Towards a Computational Approach for Collision Avoidance with Real-World Scenes

Matthias S. Keil[a], and Ángel Rodríguez-Vázquez[a]

[a]Instituto de Microelectrónica de Sevilla, Centro Nacional de Microelectrónica (CSIC), Avenida Reina Mercedes (edificio CICA), E-41012 Sevilla, Spain

ABSTRACT

In the central nervous systems of animals like pigeons and locusts, neurons were identified which signal objects approaching the animal on a direct collision course. In order to timely initiate escape behavior, these neurons must recognize a possible approach (or at least differentiate it from similar but non-threatening situations), and estimate the time-to-collision (ttc). Unraveling the neural circuitry for collision avoidance, and identifying the underlying computational principles, should thus be promising for building vision-based neuromorphic architectures, which in the near future could find applications in cars or planes. Unfortunately, a corresponding computational architecture which is able to handle real-situations (e.g. moving backgrounds, different lighting conditions) is still not available (successful collision avoidance of a robot was demonstrated only for a closed environment). Here we present two computational models for signalling impending collision. These models are parsimonious since they possess only the minimum number of computational units which are essential to reproduce corresponding biological data. Our models show robust performance in adverse situations, such as with approaching low-contrast objects, or with highly textured backgrounds. Furthermore, a condition is proposed under which the responses of our models match the so-called η-function. We finally discuss which components need to be added to our model to convert it into a full-fledged real-world-environment collision detector.

Keywords: collision avoidance, looming signalling, neurodynamics, computational model, locust, LGMD

1. INTRODUCTION

It is essential to many animal species to recognize approaching predators early enough. Considering monocular viewing, the information about an object which is available on an animal's retina is the object's subtended visual angle Θ, and its corresponding rate of expansion $\dot{\Theta}$ when the object is moving in depth. Looming-sensitive neurons could in principle use these variables to compute an imminent collision. In the locust, it has been found that the *lobula giant movement detector* (LGMD) responds vigorously to approaching objects (as opposed to receding ones).[1] The LGMD is a huge single neuron which is located in the third visual neuropile of the optic lobe (or lobula), and forms together with the *descending contralateral movement detector* (DCMD) a neural circuit which is involved in triggering escape behavior. However, the precise computation which is carried out by the LGMD is matter of ongoing debate, since there exist two contrasting views about the respective roles of *feedforward inhibition* on the one hand, and *lateral inhibition* on the other, in shaping LGMD responses.[2–4] *(i)* Feedforward inhibition collects large-field information from the retina and synapses onto a single proximal dendrite of the LGMD.[5] Excitatory inputs are provided by velocity signals from local movement detectors into the fan-shaped distal part of the LGMD dendrite. LGMD responses to stimuli which simulate an object approach could be fitted to the so-called η-*function*,[6] which is defined as $\eta(t) = C\,\dot{\Theta}\,\exp(-\alpha\Theta)$, where C and α are constants. The η-function reveals an activity peak as an object approaches at constant velocity: when the object is far away $\dot{\Theta}$ both is small and increases faster than $\exp(-\alpha\Theta)$ decreases. This gives an overall increase of $\eta(t)$. When the object is near, however, the exponential suppression dominates, and $\eta(t)$ decreases. Since the η-function peaks before collision for moderate speeds and object sizes, it was proposed that this peak provides a cue for the timing of escape behavior,[6] where the activity peak occurs earlier for larger objects, with shallower ascending-phase slopes.[7] If so, how is this computation implemented in the locust's nervous system?

Corresponding author M.S.K.
M.S.K. E-mail: matthias@imse.cnm.es
A.R.V. E-mail: angel@imse.cnm.es

Bioengineered and Bioinspired Systems, Angel Rodríguez-Vázquez, Derek Abbott, Ricardo Carmona, Editors, Proceedings of SPIE Vol. 5119 (2003) © 2003 SPIE · 0277-786X/03/$15.00

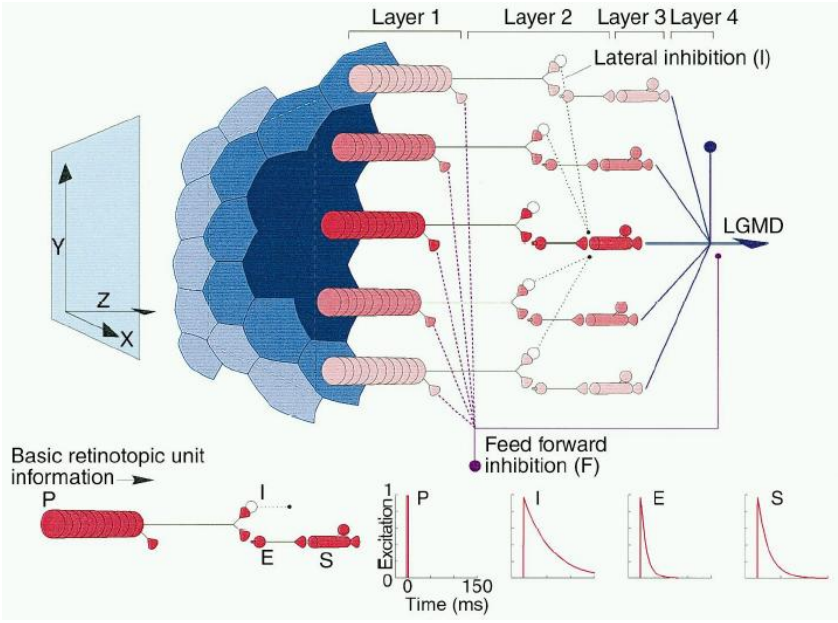

Figure 1. The Rind et al. neural network.[9,11] The network consists of three layers of retinotopically organized units as shown at the bottom of the figure. The plots show the temporal impulse response ("excitation") of I-,E-, and S-units. (Figure from Ref,[2] used with permission).

Recent evidence suggests that the multiplication of the excitatory input $\dot{\Theta}$ with the inhibitory input $\exp(-\alpha\Theta)$ is performed by logarithmically encoding both variables, then adding excitatory and inhibitory input within the LGMD, and subsequently exponentiate the sum by voltage-dependent sodium conductances at the LGMD axon.[8] The implementation of the η-function relies on postsynaptic feedforward inhibition, which represents one interpretation of LGMD responses. In contrast, another interpretation is primarily based on *(ii)* lateral inhibition, which is presynaptic to the LGMD. To compute an LGMD response which is correlated with a looming object, Rind and coworkers state that a *critical race* over the LGMD dendrites takes place between feedforward excitation and laterally spreading inhibition.[9] Lateral inhibition is elicited by feedforward inhibition in a way that only if an object's angular size on the retina grows more rapidly than inhibitory waves could propagate laterally, then excitation wins the race, and the LGMD eventually responds. In contrast, feedforward inhibition is activated only *at the end* of an object approach, and truncates LGMD responses rather abruptly.[9,10] According to this view, feedforward inhibition cannot establish an activity peak in the LGMD response. *Rind et al.* claim that it is lateral inhibition which implements the selectivity for objects moving in depth, and eliminates LGMD responses to any sudden large-scale movement. Feedforward inhibition, on the other hand, is activated with some delay when a large number of photoreceptors are excited within a short time interval, and establishes that LGMD responses are larger for approaching objects than for receding ones. Taken together, *Rind et al.* state that LGMD activity does not peak sufficiently early before collision, and therefore the LGMD responses are unlikely to implement the η-function, as described in the first part of the introduction.
In contrast to both views, we show here that our parsimonious models based on lateral inhibition indeed generate responses which look like an η-function, and examine the conditions why this similarity is observed. However, it turns out that the η-function is not approximated in a mathematical sense.

2. THE RIND ET AL. MODEL

Before we present our architectures, we give an overview over the original *Rind et al.* model[9] (figure 1), which is a variant of a model previously presented by Edwards.[12] The *Rind et al.* model was also implemented on a robotic platform, where successful collision avoidance behavior was demonstrated within a closed environment[11] (2). The model is composed of four retinotopically organized layers. Layer-1 (figure 1) represent the input into

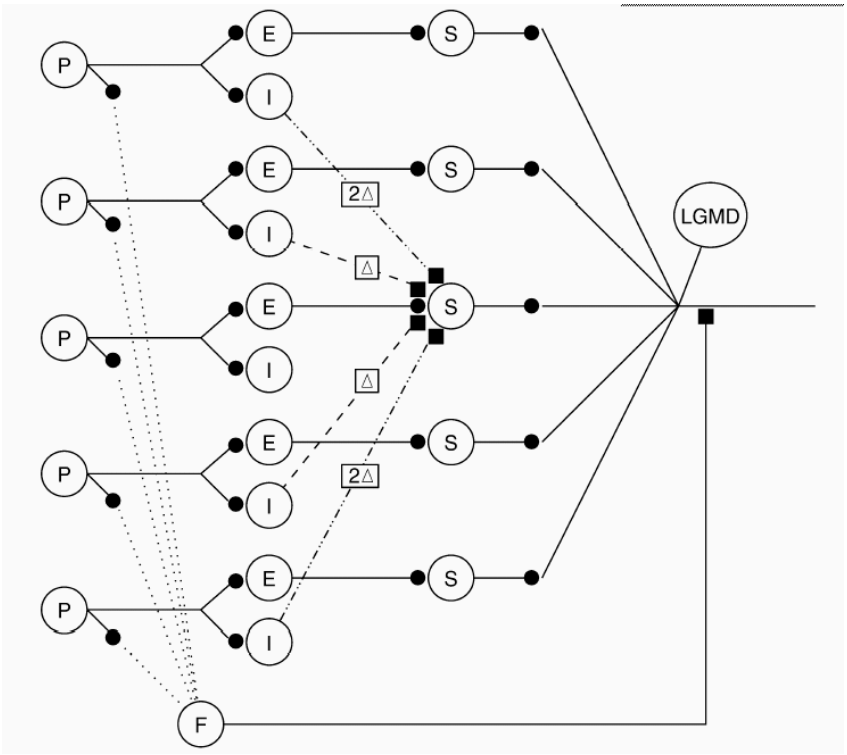

Figure 2. Lateral inhibition within the Rind et al. neural network.[9,11] This sketch visualizes the connectivity pattern and corresponding time delays of the laterally inhibitory network. (Figure from Ref,[11] used with permission).

the network, where **P-units** (figure 2) compute their output by subtracting the current luminance distribution $\mathcal{L}$ (at time t) from a previous luminance distribution (at $t - \Delta t$). Thus, these units will respond in a transient fashion if an object (assuming that $\mathcal{L}$ represents a real-world image at some time) appears on the $\mathcal{L}$-array, where the object must be darker than its respective background to elicit a corresponding OFF-response. An OFF-response is equivalently generated at the disappearance of an object being brighter than its background. The *Rind et al.* model in its present form is therefore "blind" to ON-responses, that is the appearance (disappearance) of objects which are lighter (darker) than their respective background. P-unit activity (corresponding to the thresholded state variable or potential) is simultaneously fed into an excitatory layer (which is made up of **E-units**), and an inhibitory layer (which consists of **I-units**). These layers are lumped together in layer-2 in figure 1. Layer-3 contains **S-units**, which receive direct excitatory input from the **E-units**. Activity from inhibitory units (I-units) spreads laterally with delays Δt and $2\Delta t$ before they inhibit adjacent S-units (indicated by dashed lines in figure 2). A P-unit response at a given time t and position can consequently suppress the activity of adjacent P-units at future times $t + \Delta t$ and $t + 2\Delta t$. This feature is crucial to implement "a race excitation against inhibition",[2] which increases the selectivity of the architecture to increases in the amount of edge, and acceleration of those edges on the $\mathcal{L}$-array.

The excitatory activity of all S-units converges onto a single *lobula giant movement detector* (**LGMD**) neuron in layer-4. It receives feedforward inhibition from one **F-unit**, which by-passes layer-3, and directly inhibits the LGMD neuron whenever the total number of activated photoreceptors exceeds a given threshold and photoreceptors are excited in a short time. The F-unit is important to establish selectivity for approaching over receding objects,[9] and probably to suppress translatory components in the optic flow field as they may result from ego motion.

3. A COMPUTATIONAL APPROACH TO COLLISION AVOIDANCE

Our computational model departs in a number of points from the *Rind et al.* model. First we assume that all neurons produce gradual suprathreshold responses, thus neglecting that S-units, the F-unit, and the LGMD fire action potentials or spikes. E-units integrate P-unit activity in the *Rind et al.* model, what in computational terms delays only the excitatory input of the S-units. These E-units are omitted in our architecture. Furthermore, simulations revealed that feedforward inhibition (mediated by F-units) failed to enrich the computational capabilities of our network in a significant way. Specifically, we found no evidence of how feedforward inhibition could increase the selectivity for approaching over receding objects, since it decreased all responses similarly if, for instance, an activation threshold of 50% excitation was chosen, meaning that feedforward inhibition was activated when half of the photoreceptors were activated within a short time interval. However, we only experimented with static (i.e. non-moving) backgrounds, and feedforward inhibition might turn out important to suppress self-motion induced LGMD activity, which can result from a translating background. Nevertheless, lateral inhibition in our architecture reaches the LGMD nearly instantaneously (i.e. without long delay), what renders feedforward inhibition superfluous for the experiments considered here.
Our models are now formally introduced. Each model is defined by its specific mechanism which is used to implement lateral inhibition:"inhibition by delay" or "inhibition by diffusion".

3.1. ON-OFF-neurons

ON-OFF units are known from the insect lamina of flies (e.g. Ref. 13). The membrane potential of an ON-OFF-cell at position (i, j) at time t is described by the state variable $p_{ij}(t)$, whose dynamics is governed by the following differential equation

$$\frac{dp_{ij}(t)}{dt} = g_{leak}(V_{rest} - p_{ij}) + \mathcal{L}_{ij}(t)(E_{ex} - p_{ij}) + \mathcal{L}_{ij}(t-1)(E_{in} - p_{ij}) \tag{1}$$

$g_{leak} = 50$ is the leakage conductance (or decay constant) which describes the total passive ion flow through the cell membrane. $V_{rest} = 0$ is the resting potential (or leakage reversal potential) which the cell will adopt if it does not receive any input. $E_{ex} = +1$ and $E_{in} = -1$ are the excitatory and inhibitory synaptic batteries, respectively, which confine a cell's dynamic range to $E_{in} \leq p_{ij} \leq E_{ex}$ (as long as $V_{rest} \in]E_{in}, E_{ex}[$). The values have been chosen "symmetrically" around $V_{rest} = 0$ in order to make the ON-OFF-cells equally respond to the contrast polarities "bright object on dark background" and "dark object on bright background". The ON-OFF-units replace the single-contrast-polarity P-units in the original *Rind et al.* model.. The input is provided by luminance values $\mathcal{L}(t)$ and $\mathcal{L}(t-1)$, with $0 \leq \mathcal{L} \leq 1$. In future versions of the model, the direct input by luminance values has to be replaced by an explicit photoreceptor model, in order to include sensitivity changes due to changes in ambient illumination (e.g. Ref. 14), as well as activity changes due to motion adaptation (e.g. Ref. 15). The *Rind et al.* model does not include neither of the just described mechanisms.
The activity or response of an ON-OFF-cell is given by the full-wave rectified and thresholded membrane potential

$$\tilde{p}_{ij} = \gamma \left([p_{ij} - \Theta]^+ + [-p_{ij} - \Theta]^+ \right) \tag{2}$$

with a gain factor $\gamma = 150$, a threshold $\Theta = 0.0005$, and $[\cdot]^+ \equiv \max(\cdot, 0)$.

3.2. Lateral Inhibition by Explicit Delay ("Inhibition by Delay")

This mechanism is according to *Rind et al.*,[11] and defines our first model, called "inhibition by delay". Originally, P-unit activation excites inhibitory I-units (see figure 2). In our model, we feed ON-OFF unit activity $\tilde{p}_{ij}$ into an inhibitory interneurons w_{ij}

$$\frac{dw_{ij}(t)}{dt} = g_{leak}(V_{rest} - w_{ij}) + \tilde{p}_{ij}(E_{ex} - w_{ij}) \tag{3}$$

where $g_{leak} = 10$, $V_{rest} = -0.001$, and $E_{ex} = +1$. Its activity is given by the half-wave rectified membrane potential $\tilde{w}_{ij} \equiv [w_{ij}]^+$. Lateral inhibition is implemented by convolving (symbol '$\otimes$') the response of an inhibitory neuron $\tilde{w_{ij}}$ with a nearest-neighborhood kernel $\mathcal{K}_1$ and a next-nearest-neighborhood kernel $\mathcal{K}_2$ (figure 3)

$$u_{ij}(t) = \tilde{w}_{ij}(t-1) \otimes \mathcal{K}_1 + \tilde{w}_{ij}(t-2) \otimes \mathcal{K}_2 \tag{4}$$

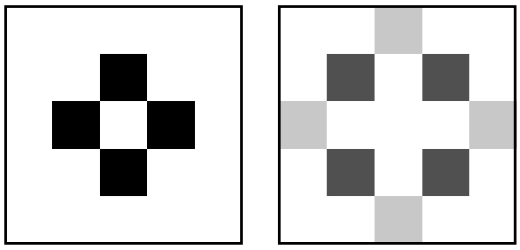

Figure 3. Convolution kernels for nearest neighbors and next-nearest neighbors. The black squares in the left image correspond to the value 0.6065, white corresponds to zero. The dark-gray squares in the right image correspond to the value 0.3679, and the light-gray squares to 0.1353. White corresponds to zero.

where the convolution kernels $\mathcal{K}_1$ and $\mathcal{K}_2$ are Gaussian-weighted (standard deviation one) annulus-like structures centered at a position (i, j).

3.3. Lateral Inhibition by Diffusion ("Inhibition by Diffusion")

Another mechanism to implement lateral inhibition is based on a diffusion layer or *syncytium*. The second model based on this mechanism shall be called "inhibition by diffusion". A diffusion layer provides an alternative biophysical mechanism to equation 4, and has the advantage that only next-neighbor interactions are required to compute its state, which makes a possible VLSI implementation less complicated.

$$\frac{ds_{ij}(t)}{dt} = g_{leak}(V_{rest} - s_{ij}) + \tilde{p}_{ij}(E_{ex} - s_{ij}) + D\,\nabla^2 s_{ij} \tag{5}$$

where $g_{leak} = 10$, $V_{rest} = -0.001$, and $E_{ex} = +1$. ∇^2 implements the Laplacian, and $D = 150$ denotes the diffusion coeffcient. The output of the syncytium is given by $\tilde{s}_{ij} \equiv [s_{ij}]^+$, and therefore g_{leak} controls the spatial extent of activity propagation. A difference between syncytium-mediated inhibition and "inhibition by delay" exists, however, due to the design of the convolution kernels (cf. figure 3). Obviously, if lateral inhibition is triggered by an event at position (i, j), then "inhibition by diffusion" inhibits this position, whereas "inhibition by delay" does not. In computational terms, nevertheless, this difference is negligible.

3.4. Excitatory neurons

Our excitatory neurons represent the S-units of the *Rind et al.* model. Excitatory neurons receive *direct* excitatory input from ON-OFF-cells, whereas in the *Rind et al.* model E-units are interposed between P-units and S-units (see figure 2). Excitatory neurons are defined as

$$\frac{dv_{ij}(t)}{dt} = g_{leak}(V_{rest} - v_{ij}) + \beta_{ex}\tilde{p_{ij}}(t)(E_{ex} - v_{ij}) + \beta_{in}\mathcal{I}_{ij}(t)(E_{in} - v_{ij}) \tag{6}$$

where $g_{leak} = 50$, $V_{rest} = -0.001$, $E_{ex} = +1$, and $E_{in} = -0.33$. $\beta_{ex} = 1$ and β_{in} are gain factors (or synaptic weights). If $\mathcal{I}_{ij}(t) \equiv u_{ij}(t)$ (equation 4), then $\beta_{in} = 100$. If $\mathcal{I}_{ij}(t) \equiv \tilde{s}_{ij}(t)$ (equation 5), then $\beta_{in} = 150$. The output is given by $\tilde{v}_{ij} \equiv [v_{ij}]^+$. The effect of increasing β_{in} is shown in figure 10.

3.5. The LGMD neuron

The LGMD neuron integrates the activity of all excitatory units according to

$$\mathcal{E} = \sum_{i,j}^{n} \tilde{v_{ij}} \tag{7}$$

In contrast to the *Rind et al.* model, our LGMD neuron does not receive feedforward inhibition. The LGMD dynamics obeys

$$\frac{dl(t)}{dt} = g_{leak}(V_{rest} - l) + \gamma_{ex}\mathcal{E}(t)(E_{ex} - l) \tag{8}$$

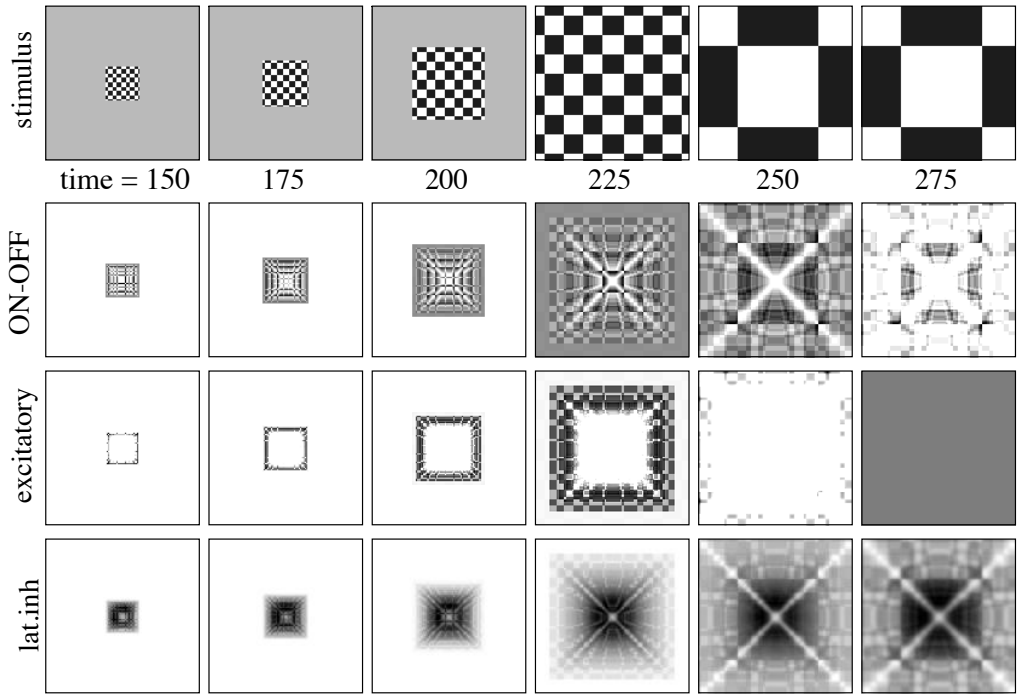

Figure 4. Temporal evolution of different model stages. A typical simulation run for a checkerboard object approaching at 50 km/h or 13.9 m/s (ttc 240 msec). Row "**stimulus**" shows the stimulus at different simulation times given in milliseconds. Row "**ON-OFF** " shows corresponding responses of ON-OFF cells (equation 2), row "**excitatory**" shows responses of the excitatory neuron (equation 6) which undergoes lateral inhibition (row "**lat.inh**") by diffusion (equation 5). In the first row, brighter colors denote higher luminance values. In all other rows, darker colors indicate higher activity. All images were individually normalized such that the full dynamical range is shown.

and the final model output corresponds to the rectified state variable $\tilde{l}(t) \equiv [l(t)]^+$. Parameter values are $g_{leak} = 50$, $V_{rest} = -0.001$, $E_{ex} = +1$, and $E_{in} = -0.33$. $\gamma_{ex} = 5 \cdot 128^2/n^2$ is a gain factors, where n is the number of rows or columns of the neuronal layers (we used $n = 128$ for input images of size 128×128).

4. METHODS

All differential equations were integrated using an explicit forth-order Runge-Kutta scheme with integration step size $\Delta t = 5$ msec. Typical matrix size was 128 $\times$ 128 pixels. Most of the simulations were performed with "inhibition by diffusion" and "inhibition by delay", in order to verify the computational equivalence of both mechanisms. Stimuli consisted of an "object" defined by a foreground luminance patch which was scaled according to physical movement in depth. The object was embedded in a background, which was held static (i.e. non-moving). The default "object" consisted of checkerboard on a uniform grey background. This implies that the mean luminance value of the overall retinal image is constant during symmetrical object movement in depth (background luminance was set equal to mean object luminance). Simulations were performed with four velocities (25 km/h, 50 km/h, 75 km/h and 100 km/h)*. Object diameter was 5 cm for figure 7, 70 cm for figure 8, and 35 cm for all other figures. Under the assumption that our stimulus was only projected onto a part of a real animal's retina, we set the visual angle simulated by our input image to 45 degrees (c.f. 100 degrees in pigeons[2]). The initial distance of the "object" was always set to 5 m, and the time-to-collision (ttc) for the four velocities is 720 msec, 360 msec, 240 msec and 180 msec, respectively. A typical simulation run with the default object/background combination is shown in figure 4.

*Corresponding to 6.9 m/s, 13.9 m/s, 20.8 m/s, 27.8 m/s.

5. RESULTS

5.1. Does Our Computational Model Approximate an η-Function?

The left plot in figure 5 visualizes the η-function[6] for the simulation parameters given in section 4. The response peaks of the η-function relative to ttc occur at -5 msec, -10 msec, -15 msec and -30 msec (i.e. before collision). Figure 6 (left) shows the corresponding LGMD responses for the "inhibition by diffusion"-model (thick curves) and the "inhibition by delay"-model (thin curves). It appears that both of our models approximate the η-function, even though neither an explicit computation of the angular velocity $\dot{\Theta}$ of the retinal image was used, nor feedforward inhibition. How is this possible? Suppose that during an object approach, at some time t_{peak} <ttc, the object cannot stimulate more retinal photoreceptors. Otherwise expressed, the angle subtended by the object $\Theta(t_{peak})$ matches the angle of the maximally excitable retinal subregion: increasing Θ further has no more effect. Since lateral inhibition is triggered by photoreceptor activity in a retinotopic way, it lags behind excitation for $t < t_{peak}$. However, at $t = t_{peak}$, no more excitation due to photoreceptor stimulation is generated, and lateral inhibition catches up with excitation. This causes the decline of LGMD activity although the object grows further until collision.

In the opposite case, when the approaching object is small enough such that $\Theta(t = ttc)$ is smaller than the maximally excitable retinal region, then the LGMD response cannot peak before ttc. Figure 7 shows that the peaks of LGMD responses shift towards to ttc for smaller objects (size 5 cm), whereas response peaks shift away from ttc for larger objects (figure 8, size 70 cm). In all, the interplay between the visual angle of an approaching object subtended on the retina at ttc, together with the maximally excitable region of the retinal photoreceptor array, determines whether or not the LGMD response peaks before ttc.

One further implication of this interplay is that increasing the strength of inhibition decreases LGMD responses, but may shift the peaks away from ttc, what is shown in figure 10. To further investigate if our LGMD responses actually approximate an η-function, we used different functions $\Theta(t)$ (figure 9). This has the effect that also $\dot{\Theta}(t)$ changes, what may cause dramatically changes in the shape of $\eta(t) = C\,\dot{\Theta}\,\exp(-\alpha\Theta)$. In addition to the "physical" approaching law $\Theta(t) = atan(d/2D(t))$ for constant velocity v (d object diameter, $D(t) = D_0 - vt$ remaining distance, D_0 initial distance), we tested our model with a linear function $\Theta(t) \propto t$ (implying $\dot{\Theta}(t)$ =const.), and an exponential function $\Theta(t) \propto \exp(\lambda t)$ (with a constant $\lambda > 0$). If our LGMD responses approximated an η-function, then their shapes should match the theoretical predicted curves in figure 5 (right). Although figure 6 (right) shows that the peaks of LGMD responses shift away from ttc, their shapes are quite different from the theoretical predicions for the linear and the exponential functions. Summarizing so far, one can say that LGMD responses in our model do not approximate an η-function in a mathematical sense, but under certain circumstances the LGMD responses may turn out to be similar to an η-function.

5.2. Receding Objects, Low Contrast Objects, and Textured Backgrounds

Figure 11 shows LGMD responses to receding objects. The model reveals a strong initial response at the begin of object recession, which is quickly suppressed by lateral inhibition. No prolonged LGMD responses are observed as a consequence of the absence of feedforward inhibition, since all retinotopic positions are instantaneously subjected to lateral inhibition. Notice, however, that our LGMD responses for receding objects are bigger in amplitude than the corresponding approaches. In the sense of *Rind et al.*, this means that our circuits show no preference for approaching over receding objects. However, the initial configuration of all simulations with receding objects was such that we started with an object at the retina. If we had done these simulations starting with an approaching object which at ttc starts to recede, then we would have observed that the LGMD responses due to recession are smaller in amplitude, since inhibition would have been present already from the approaching phase. Then, our LGMD responses would have shown this just mentioned preference. Figure 12 shows responses to an object at half contrast. Because no adaptive mechanisms were implemented at the photoreceptor level, LGMD responses drop off accordingly. If, therefore, we detected imminent collision by monitoring if LGMD activity exceeds a threshold, then we would have to chose the threshold value sufficiently low. Nevertheless, this implies that collision would be detected earlier for approaching high-contrast objects. An adaptational mechanism at the photoreceptor level could significantly reduce this problem, such that collision detection does not strongly depend on object contrast. Figure 13 shows a simulation run with a noise patch as object embedded in a noise background, where LGMD responses are also decreased, but object approach is signalled reliably.

6. DISCUSSION AND CONCLUSIONS

We presented two computational models for signalling impending collision, where each of our models relied on a different implementation of lateral inhibition. For the stimulus set which we used, the model based on "inhibition by delay" represents a computationally equivalent approach to a previously suggested one by *Rind et al.*.[9,11] However, our "inhibition by delay" model is parsimonious, since it uses only the crucial parts in order to perform the same computation. Furthermore, it works with a more universal stimulus class, since it signals approaching objects independent of contrast polarity, contrast strength and background. We compared the "inhibition by delay" model with our second model based on "inhibition by diffusion", and demonstrated that both models are computationally equivalent. This makes the "inhibition by diffusion" model interesting for a possible implementation in VLSI, since it relies exclusively on nearest-neighbor interaction, whereas "inhibition by delay" involves in addition next-nearest-neighbor interactions. We furthermore derived a condition under which our models generate LGMD responses which are similar to the η-function. This is generally the case if an approaching object stimulates a subregion of the retinal photoreceptor array which, as the object continues to approach, does not extend further. The stimulated region thus must be smaller than the object at ttc to observe an activity peak in the LGMD response. However, our model does not approximate this function in a mathematical sense, as one can verify by comparing figure 5 (right) with figure 6 (right).
Although our model robustly signals approaching objects embedded in non-moving backgrounds, the performance drops with looming or translating backgrounds, since lateral inhibition due to background movement decreases LGMD responses, and at the same time the sensitivity to approaching objects. Consequently, mechanisms for motion adaptation (e.g. Ref 15) have to be included into the model, such that stationary moving backgrounds are suppressed. Similarly, illumination adaptation, such as lateral inhibition in the retina, or gating of luminance at the photoreceptor level (e.g. Ref 14), should be included to obtain LGMD responses which are largely independent of the currently prevailing illumination conditions.

ACKNOWLEDGMENTS

This research has been partially funded by EU grant IST-2001-38097 (LOCUST). M.K. likes to thank Claire Rind (University of Newcastle) for fruitful discussions and for kindly permitting the use of figure 1 and 2. Also special thanks to E. Calderon for help in preparing the manuscript.

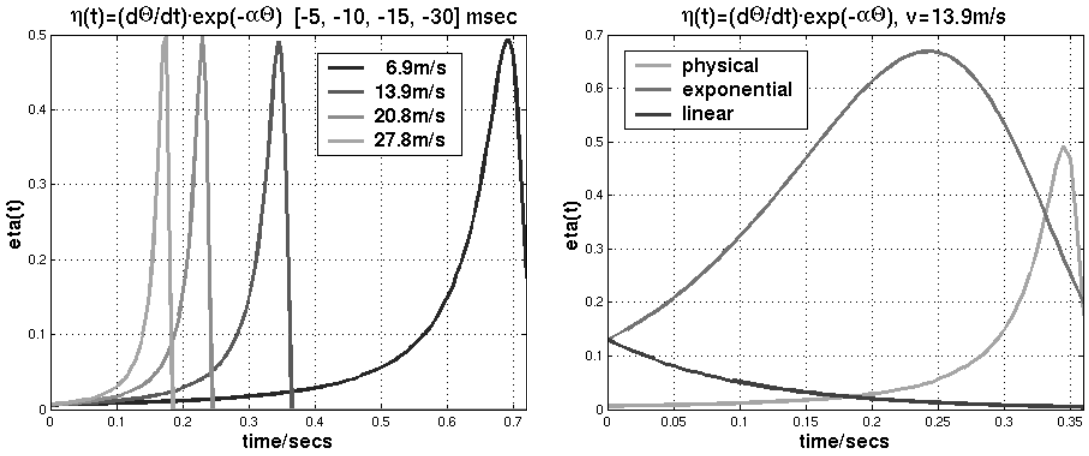

Figure 5. Theoretical predictions. The η-function ($\eta(t) = C\dot{\Theta}\exp(-\alpha\Theta)$)[6] was proposed as a model for LGMD responses. **Left:** η-function as a function of time for different approaching velocities (see section 4). The response maxima occur before ttc (see corresponding numbers in the figure headline). **Right:** η-function for different functions of angular size $\Theta(t)$ (figure 9, see also section 5.1)

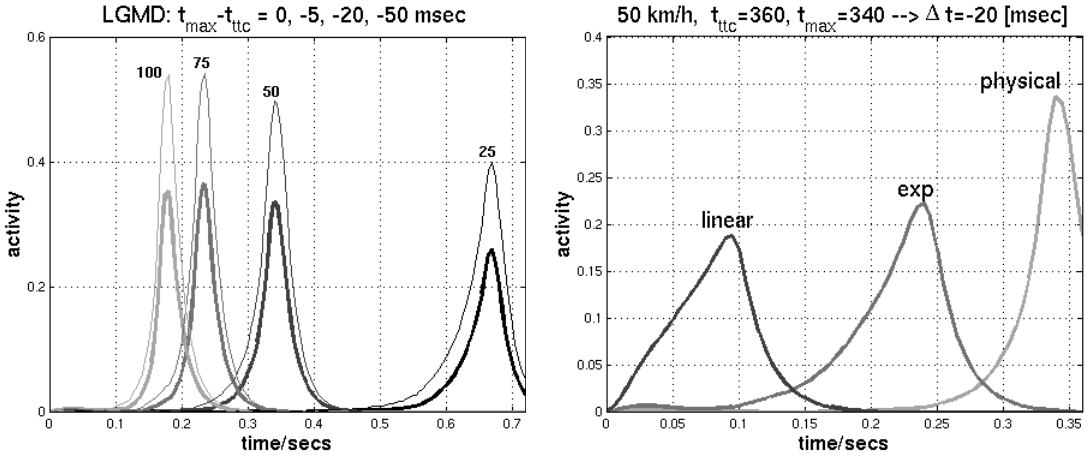

Figure 6. Model predictions. Left: LGMD responses of our model for different approaching velocities of a checkerboard object. These curves should be compared with the theoretically predicted curves of figure 5 (left). The small numbers at the curves denote object velocities in km/h. Thin curves denote "inhibition by delay", and thick curves denote "inhibition by diffusion". Peak responses for both types of inhibition coincide (see figure headline). **Right:** LGMD responses for different functions of angular size $\Theta(t)$ (see section 5.1). If the LGMD approximates an η-function, then these response curves should be similar to the theoretically predicted curves of figure 5 (right).

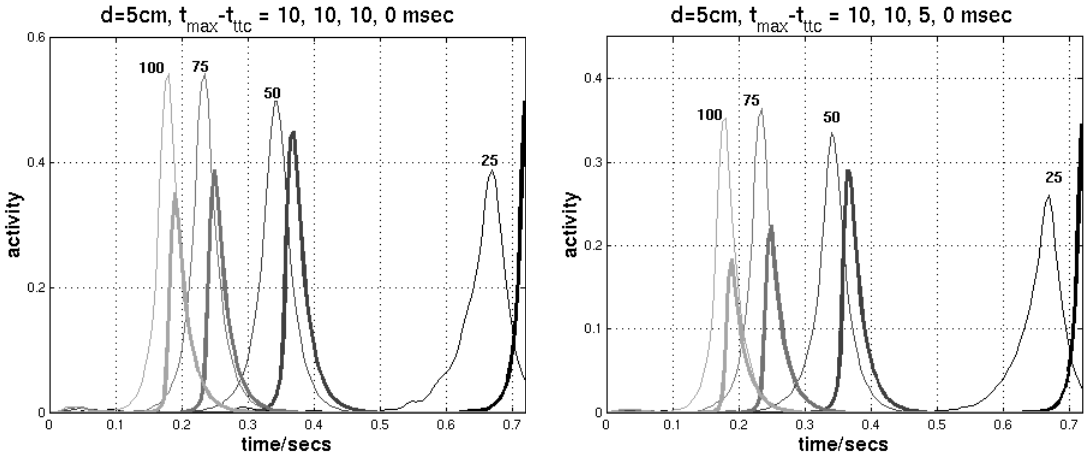

Figure 7. Smaller object size. Smaller approaching objects (diameter 5 cm, thick curves) shift response peaks towards to or even beyond ttc (see numbers in headlines) relative to bigger objects (diameter 35 cm, thin curves). **Left:** "inhibition by delay" **Right:** "inhibition by diffusion".

REFERENCES

1. F. Rind and P. Simmons, "Orthopteran DCMD neuron: a reevaluation of responses to moving objects. I. Selective responses to approaching objects," *Journal of Neurophysiology* **68**, pp. 1654–1666, 1992.
2. F. Rind and P. Simmons, "Seeing what is coming: building collision-sensitive neurones," *Trends in Neuroscience* **22**(5), pp. 215–220, 1999.
3. F. Gabbiani, G. Laurent, N. Hatsopoulos, and H. Krapp, "The many ways of building collision-sensitive neurons," *Trends in Neuroscience* **22**(5), pp. 437–438, 1999.
4. F. Rind and P. Simmons, "Reply," *Trends in Neuroscience* **22**(5), p. 438, 1999.
5. C. Rowell, M. O'Shea, and J. 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**, pp. 157–185, 1977.

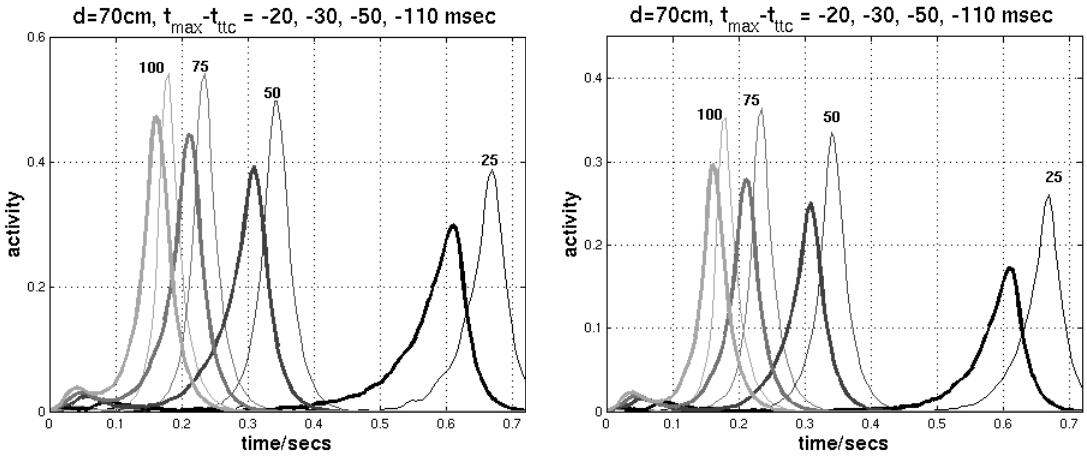

Figure 8. Bigger object size. Bigger approaching objects (diameter 70 cm, thick curves) shift response peaks away from ttc (see numbers in headlines) relative to smaller objects (diameter 35 cm, thin curves). **Left:** "inhibition by delay" **Right:** "inhibition by diffusion".

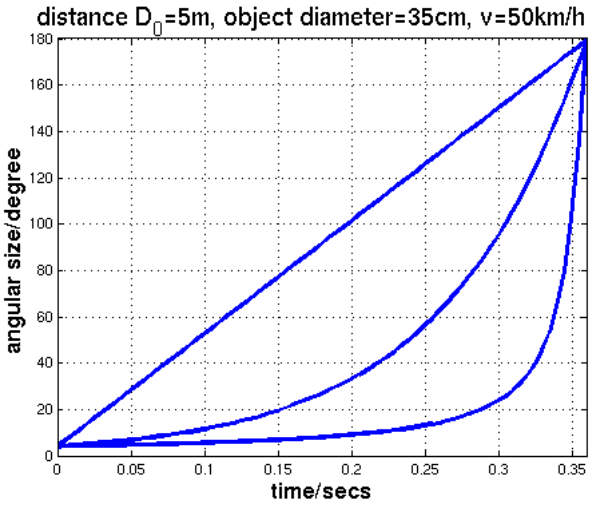

Figure 9. Different functions for the angular size Θ of an object on the retinal array. In addition to the "physical" function $\Theta(t) = atan(d/2D(t))$ (d object diameter, $D(t) = D_0 - vt$ remaining distance, D_0 initial distance, v velocity), a linear function $\Theta(t) \propto t$, and an exponential function $\Theta(t) \propto \exp(\lambda t)$ (with a constant $\lambda > 0$) are shown.

6. N. Hatsopoulos, F. Gabbiani, and G. Laurent, "Elementary computation of object approach by a wide-field visual neuron," *Science* **270**, pp. 1000–1003, 1995.
7. H. Sun and B. Frost, "Computation of different optical variables of looming objects in pigeon nucleus rotundus neurons," *Nature Neuroscience* **1**(4), pp. 296–303, 1998.
8. F. Gabbiani, H. Krapp, C. Koch, and G. Laurent, "Multiplicative computation in a visual neuron sensitive to looming," *Nature* **420**, pp. 320–324, 2002.
9. F. Rind and D. Bramwell, "Neural network based on the input organization of an indentified neuron signaling impending collision," *Journal of Neurophysiology* **75**, pp. 967–985, 1996.
10. F. Rind, "Intracellular characterization of neurons in the locust brain signaling impending collision," *Journal of Neurophysiology* **75**, pp. 986–995, 1996.
11. M. Blanchard, F. Rind, and F. Verschure, "Collision avoidance using a model of locust LGMD neuron," *Robotics and Autonomous Systems* **30**, pp. 17–38, 2000.
12. D. Edwards, "The cockroach DCMD neuron. I. lateral inhibition and the effects of light- and dark-adaptation," *Jornal of Experimental Biology* **99**, pp. 61–90, 1982.

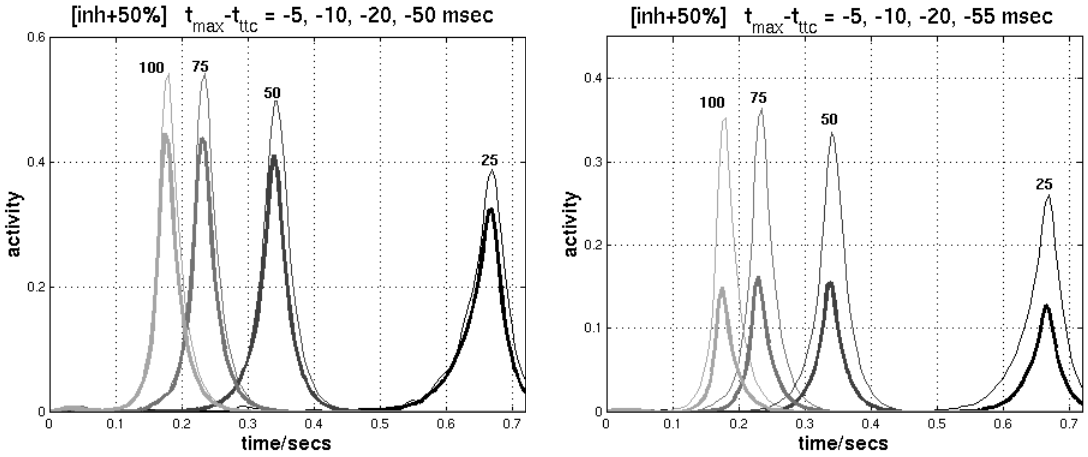

Figure 10. Increasing lateral inhibition by 50%. Increasing the strength of lateral inhibition shift LGMD response peaks away from ttc (thick curves). Thin curves correspond to default parameters (see figure 6). **Left:** "inhibition by delay" **Right:** "inhibition by diffusion".

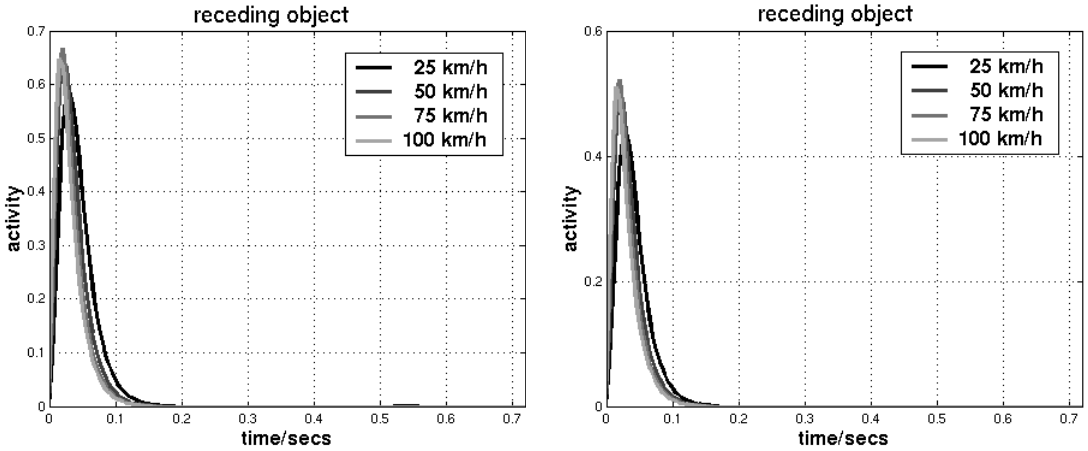

Figure 11. Receding objects. LGMD responses for receding objects with different velocities. Although no feedforward inhibition is incorporated in our model, no prolonged responses are seen. **Left:** "inhibition by delay" **Right:** "inhibition by diffusion".

13. K. Mimura, "Analysis of visual information in lamina neurons of the fly," *Journal of Comparative Physiology* **88**, pp. 335–372, 1974.
14. G. Carpenter and S. Grossberg, "Adaptation and transmitter gating in vertebrate photoreceptors," *Journal of Theoretical Biology* **1**, pp. 1–42, 1981.
15. R. Harris, D. O'Carroll, and S. Laughlin, "Adaptation and the temporal delay filter of fly motion detectors," *Vision Research* **39**, pp. 2603–2613, 1999.

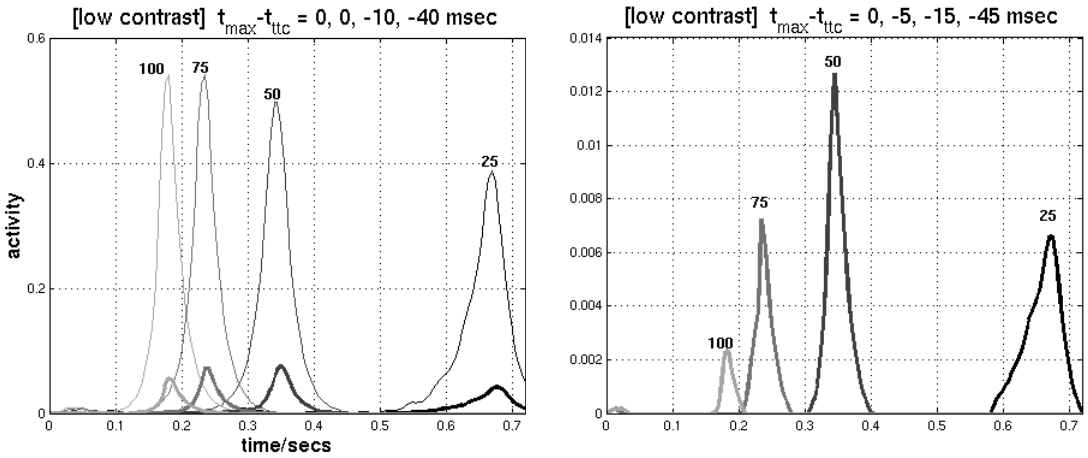

Figure 12. 50% luminance contrast. LGMD responses for a low-contrast checkerboard object (thick curves), where the contrast were half as strong as the full contrast checkerboard (thin curves). **Left:** "inhibition by delay" **Right:** "inhibition by diffusion".

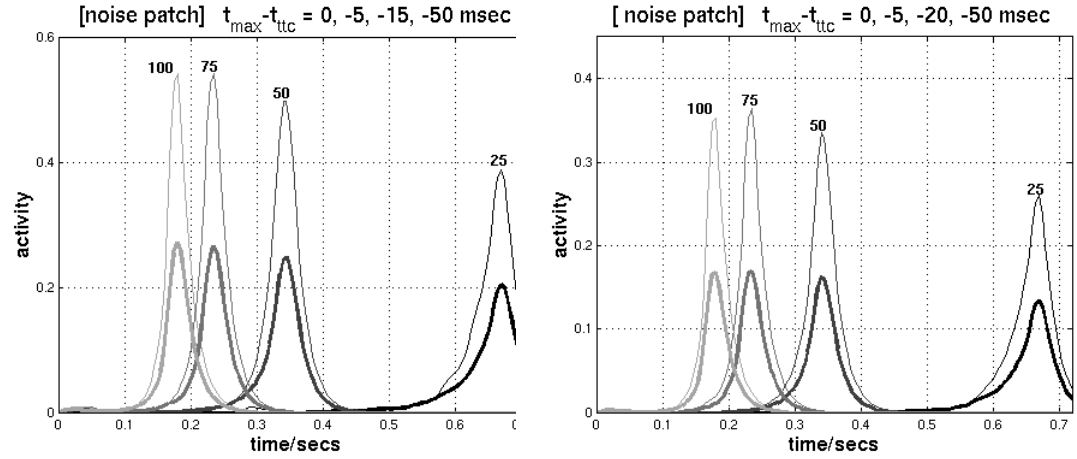

Figure 13. Noise patch on noise background. LGMD response curves to a noise object embedded in a noise background (thick curves). Thin curves correspond to default parameters (see figure 6). **Left:** "inhibition by delay" **Right:** "inhibition by diffusion".