Differences in results for experiments done on cricket sound localization when walking in arenas (usually done in the light) versus walking on treadmills (usually done in the dark) have also been attributed to the effect of visual cues¹⁶.

Comparison of cricket sound tracking with and without an attractive visual target has been examined^{16,17}. The results suggest that the cricket tends to switch between tracking the sound and the visual target, rather than using the target to help maintain the phonotactic course. That is, tracking is improved if the target and sound directions coincide, but not obviously improved if they do not. The cricket may switch between tracking the target and the sound, and if the sound ceases the cricket reverts to tracking

the target rather than using it as a guide to maintain a course in the sound source direction. Similar results were found for the interaction of polarized light and sound tracking¹⁸. Another effect of the visual target on tracking is that the cricket becomes more selective with regard to the sound parameters that induce sound tracking; i.e. the attractive visual stimulus serves as 'competition' for the phonotactic response, which hence requires a more ideal stimulus to activate it¹⁶.

By contrast it has been suggested that sound tracking with optomotor stimuli shows an additive rather than a competitive interaction of the behaviors. Bohm et. al.¹⁹ used an 'open-loop' paradigm in which the cricket was fixed in position above a treadmill, to measure its turning tendency to various stimuli. A moving visual grating normally induces rotation in the corresponding direction. When this stimulus was added while the cricket was tracking a sound, the normal turning response of the animal to sound was shifted in the corresponding direction. They conclude that the "turning tendency can be explained as the weighted sum of the two turning tendencies evoked by the two individual stimuli" and that the "the optomotor response of crickets seems to contribute continuously to the stabilization of the phonotactic course". Thus we focused on the implementation of an optomotor modulation of the phonotaxis system on the robot.

1.3. Implementation of an optomotor response

Many animals demonstrate an optomotor reflex, that is, they respond to rotation of their visual field by self-rotation, which normally serves as an auto-correction for unintended course deviations. This behavior and the underlying neurophysiology have been most extensively studied in flies e.g.²⁰⁻²⁴ but the mechanism appears to be common to many insects, including the cricket. Photoreceptors in the retina underlying the compound eye signal changes from the ambient light level. The next layer of cells, the lamina, show highpass responses and connect to the medulla. Indirect evidence suggests that in the medulla, motion detection between adjacent photoreceptors takes place. This process is described by the elementary motion detector (EMD) proposed by Hassenstein & Reichardt²⁵, which carries out delay & correlation of adjacent photoreceptor outputs. By taking the difference of units sensitive to opposite directions of motion a strongly directionally selective response is produced. Subsequently, wide-field integration across many EMD outputs appears to be implemented by a number of identified cells in the lobular plate²⁶, which thus extract optic flow²⁷ with reduced dependence on the specific visual pattern.

An analog VLSI implementation of this architecture²⁸ (see further detail below) enables a real-time, low-power, extraction of the optomotor signal from the environment. This can be directly used to control an optomotor reflex on a robot. The output is positive or negative (relative to a reference voltage) by an amount approximately proportional to the image velocity to the left or right respectively. This signals the relative rotation of the robot and the environment, which ideally should be zero. A robot with two drive wheels can simply add this signal (smoothed by a low-pass filter and multiplied an appropriate gain term) to the default motor speeds in a negative feedback loop. Details of such an implementation are described in Harrison & Koch²⁹. Our aim was to integrate this capability onto the sound-localizing robot.

1.4. Combining the behaviors

Our initial implementation of the integration of phonotactic and optomotor behaviours⁸ used a straightforward addition algorithm. That is, the two sensory systems ran in parallel and their outputs were simply summed at the level of motor speeds, using different gains. This was moderately successful in that the robot could still find a sound and showed an apparent small improvement in the directness of paths, although this was not quantified. However this depended somewhat on the starting position of the robot. If the robot started out roughly facing the sound, the optomotor response help to diminish any tendency to turn away from the sound so the robot seemed more 'locked on' to the signal. However, if the robot started facing more away from the sound, the activity of the optomotor system made it less able to locate it. This is because each turn towards the sound generated by the phonotaxis system produced an excellent visual rotation signal which the optomotor system would correct, i.e. turning the robot back away from the sound. This problem became more apparent as we tried to test the behavior more systematically.

This was an empirical demonstration of the problem first 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? It is apparent that some sort of modification of the optomotor response needs to be part of the control system underlying the intended turns. We will here briefly consider three kinds of solution that have been proposed and why we chose to implement the one we did:

Non-overlapping turning speeds: The optomotor response typically has an upper and lower bound on the velocity at which it operates, i.e. very fast or slow visual motion do not lead to a tracking reflex. It has been proposed that if the intended turns are executed rapidly (e.g. saccadic turning maneuvers in flies^{30,31}), they will cause visual motion of a speed the optomotor system does not detect, and hence will not attempt to correct. However it is unlikely that this solution applies for cricket phonotaxis, as the typical rates of rotation during sound tracking fall well within the range for which an optomotor response can be induced. Our robot which was using turning speeds chosen to resemble the animal, received a strong optomotor signal from a typical turn to sound.

Efferent copy: If a sensorimotor system is signaling a movement then it can also pass that signal, as an 'efferent' copy to the sensory systems that will be affected by that movement, enabling the subsequent processing to take it into account. In this case, a phonotaxis turn can be predicted to generate an optomotor signal, so the response to that signal can be modified. Collett³² spells out several ways in which such a signal might interact with the optomotor response. The classic form is that the expected signal is subtracted from the actual signal: thus if the movement occurs correctly, no optomotor response is invoked, but if it is subject to additional disturbance, the optomotor response will still detect this. An alternative is that 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. A third option is to make the signaled size of turn larger by an amount that will compensate for the expected optomotor response. As Collett (op cit.) demonstrates, these three schemes algorithmically all reduce to addition of the two outputs with appropriate gains: implying that our original system might have worked if we had simply scaled the gain for phonotaxis to compensate for the optomotor effect. However as Collett also notes, the schemes are not equivalent when considered at the more detailed level of the temporal dynamics of the different reflexes. This is further discussed below.

Suppression: Like the first option this suggests that the optomotor response is simply inoperative during intended turns; however it differs in postulating that this 'switching off' is active suppression mediated by 'corollary discharge' rather than a simple consequence of limits of the optomotor system. It has the disadvantage that potentially useful optomotor information is also suppressed; but the advantage of simplicity in that the size of the expected optomotor signal does not have to be predicted. This option seems particularly attractive for what might be termed 'intermittent' behaviors, where the intentional turning will not suppress the optomotor response indefinitely but rather the animal will be able to switch sufficiently frequently between the two responses to make them useful. Note also that the reported 'additive' behavior of the cricket in Bohm et al (1991) is based on averaging of its response over several minutes, thus making it difficult to determine whether the net direction is a result of switching between or actually combining the two behaviors. Critical tests for the continued activity of the optomotor system require analysis of individual turns. Analysis of this kind for several systems has shown apparent switching (e.g. in response to 'escape' signals in the locust³³ and during pursuit turns in the housefly³⁴) in some cases, and apparent summation (e.g. in hoverfly tracking³² and during 'torque spikes' in the housefly²² (but only in one direction)) in others.

In practice, we were not able to find a gain for phonotaxis that in an additive scheme would properly compensate for the induced optomotor response. This seemed to be due to the time-constant of the optomotor response - the torque signal built up during the turn and lasted a little after it, so the compensation needs to be variable, although the turn signal itself is constant. Rather than try to generate some prediction of the optomotor signal (and note precise prediction is impossible as the actual output of the optomotor system is a function of velocity, image spatial frequency content and image contrast i.e. it depends on the actual visual pattern that the robot will experience while rotating as well as the rotation velocity) we instead adopted a simple suppression scheme, in which a turn generated by the phonotaxis system temporarily turned off the optomotor reflex.

2. METHODS

2.1. Hardware

The robot based 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 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.

The auditory circuit is shown schematically in fig 1 (see ^{35,36} for further details.). 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 measured using a RMS circuit, resulting in two analog signals representing the amplitude of the two ears.

The optomotor circuit is shown schematically in fig 2. 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.

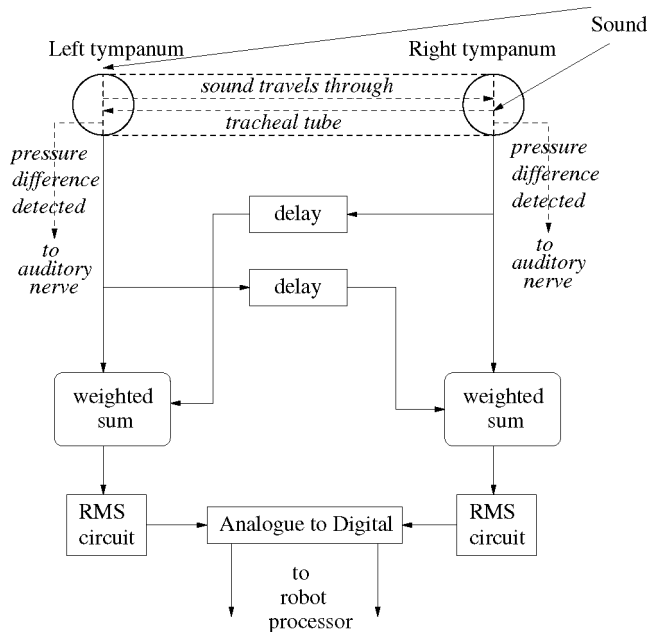


Figure 1: Schematic showing the cricket auditory system (dotted lines and italic text) and the robot auditory circuit that copies its function, i.e. the cross delay and combination of sound signals, and measurement of resulting wave amplitude, to give a response that is direction dependent

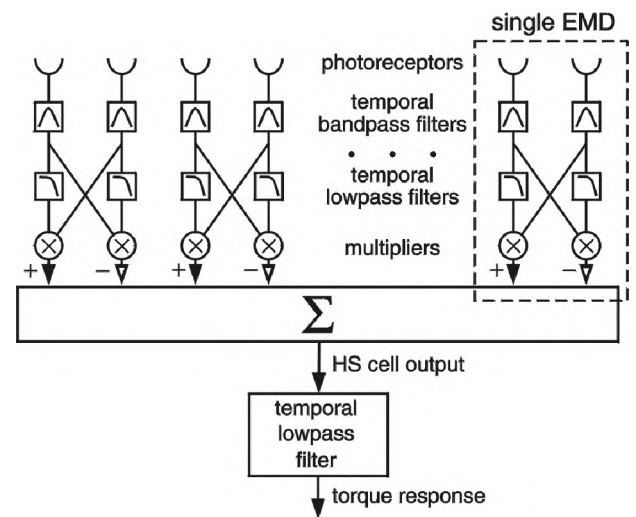


Figure 2: Schematic showing the visual system and chip design. Each photoreceptor correlates its response with the delayed (through lowpass filter) response of adjacent photoreceptor. The output is summed across the chip, and temporally filtered.

2.2. Software

The software used is a modification of that described in Webb & Scutt¹³. The main function of the original software was to simulate the cricket's auditory neurons using a simple state-based integrate-and-fire neuron model. Thus 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 which signals a turning response, and inhibits the effects of the opposite auditory-motor connection. But several successive onsets are required to excite the motor neuron above threshold.

Spikes in the motor neurons represent signals to turn to the right or left. In the original model, the turn was implemented by altering the speeds of the two motors appropriately for a fixed time. Effectively, a left spike sets a *left_turn* variable to one, which is reset to 0 after a fixed time (around 400 ms). If another left spike occurs in that time, the time is extended. While *left_turn* = 1, the left wheel speed is set to zero or backwards, so the robot turns left. The same mechanism applies for right spikes. If a right spike occurs during a left turn (or vice versa), the current turn is stopped and the opposite turn started. To integrate this response with the optomotor signal the *left_turn* and *right_turn* variables were combined as:

$$ears_signal = right_turn - left_turn$$

i.e. *ears_signal* would equal +1 for a right turn, -1 for a left turn, or 0 for no turn.

The output of the optomotor circuit was differenced from a reference value measured before the robot was moved at the beginning of each trial. Thus a positive *opto_signal* indicated rightwards visual motion and a negative *opto_signal* leftwards visual motion. Generally the value ranged between +/- 100. By multiplying this value by a fixed gain term, the turning rate of the robot could be set to follow this motion, thus compensating for the rotation. To implement the suppression, this gain was set to 0 whenever the *ears_signal* was non-zero.

The final motor output was thus calculated as:

$$left_speed = base_speed + opto_gain * opto_signal + ears_gain * ears_signal$$

$$right_speed = base_speed - opto_gain * opto_signal - ears_gain * ears_signal$$

The base speed used in the following experiments was 20, which is equivalent to moving at about 10 cm/sec. The ear_gain was set at 20, so that the robot would respond to a turn signalled from the ears by stopping one wheel and doubling the speed of the other. The optomotor gain was set at 0.14, so that the robot would slow down one wheel and speed up the other by an amount proportional to the visual velocity signal. Alternative gain settings were tried but not systematically explored, as these values seemed to give adequate behavior.

Some additional minor modifications to the program concerned the transfer of the phonotaxis behavior to the Koala from the previously used Khepera. It was necessary to write a new interface to the ears circuit. More significantly, the larger distances covered by the larger robot required it to work successfully over a wider range of amplitudes of the auditory signal. The auditory circuit was able to encode this range but the neural simulation encountered saturation problems. Hence we introduced an approximate logarithmic compression of the *ears_signal* as a prior process to the neural summation.

2.3. Experimental methods

The robot behavior was tested in the normal lab environment. 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 3 Hz. This sound was simulated using a customized program under Linux, writing directly to the sound card. The speaker was placed on the floor of the lab, about one meter from the top wall. No special soundproofing or other controls for noise or echoes were used. The visual stimulus was simply the lab furniture.

We carried out trials with the robot starting in one of three positions: from near the center of the room, facing the speaker from about 250cm; and from half-way down each side wall, facing the opposite wall, about 200cm from the speaker (see track figures). The robot was stopped (manually) if it was about to hit the speaker or else one of the lab walls; a successful trial was counted as the center of the robot being within 45cm of the speaker at the point it was stopped.

The tracks were recorded using the shaft encoders of the robot, which are sufficiently accurate on the Koala for dead reckoning over the short paths we were using. The number of shaft encoder counts for the left and right wheels for each successive point recorded (usually every 100ms) was converted to x-y position data as follows:

$$x_{i+1} = x_i + distance_i * \cos(\phi_i)$$

$$y_{i+1} = y_i - distance_i * \sin(\phi_i)$$

where:

$$distance_i = \frac{shaftencoder * (leftcount_i + rightcount_i)}{2}$$

$$\varphi_i = \varphi_{i-1} + \frac{shaftencoder * (leftcount_i - rightcount_i)}{2\pi * wheelbase}$$

$shaftencoder = 0.004374$ (centimeters distance per shaft-encoder count, determined by calibration)

$wheelbase = 28.5$ (centimeters distance between the left and right wheels)

x_1, y_1, φ_1 are specified according to the starting position, where the origin ($x=0, y=0$) is the speaker position and $\varphi=0$ is the direction of the positive x-axis.

Using this data, we also calculated several measures to compare how directly the robot had approached the sound, including:

- $heading_i = \arctan\left(\frac{-x_i}{-y_i}\right) - \varphi_i$ (the heading of the robot relative to the speaker)
- the standard deviation of the $heading$ for the entire track, i.e. how much the robot varied its heading.
- $length = \sum_i distance_i$ i.e. the total distance covered by the robot in the track
- the normalized mean vector for the track, calculated as follows:

$$magnitude = \bar{x}^2 + \bar{y}^2 \quad angle = \arctan\left(\frac{\bar{y}}{\bar{x}}\right) \quad \text{where } \bar{x} = \frac{\sum distance_i * \cos(heading_i)}{length} \quad \bar{y} = \frac{\sum -distance_i * \sin(heading_i)}{length}$$

The angle of the normalized mean vector indicates the average heading of the robot relative to the speaker, 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. To look for a significant difference in the directness of tracks we adapted the measure used by Schul³⁷ for crickets tracking on a treadmill. This combines the cosine of the angle of the mean vector (which varies from 1 to 0 as the robot deviates from going towards the speaker) and the length of the mean vector (which varies from 1 to 0 as the robot deviates more around the mean angle) and also (in Schul's paper) a measure of the straight-line distance covered by the animal while walking on a treadmill for a given length of time, divided by the maximum distance. Because our robot was instead covering a fixed distance in a variable amount of time we substituted the measure:

$$tracktime = \frac{\text{minimum time to do trial}}{\text{actual time to do trial}} \quad (\text{which varies from 1 to 0 as the robot takes longer to finish the trial})$$

as the third component. Overall directness is then scored as:

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

The more direct the track to the speaker, the closer D should approach to 1. Both standard deviation of the heading (H) and D were used in statistical tests for mean differences.

3. RESULTS

3.1. Phonotaxis only

With just the phonotaxis system operating (figure 3), the robot displayed the characteristic (of crickets) zigzag tracks to reach the sound source, when started either directly opposite the speaker (12 trials) or at either side of the room (7 trials each). The two failures to reach the speaker seemed principally due to adverse echoes. Essentially, this confirmed that the mechanism previously used with the Khepera had been successfully transferred to the larger robot. However, we did not carry out any tests of the selectivity of the phonotaxis system.

3.2. Adding the optomotor response

Next we added the optomotor response, as described in the Methods, and ran the same sequence of trials. In this case the robot never failed to find the speaker, and the tracks, on visual inspection, look a little more direct (figure 4). A two-way ANOVA on the standard deviation of the headings (H) found no significant interaction, no effect of starting position and a marginally significant improvement with the optomotor response ($p=0.052$, $F=3.970$, $d.f.=1$). Similar results were found in a two-way ANOVA on the directness measure (D) - in this case the optomotor improvement was not significant ($p=0.102$, $F=2.775$, $d.f.=1$) but it should be noted there was a significant difference in the variance of D across the conditions (Levine's test, $p=0.003$, $F=5.221$, $d.f.=3$) due to a much wider range of values for paths without optomotor control (see error bars in figure 7). Thus adding the optomotor response was not detrimental to phonotaxis, and may have made it slightly more reliable, but these effects were marginal. We did not get, for example, the large reduction in range of heading angles reported by Weber et. al.¹⁵ for crickets walking in the light vs. the dark.

The lack of obvious improvement can be mostly attributed to the fact that phonotaxis control was already fairly good at getting the robot directly to the sound source, because there was little else to put it off-course once it had turned to face the right direction. The main function of an optomotor response is to maintain a straight line heading when subject to disturbance or drift. The Koala robot when instructed to move in a straight line on a flat floor has little difficulty maintaining a straight course and hence little to gain from an optomotor reflex. In fact it is noticeable that the tracks from the sides tend to be more gently curved towards the sound when the optomotor response is included, whereas without it, the robot sometimes turns very directly towards the sound, but sometimes makes substantial deviations. The optomotor response mostly seems to be acting to limit the size of the turns - as before, this is useful if the robot is facing the sound but can be detrimental if it is trying to turn through a reasonably large angle to face it.




Figure 3: Tracks of the robot using phonotaxis only to travel from three different starting positions towards the sound source (at '+'). It tends to 'zigzag' towards the sound, and misses it twice.




Figure 4: Tracks of the robot using phonotaxis and the optomotor reflex to travel from three different starting positions towards the sound source. The tracks are marginally more direct than in Figure 3.

3.3. Correcting for motor bias

We next added a 'disturbance' to the robot's normal motor behavior. This was a systematic bias in its base motor speeds, such that the left motor would run 20% faster than the right. Most crickets walking on a treadmill show some systematic directional bias³⁷; and crickets in their natural environment often have motor asymmetries, as well as having to deal with environmentally caused deviations. With this bias, the robot without phonotactic or optomotor control would tend to turn to the right. With optomotor control (and no sound), it was able to correct for the bias to produce an approximately straight path (this replicated the experiments described in²⁹).

Figure 5 illustrates 12 sound-tracking trials with the bias and without the optomotor reflex. The robot still reaches the sound in half the trials, but the tendency to head to the right of the speaker is clear. In one case it loses the sound altogether; in another it makes a complete circle in the bias direction before moving to the sound. When the optomotor response is added (figure 6) there is a clear improvement in the behavior. A t-test on the difference in heading H shows a highly significant improvement for the optomotor condition ($p=0.001$, $t=3.625$, $d.f. =22$) and similar results pertain to the directness measure D ($p=0.012$, $t=2.727$, $d.f. =22$). In fact the tracks with optomotor correction become comparable to those in the original, no bias conditions. There was no significant difference for H or D when comparing the optomotor tracks with and without the motor bias.

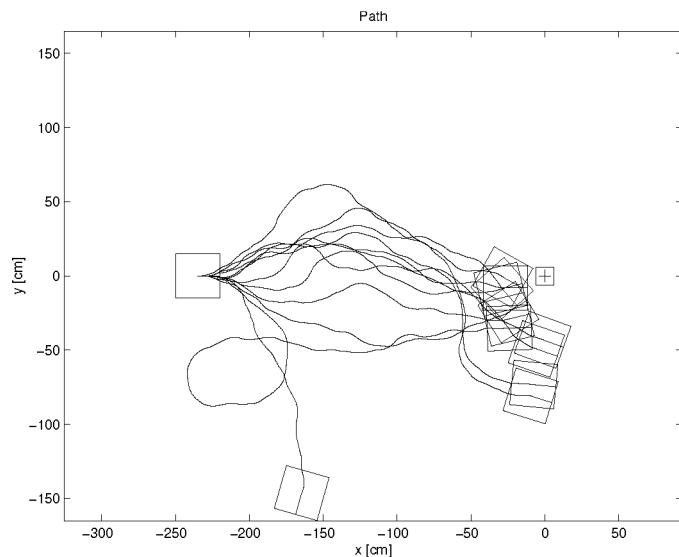


Figure 5: Tracks of the robot when it has a 20% lag induced between its motors with phonotaxis only. The robot tends to miss the speaker by heading too far to the right

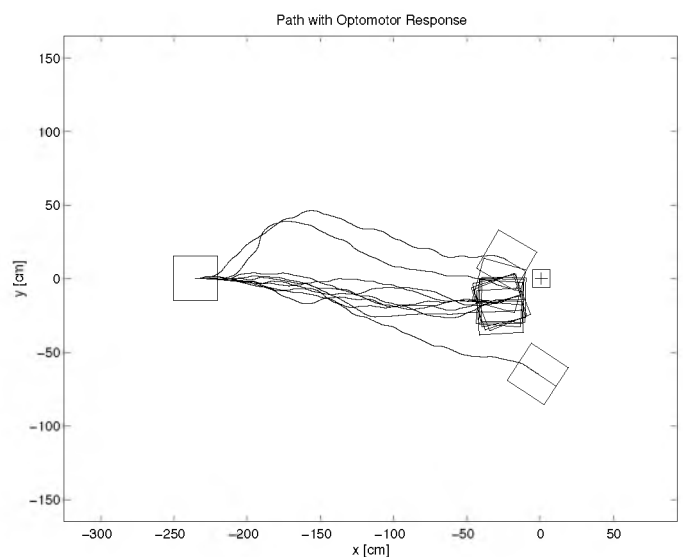


Figure 6: Tracks of the robot when it has a 20% lag induced between its motors with the optomotor response added. The optomotor response substantially improves the behavior

3.4. Summary

These results are summarized in figures 7&8 which show the mean and standard deviation of the measures H and D for the six experimental conditions described: tracking from directly in front, or from the sides, or with a motor bias; with and without the optomotor reflex. It can be seen that the optomotor reflex always improves tracking (reducing heading variance and 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. Adding the bias makes tracking significantly worse without the optomotor response, but does not affect it with optomotor response. The optomotor conditions also show substantially less variance than the phonotaxis only conditions.

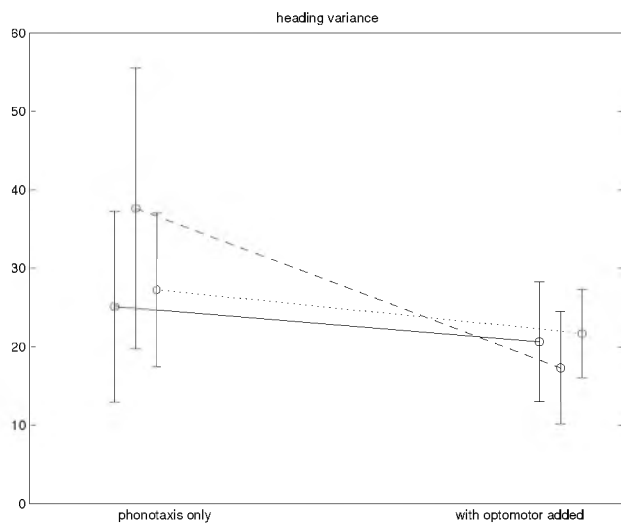


Figure 7: Variance of the heading relative to the speaker (means and standard deviation error bars). Solid line, tracks from center with no bias, dotted line, tracks from the sides with no bias, dashed line, tracks from center with 20% motor bias. Left with phonotaxis only, right with the optomotor response added.




Figure 8: "Directness" of the tracks, coding as in figure 7. The optomotor response substantially improves the behavior when the motor bias is imposed. It also decreases the variance of the tracks.

4. DISCUSSION

The scheme we have investigated here for the integration of visual and auditory reflexes on a robot model of the cricket is a very simple one. Essentially it works by allowing one behavior, when active, to inhibit the other. We have shown that this mode for adding the optomotor response to phonotaxis:

- Is consistent with biological evidence, although other solutions are also possible
- Does not interfere with normal phonotaxis under good conditions
- Improves phonotaxis significantly under conditions of degraded motor control

While continued work in this direction - stepwise introduction of a more varied and flexible range of sensorimotor behavior - will undoubtedly require more complex integration mechanisms in the future, even at this preliminary level some interesting issues have been raised.

One is the importance of matching the task context of the robot and the animal. Most insects have an optomotor response, but many robots do not need one, as they are not subject to equivalent random or systematic course perturbations. We found it necessary to introduce an arbitrary perturbation to illuminate the effectiveness of optomotor control. However certain robot tasks are more likely to inherently contain these irregularities - in particular, tasks that involve movement in natural rather than man-made environments. Flying or swimming robots will be subject to unpredictable air or water currents, but even terrestrial robots navigating over uneven terrain with varied surface traction are unlikely to be able to rely on purely feedforward control to maintain a steady direction. Outdoor phonotaxis is one current target of our research, and optomotor control is likely to be a useful component of achieving this, but it is anticipated that additional sensing will also be required.

Another issue is the importance of the 'implementation' level in deciding which scheme for combining behaviors is in fact most viable. As mentioned above, "adding in" an efferent signal to modify the optomotor response has rather different consequences if implemented in different ways, largely because of the difference in the temporal dynamics of the various control systems: for example instabilities are harder to avoid in a scheme that adds a proportional increase to the gain of the intended turn than in a scheme that subtracts a proportional signal from the optomotor signal³². Even in the simpler case of suppression there are different options for the interaction that have different consequences. We used an inhibition of the output of the optomotor system, but this allowed the integration of the signal to continue. A consequence was that when the phonotactic turn ended, the optomotor system often reacted to the residual signal and made a small correction. We could have attempted to extend the

inhibition to cover this period. Alternatively we could have had the inhibition occur at the level of input to the integrator, so the rotation during the phonotactic turn would not contribute to it. However in our current configuration, the integration takes place in hardware so this would have necessitated implementing an external switch.

We are also interested in translating the integration algorithms into models based on realistic neural mechanisms. There are several ways in which 'inhibition' can be neurally implemented and again these methods may have different behavioral effects. There is some physiological evidence³⁸ regarding multimodal neurons in the cricket brain that respond to both calling song and optomotor stimuli, which also show the interesting characteristic of differential responses when the animal is walking or stationary. Understanding the functioning of these biological circuits, of interest in itself, may also provide novel ideas for engineering sensorimotor integration.

ACKNOWLEDGEMENTS

This work was supported by a grant from the Gatsby Charitable Foundation. Thanks also to Richard Reeve who assisted with some of the programming and experiments.

REFERENCES

1. R. Möller, D. Lambrinos, R. Pfeifer, T. Labhart, and R. Wehner, "Modeling Ant Navigation with an autonomous agent," in: *From Animals to Animats 5*, MIT Press, Cambridge, Mass., 1998.
2. F. Grasso, T. Consi, D. Mountain, and J. Atema, "Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: Progress and challenges," *Robotics and Autonomous Systems* **30**, pp. 115-131, 2000.
3. N. Burgess, J. G. Donnett, and J. O'Keefe, "Using a mobile robot to test a model of the rat hippocampus," *Connection Science* **10**, pp. 291-300, 1998.
4. M. Blanchard, F. C. Rind, and P. F. M. J. Verschure, "Collision avoidance using a model of the locust LGMD neuron," *Robotics and Autonomous Systems* **30**, pp. 17-38, 2000.
5. S. Huber, M. O. Franz, and H. H. Bulthoff, "On robots and flies: modeling the visual orientation behaviour of flies," *Robotics and Autonomous Systems* **29**, pp. 227-242, 1999.
6. T.M. Morse, T.C. Ferree, S.R. Lockery, "Robust spatial navigation in a robot inspired by chemotaxis in *Caenorhabditis elegans*," *Adaptive Behaviour* **6**, pp. 393-410, 1998.
7. R. Wehner, "Matched filters - neural models of the external world," *Journal of Comparative Physiology A* **161**, pp. 511-531, 1987.
8. B. Webb and R. Harrison, "Eyes and ears: combining sensory motor systems modelled on insect physiology," in: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 3913-3918, 2000.
9. A. Michelsen, A. V. Popov, and B. Lewis, "Physics of directional hearing in the cricket *Gryllus bimaculatus*," *Journal of Comparative Physiology A* **175**, pp. 153-164, 1994.
10. F. Huber, "Behavior and neurobiology of acoustically oriented insects," *Naturwissenschaften* **79**, pp. 393-406, 1992.
11. J. Thorson, T. Weber, and F. Huber, "Auditory behaviour of the cricket II Simplicity of calling-song recognition in *Gryllus* and anomalous phonotaxis at abnormal carrier frequencies," *Journal of Comparative Physiology A* **146**, pp. 361-378, 1982.
12. J.A. Doherty, "Phonotaxis in the cricket, *Gryllus bimaculatus* De Geer: comparisons of choice and no choice paradigms," *Journal of Comparative Physiology A*, **157**, pp. 279-289, 1985.
13. B. Webb, and T. Scutt, "A simple latency dependent spiking neuron model of cricket phonotaxis," *Biological Cybernetics* **82**, pp. 247-269, 2000.
14. J. Stabel, G. Wendler, and H. Scharstein, "Cricket phonotaxis: localization depends on recognition of the calling song pattern," *Journal of Comparative Physiology A* **165**, pp. 165-177, 1989.
15. 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**, pp. 215-232, 1981.
16. J. Stout, G. Atkins, T. Weber, F. Huber, "The effect of visual input on calling song attractiveness for female *Acheta domesticus*," *Physiological Entomology* **12**, pp. 135-140, 1987.
17. T. Weber, G. Atkins, J. Stout, F. Huber, "Female *acheta-domesticus* track acoustical and visual targets with different walking modes," *Physiological Entomology* **12**, pp. 141-147, 1987.

18. T. Weber, "Phonotaxis and visual orientation in *Gryllus Campestris*: Behavioural experiments," in: *Sensory Systems and Communication in Arthropods*. Basel: Birkhauser Verlag; 1990.
19. H. Bohm, K. Schildberger, F. Huber, "Visual and acoustic course control in the cricket *gryllus-bimaculatus*," *Journal of Experimental Biology* **159**, pp. 235-248, 1991.
20. K.G. Gotz, "The optomotor equilibrium of the *Drosophila* navigation system," *Journal of Comparative Physiology* **99**, pp. 187-210, 1975.
21. W. Reichardt, and T. Poggio, "Visual control in the fly. A quantitative analysis," *Quarterly Review of Biophysics* **9**, pp. 311-375, 1976.
22. M. Heisenberg, R. Wolf, "Reafferent control of optomotor yaw torque in *drosophila-melanogaster*," *Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology* **163**, pp. 373-388, 1988.
23. M. Egelhaaf, "Visual afferences to flight steering muscles controlling optomotor responses of the fly," *Journal of Comparative Physiology A* **165**, pp. 719-730, 1989.
24. M. Egelhaaf, A. Borst, W. Reichardt, "The nonlinear mechanism of direction selectivity in the fly motion detection system," *Naturwissenschaften* **76**, pp. 32-35, 1989.
25. B. Hassenstein, and W. Reichardt, "Systemtheoretische Analyse der Zeit-, Reihenfolgen-, und Vorzeichenbewertung bei der Bewungsperezeption des Rüsselkäfers," *Chlorophanus. Z. Naturforschung* **11b**, pp. 513-524, 1956.
26. K. Hausen, "Motion sensitive interneurons in the optomotor system of the fly .2. The horizontal cells - receptive-field organization and response characteristics," *Biological Cybernetics* **46**, pp. 67-79, 1982.
27. H.G. Krapp, and R. Hengstenberg, "Estimation of self-motion by optic flow in single visual interneurons," *Nature* **384**, pp. 463-466, 1996.
28. R.R. Harrison, C. Koch, An analog VLSI model of the fly elementary motion detector . in: *Advances in Neural Information Processing Systems 10*, pp. 880-886, MIT Press, Cambridge, MA, 1998.
29. R.R. Harrison, and C. Koch, "A robust analog VLSI motion sensor based on the visual system of the fly," *Autonomous Robotics* **7**, pp. 211-224, 1999.
30. H. Wagner, "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**, pp. 581-595, 1986.
31. M. Egelhaaf, "Dynamic properties of two control systems underlying visually guided turning in house-flies," *Journal of Comaparive Physiology* **161**, pp. 777-783, 1987.
32. T. Collett, "Angular tracking and the optomotor response: an analysis of visual reflex interaction in a hoverfly," *Journal of Comparative Physiology* **140**, pp. 145-158, 1980.
33. D. Robert, C.H.F. Rowell, "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**, pp. 53-62, 1992.
34. M.S. Srinivasan, and G. D. Bernard, "The pursuit response of the housefly and its interaction with the optomotor response," *Journal of Comparative Physiology* **115**, pp. 101-117, 1977.
35. H. H. Lund, B. Webb, and J. Hallam, "A robot attracted to the cricket species *Gryllus Bimaculatus*," in: *Fourth European Conference on Artificial Life*, MIT Press, Cambridge MA, 1997.
36. 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**, pp. 95-107, 1998.
37. J.L. Schul, "Song recognition by temporal cues in a group of closely related bushcricket species (genus *Tettigonia*)," *Journal of Comparative Physiology A* **183**, pp. 401-410, 1998.
38. H. Bohm, K. Schildberger, "Brain neurons involved in the control of walking in the cricket *gryllus-bimaculatus*," *Journal of Experimental Biology* **166**, pp. 113-130, 1992.