Evolving Predictive Visual Motion Detectors

Jonas Ruesch and Alexandre Bernardino

***Abstract*—The geometrical organization of a visual sensor is of major importance for the later processing of sensed stimuli. We present an approach to evolve artificial visual detectors which adapt their size and orientation according to the experienced sensory stimulation. The criterion for the introduced optimization method is given by a Reichardt correlation measure on the input signal. Under the described conditions, the visual receptors organize their spatial arrangement following the average luminance flow recorded by the sensor over time.**

***Index Terms*—Sensor morphology, Reichardt correlation, elementary motion detection, visual receptor distribution, self-organization, optimization.**

I. Introduction

Since the first light sensitive receptors appeared in primitive animals, presumably 550 million years ago, visual sensors have evolved into highly complex perception systems occupying big areas of the human brain. As has been shown in embodied artificial intelligence research, the structure of the information flow in a sensory system is essential for later processing [1], [2]. This structure is determined conjointly by the agent's behavior and the organization of the sensor.

In this paper we address the aspect of evolving a geometrical sensor layout given the particular stochastic properties of an input signal. In nature many examples can be found where visual sensors have evolved in accordance to an animal's behavior. For instance the human retina was shown to have an optimal spatial organization for tracking tasks [3]. Another example is the housefly's compound eye, whose photo-receptors are distributed such as to simplify distance estimation from motion parallax [4], [5]. An interesting study in [6] has recorded images from a camera mounted on a cat's head and has found that the video stream presents clear statistical regularities. For example a predominance of horizontal and vertical orientations over oblique contours was found. These statistics seem to influence the organization of retinal ganglion cells and their dendrites, that are preferentially arranged along the vertical and horizontal meridians.

The mechanism that drives the emergence of such sensory organization has been subject of research by several authors [7]–[9]. Simulations have shown that sensor structures can evolve in direct response to the structure of the visual stimulus [10]. In that work a Kohonen learning rule is used to adapt a network of receptors. The network is adapted according to self-induced transformations of previous receptor locations. It is shown that the network converges to a retina-like spatial organization with efficient prediction of the expected transformation. Thus the sensor morphology adapts to the signal's statistical properties in order to simplify motion prediction. Clippingdale's work [10] addresses sensor structure adaptation assuming known motions. However, motion estimation in visual sensors is by no means trivial and a large body of research has been dedicated to this issue. Energy-based models [11] and Reichardt-detectors [12], [13] have a clear biological motivation and have been shown to outperform linear gradient schemes under varying spatial frequencies [14]. In our work we follow Clippingdale's paradigm but use Reichardt type detectors that are adapted incrementally according to the statistics of a two-dimensional image stream. We propose an on-line algorithm to optimize the detector's layout to reflect the average intensity displacements. For this purpose we furthermore study some aspects of the strongly non-linear properties of Reichardt-detectors.

In related work [15] a one-dimensional array of receptors is used for obstacle avoidance in a mobile robot. The morphology of the sensor is evolved using an genetic algorithm where each solution codes for a different sensor geometry. On the contrary, we adopt an ontogenetic approach to organize two-dimensional detectors which continue to adapt if the characteristics of the input stimulus change over time.

In the next section we briefly review the Reichardt Correlation Model. The proposed optimization is introduced in section III and accompanied by an analysis of the search landscape. Demonstrations and results for one- and two-dimensional sensors are presented in section IV.

II. Measuring Intensity Displacements

Motion directly defined by intensity changes is called *primary motion* or *Fourier motion* and can be predicted from the spatio-temporal Fourier transform [14]. The Reichardt-detector described in the next section is suitable and biological plausible to measure direction and velocity of this type of motion.

A. A Short Review on the Reichardt-Detector

Originally proposed by W. Reichardt, the so-called *elementary motion detector* of the correlation type (EMD) detects the displacement of a luminance distribution in a given direction. This particular mechanism is also referred to as the *Reichardt-detector* and was developed based on the analysis of insect behavior [12], [13], [16]. The detector consists of two input elements located at two different spatial locations $\mathbf{A}$ and $\mathbf{B}$ (Fig. 1). Each of these receptors has a receptive field filtering the input intensity $I(\mathbf{x}, t)$ spatially with $\mathcal{F}(I(\mathbf{x}, t))$.

To compute the response of a Reichardt-detector, a first correlation measure is computed by multiplying the resulting signal at location $\mathbf{B}$ with the temporally filtered signal at

J. Ruesch and A. Bernardino are with the Institute for Systems and Robotics, Instituto Superior Técnico, Lisbon, Portugal, e-mail: {jruesch,alex}@isr.ist.utl.pt

978-1-4244-4118-1/09/$25.00 ©2009 IEEE

location $\mathbf{A}$. By repeating this step in a reciprocal way, a second correlation measure is obtained. In a next step, these two results are subtracted from each other yielding the actual correlation measure related to intensity displacement:

$$C_{AB} = \mathcal{F}(I_B) \cdot \widetilde{\mathcal{F}}(I_A) - \mathcal{F}(I_A) \cdot \widetilde{\mathcal{F}}(I_B), \quad (1)$$

where $\widetilde{\mathcal{F}}$ computes the temporally filtered response. Note that the subtraction step eliminates any output resulting from temporal stationary intensity signals and that directional sensitivity to intensity displacements is conserved by this operation. The sign of the output indicates the orientation of the detected change. The response of this detector is influenced on the one hand by the receptor locations $\mathbf{A}$ and $\mathbf{B}$, and on the other hand by the time constant T of the temporal filter in use. For visual receptors in insects, the temporally delayed signal is usually assumed to be low-pass filtered. An additional parameter influencing the correlation measure is the size σ of the receptive field, which affects the spatial frequency resolution of the detector.

B. Implementation

Each receptor samples an intensity value from an underlying pixel array using a gaussian filter with a kernel size σ. Receptors can be placed with sub-pixel accuracy and interpolate intensity values linearly between pixels. Pixels can have assigned a value in the range of [0,1].

In order to subsequently express a correlation measure depending on the receptor locations and the location $\mathbf{x}$ of a detector we rewrite (1) as:

$$C_{\mathbf{x}}(\mathbf{d}, t) = \mathcal{F}(I(\mathbf{x} + \mathbf{d}/2, t)) \cdot \widetilde{\mathcal{F}}(I(\mathbf{x} - \mathbf{d}/2, t) - \mathcal{F}(I(\mathbf{x} - \mathbf{d}/2, t) \cdot \widetilde{\mathcal{F}}(I(\mathbf{x} + \mathbf{d}/2, t), \quad (2)$$

where the vector $\mathbf{d} = \mathbf{B} - \mathbf{A}$ describes the distance between the receptors for the detector at location $\mathbf{x}$.

For our analysis we define $\widetilde{\mathcal{F}}$ as a time delay with period $T = 1$. Although such simple time delays are not commonly found in nervous systems, this model corresponds qualitatively to motion detectors found in flies [16].

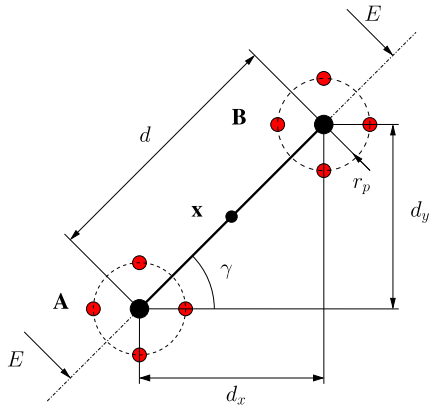

Fig. 1. A Reichardt-detector located at position $\mathbf{x}$ with receptors at locations $\mathbf{A}$ and $\mathbf{B}$ at a distance d from each other. The four marked positions at distance r_p from $\mathbf{A}$ and $\mathbf{B}$ indicate the perturbed receptor locations used in the gradient estimation step of the optimization algorithm.

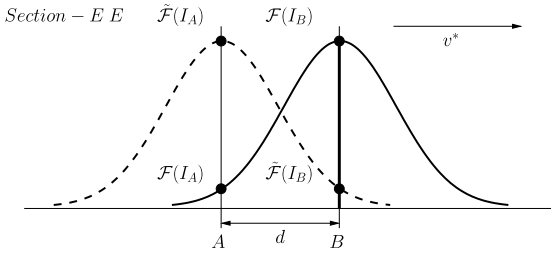

Fig. 2. Intensity measured for each spatial location (x-axis) while an intensity pulse is moving from left to right. Each location filters intensity with a Gaussian receptive field. The situation corresponds to the response to an intensity pulse passing at the characteristic velocity v^*. The pulse moves at a distance per time step of exactly d and is currently located at location B. The dashed curve represents the time-delayed response.

Of major importance for our work is that implicitly, the distance $||\mathbf{d}||$ and the time constant T define a velocity for which a signal traveling in the direction of the detector evokes a maximum response (see Fig. 2). This characteristic velocity v^* increases if the the distance between $\mathbf{A}$ and $\mathbf{B}$ is enlarged or if the time constant T is decreased and vice versa. As we keep T constant, the vector $\mathbf{d}$ only defines the characteristic velocity $\mathbf{v}^*$. The receptive field of each sensor is modeled by a Gaussian of size σ.

III. Optimization Algorithm

Our objective is to place the receptor pairs of each Reichardt-detector in a way such that the maximum correlation is recorded over time. If we consider the detectors to be placed at fixed locations $\mathbf{x}$, the maximization of the total correlation over time and space can be done by optimizing each individual detector independently. In the discussion of the computational algorithms we consider that Reichardt-detectors are sampled at the rate of the incoming images, thus t is considered, from now on, a discrete time variable.

To find the optimum receptor configuration for a detector at location $\mathbf{x}$ we want to find a receptor distance $\mathbf{d}$ which maximizes the expected value of the correlation measure:

$$\arg\max_{\mathbf{d}} E\left[C_{\mathbf{x}}(\mathbf{d}, t)\right], \quad (3)$$

As long as we keep $\mathbf{d}$ constant, C is a stationary stochastic process and a good estimate of the expected value in (3) is given by the sample mean:

$$\overline{C}_{\mathbf{x}}(\mathbf{d}, t) = \frac{1}{t} \sum_{\tau=0}^{t} C_{\mathbf{x}}(\mathbf{d}, \tau) \quad (4)$$

For the purpose of finding a $\mathbf{d}$ which maximizes $\overline{C}_{\mathbf{x}}(\mathbf{d}, t)$ we employ a stochastic gradient descent method where the respective gradient is:

$$\frac{\partial}{\partial \mathbf{d}} \overline{C}_{\mathbf{x}}(\mathbf{d}, t) = \frac{1}{t} \left(\sum_{\tau=0}^{t} \frac{\partial}{\partial \mathbf{d}} C_{\mathbf{x}}(\mathbf{d}, t) \right) \quad (5)$$

Though, because the statistical properties of $\overline{C}_{\mathbf{x}}$ changes as soon as $\mathbf{d}$ is changed, we would need to wait an infinite amount

of time so that $\overline{C}_{\mathbf{x}}$ adapts to the changes in the input statistics. On the other end we could use the approximation:

$$\frac{\partial}{\partial \mathbf{d}}\overline{C}_{\mathbf{x}}(\mathbf{d},t) \approx \frac{\partial}{\partial \mathbf{d}}C_{\mathbf{x}}(\mathbf{d},t) \tag{6}$$

However, using such a coarse approximation of $\partial\overline{C}_{\mathbf{x}}/\partial\mathbf{d}$ proofed to be too noisy to update $\mathbf{d}$. We therefore opt for using an exponentially smoothed moving average and estimate $\overline{C}_{\mathbf{x}}$ using the update rule:

$$\frac{\partial}{\partial \mathbf{d}}\overline{C}_{\mathbf{x}}(t,\mathbf{d}) = (1-\alpha)\frac{\partial}{\partial \mathbf{d}}C_{\mathbf{x}}(t,\mathbf{d}) + \alpha\frac{\partial}{\partial \mathbf{d}}\overline{C}_{\mathbf{x}}(t-1,\mathbf{d}-1), \tag{7}$$

where α defines the time window over which $\overline{C}_{\mathbf{x}}$ is averaged. The final update of the receptor positions can then be done as:

$$\mathbf{d}(t) = \beta\frac{\partial}{\partial \mathbf{d}}\overline{C}_{\mathbf{x}}(t-1). \tag{8}$$

Note, as the update of the estimation of $\partial C_{\mathbf{x}}(\mathbf{d},t)/\partial\mathbf{d}$ and the update of $\mathbf{d}$ are done concurrently at each time step, the parameters α and β are interdependent. This interdependency is discussed later.

A. Receptor Perturbations

To compute $\partial C_{\mathbf{x}}(\mathbf{d},t)/\partial\mathbf{d}$ we use a finite difference approximation where the correlation measures for the differences are obtained by perturbing the receptors in every dimension by a distance of r_p in positive and negative direction. The current gradient is then calculated by subtracting correlation measures from two opposing perturbed configurations in each dimension.

Applying this scheme to the two-dimensional case, we compute for each receptor the correlation values C_{NS}, C_{SN}, C_{WE}, and C_{EW} using the receptor pairings $\mathbf{A}_S$-$\mathbf{B}_N$, $\mathbf{A}_N$-$\mathbf{B}_S$, $\mathbf{A}_W$-$\mathbf{B}_E$, and $\mathbf{A}_E$-$\mathbf{B}_W$. The indices code for receptor perturbations in direction north, west, south and east of the receptor's original position (see Fig. 1). The gradient $\partial C_{\mathbf{x}}(\mathbf{d})/\partial\mathbf{d}$ for the current time step is then approximated as:

$$\frac{\partial}{\partial \mathbf{d}}C_{\mathbf{x}}(\mathbf{d}) = \begin{pmatrix} (C_{NS}(\mathbf{d}) - C_{SN}(\mathbf{d}))/2r_p \\ (C_{WE}(\mathbf{d}) - C_{EW}(\mathbf{d}))/2r_p. \end{pmatrix} \tag{9}$$

In this work, the receptor perturbations are done during one time step. Referring to a real sensor in an artificial or biological system, the analogon to such a mechanism could either happen at a faster timescale or be provided by a number of coupled receptors. Furthermore, we want a detector at location $\mathbf{x}$ with direction γ to give the same response as a detector at location $\mathbf{x}$ with direction $\gamma + 180\,\mathrm{deg}$. We therefore use the absolute value of the correlation response to compute $\partial C_{\mathbf{x}}(\mathbf{d},t)/\partial\mathbf{d}$. This in fact introduces a second optimum $\mathbf{d}_2$ which equals $\mathbf{d}_1$ mirrored at the origin.

B. Optimization Parameters

Adjustable parameters for the described optimization are the gradient averaging rate α, the optimization update rate β, the perturbation radius r_p and the σ of the input filter.

For the parameter α, two things have to be taken into account:

1) Considering $\mathbf{d}$ constant, α corresponds to a learning rate for $\partial\overline{C}_{\mathbf{x}}/\partial\mathbf{d}$, or seen from a different perspective, α is related to the bandwidth of a temporal filter which must be tuned to attenuate the noise in the instantaneous gradient calculations.
2) If $\mathbf{d}$ is changing, α also corresponds to a momentum term as used in common acceleration techniques in gradient based optimization.

So, when choosing the parameter α, we must primarily find a value specifying a time window which is large enough for having a good estimate of $\partial\overline{C}_{\mathbf{x}}/\partial\mathbf{d}$ and at the same time, we have to keep in mind that by increasing α we increase the momentum term which can lead to overshooting and oscillation when optimizing $\mathbf{d}$. Because we want to estimate the average gradient over long time courses, we are forced to choose a relatively high α and therefore the momentum of the gradient descent is necessarily high. To keep this momentum within reasonable bounds we should nevertheless choose α as low as possible. How low α can be chosen depends on the observed signal variability. If the characteristic patterns of the input signal are repeated in short time periods, then α can be smaller. The values chosen for the actual experiments are described in section IV.

Because of the introduced momentum, care has to be taken as well when choosing the update rate β in relation to α because a high adaptation rate combined with a high momentum leads to the previously mentioned behavior of overshoot and oscillation. In other words, we need a β small enough for the gradient to have time to be updated sufficiently accurate for a new distance $\mathbf{d}$.

Choosing the perturbation radius r_p is less tricky. The bounds for r_p are constrained by the curvature of the search landscape. As the highest curvature always occurs close to the maximum, as can be seen from Fig. 4, we just need to choose r_p small enough to estimate the gradient around the optimum with a desired accuracy.

The effects of changing parameter σ are described in the following section.

C. Search Landscape

Before running the optimization, we briefly analyze the search landscape on which it operates. This will provide us with an insight on the relationship between receptor distances for which results a maximum correlation response and the corresponding intensity displacement velocities.

For the one-dimensional case the correlation response for a moving intensity pulse plotted against the receptor distance and the signal velocity is shown in Fig. 3. The receptors have a receptive field of size $\sigma = 2.5$. Note that for every distance and velocity configuration the plot actually shows the response value for the point in time when the detector reacts with maximum response to the passing intensity pulse. For different points in time the plot remains the same up to a scale factor. Of importance for our purpose is the approximately one-to-one relationship between distances and velocities for distances > 5 which maximize the correlation response. Due to the drop in response close to the origin, the receptor distances found by

the optimization have a minimum bound. One can envisage this relationship by looking at Fig. 2: When decreasing the signal velocity and decreasing at the same time the measuring distance, then there will be a point where the correlation response starts to decrease because the result of $\mathcal{F}(I_A)\cdot\widetilde{\mathcal{F}}(I_B)$ approaches $\mathcal{F}(I_B)\cdot\widetilde{\mathcal{F}}(I_A)$ before $\mathcal{F}(I_B)$ and $\widetilde{\mathcal{F}}(I_A)$ reach their peak response. The distance d for which this happens depends on the σ used in the pre-processing of the input signal and equals approximately 2σ. Fig. 4 shows two cut-sections of Fig. 3 for two different velocities. These are the profiles on which the gradient descent is performed for a fixed velocity. Measures in pixels relate to the resolution of the observed signal. To see the changes to the correlation response for a different σ compare these plots with Fig. 6(d).

A plot for the two-dimensional case can be seen in Fig. 5. Again, the correlation response is maximized for receptor distances d_x and d_y which match the given intensity displacement velocity $v_x = 6$ and $v_y = 4$.

Apart from defining a lower bound for d, the σ of the pre-processing filter has another important impact: If one wants to prevent aliasing, then the width of the receptive field should be twice the distance d between the receptors. This relation is in accordance with the Nyquist sampling theorem. However, as can be seen from Fig. 3, the proposed optimization requires a σ which does not fulfill this condition and therefore aliasing can occur. This means a motion detector can be stimulated not only by a single moving edge, but by different portions of the intensity distribution moving over it. The effect of this related to our optimization is that receptor pairs initialized at a distance far away from the optimum can converge to a different distance which corresponds to maximum correlation in case of aliasing.

IV. Results

To test the proposed algorithm we present first an experiment for a one-dimensional d. Subsequent results relate to two-dimensional input signals.

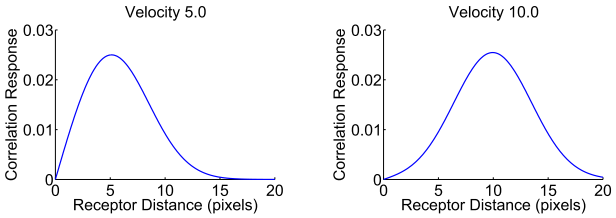

Fig. 4. The correlation response plotted against the receptor distance d for two different intensity displacement velocities (cut-sections of Fig 3).

A. One-Dimensional Case

In this setup 20 detectors are uniformly distributed over a one-dimensional input signal of size 320 pixels. Each detector is initialized with its receptors at a distance of 8 pixels. The receptive field of each receptor has size $\sigma = 1.5$. During the simulation an intensity pulse moves on the input image in positive direction according to a sinusoidal velocity profile with a mean displacement of 8 pixels and an amplitude of 4 pixels while the period 2π of the sine extends over 160 pixels (see Fig. 6(b)).

Fig. 6(e) and 6(f) show the results of a simulation lasting $10'000$ time steps. Remember that receptors can move and measure intensity with sub-pixel accuracy using linear interpolation. A slight overshoot can be observed for some distance values. The noise in the gradient estimation is dependent on the parameter α. The parameters used for this simulation are summarized in TABLE I. In Fig.6(d) a plot of the search landscape according to the used σ is shown. On the dashed line the initial distances are marked. The arrows point in the direction the distances change to reach the optimum on the bold red line.

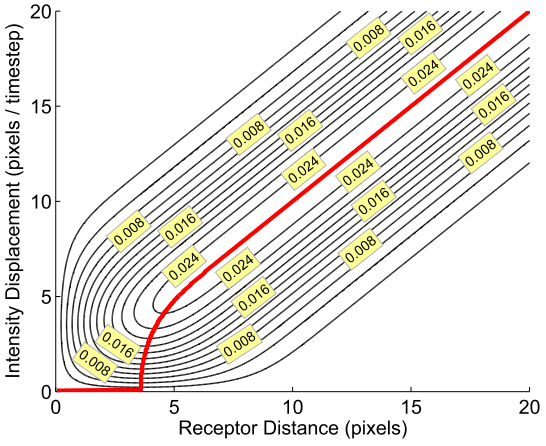

Fig. 3. The correlation response plotted against the receptor distance (x-axis) and the signal velocity (y-axis). The correlation is computed assuming an intensity pulse filtered by receptive fields of size $\sigma = 2.5$. The curve in bold red indicates for each velocity value the receptor distance for which the maximum correlation is measured. Values in pixels relate to the resolution of the input signal.

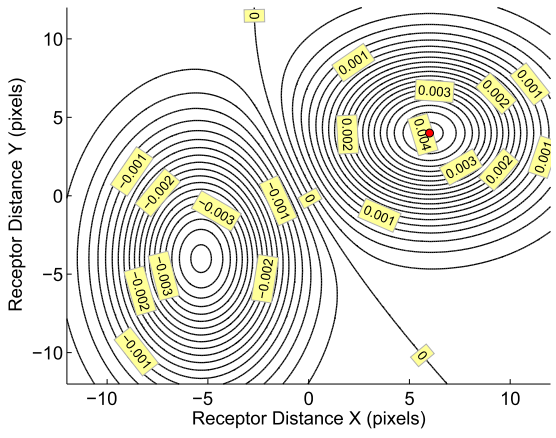

Fig. 5. The correlation response of an intensity pulse moving in two-dimensional space in x-direction at a velocity of 6 pixels per time step and in y-direction at 4 pixels per time step. The correlation response is plotted against varying horizontal and vertical receptor distances. The dot in bold red marks a receptor with distances $d_x = 6$ and $d_y = 4$ which maximizes the correlation response.

TABLE I
PARAMETER SETTINGS

Parameter		one-dimensional	two-dimensional
Perturbation radius	r_p	0.5pixels	0.5 pixels
momentum, (1 - learning rate)	α	$1 - 2 \cdot 10^{-3}$	$1 - 5 \cdot 10^{-4}$
Receptor position update rate	β	0.1	1.0
Receptive field size	σ	1.5 pixels	1.5 pixels

B. Two-Dimensional Case

The results for a two-dimensional input image were obtained by generating a pattern of 80 intensity pulses at random locations on a 96 by 96 pixel array. The initial configuration consists of 64 detectors distributed as shown in Fig. 7(b) with receptors at a distance of 3 pixels randomly oriented. The receptive field of each receptor has a size of 1.5 pixels. A typical input pattern as filtered by the receptors is shown in Fig. 7(a). To prevent a particular pattern to introduce a bias, a random pattern is regenerated every 100 time steps. The response from two different motion patterns are presented here. The final detector configurations can be seen in Fig. 7(e) and Fig. 7(f): i) The random pattern is rotated at an angle of $15\,\mathrm{deg}$ per time step for $10'000$ time steps; ii) The intensity pulses are repeatedly scaled to a distance of 0.875 and 1.125 times their original distance from the center of the image.

V. DISCUSSION AND FUTURE WORK

In the two-dimensional experiments a minor number of receptors did not converge to the optimum distance after $10'000$ time steps. As can be seen from the plots in Fig. 7, these are mainly detectors which have their receptors located at a $\mathbf{d}$ perpendicular to the optimum $\mathbf{d}$. Referring to the search landscape these detectors are located on the contour curve of level 0 in Fig. 5. To explain this, we have to take into account that the motion patterns we generate exert a displacement velocity in exactly one direction at each location. Depending on the distance of a receptor pair, this displacement cannot be detected by receptors with a $\mathbf{d}$ perpendicular to this direction.

Another observation concerns a few detectors which show a receptor distance which is too large to match the intensity displacement at that location. These detectors suffer from the effect of aliasing described in section III-C.

In future work we expect both of the above described flaws to diminish by introducing an inter-unit coupling. Coupling forces between detectors could lead to a smoothing effect eliminating outliers. Furthermore, by connecting neighboring units the organization of a coherent network of motion detectors of different sizes becomes possible. Some work is required related to the optimization, as with this extension the assumption of a fixed location for each detector must be abandoned.

Eventually we will investigate how such detector-networks organize when observing image streams similar to the ones sensed by organisms moving in their natural environment; e.g. images corresponding to the perspective of a fly during flight, or the image stream resulting from a saccading sequence of a human.

ACKNOWLEDGMENT

This work was supported by the European Commission, Project IST-004370 RobotCub, and by the Portuguese Government – Fundação para a Ciência e Tecnologia (ISR/IST plurianual funding) through the POS_Conhecimento Program that includes FEDER funds, and through project BIO-LOOK, PTDC / EEA-ACR / 71032 / 2006.

REFERENCES

[1] Pfeifer, R., Lungarella, M., Sporns, O., Kuniyoshi, Y. On the Information Theoretic Implications of Embodiment: Principles and Methods, *50 Years of Artificial Intelligence*, Berlin, Germany, December ,2007, pp76-86.

[2] Lungarella, M., Sporns, O. Mapping information flow in sensorimotor networks, *PLoS Computational Biology*, Vol. 2, Oct., 2006, pp. 1301-1312.

[3] Vincze, M., Weiman, C.F.R. General relationship for optimal tracking performance, *Proc. SPIE*, Vol. 2904, Nov., 1996, pp. 402-412.

[4] Van der Zwaan, S., Santos-Victor, J. An Insect Inspired Visual Sensor for the Autonomous Navigation of a Mobile Robot, Proc. 7th Int. Symp. on Intelligent Robotic Systems - SIRS99, Coimbra, Portugal, Jul., 1999.

[5] Franceschini, N., Pichon, J. M., Blanes, C. From Insect Vision to Robot Vision, *Philosophical Transactions: Biological Sciences*, Vol. 337, No. 1281, Sep., 1992, pp. 283-294.

[6] Betsch, B. Y., Einhäuser, W., Körding, K. P., König, P. The world from a Cat's Perspective Statistics of Natural Videos, *Biological Cybernetics*, Vol. 90, No. 1, Jan., 2004, pp. 41-50.

[7] Linsker, R. Self-Organization in a Perceptual Network, *IEEE Computer*, Vol. 21, No. 3, Mar., 1988, pp. 105-117.

[8] Schwartz, E. L. Computational Anatomy and Functional Architecture of Striate Cortex: A Spatial Mapping Approach to Perceptual Coding, *Vision Research*, Vol. 20, 1980, pp. 645-669.

[9] Hubel, D. H., Wiesel, T. N. Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex, *J. Physiol. Lond.*, Vol. 160, Jan., 1962, pp. 106-154.

[10] Clippingdale S.m Wilson R. Self-similar Neural Networks Based on a Kohonen Learning Rule, *Neural Networks*, Vol. 9, No. 5, Jul., 1996, pp. 747-763.

[11] Adelson, E.H., Bergen, J.R. Spatiotemporal energy models for the perception of motion, *J Opt Soc Am A*, Vol. 2, No. 2, Feb., 1985, pp. 284-299.

[12] Reichardt, W. "Autokorrelationsauswertung als Funktionsprinzip des Zentralnervensystems", Z. Naturforsch. 12b, 1957, p447-457.

[13] Reichardt, W. Autocorrelation, a principle for evaluation of sensory information by the central nervous system, *Principles of Sensory Communications*, New York, 1961, pp. 303-317.

[14] Zanker, J.M. On the Elementary Mechanism Underlying Secondary Motion Processing, *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*,Vol. 351, 1996, pp. 1725-1736.

[15] Lichtensteiger, L., Eggenberger, P. Evolving the Morphology of a Compound Eye on a Robot, Proc. Third European Workshop on Advanced Mobile Robots (Eurobot '99), Piscataway, NJ, USA, 1999, p.127-34.

[16] Srinivasan, M. V., Poteser, M., Kral, K. Motion Detection in Insect Orientation and Navigation, *Vision Research*, Vol. 39, 1999, pp. 2749-2766.

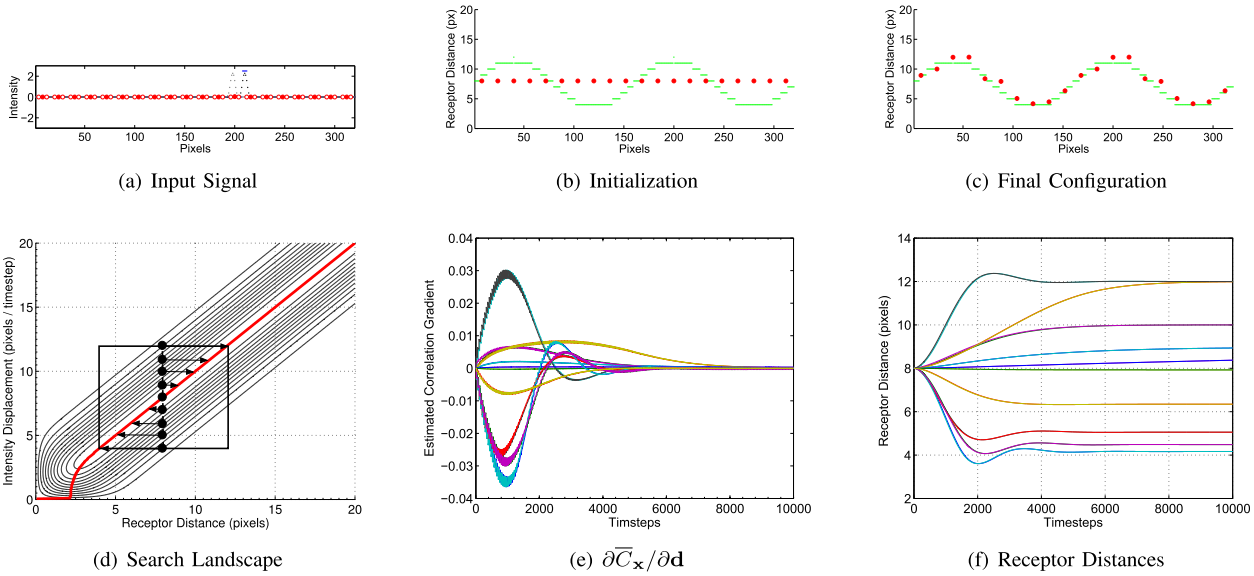

(a) Input Signal (b) Initialization (c) Final Configuration

(d) Search Landscape (e) $\partial\overline{C}_{\mathbf{x}}/\partial\mathbf{d}$ (f) Receptor Distances

Fig. 6. Results for the one-dimensional case: (a) An intensity pulse and its delayed response travelling from left to right over 320 pixels. (b) The initial receptor distances for 20 detectors (red dots) and the intensity pulse displacement profile (green dots). (c) The final receptor distances after $10'000$ iterations. (d) Search landscape for a receptive field of size $\sigma = 1.5$. The dots on the dashed line mark the initial receptor distances within the observed intensity displacement velocity range. The arrows indicate the changes during the optimization towards the optimum value on the bold red line. Note that each receptor experiences one particular velocity and therefore the gradient points in a horizontal direction in this plot. (e) The receptor distances plotted against time. (f) $\partial\overline{C}_{\mathbf{x}}/\partial\mathbf{d}$ plotted against time.

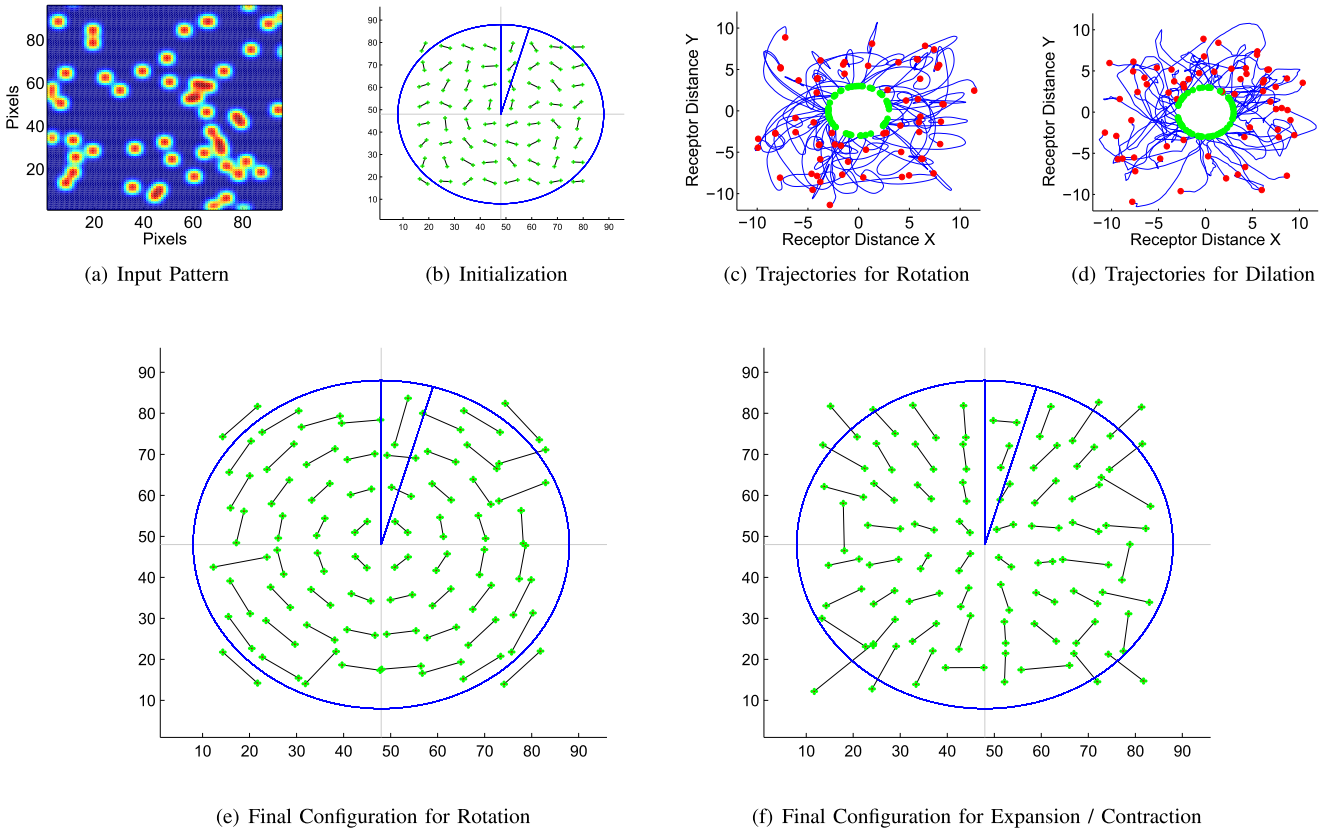

(a) Input Pattern (b) Initialization (c) Trajectories for Rotation (d) Trajectories for Dilation

(e) Final Configuration for Rotation (f) Final Configuration for Expansion / Contraction

Fig. 7. Results for the two-dimensional case: (a) A typical input pattern of 80 intensity pulses distributed randomly on an image of 96 by 96 pixels. (b) The initial distribution of 64 detectors initialized with receptors at a distance of 3 pixels at random orientation. (c), (d) The evolution of the receptor distances over $10'000$ timesteps. (e), (f) The final configurations for the 64 receptor pairs after observing a rotating, respectively a scaled input pattern for $10'000$ time steps.