

Obstacle avoidance with LGMD neuron: towards a neuromorphic UAV implementation.

Llewyn Salt

School of Information Technology
and Electrical Engineering
The University of Queensland, Australia
Email: llewyn.salt@uqconnect.edu.au

Giacomo Indiveri

Institute of Neuroinformatics
University and ETH Zurich
Zurich, Switzerland
Email: giacomo@ini.uzh.ch

Yulia Sandamirskaya

Institute of Neuroinformatics
University and ETH Zurich
Zurich, Switzerland
Email: ysandamirskaya@ini.uzh.ch

Abstract—We present a neuromorphic adaptation of a spiking neural network model of the locust Lobula Giant Movement Detector (LGMD), which detects objects increasing in size in the field of vision (looming) and can be used to facilitate obstacle avoidance in robotic applications. Our model is constrained by the parameters of a mixed signal analog-digital neuromorphic device, developed by our group, and is driven by the output of a neuromorphic vision sensor. We demonstrate the performance of the model and how it may be used for obstacle avoidance on an unmanned aerial vehicle (UAV).

I. INTRODUCTION

Biological systems can be a source of inspiration when solving tasks in which robots have to respond to sensory events quickly, reliably, and efficiently in real-world environments. One such biological system is the locust Lobula Giant Movement Detector (LGMD) present in the locust nervous system. The LGMD allows the locust to escape rapidly approaching predators by responding to *looming* stimuli, ignoring other types of visual patterns. A looming stimulus is one that expands over the field of view. When activated, the LGMD neuron stimulates the descending contra-lateral movement detector neurons (DCMDs). These in turn stimulate leg and wing motor neurons that trigger an escape response.

Simplified models of the LGMD neuron have been used previously in robotic applications for obstacle avoidance tasks. Santer et al. demonstrated the ability of the LGMD model to detect simple looming stimuli [1]. More recent models demonstrated robust collision avoidance by combining the LGMD and DCMD neuron outputs [2]. In [2] the authors demonstrated the ability of a robotic vehicle to avoid rolling balls at different speeds. These LGMD models were based on the simulation of spiking neural networks (SNNs) driven by standard frame-based cameras [1]–[3]. However, frame-based cameras are not ideal inputs for SNNs. To take full advantage of the spiking nature of these networks, a spike-based vision sensor can be used.

Neuromorphic vision sensors, such as the Dynamic Vision Sensor (DVS) [4], provide spike-like (event-based) responses. A DVS communicates positive or negative luminance changes for pixels as soon as they occur. It detects events with microsecond accuracy and does not suffer from blurring artefacts with fast changes in the visual scene. It is therefore an excellent candidate for use with neuronal controllers in high-speed

applications, e.g. for obstacle avoidance on a UAV¹.

In this paper, we propose an obstacle-avoidance controller for a UAV using the output of a DVS and an LGMD-DCMD model. For such a high-speed application, standard computer simulations of the SNN-based models of the LGMD neuron are too time consuming and would require off-board processing to enable real-time functionality. For this reason, we target a model that can be readily implemented in custom hardware neuromorphic processors [5]. These neuromorphic processors represent a shift away from the von Neumann architecture to a parallel event-based architecture in which memory and computation are co-localized [5]. Given their massively parallel processing abilities, they also allow a real-time, low-power implementation of the looming detector model we propose, suitable for real-time control of a UAV.

We present a behavioural simulation of the LGMD and the DCMD neurons that has been adapted to match the features and restrictions of the neuromorphic hardware. Specifically, we matched the simulation equations and parameters to the circuits present on the mixed signal analog/digital neuromorphic device presented in [6]. We demonstrate the functionality of our simulated LGMD model using visual stimuli in a realistic scenario, collected from a DVS mounted on a moving quadrotor UAV.

II. HARDWARE

A. Dynamic Vision Sensor (DVS)

Silicon retinas are an attempt to break away from the frame-based nature of standard cameras. The biological neural systems don't compute with frames but rather each receptor in the retina sends a spike to the visual cortex when its activity level exceeds a threshold. This preprocessing on the sensor level prevents excess or redundant information to be communicated to the cortex. The address event representation (AER) protocol of the DVS is the realisation of this brain-like spike-based communication. Each pixel of the DVS computes a normalised time derivative of light intensity [4]. If a change in intensity is detected, the pixel communicates its coordinate and the direction in which the intensity has changed. The DVS has a wide operating range in terms of light intensity, low output data rate, and microsecond latency [7]. Fig. 1b shows an example output of the DVS when presented with a circular looming stimulus.

¹Unmanned Aerial Vehicle

B. QUAV




Fig. 1: QUAV and example DVS output

A quadrotor UAV (QUAV) platform was built and used to collect DVS data (Fig. 1a). The QUAV was custom made using off the shelf parts. It was designed to be fairly compact so that it could be used both in and outdoors. The drive-train consisted of 20A electronic speed controllers (ESCs) and Cobra2206-2100kv motors so that the QUAV could use 3S or 4S lithium polymer batteries. With 6045 propellers and a 4S battery, the motors can provide 1.1kgs of thrust each at 100% power output.

III. BUILDING AN LGMD MODEL

We constructed the LGMD model based on descriptions of the locust's neuron, available in [1]. The LGMD model consists of a photoreceptor (P), a summing (S), and an LGMD neuron layer. The neurons in these layers are typically modelled as integrate and fire neurons: they sum up the inputs and spike if the membrane potential exceeds a threshold [1]. These three layers are connected by intermediate excitatory (E), inhibitory (I), and feed-forward (F) connections, which are modelled as linear threshold inter-neurons.

The feed-forward neurons (F) are intended to inhibit translational motion. The inhibitory connections (I) from the photoreceptor to the summing layer inhibit non-looming stimuli. The weights of the inhibitory connections are assigned based on their distance from the central excitatory neuron. This connection configuration spans the P layer like a kernel. This model has been used previously to study collision avoidance on robotic vehicles [1], [2].

The novel aspect of this work lies in the constraints of the neuromorphic processor computing platforms that we have available in our group. These constraints are:

- The neuromorphic computing platform currently available has a total of at most 9k neurons.
- Each neuron has a fan-in of only 64 non-zero (programmable) synapses.
- Each neuron has a fan-out of 4000 units, but subdivided into four distinct clusters.
- Each synaptic weight can assume one of four possible analog values (two positive and two negative ones).




Fig. 2: (a) The neuromorphic LGMD model. Black lines: excitatory connections; red lines: inhibitory connections; dashed lines: slower inhibitions; solid lines: faster inhibitions. (b) The network topology for converting detected looms (LGMD neurons) into steering directions (DCMD neurons). Solid lines: excitatory connections; dashed lines: inhibitory connections.

- Excitatory and inhibitory synapse low pass filter circuits have shared parameters (same dynamics).
- There are two sets of parameters for choosing the type of synapse dynamics (fast or slow).

A. Neuromorphic LGMD Model

The LGMD model introduced in the literature is not suitable for neuromorphic devices that have a limit on how many synapses can be connected to a single neuron. Intermediate layers between the LGMD neuron and the S and P layers needed to be added to make the network compatible with a neuromorphic processor (see Fig. 2a). Each 5×5 neurons in the S layer in our model connect to each neuron in the intermediate S (IS) layer which consists of 6×6 neurons. Similarly, each 8×8 neurons in the P layer connect to each neuron in the intermediate P (IP) layer which consists of 4×4 neurons. These layers are connected to the LGMD neuron with an excitatory connection from the IS layer and an inhibitory connection from the IP layer.

B. Adding Direction to the LGMD Model

The LGMD model can only signal whether or not an object in the field of view is looming. This doesn't provide any information about the direction that the stimulus is coming from, which is required for effective obstacle avoidance.

We split the field of view (FoV) into halves twice, across the centre of the FoV both horizontally and vertically (top/bottom and left/right) and implement an LGMD neuron for each section: top, bottom, left, and right. This combination allows the use of the same network parameters as the whole FoV implementation. Critically, this set-up doesn't add any additional neurons, which is important to adhere to the neuromorphic processor constraints.

C. Translating the LGMD Output to Motor Response

To translate the directional LGMD model outputs into a decision about an obstacle avoidance manoeuvre, we add four

DCMD neurons: up, down, left, and right. These neurons are linked by excitatory synapses to their counterpart LGMD neurons: up to bottom, down to up, left to right, and right to left. They are also linked by inhibitory synapses to the LGMD neuron that would evoke the opposite response: up to top, down to bottom, left to left, and right to right. Fig. 2b shows the excitatory and inhibitory connections between the LGMD and DCMD neurons. This connectivity makes the controller similar to the steering wheel model presented in [2]. The steering commands also describe the amount the QUAUV needs to turn to avoid the looming stimulus.

D. Accounting for the Size of Neuromorphic Hardware

To account for limitations of the neuromorphic hardware, the P layer was not modelled as the DVS128 pixels directly, but the 128×128 pixels were fed into a 32×32 P layer. This reduces the number of neurons to the size of a single neuromorphic processor chip. Down-sampling the P layer acts as a low-pass filter on the DVS events and improves the performance of the model. Furthermore, the inhibitory connections from the P to the S layer were made using a ring-shaped kernel with two levels of weight values as in [2]. The S layer was reduced to 30×30 . The excitatory connections were made by connecting the centre of the kernel in the P layer to the neuron directly behind it in the S layer.

IV. EXPERIMENTS AND DATA COLLECTION

The test data was collected on the QUAUV to validate its performance in a controlled, but realistic environment. For safety during data collection, the QUAUV was fixed by a string on each arm. These four strings could be moved independently to restrict X,Y, and Z motion. To collect the simple stimuli the QUAUV was set to hover in front of a computer generated scene which consisted of a black shape on a white background translating and moving at various speeds. For complicated stimuli the QUAUV would either hover or fly towards a ball, a hand, or a cup. A damping plate was used to reduce vibrations generated by the QUAUV.

The stimuli were recorded using the jaERViewer. Fig. 1b shows the output of the DVS128 when recording the looming circle stimulus. This displays events that were accumulated over a period of $20 \mu s$. The salt and pepper noise poses a challenge to the looming detector, especially as it tends to increase for a moving camera. However, the low-pass filtering capabilities of the modified P layer helped with this noise. jaERViewer was used to label the event streams as either looming or non-looming. These labeled event streams were used to evaluate the model.

V. RESULTS

This section demonstrates the performance of the modified LGMD model, which was tested on various translating and looming stimuli of different speeds, shapes, and sizes.

A. Analysis of LGMD Layers

Fig. 3 shows raster plots of the neuromorphic LGMD model. The DVS output, the P, S, IS, IP, and LGMD layers of the model are shown. The membrane potential of the LGMD





Fig. 3: Raster plots of the LGMD model layers. Last row: activation of the LGMD neuron.

neuron over time is also shown for two looming motions of a dark circle against light background.

Fig. 3 shows that the input is noisy and that the P layer reduces the noise of the input. This is due to it acting as a low-pass filter in addition to its down sampling functionality. It is evident in the S layer that the receding component of the input from the P layer is more inhibited than the looming component. This is key to the success of the looming network. The additional intermediate layers (IP, IS) act as additional low-pass filters. The IP layer reduces the number of spikes that would inhibit the LGMD neuron when compared to the LGMD model in [1]. The translational component in the IP layer is similar to that in the S layer, which is important because this is how the feed-forward inhibition inhibits translational motion. The LGMD neuron output spikes during looming and not during non-looming stimulus.

B. Analysis of Performance for Different Stimuli

We have tested the model with stimuli of differing complexity, two of which are shown in Fig. 4. We used the following stimuli – comp: composite translating and looming stimulus shown in Fig. 4a; cSlow/cFast: a slow/fast looming circle; sSlow/sFast: a slow/fast looming square; ball: a rolling ball; cup: a QUAUV flying towards a suspended cup shown in Fig. 4b; and hand: a looming hand in front of a hovering QUAUV.



(a) Black circle on a white background translating and looming at increasing speeds.

(b) QUAUV flying towards a cup suspended in front of it with a white background.

Fig. 4: Activity (raster plots) of the P layer of the LGMD network over time for two exemplary inputs. Blue are the looming times.

We used *accuracy* to determine the ability of the network to detect looming and suppress translational and receding stimuli. The LGMD network was defined to have detected a looming stimulus if the output neuron's spike rate exceeded a fixed value (11 or more spikes in 10ms). We used differential evolution to find the parameters of the LGMD network that maximised accuracy [8]. Table I shows the performance

Stimulus	Model	Accuracy	Sensitivity	Precision	Specificity
comp	LGMD	0.80	0.80	0.80	0.80
	NLGMD	0.90	1.00	0.83	0.80
cSlow	LGMD	1.00	1.00	1.00	1.00
	NLGMD	0.80	0.60	1.00	1.00
cFast	LGMD	0.44	0.00	0.00	1.00
	NLGMD	1.00	1.00	1.00	1.00
sSlow	LGMD	1.00	1.00	1.00	1.00
	NLGMD	1.00	1.00	1.00	1.00
sFast	LGMD	0.50	0.00	0.00	1.00
	NLGMD	1.00	1.00	1.00	1.00
ball	LGMD	0.33	0.33	1.00	0.00
	NLGMD	0.66	0.66	1.00	0.00
cup	LGMD	0.70	1.00	0.62	0.40
	NLGMD	0.70	1.00	0.62	0.40
hand	LGMD	0.50	1.00	0.50	0.00
	NLGMD	0.50	1.00	0.50	0.00

TABLE I: Performance metrics of the performance of the LGMD models for different looming stimuli

metrics of our modified LGMD (NLGMD) model and the original LGMD model [1] which used the same P layer size for fairness of comparison. Our results show that the modified LGMD model is capable of detecting looming stimuli while filtering out translational movement as well as or better than the original model on all but the slow circular looming stimulus (the added IS and IP layers excessively filter the already weak DVS output from a slow stimulus reducing sensitivity). The difference in the model's performance in all cases is the sensitivity. The specificity remains the same. Both models had the same ability to identify non-looming stimuli but different abilities to identify looming stimuli.

Fitting the model to neuromorphic hardware only decreased the accuracy of the model on particularly sparse inputs, but increased its performance on stimuli of varying shapes or at high speed.

C. Using the LGMD and DCMD Neurons to Control the QUAV

To enable reactive directional control, the model will provide continuous input to the motors depending on the number of spikes output by the DCMD motor neurons. Fig. 5 shows activity of the directional LGMD and DCMD neurons when an object is looming on the bottom left side of the FoV. The bottom and left LGMD neurons spike more than their counterparts. These LGMD neurons excite the respective DCMD neurons, while the inhibitions mostly suppress the other responses. Based on the activity of the DCMD neurons, the QUAV would correctly steer to go right and slightly up.

VI. CONCLUSIONS

We presented a modified implementation of the locust LGMD model for use on a neuromorphic processor. The network detects looming stimuli and can be used in robotic controllers to facilitate obstacle avoidance. Our implementation




Fig. 5: Exemplary output of the four LGMD neurons and the respective DCMD neurons for a looming stimulus (green lines are ground-truth looms).

uses DVS as input to the model and is tailored to implementation in a neuromorphic device. We have tested the model on input, recorded with a DVS mounted on a QUAV. The performance of the model is superior to the original model and can be further improved by using another bio-inspired vision system – the fly elementary motion detector (EMD) neuron, which detects the direction in which stimuli are propagating across the field of view. The EMD neurons can be combined with the LGMD to further suppress the translational motion in order to improve detection of the looming objects.

ACKNOWLEDGMENT

Supported by EU H2020-MSCA-IF-2015 grant 707373 ECogNet and EU ERC-2010-StG 20091028 grant 257219 NeuroP. We would also like to thank Erin Oliver for her pro bono editing assistance.

REFERENCES

- [1] R. D. Santer, R. Stafford, and F. C. Rind, "Retinally-generated saccadic suppression of a locust looming-detector neuron: investigations using a robot locust," *Journal of The Royal Society Interface*, vol. 1, no. 1, pp. 61–77, 2004.
- [2] S. Yue, R. D. Santer, Y. Yamawaki, and F. C. Rind, "Reactive direction control for a mobile robot: a locust-like control of escape direction emerges when a bilateral pair of model locust visual neurons are integrated," *Autonomous Robots*, vol. 28, no. 2, pp. 151–167, 2010.
- [3] R. Stafford, R. D. Santer, and F. C. Rind, "A bio-inspired visual collision detection mechanism for cars: combining insect inspired neurons to create a robust system," *BioSystems*, vol. 87, no. 2, pp. 164–171, 2007.
- [4] S. C. Liu and T. Delbruck, "Neuromorphic sensory systems," *Current Opinion in Neurobiology*, vol. 20, no. 3, pp. 288–295, 2010.
- [5] E. Chicca, F. Stefanini, C. Bartolozzi, and G. Indiveri, "Neuromorphic electronic circuits for building autonomous cognitive systems," *Proceedings of the IEEE*, vol. 102, no. 9, pp. 1367–1388, 2014.
- [6] G. Indiveri, F. Corradi, and N. Qiao, "Neuromorphic architectures for spiking deep neural networks," in *2015 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 2015, pp. 4–2.
- [7] T. Serrano-Gotarredona and B. Linares-Barranco, "A 128×128 1.5% contrast sensitivity 0.9% FPN 3 μs latency 4 mW asynchronous frame-free dynamic vision sensor using transimpedance preamplifiers," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 3, pp. 827–838, 2013.
- [8] L. Salt, "Optimising a Neuromorphic Locust Looming Detector for UAV Obstacle Avoidance," Master's thesis, The University of Queensland, 2016.