

Collision avoidance using a model of the locust LGMD neuron

Mark Blanchard^{a,*}, F. Claire Rind^{b,2}, Paul F.M.J. Verschure^{a,1}

^a Institute of Neuroinformatics, ETH-University Zurich, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland

^b Department of Neurobiology, Medical School, University of Newcastle upon Tyne, Framlington Place, Newcastle upon Tyne, NE2 4HH, UK

Abstract

The lobula giant movement detector (LGMD) system in the locust responds selectively to objects approaching the animal on a collision course. In earlier work we have presented a neural network model based on the LGMD system which shared this preference for approaching objects.

We have extended this model in order to evaluate its responses in a real-world environment using a miniature mobile robot. This extended model shows reliable obstacle detection over an eight-fold range of speeds, and raises interesting questions about basic properties of the biological system. ©2000 Elsevier Science B.V. All rights reserved.

Keywords: Insect vision; Lobula giant movement detector (LGMD); Mobile robot; Obstacle detection; Collision avoidance

1. Introduction

For animals the ability to detect approaching objects is important for survival, serving both to prevent collisions as the animal moves and also to avoid capture by predators [30]. While the fate of a mobile robot is unlikely to involve being eaten the ability to avoid collisions is equally important. Traditionally robotic technology has involved active sensors, such as ultrasound and infra-red devices, or high-precision sensors, such as laser scanners, for the detection of obstacles [7]. In biology, however, many examples are found of systems which rely on visual information to accomplish this task, and these have been successfully applied to real-world robotic control tasks [8,49]. Neurons tuned

to respond to approaching objects have been identified in species as diverse as pigeons [42] and locusts [31,38].

The lobula giant movement detector (LGMD), a large visual interneuron in the optic lobe of the locust [24], is one such neuron. Originally the LGMD was thought to be tuned to detect novel movement of small objects [35]. However, more recent work has shown that the LGMD responds most strongly to approaching objects [31,38] and that it is tightly tuned to objects approaching the animal on a direct collision course [14]. Receding objects produce little or no response [31]. Approaching objects are distinguished by the LGMD using the increasing speed of edge movement and increasing length of the edges [40]. The response dynamics during the approach of an object are the subject of two conflicting reports, one showing that the spike rate increases continuously during approach [31], the other showing that the spike rate may peak before collision [9,12]. At the time of writing, this conflict is unresolved.

* Corresponding author. Tel.: +41-1-635-30-70; fax: +41-1-635-30-53.

E-mail addresses: jmb@ini.phys.ethz.ch (M. Blanchard), claire.rind@ncl.ac.uk (F.C. Rind), pfmjv@ini.phys.ethz.ch (P.F.M.J. Verschure).

¹ Supported by SPP, Swiss National Science Foundation.

² Supported by a Royal Society Research Fellowship.

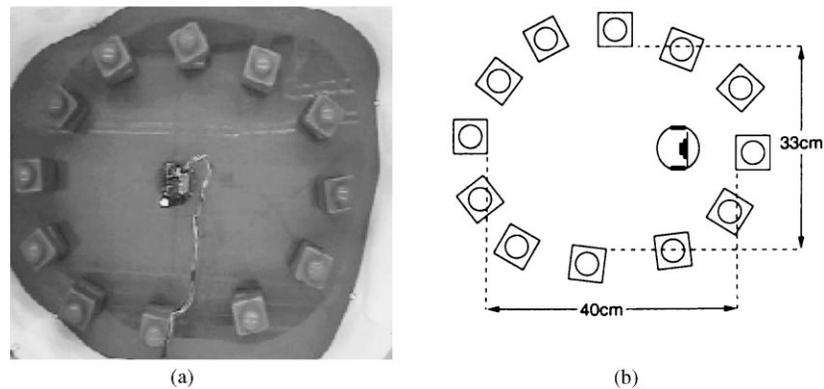


Fig. 1. Experimental environment. (a) Robot's arena as seen by the overhead camera. (b) Dimensions of the arena.

The LGMD is believed to play a role in triggering escape jumps and steering responses during flight. This belief is supported by the strong connection of the LGMD to the descending contralateral movement detector (DCMD) neuron [23] which in turn makes connections with interneurons and motoneurons in the thoracic ganglia [5,39]. Responses in the DCMD match responses in the LGMD one-for-one [23,27].

A neural network model of the input circuitry of the LGMD was developed by Rind and Bramwell [29]. Based closely on the anatomy and physiology of the optic lobe, the network comprised three principal layers:

- the input photoreceptive layer, which responded to changes in the image in order to detect the edges of moving objects;
- a processing layer, where excitation passed retinotopically through the network while delayed inhibition spread laterally;
- the output layer, which represented the LGMD, where the excitation and inhibition were combined.

In addition, a feed-forward pathway inhibited the output of the model LGMD during large changes in the image, such as those caused by ego-motion. This model displayed the same preference for approaching objects as the LGMD and revealed that, at least for simple stimuli, this preference results from a critical race between the excitation produced by the movement of an object's edges and the delayed lateral flow of inhibition within the network. This model was subsequently extended to allow the network to respond to textured stimuli [3].

In this paper we evaluate the behaviour of the LGMD network in a real-world environment. We show that robust obstacle detection can be achieved using an insect-based solution which relies on vision. Our results have interesting implications for the understanding of the biological LGMD system and demonstrate the potential of our LGMD model for real-world obstacle detection applications.

2. Methods

2.1. Experimental apparatus

Experiments were conducted in the environment shown in Fig. 1. This comprised of small stacks of Duplo blocks of various colours (red, green or blue) within a white outer wall. At the widest point the blocks were separated by 40 cm and at the narrowest point by 33 cm (Fig. 1(b)). The spaces between blocks were approximately 5 cm. The floor of the environment was clear Perspex over a gray sheet of paper.

A Khepera mobile robot (K-Team A.G., Lausanne, Switzerland) was fitted with a monochrome pinhole CCD camera (K2D-BW-PAL, K-Team A.G., Lausanne, Switzerland) which gave a field-of-view of approximately 60°. The camera was angled downwards by 5° to exclude views over the wall of the environment. The views seen via this camera are shown in Fig. 2.

A colour CCD camera (K2D-C-PAL, K-Team A.G., Lausanne, Switzerland) was mounted above the cen-

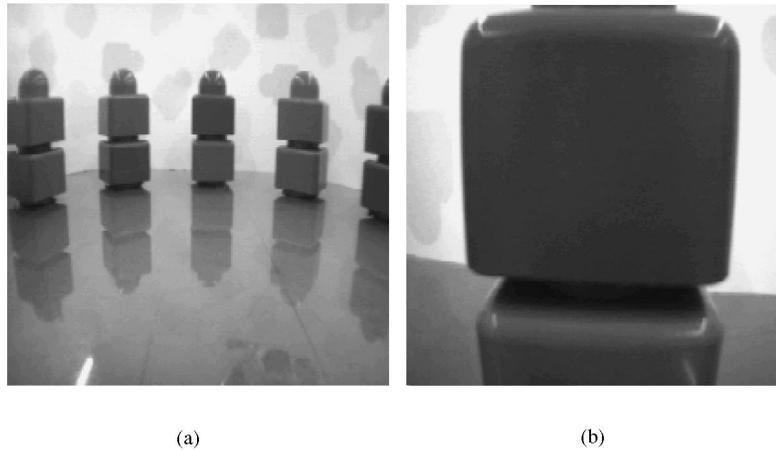


Fig. 2. View of the environment as seen from the camera mounted on the robot: (a) across the full width of the environment; (b) 2 cm from a block. Notice that the images include significant shadows and reflections.

tre of the environment and faced directly downwards. The images from this camera were used to track the position of the robot during our experiments. Fig. 1(a) shows the view from this camera.

2.2. Simulation system

We performed our experiments using the simulation software IQR421, formerly known as Xmorph [43]. IQR421 provides a graphical programming environment for the specification, analysis, and documentation of large-scale heterogeneous neural models which can be interfaced to external devices. Real-time robot control can be achieved using a distributed computing environment based on the TCP/IP protocol. This distributed system allows both the computational load and specialised hardware to be spread among several machines.

Processes were run on three Pentium II (two 450 MHz and one 333 MHz) PCs connected via an ethernet. All the machines used the Red Hat Linux 5.1 operating system. Video frame grabbers (Mediavision ProMovie Studio) were fitted into each 450 MHz machine which provided 210×210 pixel images. The robot was connected to the serial port on the 333 MHz machine.

For these experiments we constructed the simulation system shown in Fig. 3. The simulation system comprised of four processes, the LGMD model,

the robot control system, the IQR421 tracker module TraX and the IQR421 graphical user interface. In total, our simulation comprised over 1800 cells and 6000 synapses, and ran at approximately 16 timesteps per second. The LGMD, robot control and TraX processes are described below. Equations for the cell types used and the parameter values are given in Appendix A.

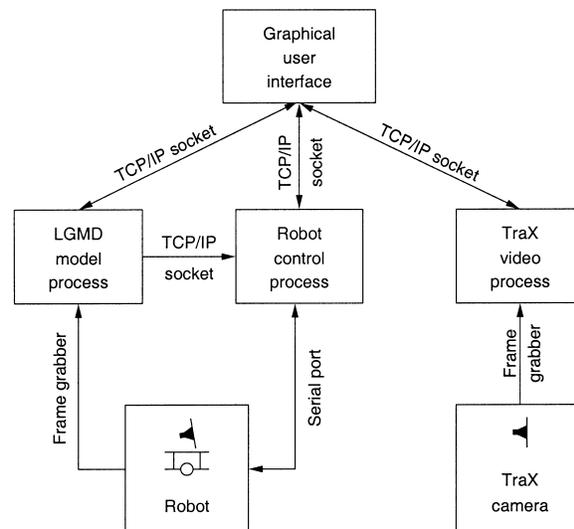
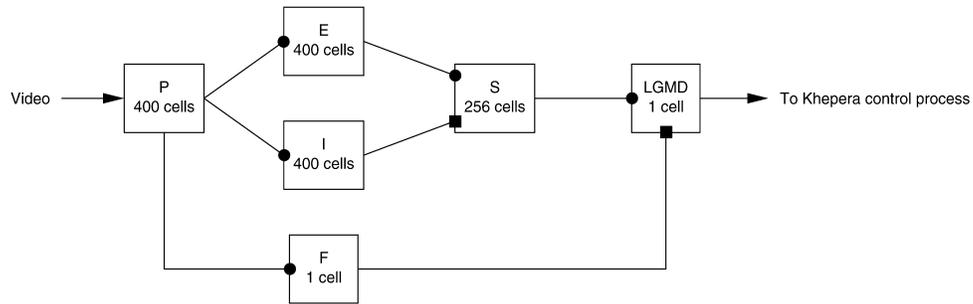
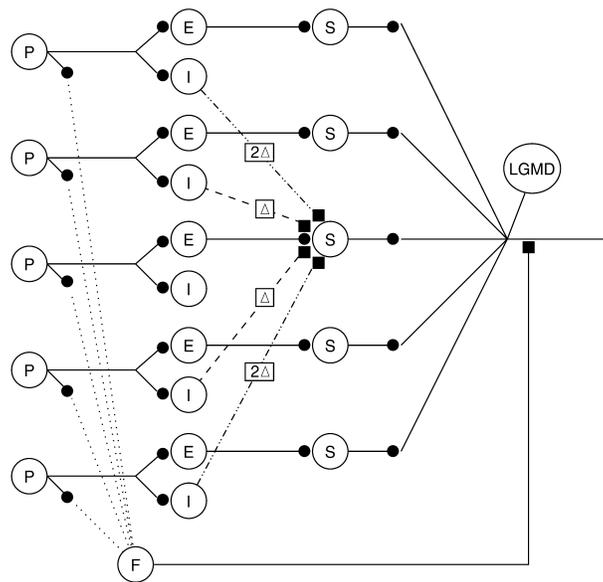


Fig. 3. IQR421 simulation system used for the experiments. The processes were distributed in order to spread the input/output load.



(a)



(b)

Fig. 4. Circuit for LGMD model, which comprised approximately 1500 cells and 4600 synapses. (a) Cell groups used in the model (P, photoreceptive; E, excitatory; I, inhibitory; S, summing; F, feed-forward; LGMD) and the global connections between these groups. The P cell group received the input from the video camera. (b) Detail of the connections between specific cells within the model. The S cells received strong excitation from their topographically aligned E cell without delay, and weaker delayed inhibition from topographically neighbouring I cells. Circles, excitatory connections; squares, inhibitory connections; Δ , delay of one timestep; solidity of lines represents connection strengths.

2.2.1. LGMD model process

The LGMD model was based on the input organisation of the biological system. Each element within the network represented a cell type found in the locust eye. An overview of this model is given here, full descriptions are given in [3,29].

The model comprised of the four groups of cells (P, photoreceptive; E, excitatory; I, inhibitory and S, summing) and two single cells (F, feed-forward inhibition and LGMD) used in the original model (Fig. 4(a)). The detailed connections (Fig. 4(b)) are described below.

Input to the model, via the 20×20 group of photoreceptive P cells, was taken from the difference between successive images derived from the robot-mounted camera, which were sampled at 16 frames per second. The P cells were integrate-and-fire cells which responded to edges crossing the receptive field of the cell with discrete spikes (impulses). The number of cells was chosen to reduce the resolution of the camera image to represent approximately that of the locust eye.

Output from the P cells passed to two groups, the excitatory E cells and the inhibitory I cells, and the single feed-forward inhibitory F cell. The P-to-E and P-to-I groups were connected topographically with each P cell connected to one E and one I cell. On the contrary, the excitatory projections from the P cell group to the single F cell connected the central 256 P cells to the F cell equally. This gave a measure of the total change within the image, and was used to inhibit the LGMD during large changes in the image. Only the central 256 P cells were used in the P-to-F connections in order to eliminate edge effects. This strategy was also used for the E-to-S and I-to-S connections.

The linear threshold E and I cells integrated their inputs and passed suprathreshold values directly to their outputs, approximating the graded potential neurons observed in insect visual systems. The output from these cells passed to the S cells. The excitatory E-to-S connection linked the central 256 cells in each group topographically. However, in the inhibitory I-to-S connection the central 256 S cells each received input from a ring of 12 I cells in the nearest and next-nearest neighbouring positions, forming a lateral inhibitory network within the model. The weights and delays of the synapses within these connections depended on the distance between the I and S cells, with the weight decreasing and the delay increasing as the distance increased.

The integrate-and-fire S cells received excitatory input from the E cells and inhibitory input from the I cells. The gains of these two inputs were independent, and the output from the S cells was computed using the difference between the total excitation and the total inhibition received.

Finally, input to the single integrate-and-fire LGMD cell was from the S cells and the F cell. The S-to-LGMD connection connected the cells in the S group to the LGMD equally. Input from the S cells

was excitatory, input from the F cell was inhibitory. The cell produced a pulse of activity whenever the integrated input exceeded a preset threshold, after which the membrane potential of the cell was reset.

2.2.2. Robot control process

The robot control process was derived from our previous work called Distributed Adaptive Control (DAC, a study of synthetic models of learning and problem solving, see [44,45] for an overview). It consisted of a so-called reactive control structure capable of controlling the robot using only the input from the infra-red sensors. The input from the LGMD was used to trigger avoidance actions, a turn to the left, when it responded to obstacles. Fig. 5 shows the structure of the circuit used.

This control structure provided three specific behaviours:

- basic exploratory activity (translation);
- simple collision avoidance responses using the infra-red sensors (rotation);
- avoidance of more distant obstacles, triggered by the responses LGMD (rotation);

All sensory inputs converged onto the MotorOut cell group, which comprised an array of 100 (10×10) linear threshold cells. This group represented a motor map with the positions of the cells producing specific motor responses (Fig. 6). Cells in the upper half of the array produced forward motion, those in the lower half produced backward motion, and cells to the left and right produced turning responses in their respective directions. The motor response was selected on a winner-take-all basis with the cell with greatest activity defining the movement at each timestep.

The infra-red sensors were mapped onto the Coll group (Fig. 7), which consisted of eight linear threshold neurons. When any of these cells was active, avoidance reactions were triggered via the Reflex group of linear threshold cells, which connected one-to-one with the MotorOut group. Predefined connection patterns caused collisions on the left to produce turns to the right, and vice versa.

Exploratory activity was produced by the Noise group of random spike cells. The overall activity in this group was summed in the Explore group and resulted in forward motion by activating a specific cell in the Reflex group. During avoidance reactions triggered by

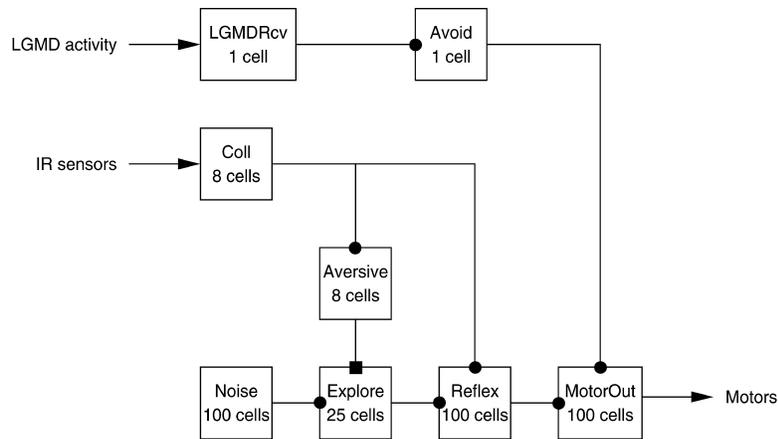


Fig. 5. Circuit of reactive control structure, which comprised 360 cells and over 1400 synapses. Input from the robot's infra-red sensors and the LGMD was received by the 8 Coll cells and the LGMDRcv cell respectively. The robot motors were controlled via the MotorOut group, which represented a map of possible motor states. Functional details are given in the text. Circles, excitatory connections; squares, inhibitory connections.

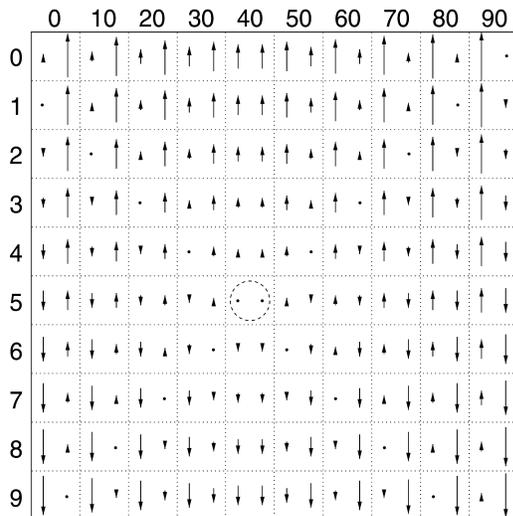


Fig. 6. Mapping of activity in the MotorOut group onto motor actions. The arrows show the speed and direction of the robot's two motors, with longer arrows denoting higher speeds. The cell with the highest response triggered the corresponding pattern.

the infra-red sensors, the Explore group was inhibited in order to suppress forward motion, ensuring that the avoidance reaction dominated in the MotorOut group.

LGMD responses were added into this control structure via two additional cells not present in previous DAC systems, LGMDRcv and Avoid. The activity from the LGMD was integrated in the LGMDRcv integrate-and-fire cell, which produced a single spike

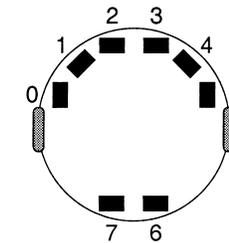


Fig. 7. Mapping of the robot's infra-red sensors onto the Coll cell group, viewed from above the robot. The numbers show the cell indices.

when an obstacle was detected by the LGMD circuit. This spike then triggered a pulse in the Avoid cell which connected directly onto MotorOut cell 4, which produced strong anti-clockwise turns. This simple avoidance strategy, which was not intended to replicate any specific biological circuits, was sufficient for the present study where the emphasis was on the responses of the LGMD model.

It was necessary to tune the responses of the LGMDRcv cell according to the speed of the robot. At the lower speeds tested (1.5–7.5 cm/s) several LGMD spikes were required in rapid succession in order to produce each LGMDRcv spike. However, at the highest speeds (10.0 and 12.5 cm/s), the LGMDRcv cell was tuned to produce a spike when only a single LGMD spike was received. The need for this retuning is discussed later.

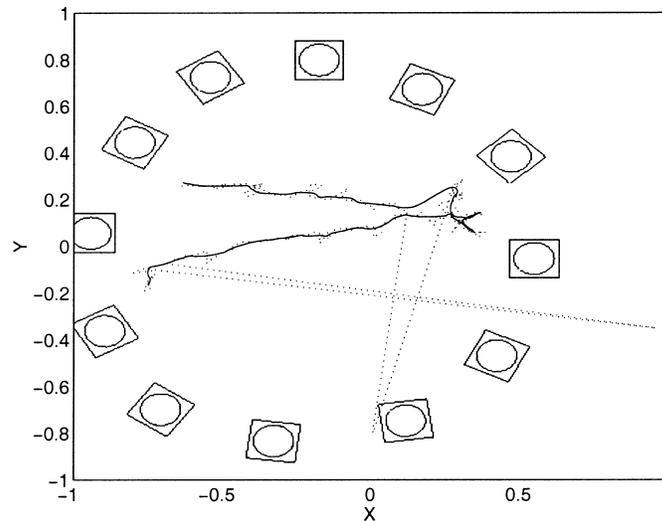


Fig. 8. Comparison between raw and edited position data, illustrating the removal of spurious jumps in position. The data shown was taken from the 5 cm/s experiment. Dotted line, raw data; solid line, edited data.

2.2.3. TraX process

The position of the robot during the experiments was monitored using IQR421's internal tracking system TraX. The position of the brightest moving point within the images from the overhead camera was found and mapped onto a normalised Cartesian coordinate system where $(-1, -1)$ corresponded to the bottom left corner of the image and $(1, 1)$ represented the top right corner. To ease the recognition of the robot, and hence minimise noise within the signal, a small lamp was attached to the left side of the robot (visible in Fig. 1(a)). The offset position of this lamp resulted in a slight error in the position recorded, particularly during rotation of the robot when the tracked position appeared to translate.

2.3. Protocol

At the beginning of each experiment the robot was placed in approximately the same position (indicated in Fig. 1(b)) with its centre 10 cm from the right-hand blocks, facing leftwards towards the distant blocks. Data was recorded for 1000 timesteps as the robot moved, which took on average 62 s. Six different speeds of movement were tested: 1.5, 2.5, 5, 7.5, 10, 12.5 cm/s.

Experiments were performed under normal fluorescent room lighting. Only one level of illumination was

used due to the limitations imposed by the CCD camera, which did not provide the constant contrast responses characteristic of insect visual systems [18].

The raw position data recorded by TraX was edited to remove spurious jumps caused by other objects (for example, the wall of the environment or the robot's cable) being tracked. Jumps of greater than 0.1 (5% of the range) during a single timestep were assumed to be errors, and missing data values were interpolated linearly from the position values immediately before and after the error. The data was then smoothed with a sliding average of width 11 timesteps. Fig. 8, which shows a comparison between raw and edited position data, illustrates that we were able to track the position of the robot robustly.

3. Results

The directional selectivity of both the LGMD neuron and our neural network model is illustrated in Fig. 9. During the approach of an object the spike rate from the LGMD increases whereas there is only a brief response to receding objects. This response is found over a wide range of speeds of movement [28] and is tightly coupled to objects moving on a direct collision course [14].

In our present study we ran two series of experiments. In the first set of experiments, using only reac-

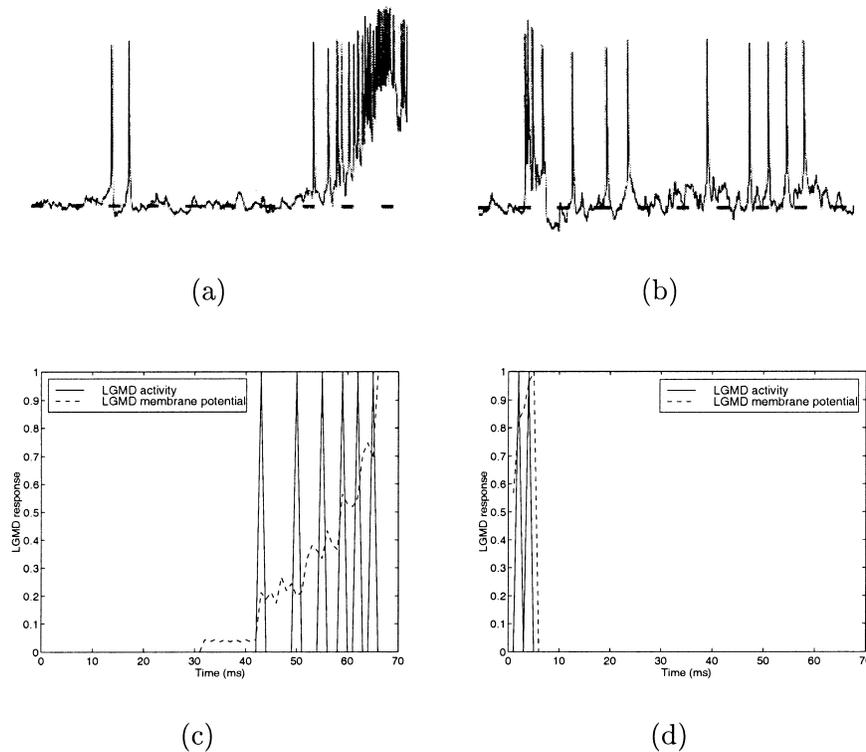


Fig. 9. Example responses to looming (a,c) and receding (b,d) objects of the LGMD neuron (a,b) and the LGMD model (c,d). For (a,b), the responses of the neuron to a $100\text{ mm} \times 75\text{ mm}$ object moving at 5 m/s were recorded using a sharp intracellular electrode. Each record shows 1.1 s of data. The traces in (c,d) show the spike rate of the model LGMD during movement of a 70 mm square object at 6 m/s . The responses of the original model (shown as dashed lines) have been converted into spikes (solid lines) to simplify comparison with the physiological record. Neuron responses from [28], model responses modified from [3].

tive infra-red based control for the robot, we recorded the responses of the LGMD model in isolation. In the second, the LGMD responses were used to control the robot. The parameters of the LGMD model, which are given in Appendix A, were modified to match the responses to the speed of movement of the robot, a strategy also observed in insect visual systems [22].

3.1. Reactive control

Example responses of the LGMD model at two translational speeds are shown in Fig. 10. As the robot approached the blocks in its environment the LGMD responded with trains of spikes which tended to increase in rate during the approach. Notice, however, that the timing of the spikes within these responses is less predictable than those observed in both the bio-

logical system and our previous simulations: for example, the two responses in Fig. 10(b) show very different patterns of activity during approach. We attribute this to the more complex nature of the stimuli used here, e.g. blocks of different colours, and noticeable shadows and reflections (see Fig. 2).

Fig. 11 shows the responses within the model during approach. When the robot was far from the blocks, the movement of edges was slow and the lateral inhibition from the I cells prevented responses from the S cells. However, just before collision the edges of the objects moved more quickly than the lateral inhibition and the S cells were active.

Single approach responses for all speeds tested are shown in Fig. 12. These traces show that the theory that LGMD spike rate increases with speed of motion and increases until collision, as discussed in the introduction, is not necessarily true when the variability of

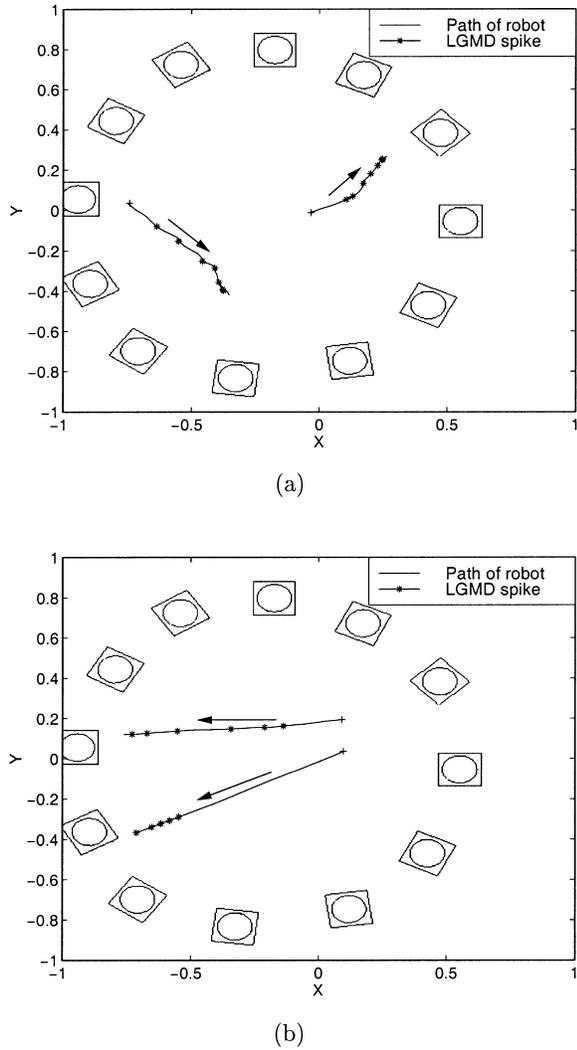


Fig. 10. Example responses of the LGMD when the robot moved under IR control at (a) 2.5 cm/s, (b) 10 cm/s. LGMD spikes are shown as *, the starting position of each trace is shown with a +, and the direction of motion is indicated.

real images is considered. This would affect the ability of the LGMD to trigger escape responses at specific times during approach. The traces for the four lowest speeds tested show the peak spike rate increased with speed. The peak rate during motion at 7.5 cm/s was close to the end of the response, but the peak occurred near the beginning of the response at 5 cm/s. At the two highest speeds the spike rate was lower than at either 5 or 7.5 cm/s. This specific result, however,

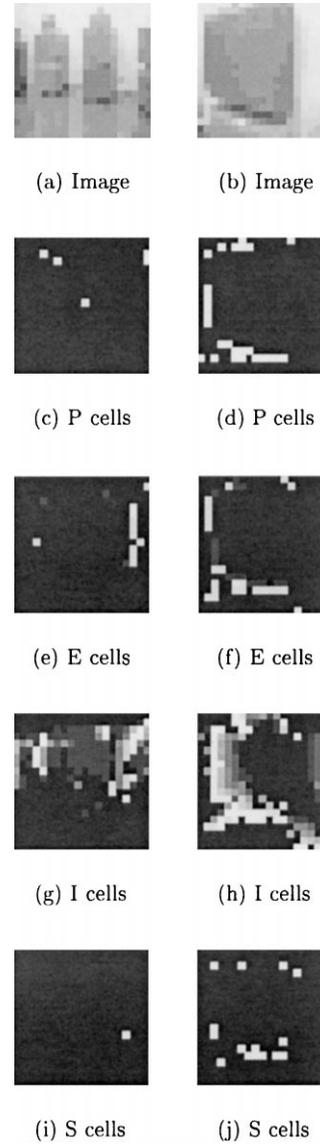


Fig. 11. Responses within the different cell groups during approach at 7.5 cm/s. Plots (a,c,e,g,i) show the responses when the robot was far from the blocks. Plots (b,d,f,h,j) show the responses just prior to a collision. In all plots, brighter squares represent higher activity. The time constant of the S cells was increased to aid interpretation of these static images.

could be an artifact of the limited rate of video capture and cannot at this point be interpreted as an inherent property of the LGMD model.

During motion which covered the whole visual field, inhibition from the I cells reduced the S cell responses.

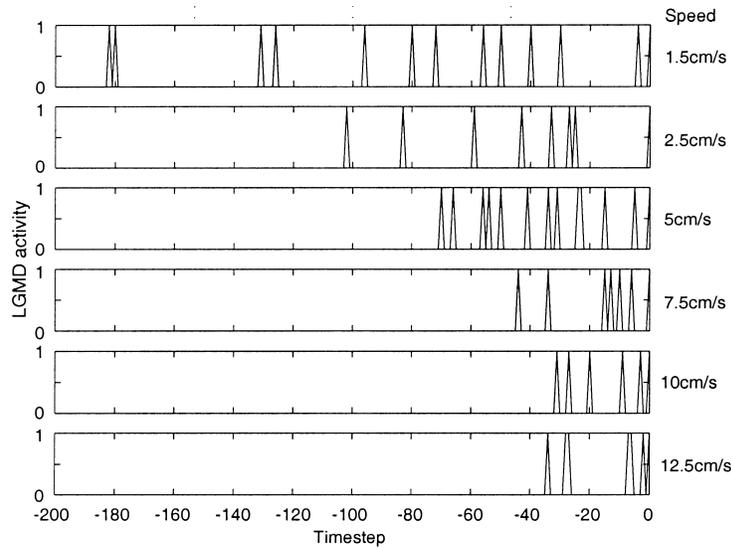


Fig. 12. LGMD responses during a single approach for all six speeds tested. Responses are aligned at the time of the last spike, which is taken as zero. From the top, the traces show responses during motion at: 1.5, 2.5, 5, 7.5, 10 and 12.5 cm/s.

This is shown in Fig. 13 where the excitatory input, inhibitory input and output activity of the S cell group are compared. During forward motion (Fig. 13(a)) the S cell activity followed the excitation closely with the inhibition having little effect. However, during rotation (Fig. 13(b)), which occurred at the same motor speed as forward motion but with one motor reversed, visual features flow equally over all parts of the image. This increased optic flow produced greater activity in all cell groups, and the inhibition reduced strongly the activity in the S cell group, and consequently the input to the LGMD. Note the relatively long decay time of the inhibition, which persisted after rotation ceased.

A second mechanism which controlled LGMD responses during whole-field motion involved the summed P cell activity inhibiting the LGMD via the F cell (Fig. 14). The F cell received excitation from all P cells but, due to its high threshold, was usually inactive during forward motion. However, during rotation, which produced activity in a large number of P cells, the F cell was activated strongly. This activity reduced the membrane potential of the LGMD.

3.2. LGMD control

Fig. 15 illustrates the behaviour of the robot when the LGMD response was used to avoid collisions. At

Table 1

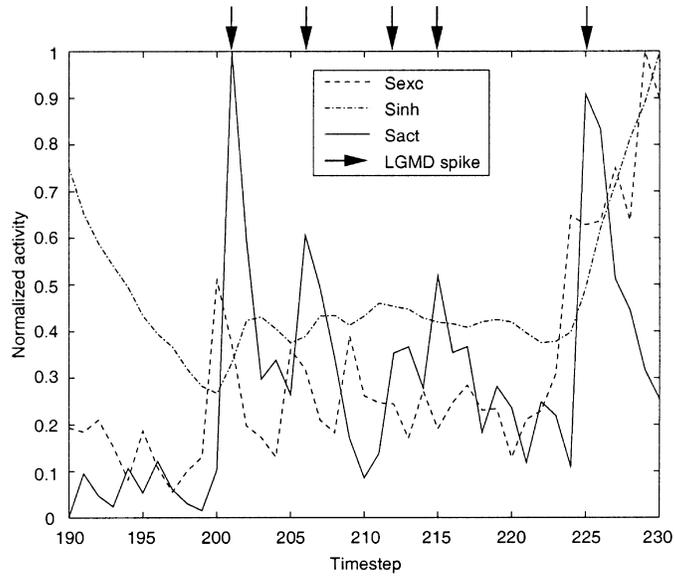
Percentage of avoidance reactions produced by the LGMD based reactive control structure relative to total events (events were defined as avoidance reactions + collisions)

Speed (cm/s)	Avoidance reactions (%)
1.5	90
2.5	91
5.0	81
7.5	69
10.0	88
12.5	92

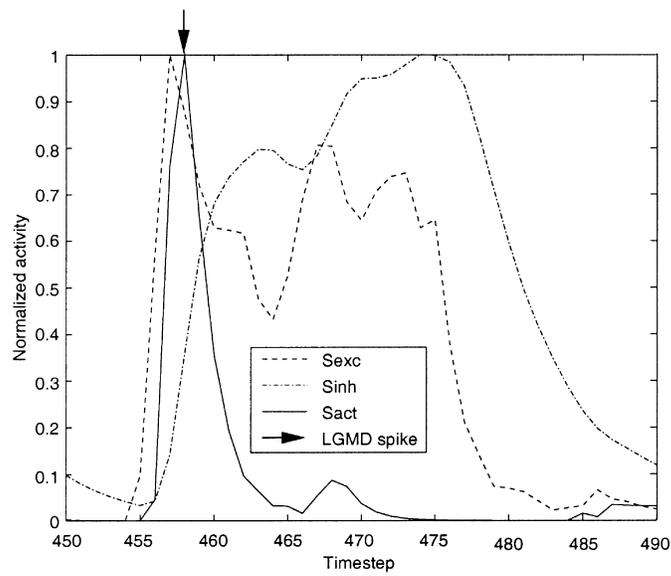
the lower speed (Fig. 15(a)) the input from the LGMD was integrated and several spikes were required to produce an avoidance reaction, whereas at the higher speed (Fig. 15(b)) the parameters of the LGMDRcv cell in the robot control circuit were modified to generate avoidance behaviour in response to individual LGMD spikes.

The performance of the LGMD control structure is shown in Table 1. The effectiveness of control via the LGMD was always better than 69% and for half of the speeds tested it was over 90%.

Most of the collisions which occurred while the robot was under LGMD control were due to the inhibition to the S cell group decaying slowly after the robot rotated. This is illustrated in Fig. 16 which shows



(a)



(b)

Fig. 13. Comparison between the normalised excitatory input, inhibitory input and output activity of the S cell group during (a) forward motion and (b) rotation, when the S group was inhibited and the LGMD produced no spikes. The data is taken from the 7.5 cm/s experiment. Dashed line, excitatory input to S cells; dot-dash line, inhibitory input to S cells; solid line, S cell group output activity; arrows, LGMD spikes.

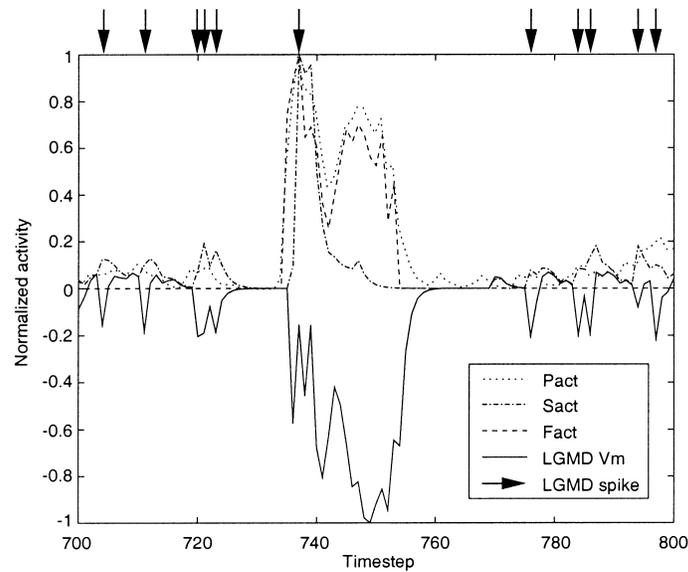


Fig. 14. Effect of the F cell during motion at 7.5 cm/s. The robot moved forward during the first 40 timesteps, then rotated for the next 40 steps before resuming forward motion. Dotted line, P cell group activity; dot-dash line, S cell group activity; dashed line, F cell group activity; solid line, LGMD membrane potential; arrows, LGMD spikes. Responses of P, S and F groups are normalised to maximum of P, S and F cell activity, LGMD membrane potential is normalised to its own minimum value.

both the activity within the LGMD model and the behaviour of the robot. Particularly at high translation speeds (at the highest speed tested the robot crossed its environment in approximately 3 s) the time between an avoidance reaction and the robot's next encounter with an obstacle could be short. In this case the activity of the I cells had insufficient time to decay, inhibiting the S cells and in turn suppressing LGMD responses. It will be interesting to challenge the locust using similar stimuli to see whether this response is observed in the neuron.

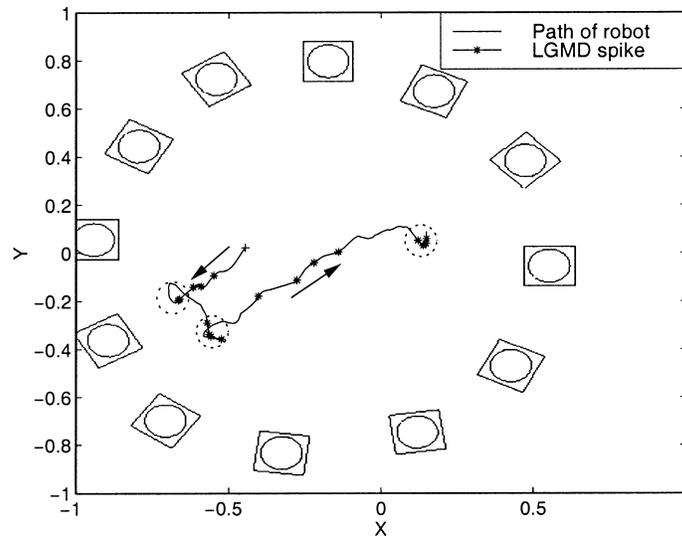
4. Discussion

In this project we evaluated a biologically based model of the locust LGMD in behavioural terms using a mobile robot. We found that the model responded reliably as the robot approached obstacles in its path over an eight-fold range of speeds. This illustrates that robust obstacle avoidance can be achieved using an insect based solution which relies on vision. Our experiments provided new insights into the dynamic

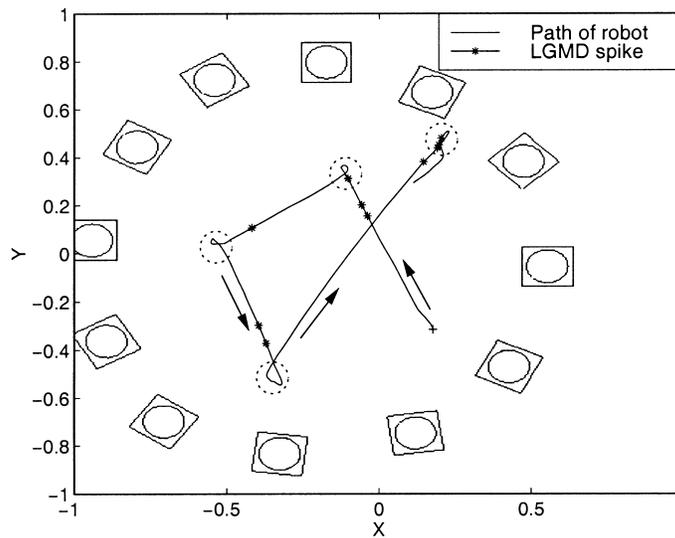
responses of our model, which will be investigated in future experiments using both the model and the neuron.

4.1. LGMD model

This is the first time that our LGMD model has been challenged with complex real-world stimuli, and we have shown that it is able to detect approaching objects reliably. The need to tune the model's parameters for the speed of motion suggests new interpretations of existing features of the LGMD system. For example, it was necessary to increase the P cell threshold for higher speeds of motion. Although we were only able to test the model at one level of luminance due to the variable contrast gain provided by our camera (see methods above), we suggest that this tuning may correlate with the habituation observed in the input neurons to the LGMD [35]. It is known that stimulation of the LGMD in one region of the eye produces responses which decrease with repetition of the stimulus. This habituation could act to stabilise the amount of excitation received by the LGMD due to increased



(a)

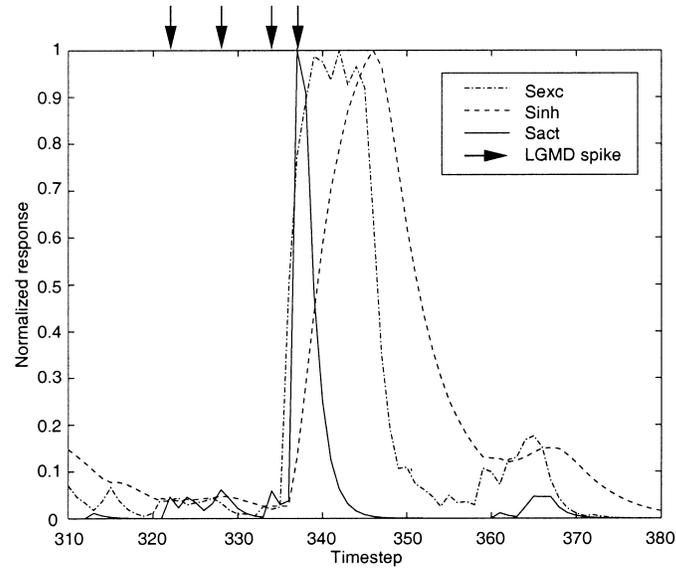


(b)

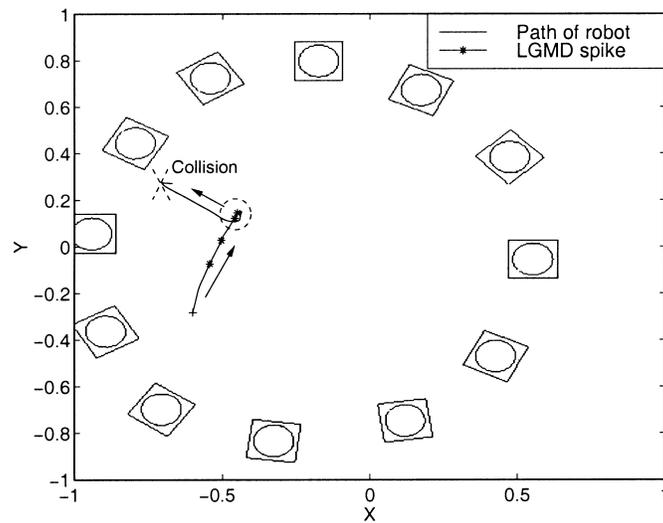
Fig. 15. Example responses of the LGMD when the robot moved under LGMD control at: (a) 2.5 cm/s; (b) 10 cm/s. LGMD spikes are shown as *, the starting position of each trace is shown with a +, and the direction of motion is indicated. Dashed circles indicate the positions at which avoidance reactions occurred.

optic flow. This is reminiscent of an automatic gain control system. Habituation of responses is seen in transient cells which are believed to provide the input to the LGMD [25]. In our model, addition of an

adapting threshold to the P cells may be sufficient to incorporate the effect of habituation. Alternatively, the model could be extended to include the earlier steps of visual processing which are only approximated by



(a)



(b)

Fig. 16. Collision due to prolonged decay of inhibition to S cells after rotation. (a) Activity within the model. Dotted line, excitatory input to S cells; dashed line, inhibitory input to S cells; solid line, S cell group output activity; arrows, LGMD spikes. (b) Corresponding behaviour of the robot. Solid line, path of robot; *, LGMD spikes; dashed circle, avoidance reaction; dashed X, collision. Speed of experiment was 7.5 cm/s, robot rotated between timesteps 335 and 350.

our P cells. A recent model of ON–OFF units in the fly visual system [37], which incorporated synaptic depression, could produce a similar effect.

There is evidence that the sensitivity of these circuits can be restored by neuromodulators released by particular neurons projecting from the brain back into the optic lobe. Octopamine was found to reduce the habituation seen in the LGMD during responses to repeated visual stimuli [2,41]. During motion, a graded release of octopamine could regulate the gain within the LGMD system according to the speed of flight. Presently, the dynamics of octopamine release require further investigation.

Most existing records of LGMD responses were obtained using simple stimuli comprising an object of constant illumination moving smoothly against an evenly illuminated background. The interpretation of these responses as either a continuous rise in spike rate during object approach [31,32] or as a peak before collision which codes the angular size of the object [12,16], relied on the LGMD seeing a single unambiguous object. Our experiments show that the responses of our model LGMD do not increase predictably to more complex real-world scenes, where the edges of objects may be occluded, in shadow, etc. This suggests that the responses of the LGMD may be poorly suited for triggering escape behaviour at a specific time during the approach of an object, such as when the object subtends a specific angular size [34] or at a constant time before collision [10]. Even though we used a simplified visual environment, the responses of the model could be strongly asymmetric during approaches towards the blocks. This implies that the behaviour of the LGMD is difficult to predict from its responses to simple stimuli and that an extensive study of LGMD responses to complex stimuli should be conducted. We shall pursue this goal using both our robot-based model and physiological experiments.

The suppressed responses of the model following rotation of the robot, caused by the prolonged inhibition of the S cells, enforced the tuning of the motor responses in order to trigger avoidance behaviour. However, this was not sufficient to prevent all collisions at the highest speeds tested. It is not known whether this prolonged inhibition is also present in the locust. We will perform further experiments on the locust using similar stimuli in order to resolve this question. A

modification of the model, moving the target of lateral inhibition from the S cells to the P cells (negative feedback), may reduce this problem by restricting the build-up of activity in the I cells. The origin of this proposed modification is recent electron-microscope analysis of the inputs to the LGMD [33]. This analysis revealed a novel arrangement of synapses around the dendrites of the LGMD capable of mediating both excitation of the LGMD and inhibition of the presynaptic cells using a single neurotransmitter.

LGMD responses also show pronounced inhibition as the animal prepares to jump or move its head rapidly [50]. This saccadic suppression is generated centrally, but the site within the LGMD input pathway at which this inhibition acts has not been identified experimentally. Inhibition of this form, acting at the P cells in our model, would limit activity in the network during rotation but would be ineffective during involuntary motion.

4.2. Robot control structure

It has become customary to relate reactive control systems to the so-called subsumption architecture [4], which was based on earlier work on hierarchical task decomposition by Albus [1]. In this case, control architectures are described as consisting of multiple processing layers, each of which contains input, processing and output components. In addition, specific preference relationships between the processing layers are predefined. It is important to point out, however, that this proposal is neutral towards the content that the designer assigns to these processing layers [44]. For instance, in the original proposal one layer of a hypothetical subsumption architecture contained a hypothetical expert system.

Loosely, one could find hierarchical task decompositions at any level in a neural system, from the signal transduction performed by single synapses to overall relationships between sensory and motor events of the behaving system. An important aspect one loses in this perspective, however, is the extent to which the performance of the system depends on the continuous feed-forward and recurrent integration of activity within and between multiple processing stages. Hence, although we appreciate the use of this perspective in the analysis and engineering of control systems, at present it does not seem to be a concept which pro-

vides additional insight into the invertebrate system we study.

4.3. Comparison with other biologically-based robotics research

Our model was based closely on the anatomy and physiology of the locust LGMD circuit. This close relationship with biology is shared with studies of chemotaxis in ants [48]. A Lego robot equipped with two gas sensors and a simple on-board computer reproduced the trail-following behaviour observed in ants using a small network of neurons. However, unlike our LGMD model, the neural circuit implemented was not based on identified neurons. Instead, data from a number of insect species was combined and no specific predictions could be made about the exact neural circuit used by the ant.

In other biologically-based robotics studies, more abstract models of neural processing have been implemented on robots. The polarized light compass used by the desert ant *Cataglyphis* has been studied using Sahabot [17] and Sahabot 2 [21]. Input to the compass was obtained using three polarised light sensors modelled after the POL neurons found in the optic lobe of the cricket. Subsequent processing, however, was performed mathematically as the neural circuitry was unknown. The later Sahabot 2 combined the polarisation compass with algorithms for path integration and landmark learning. Abstract models were also used to test strategies for phonotaxis in crickets [47,48] and navigation in bees [36,49]. In both cases, the neural circuitry controlling these behaviours in the animal had not been identified.

Fly visual motion detection has been successfully applied to robot navigation (for a review, see [8]). Based on the Reichardt correlation-type elementary movement detector [26], a model of the fly's eye was constructed using an analogue design technique similar to neuromorphic engineering [6]. The optic flow detected by this "eye" was used to steer a robot through a cluttered environment successfully. This analogue modelling approach reproduces neural processing more faithfully than computer simulations, but makes it more difficult (if not impossible) to observe the internal states of a model. On the contrary, our simulation system, IQR421, allows us to see the

states of all the cells in our model in real-time during a simulation.

4.4. Application in robotic control tasks

Our results suggest that the implementation of the basic LGMD circuitry using a VLSI, or neuromorphic, technology is feasible [46]. By adding the circuitry of this model into a silicon retina [20], a smart sensor for visually based obstacle detection could be developed. Neuromorphic sensors based on models of insect motion detection have already been designed [11,19] and applied to robotic control tasks [13,15]. Development of a smart sensor based on the LGMD would provide a robotics engineer with a compact, safe, low-power solution for obstacle detection using vision, eliminating the need to use traditional image processing methods or active sensing (e.g. laser scanners).

4.5. Conclusions

In our investigations we attempt to develop biologically realistic models of sensory processing which can be applied to mobile platforms. The LGMD circuit investigated in the present study is a step in this direction. This approach illustrates the synergistic relationship between basic neuroscientific research, modelling studies using digital simulation and real-world devices, and neuromorphic engineering.

We have demonstrated that our model of the LGMD system of the locust detected approaching obstacles robustly in a complex visual environment, but that the responses were less predictable than those obtained using simple stimuli. We will continue to investigate the responses of both the neuron and our robot-based model to complex stimuli in order to determine the behavioural properties of the LGMD and assess possible applications in robot control systems. Although the precision of visual systems such as the LGMD is limited when compared with traditional robotic range sensors, their widespread use in diverse animal species suggests that such systems can become valuable tools for robotics engineers.

Acknowledgements

The authors wish to thank the anonymous referees for their constructive comments.

Appendix A

A.1. Cell types

Three types of cell were used: linear threshold, integrate-and-fire and random spike.

A.1.1. Linear threshold

This cell type was used to model graded potential neurons. The membrane potential of a linear threshold cell i at time $t + 1$, $v_i(t + 1)$, was given by

$$v_i(t + 1) = p_i v_i(t) + g_i^{\text{Exc}} \sum_{j=1}^m w_{ij} a_j(t - \delta_{ij}) - g_i^{\text{Inh}} \sum_{k=1}^n w_{ik} a_k(t - \delta_{ik}), \quad (\text{A.1})$$

where $p_i \in \{0, 1\}$ was the persistence of the membrane potential, g_i^{Exc} and g_i^{Inh} were the gains of the excitatory and inhibitory inputs, respectively, m was the number of excitatory inputs, n was the number of inhibitory inputs, w_{ij} and w_{ik} were the strengths of the synaptic connections between cells i and j and i and k respectively, a_j and a_k were the output activities of cells j and k , respectively, and $\delta_{ij} \geq 0$ and $\delta_{ik} \geq 0$ were the delays along the connections between cells i and j and cells i and k respectively.

The output activity of cell i at time $t + 1$, $a_i(t + 1)$, was given by

$$a_i(t + 1) = \begin{cases} v_i(t + 1) & \text{for } v_i(t + 1) \geq \theta, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A.2})$$

where θ was the membrane potential threshold.

A.1.2. Integrate-and-fire

The dynamics of the membrane potential of the integrate-and-fire cell was identical to that of the linear threshold cell described above (Eq. (A.1)). However, an integrate-and-fire cell produced a discrete spike whenever its membrane potential reached a threshold after which the membrane potential was hyperpolarised. The activity of integrate-and-fire cell i at time $t + 1$, $a_i(t + 1)$, was given by

$$a_i(t + 1) = \begin{cases} \beta & \text{for } v_i(t + 1) \geq \theta, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A.3})$$

where β was the amplitude of the output spikes and θ was the threshold membrane potential.

After producing a spike at time $t + 1$, the membrane potential of cell i was hyperpolarised such that

$$v_i'(t + 1) = v_i(t + 1) - \alpha, \quad (\text{A.4})$$

where $v_i'(t + 1)$ was the membrane potential after hyperpolarisation and α was the amplitude of the hyperpolarisation.

A.1.3. Random spike

This cell type generated random spiking activity with adjustable spiking probability. It received no input and had no membrane potential. The output activity of neuron i at time $t + 1$, $a_i(t + 1)$, was given by

$$a_i(t + 1) = \begin{cases} \beta & \text{with probability } P, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A.5})$$

where β was the amplitude of output spikes and P was the probability that the cell would be active at any given time.

A.2. Connection mapping

Connections were defined around topographically aligned cells, cells in the same relative positions within their groups. A presynaptic cell with position x_{pre} and y_{pre} was topographically aligned with a postsynaptic cell with position x_{post} and y_{post} when

$$\begin{aligned} x_{\text{pre}} &= (x_{\text{post}}/w_{\text{post}}) \times w_{\text{pre}}, \\ y_{\text{pre}} &= (y_{\text{post}}/h_{\text{post}}) \times h_{\text{pre}}, \end{aligned} \quad (\text{A.6})$$

where w_{pre} and w_{post} were the widths of the pre- and postsynaptic groups, respectively, and h_{pre} and h_{post} were the heights of the pre- and postsynaptic groups, respectively. The values of x_{pre} and y_{pre} were rounded to the nearest integer. Fig. A.1, shows an example.

A.3. LGMD model process

Tables A.1 and A.2 show the parameters and group sizes used in the LGMD model, which comprised approximately 1500 cells. The P cell was a special case: its input was the difference between two successive

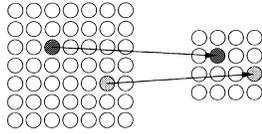


Fig. A.1. Topographic connections between a presynaptic group of 49 cells and a postsynaptic group of 16 cells.

Table A.1

Parameters used in LGMD model for a robot speed of 5 cm/s^a

Name	Type	g^{Exc}	g^{Inh}	Θ	β	p	α
P Cell	I&F	–	–	0.3*	1	0.4	0.5
E Cell	LinTh	0.6	0	0	–	0.1	–
I Cell	LinTh	0.2	0	0	–	0.8	–
S Cell	I&F	1	1	0.5	1	0.4	0.5
F Cell	LinTh	0.2	0	0.15	–	0.1	–
LGMD Cell	I&F	2	5	0.25	1	0.4	0.25

^aValues labelled * were dependent upon the robot speed (see Table A.8). LinTh, linear threshold; I&F, integrate-and-fire; g^{Exc} , excitatory input gain; g^{Inh} , inhibitory input gain; Θ , threshold; β , output spike height; p , membrane potential persistence; α , membrane potential reset (after hyperpolarisation).

Table A.2

Cell groups used in LGMD model

Name	Size	Cell type
P	20 × 20	P Cell
E	20 × 20	E Cell
I	20 × 20	I Cell
S	20 × 20	S Cell
F	1	F Cell
LGMD	1	LGMD Cell

images from the monochrome camera mounted on the robot. For this cell type, Eq. (A.1) became

$$v_i(t+1) = p_i v_i(t) + x_i(t+1) - x_i(t), \quad (\text{A.7})$$

Table A.3

Connection types used in the LGMD model^a

Name	Type	Arborisation	N_{pre}	N_{post}	w	δ
P-to-E	+	1 : 1	400	400	1	0
P-to-I	+	1 : 1	400	400	1	0
E-to-S	+	1 : 1	400	400	1	0
I-to-S	–			(see Fig. A.2)		
S-to-LGMD	+	256 : 1	256	1	0.04	0
P-to-F	+	256 : 1	256	1	0.04	0
F-to-LGMD	–	1 : 1	1	1	1	1

^aTypes: +, excitatory; –, inhibitory; N_{pre} , number of presynaptic cells; N_{post} , number of postsynaptic cells; w , strength of connection between each pair of pre-and postsynaptic cells; δ , delay of connection between each pair of pre-and postsynaptic cells.

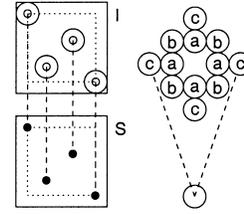


Fig. A.2. Connection between I and S groups in the LGMD model circuit. Each of the 256 central cells in S was connected to a ring of 12 cells in I to produce lateral inhibition, with the remaining S cells not connected. The weights and delays were related to the relative positions of the pre- and postsynaptic cells. The strength, w , and delay, δ , parameters used were: a, $w = 0.4$, $\delta = 1$; b, $w = 0.32$, $\delta = 1$; c, $w = 0.2$, $\delta = 2$.

where $x_i(t+1)$ and $x_i(t)$ were the current and previous states of the pixel which was topographically aligned with cell i . The output activity was calculated using Eq. (A.3) as normal.

There were seven connection types, which produced 4600 synapses, used in the model. The parameters for these connection types are given in Table A.3 and Fig. A.2.

A.4. Robot control process

Tables A.4 and A.5 show the parameters used in the robot control circuit, which comprised 360 cells.

There were eight connection types, which produced over 1400 synapses, used in the model. The parameters for these connection types are given in Tables A.6 and A.7.

Table A.4
Parameters used in the robot control circuit for a robot speed of 5 cm/s^a

Name	Type	g^{Exc}	g^{Inh}	Θ	β	P	p	α
AvCell	LinTh	0.25	0	0.05	0	–	0.7	–
AvoidCell	I&F	1	0	0.05	5	–	0.75	0.05
CollCell	LinTh	1	0	0.54	1	–	0.55	–
ExpCell	LinTh	0.5	0.25	0.25	1	–	0.7	–
LGMDRcvCell	I&F	0.75*	0	1*	1	–	0.9	2.0
MOCcell	LinTh	1	0	0.3	0	–	0.5	–
NoiseCell	RandSpk	0	0	0	1	0.15	0.85	–
ReflexCell	LinTh	0.55	5	0.5	0	–	0.85	–

^a Values labelled * were dependent upon the robot speed (see Table A.8). LinTh, linear threshold; I&F, integrate-and-fire; g^{Exc} , excitatory input gain; g^{Inh} , inhibitory input gain; Θ , threshold; β , output spike height; P , output probability; p , membrane potential persistence; α , membrane potential reset (after hyperpolarisation).

Table A.5
Cell groups used in robot control circuit

Name	Size	Cell type
Aversive	8	AvCell
Avoid	1	AvoidCell
Coll	8	CollCell
Explore	5 × 5	ExpCell
LGMDRcv	1	LGMDRcvCell
MotorOut	10 × 10	MOCcell
Noise	10 × 10	NoiseCell
Reflex	10 × 10	ReflexCell

Table A.6
Connection types used in the robot control circuit^a

Name	Type	Arborisation	N_{pre}	N_{post}	w	δ
Noise-to-Explore	+	$\mathcal{R}(12, 100) : 1$	100	25	$\mathcal{R}(0.05, 0.1)$	0
Explore-to-Reflex	+	25 : 1	25	1 (42)	0.25	0
Reflex-to-MotorOut	+	1 : 1	100	100	1	0
Coll-to-Aversive	+	1 : 1	8	8	1	0
Aversive-to-Explore	–	25 : 1	8	25	1	0
Coll-to-Reflex	+			(see Table A.7)		
LGMDRcv-to-Avoid	+	1:1	1	1	1	0
Avoid-to-MotorOut	+	1 : 1	1	1 (4)	1	0

^aTypes: +, excitatory; –, inhibitory. N_{pre} , number of presynaptic cells; N_{post} , number of postsynaptic cells; numbers in parentheses give specific cell IDs. w , strength of connection between each pair of pre-and postsynaptic cells; δ , delay of connection between each pair of pre-and postsynaptic cells. $\mathcal{R}(x, y)$ denotes a random number between x and y .

Table A.7

Specific connections between the Coll and Reflex cell groups in the robot control circuit (w , strength of connection; δ , delay of connection)

Coll cell ID	Reflex cell ID	w	δ
0	64	2	0
1	64	2	0
2	74	2	0
3	14	2	0
4	24	2	0
5	24	2	0
6	41	2.5	0
7	41	2.5	0

Table A.8

Variation of cell parameters with robot speed

Speed (cm/s)	P Cell θ	LGMDRcvCell g^{Exc}	LGMDRcvCell θ
1.5	0.20	0.75	1.0
2.5	0.25	0.75	1.0
5.0	0.30	0.75	1.0
7.5	0.30	0.75	1.0
10.0	0.42	1.0	0.9
12.5	0.50	1.0	0.9

A.5. Variation of cell parameters

Table A.8 gives the parameter changes required to tune the LGMD circuit for different speeds of motion.

References

- [1] J.S. Albus, A.J. Barbera, M.L. Fitzgerald, R.N. Nagel, G.J. VanderBrug, T.E. Wheatley, A measurement and control model for adaptive robots, in: Proceedings of the Tenth International Symposium on Industrial Robots, SME, Dearborn, MI, 1979.
- [2] J.P. Bacon, K.S. Thompson, M. Stern, Identified octopaminergic neurons provide an arousal mechanism in the locust brain, *Journal of Neurophysiology* 74 (1995) 2739–2743.
- [3] J.M. Blanchard, Collision avoidance: A biologically inspired neural network for the detection of approaching objects, Ph.D. Thesis, Faculty of Medicine, University of Newcastle upon Tyne, UK, 1998.
- [4] R.A. Brooks, A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation* RA 2 (1) (1986) 14–23.
- [5] M. Burrows, C.H.F. Rowell, Connections between descending visual interneurons and metathoracic motoneurons in the locust, *Journal of Comparative Physiology* 85 (1973) 221–234.
- [6] R. Douglas, M. Mahowald, C. Mead, Neuromorphic analogue VLSI, *Annual Review of Neuroscience* 18 (1995) 255–281.
- [7] H.R. Everett, *Sensors for Mobile Robots: Theory and Application*, AK Peters, Wellesley, MA, 1995.
- [8] N. Franceschini, J.M. Pichon, C. Blanes, Bionics of visuo-motor control, in: T. Gomi (Ed.), *Evolutionary Robotics: From Intelligent Robots to Artificial Life*, AAI Books, Ottawa, 1997, pp. 49–67.
- [9] F. Gabbiani, H.G. Krapp, G. Laurent, Computation of object approach by a wide-field motion-sensitive neuron, *Journal of Neuroscience* 19 (3) (1999) 1122–1141.
- [10] J.J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, MA, 1979.
- [11] R.R. Harrison, C. Koch, An analog VLSI model of the fly elementary motion detector, in: M.I. Jordan, M.J. Kearns, S.A. Solla (Eds.), *Advances in Neural Information Processing Systems*, Vol. 10, The MIT Press, Cambridge, MA, 1998.
- [12] N. Hatsopoulos, F. Gabbiani, G. Laurent, Elementary computation of object approach by a wide-field visual neuron, *Science* 270 (1995) 1000–1003.
- [13] G. Indiveri, P. Verschure, Autonomous vehicle guidance using analog VLSI neuromorphic sensors, in: W. Gerstner, A. Germond, M. Hasler, J.-D. Nicoud (Eds.), *Proceedings of the Artificial Neural Networks — ICANN'97*, Lausanne, Switzerland, Springer, Berlin, 1997, pp. 811–816.
- [14] S.J. Judge, F.C. Rind, The locust DCMD, a movement-detecting neurone tightly tuned to collision trajectories, *Journal of Experimental Biology* 200 (1997) 2209–2216.
- [15] J. Kramer, G. Indiveri, Neuromorphic vision sensors and preprocessors in system applications, in: *Proceedings of the Second International Conference on Advanced Focal Plane Arrays and Electronic Cameras (AFPAEC'98)*, 1998.
- [16] H.G. Krapp, F. Gabbiani, C. Koch, G. Laurent, Neuronal multiplication in the locust visual system, in: N. Elsner, R. Wehner (Eds.), *Proceedings of the 26th Göttingen Neurobiology Conference*, Vol. II, Georg Thieme, Stuttgart, 1998, p. 410.
- [17] D. Lambrinos, M. Maris, H. Kobayashi, T. Labhart, R. Pfeifer, R. Wehner, An autonomous agent navigating with a polarized light compass, *Adaptive Behaviour* 6 (1) (1997) 131–161.
- [18] S.B. Laughlin, Matching coding, circuits, cells, and molecules to signals — General principles of retinal design in the fly's eye, *Progress in Retinal and Eye Research* 13 (1994) 165–196.
- [19] S. Liu, Silicon model of motion adaptation in the fly visual system, in: *Proceedings of the Third UCSD–Caltech Joint Symposium on Neural Computation*, 1996, pp. 1–10.
- [20] M.A. Mahowald, C. Mead, The silicon retina, *Scientific American* 264 (1991) 76–82.

- [21] R. Möller, D. Lambrinos, R. Pfeifer, T. Labhart, R. Wehner, Modeling ant navigation with an autonomous agent, in: R. Pfeifer, B. Blumberg, J.-A. Meyer, S.M. Wilson (Eds.), *Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior, From Animals to Animats*, Vol. 5, MIT Press, Cambridge, MA, 1998, pp. 185–194.
- [22] D.C. O’Carroll, N.J. Bidwell, S.B. Laughlin, E.J. Warrant, Insect motion detectors matched to visual ecology, *Nature* 382 (1996) 63–66.
- [23] M. O’Shea, C.H.F. Rowell, J.L.D. Williams, The anatomy of a locust visual interneurone: The descending contralateral movement detector, *Journal of Experimental Biology* 60 (1974) 1–12.
- [24] M. O’Shea, J.L.D. Williams, The anatomy and output connection of a locust visual interneurone: The lobular giant movement detector (LGMD) neurone, *Journal of Comparative Physiology* 91 (1974) 257–266.
- [25] D. Osorio, Mechanisms of early visual processing in the medulla of the locust optic lobe: How self-inhibition, spatial-pooling, and signal rectification contribute to the properties of transient cells, *Visual Neuroscience* 7 (1991) 345–355.
- [26] W. Reichardt, M. Egelhaaf, Properties of individual movement detectors as derived from behavioural experiments on the visual system of the fly, *Biological Cybernetics* 58 (1988) 287–294.
- [27] F.C. Rind, A chemical synapse between two motion detecting neurones in the locust brain, *Journal of Experimental Biology* 110 (1984) 143–167.
- [28] F.C. Rind, Intracellular characterization of neurons in the locust brain signaling impending collision, *Journal of Neurophysiology* 75 (1996) 986–995.
- [29] F.C. Rind, D.I. Bramwell, Neural network based on the input organization of an identified neuron signaling impending collision, *Journal of Neurophysiology* 75 (1996) 967–985.
- [30] F.C. Rind, P.J. Simmons, Seeing what is coming: Building collision sensitive neurons, *Trends in Neurosciences*, 22 (1999) 215–220.
- [31] F.C. Rind, P.J. Simmons, Orthopteran DCMD neuron: A reevaluation of responses to moving objects. I. Selective responses to approaching objects, *Journal of Neurophysiology* 68 (1992) 1654–1666.
- [32] F.C. Rind, P.J. Simmons, Signaling of object approach by the DCMD neuron of the locust, *Journal of Neurophysiology* 77 (1997) 1029–1033.
- [33] F.C. Rind, P.J. Simmons, Local circuit for the computation of object approach by an identified visual neuron in the locust, *Journal of Comparative Neurology* 395 (3) (1998) 405–415.
- [34] R.M. Robertson, A.G. Johnson, Retinal image size triggers obstacle avoidance in flying locusts, *Naturwissenschaften* 80 (1993) 176–178.
- [35] C.H.F. Rowell, M. O’Shea, J.L. Williams, The neuronal basis of a sensory analyser, the acridid movement detector system. IV. The preference for small field stimuli, *Journal of Experimental Biology* 68 (1977) 157–185.
- [36] J. Santos-Victor, G. Sandini, F. Curotto, S. Garibaldi, Divergent stereo in autonomous navigation: From bees to robots, *International Journal of Computer Vision* 14 (1995) 159–177.
- [37] M. Sarikaya, W. Wang, H. Ögmen, Neural network model of on-off units in the fly visual system: Simulations of dynamic behavior, *Biological Cybernetics* 78 (1998) 399–412.
- [38] G.R. Schlotterer, Response of the locust descending contralateral movement detector neuron to rapidly approaching and withdrawing visual stimuli, *Canadian Journal of Zoology* 55 (1977) 1372–1376.
- [39] P. Simmons, Connexions between a movement-detecting visual interneurone and flight motoneurons of a locust, *Journal of Experimental Biology* 86 (1980) 87–97.
- [40] P.J. Simmons, F.C. Rind, Orthopteran DCMD neuron: A reevaluation of responses to moving objects. II. Critical cues for detecting approaching objects, *Journal of Neurophysiology* 68 (1992) 1667–1682.
- [41] M. Stern, K.S.J. Thompson, P. Zhou, D.G. Watson, J.M. Midgley, M. Gewecke, J.P. Bacon, Octopaminergic neurons in the locust brain: Morphological biochemical and electrophysiological characterization of potential modulators of the visual system, *Journal of Comparative Physiology A* 177 (1995) 611–625.
- [42] H. Sun, B.J. Frost, Computation of different optical variables of looming objects in pigeon nucleus rotundus neurons, *Nature Neuroscience* 1 (4) (1998) 296–303.
- [43] P.F.M.J. Verschure, Xmorph: A software tool for the synthesis and analysis of neural systems, Technical Report, Institute of Neuroinformatics, ETH-University Zurich, June 1997.
- [44] P.F.M.J. Verschure, Synthetic epistemology: The acquisition, retention, and expression of knowledge in natural and synthetic systems, in: *Proceedings of the IEEE World Conference on Computational Intelligence, WCCI’98*, Anchorage, 1998, pp. 147–153.
- [45] P.F.M.J. Verschure, T. Voegtlin, A bottom up approach towards the acquisition and expression of sequential representations applied to a behaving real-world device: Distributed Adaptive Control III, *Neural Networks* 11 (1998) 1531–1549.
- [46] A. Watson, Neuromorphic engineering: Why can’t a computer be more like a brain, *Science* 277 (1997) 1934–1935.
- [47] B. Webb, Robots, crickets and ants: Models of neural control of chemotaxis and phonotaxis, *Neural Networks* 11 (1998) 1479–1496.
- [48] B. Webb, A cricket robot, *Scientific American* 275 (1996) 94–99.
- [49] K. Weber, S. Venkatesh, M.V. Srinivasan, Insect inspired behaviours for the autonomous control of mobile robots, in: M.V. Srinivasan, S. Venkatesh (Eds.), *From Living Eyes to Seeing Machines*, Vol. 11, Oxford University Press, Oxford, UK, 1997, pp. 226–248.
- [50] M. Zaretsky, Quantitative measurements of centrally and retinally generated saccadic suppression in a locust movement detector neurone, *Journal of Physiology* 328 (1982) 521–533.



Mark Blanchard is a postdoctoral researcher at the Institute of Neuroinformatics, University-ETH Zurich, Switzerland. He received both his first degree (Microelectronics and Software Engineering, 1993) and his doctorate (Neurobiology, 1998) at the University of Newcastle upon Tyne, UK. His current research uses mobile robots to study aspects of insect vision which cannot be studied by traditional physiological and simulation methods.



Claire Rind received her B.Sc. in Animal Physiology at Canterbury University, New Zealand, and both an MA and Ph.D. at Cambridge University, UK. She is currently a Royal Society Research Fellow at the University of Newcastle upon Tyne. Her research interests are the neural mechanisms used by insects to detect motion and the application of these mechanisms into artificial systems.



Paul F.M.J. Verschure, born 1962, is a group leader at the Institute of Neuroinformatics, ETH-University Zurich, Switzerland. He received both his MA and Ph.D. in Psychology. His scientific aim is to find a grounding for psychological explanations in neuroscience through the use of synthetic methods. He has pursued his research working at different institutes in the US (Neurosciences Institute and The Salk Institute, both in San Diego) and Europe (University of Amsterdam and University of Zurich). He works on biologically realistic models of perception, learning, and problem solving which are applied to robots and on the software tools which support this research in synthetic epistemology, IQR421. His work has made significant contributions to the developing field of, so-called, new artificial intelligence. In addition he applies these models in the domain of art and technology.