

A Biologically Inspired Computational Vision Front-end based on a Self-Organised Pseudo-Randomly Tessellated Artificial Retina

Sumitha Balasuriya and Paul Siebert

Department of Computing Science

University of Glasgow

Glasgow G12 8QQ

E-mail: sumitha@dcs.gla.ac.uk, psiebert@dcs.gla.ac.uk

Abstract— This paper considers the construction of a biologically inspired front-end for computer vision based on an artificial retina ‘pyramid’ with a self-organised pseudo-randomly tessellated receptive field tessellation. The organisation of photoreceptors and receptive fields in biological retinae locally resembles a hexagonal mosaic, whereas globally these are organised with a very densely tessellated central foveal region which seamlessly merges into an increasingly sparsely tessellated periphery. In contrast, conventional computer vision approaches use a rectilinear sampling tessellation which samples the whole field of view with uniform density.

Scale-space interest points which are suitable for higher level attention and reasoning tasks are efficiently extracted by our vision front-end by performing hierarchical feature extraction on the pseudo-randomly spaced visual information. All operations were conducted on a geometrically irregular foveated representation (data structure for visual information) which is radically different to the uniform rectilinear arrays used in conventional computer vision.

I. INTRODUCTION

In this paper the authors describe the construction of a generic space-variant computational vision front-end capable of identifying locations of interest on an image suitable for the extraction of feature information. The authors’ approach differs from that used previously by using a biologically inspired methodology to extract visual information from an image. A multi-resolution ‘pyramid’ of artificial retinae with independently self-organised receptive field tessellations fixates upon and samples visual information from an image. The tessellation of the artificial retinae resembles the tessellation of retinal receptive fields and photoreceptors in the human retina [1], with a local hexagonal grouping which is very dense in the central foveal region of the retina and becomes increasingly sparse in the surrounding space-variant periphery.

A space-variant sampling of visual information using a foveated sensor reduces the combinatorial explosion of information and massive computational processing load associated with vision. In human retinae, the highest acuity region in the foveola has a diameter of about $1\frac{1}{2}$ degrees around the point of fixation in our field-of-view. This corresponds to about 1cm at arms length. The rest of our field-of-view is sampled at reduced acuity by the rest of the fovea

(with a diameter of 5 degrees) and at increasingly reduced detail in the large periphery (diameter 150 degrees). This reduction in sampling density with eccentricity isn’t just because of the biological difficulty of tightly packing sensor elements in our retinae. The visual information extracted by our retinae undergoes extensive processing in the visual cortex. Based on the biological computational machinery dedicated to the human fovea, our brains would have to weigh about 60kgs if we were to process our whole field-of-view at foveal resolution!

Most machine approaches to computer vision tend to rely on uniform, rectilinear information sampling approaches and internal representations. The whole field-of-view is evaluated at the same resolution while processing resources and the sampling machinery are not sparingly conserved. The internal representation of visual information in these machine systems is almost invariably a pixel array which contains the responses of sampling elements or the results after filtering operations.

There is a discrepancy between the sampling and representations found within biological vision systems and those available to computational machine approaches to vision. This paper demonstrates that it is viable to construct a computational artificial vision front-end that is capable of extracting useful visual information and which is based on the sampling of a biologically-inspired artificial retina with a tessellation similar to that found in human retinae.

II. BACKGROUND

Most artificial retinae that sample images in modern computer vision are based on an analytic retino-cortical transform that projects locations in retinal space (image space) to an associated ‘cortical space’ which is biased to the foveal region near the point of fixation. These transforms are motivated by the topography of the projection of retinal afferents to the LGN and visual cortex in biology. Schwartz’s complex-log transform [2] creates a retinotopic cortical space image by taking the complex logarithm of the image pixel coordinates. While almost all machine retinae that sample images are based on this model it contains severe limitations. There is a singularity in the centre of the retina and the model over-samples visual information on the image near the point of

fixation resulting in a highly correlated foveal area in the cortical image. While other retino-cortical transforms have been proposed [3, 4] none quantitatively calculate the actual receptive field locations (tessellation) of an artificial retina. The work is based on the projection of the radial component of retinal coordinates to cortical space and does not consider the angular relationships between the (polar) coordinates of the retinal plane when performing the projection. No analytic approach or geometric mapping that can describe the gradual change in topography of the retina between a uniform fovea and space-variant periphery transforming the lateral and radial coordinate of a receptive field or image pixel to a cortical space has been reported in the literature to date.

Researchers [5, 6] have also attempted to process the central foveal region separately to avoid the central singularity and over-sampling. Sandini et. al. [7] even created hardware implementations of such retinae. The central region comprises of a uniform density mosaic while the periphery region is space-variant. However, using this approach results in a discontinuity between the foveal and periphery in the retinal sampling and internal visual representation. Thus far researchers have failed to report computing plausible retinae based on an analytical transform that can sample an image using a single space-variant retinal topology and map locations in the field of view to a continuous cortical layer.

The authors question the tractability of finding a continuous space-variant retinal tessellation using an analytic perspective that meets the constraints of a continuous uniform fovea and space-variant periphery sampling regimen. Therefore in this paper the authors determined the retinal receptive field locations using self-organisation.

III. RETINA TESSELLATION

A retinal tessellation with a uniform foveal region which seamlessly merges into a space-variant periphery was created using Self-Similar Neural Networks [8] self-organisation. This approach resembles Kohonen's Self-Organising Feature Map [9] but differs by deriving the input stimulus by applying a given sequence of transformations to the current network of weights. In this paper, a two-dimensional point $x_i \in \mathcal{R}^2$ represents each retinal neuron or network weight. The network is initialized by an arbitrary random tessellation of the required number of retinal neuron positions which will eventually self-organise into the retinal tessellation that is required for a foveated retina. For a retina with N neurons, the input stimulus $y_i(n)$ at iteration n is calculated by the following, where $x_i(n-1)$ is the i^{th} neuron at iteration $n-1$ and $1 \leq i \leq N$.

$$y_i(n) = T(n) x_i(n-1) \quad (1)$$

The final global configuration of the network is governed by the composite transform T . The following sequence of transforms T was used to make the resultant network converge onto a foveated space-variant retinal tessellation.

- 1) A rotation of the network with a random transformation angle between 0 and 2π .

- 2) A dilation (increase in eccentricity) comprising of the exponent of a dilation factor which is random between 0 and $\log(8)$.
- 3) A random translation of the network weights with a value between 0 and f .

As the input stimuli $y_i(n)$ could be transformed beyond the domain of the network, those points which lie outside the bounds of the retina were culled before the network weights $x_j(n-1)$ were stimulated to calculate $x_j(n)$. The following learning rule was used to stimulate the network.

$$x_j(n) = x_j(n-1) + \alpha(n) \sum_{i \in \Lambda_j(n)} (y_i(n) - x_j(n-1)) \quad (2)$$

$$\Lambda_j(n) = \left\{ i : |y_i(n) - x_j(n-1)| < |y_i(n) - x_k(n-1)|, k \neq j \right\} \quad (3)$$

$\Lambda_j(n)$ contains the indices to the input stimuli $y_i(n)$ to which $x_j(n-1)$ is the closest network vector.

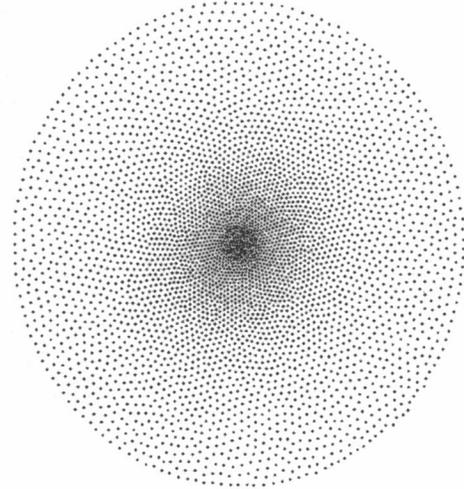


Fig. 1. Self-organised retina receptive field tessellation with 4096 nodes

Unlike complex-log transform based retinae, the retina tessellation resulting from self-organisation has a uniform foveal sampling density (Fig. 1 and Fig. 2). The above retina tessellation was implemented using a composite of horizontal ($f = 0.2$), vertical ($f = 0.2$) and radial ($f = 0.066$) translations in the composite transform T to create a tessellation with a large uniform isometric foveal region and the learning rate was annealed for 20000 iterations.

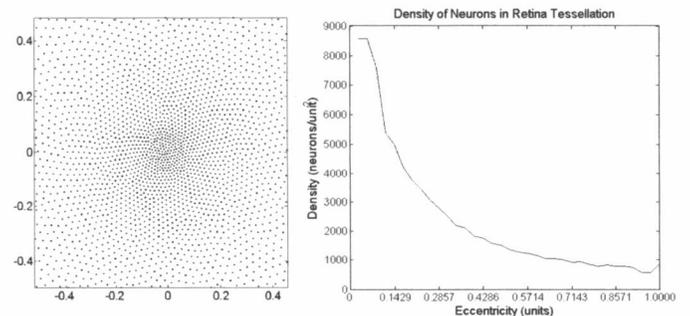


Fig. 2. Magnified view of the foveal and receptive field density. The retina tessellation has a radius of one unit.

IV. RECEPTIVE FIELDS

A retinal tessellation cannot sample visual information. Visual information must not be gathered by point sampling the locations indicated by the tessellation. Such an approach would cause aliasing in the extracted visual information. Instead neurons need to extract information using overlapping receptive fields. In this paper receptive fields with Gaussian support profiles were used initially sample visual information from an image because the wide, low-pass frequency support of a Gaussian would reduce aliasing of the retinal sampling. A circularly symmetric un-normalized two-dimensional Gaussian support $G(x,y)$ with standard deviation σ is as follows,

$$G(x,y,\sigma,X,Y) = \exp\left(-\frac{(x-X)^2+(y-Y)^2}{2\sigma^2}\right) \quad (4)$$

(X, Y) are the horizontal and vertical sub-pixel offsets in placing the retinal receptive field on the integer locations of the image. The receptive field is therefore placed on the digital image with sub-pixel precision. All the input weights of the retinal neuron's simple receptive field were normalized to sum to unity. The standard deviation σ of the support region was chosen such that

$$\sigma = d/6 \quad (5)$$

where d is the diameter of the receptive field. It can be argued that the size of the receptive field is related to local node density. Receptive fields in the periphery of the retina have to be much larger than those in the fovea. The following was used to determine the receptive field diameter d_i of retinal neuron i

$$d_i = \frac{s}{k} \sum_{j=2}^k A_{i,j} \quad (6)$$

where $A_{i,j}$ is the sorted Euclidean distance matrix of the retinal tessellation, k is the neighbourhood size for determining node density and s is a scaling constant. A value of $s = 4$ was used to create overlap of receptive fields and suppress aliasing.

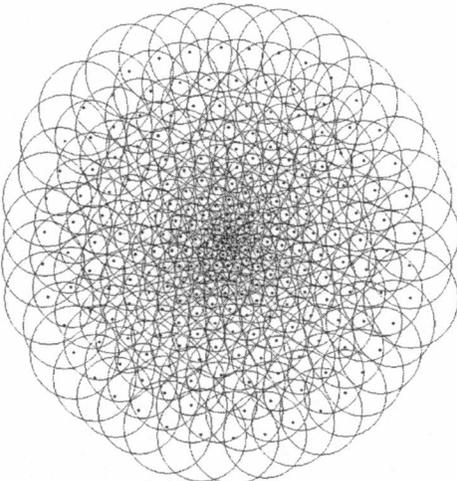


Fig. 3. Overlapping receptive fields on a 1024 node retina tessellation. The indicated receptive field area subsumes 95% of its support region.

The response of the receptive field was generated by multiplying the underlying image pixels with coefficients of the receptive field $G(x,y,\sigma,X,Y)$. As the implemented retina was not based on a retino-cortical transform there isn't a coherent cortical space structure to store and represent the retinal responses. However it is useful to represent retinotopic relationships as information extracted in adjacent areas in the field of view tend to be processed together in higher level feature extraction operations. In this paper a graph structure (Fig. 4) extending the work by Wallace et. al. [10] was used to define higher level receptive fields in the processing hierarchy.

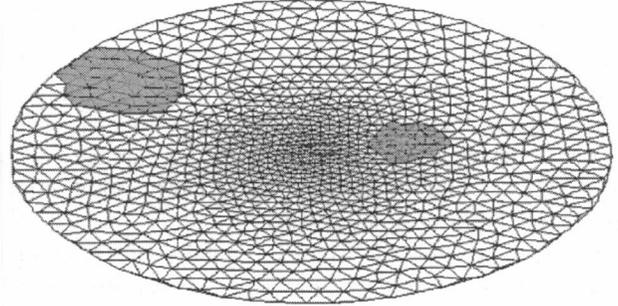


Fig. 4. Cortical graph for a retina with 1024 nodes generated using Delaunay triangulation. A parafoveal and a peripheral receptive field are indicated.

The response of a higher order receptive field in a feature extraction hierarchy $O(c)$ on node c on a layer with N neurons is given by

$$O(c) = \sum_i R(i) W_c(v_i) \text{ where } 1 \leq i \leq M, 1 \leq c \leq N \quad (7)$$

Here $R(i)$ is the response of the receptive field on node v_i of the immediately lower layer in the feature extraction hierarchy with M nodes and W_c is the coefficient of the higher layer receptive field for the afferent from lower layer node v_i .

The Laplacian of Gaussian (LoG) [11] is of interest in biologically inspired computer vision due to its resemblance to the profile of biological retinal ganglion cell receptive fields. As a LoG filter is a spatial frequency band pass filter, the Gaussian responses from the artificial retina can be processed by neurons with a LoG support profile to detect contrast in the extracted visual information. The following equation was used to compute the receptive field coefficients $W_c(v_i)$ for a LoG neuron.

$$L(r,\sigma) = \frac{(r^2 - 2\sigma^2)}{2\pi\sigma^6} e^{-\frac{r^2}{2\sigma^2}} \quad (8)$$

Here r is the retinotopic distance to an afferent node from the LoG neuron's centre. The response of the neuron is computed using Eqn. 7.

V. RETINA PYRAMID

Although the space-variant density of the retinal tessellation results in higher spatial acuity in fovea than in the periphery of the field-of-view, a single LoG neuron layer is restricted to extracting visual information at a single narrow band-pass frequency range from a particular location in the scene from a

single retinal fixation. Since the LoG is a band-pass filter, in this paper Gaussian neuronal sampling is used for multi-resolution feature extraction. While it is possible to self-organise not only Gaussian neuronal retinal layer positions but also multi-resolution support regions, extracting higher level features from the responses of the resulting self-organised multi-scale sampling mechanism is not arbitrary. In this paper the standard computer vision Gaussian pyramid [12] was extended to be able to operate on an retina-based irregular information representation. Only the finest scale retina layer actually samples image. All others sub-sample the immediately finer Gaussian retina layer. Each sub-sample operation of the Gaussian retina pyramid necessitates the pre-computing of unique receptive field coefficients over an irregular representation as indicated in Section IV (in contrast to the regular rectilinear arrays normally found in computer vision). However the computational advantage of sub-sampling visual information in a Gaussian pyramid of retinae far exceeds this processing penalty.

The pyramid structure, which the authors refer to as a retina pyramid, has layers which approximate octave separation. The Gaussian retina pyramid illustrated in Fig. 5 contains 8192, 4096, 1024 and 256 nodes.

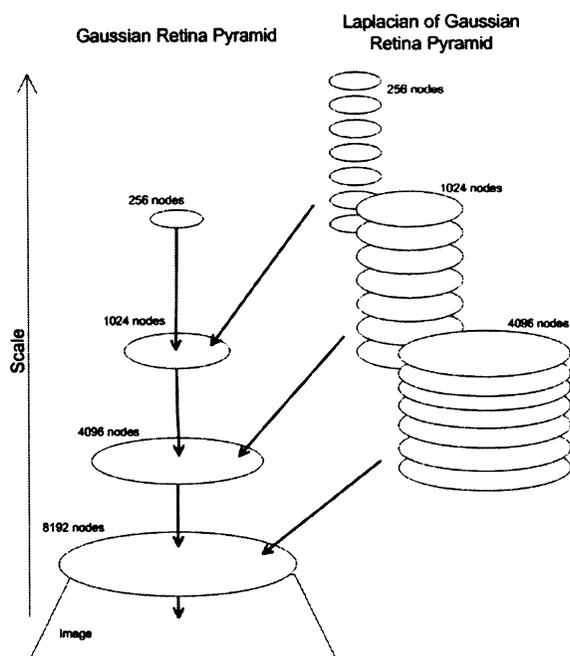


Fig. 5. Gaussian and associated Laplacian of Gaussian retina pyramid. Dark arrows indicate sampling or feature extraction.

Instead of processing each Gaussian retina layer with a single LoG retina layer, the granularity of the LoG sampling across scale was increased by sampling the Gaussian efferents using LoG retina layers with receptive fields having different support region standard deviations σ_n . The standard deviation of the n^{th} LoG layer in an octave 'stack' is given by

$$\sigma_n = \sigma_o s^{n-1} \quad (9)$$

where σ_o is the standard deviation of the finest LoG layer in the octave and s is the increase in standard deviation of the

LoG receptive fields between layers.

VI. LOG SCALE-SPACE EXTREMA

The identification of stable interest points is an important step in modern machine vision. The local visual information extracted at interest point locations is more invariant and robust than global information extracted from the whole object or scene [13]. This reduces the ambiguity and increases the discrimination of the vision system in performing tasks such as recognition and visual search. The detection of interest points based on discrete LoG scale-space extrema enables the extraction of features not only at stable salient spatial locations in the image but also at the characteristic scale of a particular salient feature [14, 15].

The LoG layers within each octave of the LoG retina pyramid detect contrast, yet are not suitable for the detection of discrete scale-space extrema because the response amplitude of a LoG generally decreases with scale. Lindeberg [16] and Mikolajczyk [14] showed that the scale normalised derivative D of order m centred on v_c can be given by

$$D_m(v_c, \sigma) = \sigma^m L_m(v_c, \sigma) \quad (10)$$

The σ^m term helps to somewhat normalise the image derivative's L_m response to scale ($m=2$ for LoG). In our experiments the authors found that even this normalisation was not sufficient to generate scale space extrema equitably on LoG layers within an octave of the retina pyramid. Therefore each (unique) LoG receptive field in the pyramid was independently normalised by the response of that receptive field when the vision front-end was fixated upon random dot stimuli. If \bar{D}_2 contains the mean response to random dot stimuli, the normalised responses \tilde{D}_2 is given by

$$\tilde{D}_2(v_c, \sigma) = \frac{D_2(v_c, \sigma)}{\bar{D}_2(v_c, \sigma)} \quad (11)$$

Using the normalised LoG responses \tilde{D}_2 , it is possible to detect extrema in discrete LoG scale-space. A classification of a LoG receptive field (v_c, σ^j) in the pyramid centred at v_c and LoG standard deviation spread σ^j is given by the following

$$(v_c, \sigma^j) = \begin{cases} \text{extrema} : (v_c, \sigma^j) > \forall (v_k, \sigma^j), j = i-1 \dots i+1 : v_k \in \mathbb{N}(v_c) \\ \text{extrema} : (v_c, \sigma^j) < \forall (v_k, \sigma^j), j = i-1 \dots i+1 : v_k \in \mathbb{N}(v_c) \\ \text{not extrema} : \text{otherwise} \end{cases} \quad (12)$$

where node v_k is a neighbour of node v_c in space and scale(σ_n) in the LoG retina pyramid.

Since for the detection of extrema in LoG responses across frequency requires an additional two LoG retina layers, the scaling factor s (Eqn. 9) for a LoG retina pyramid with n layers per octave is

$$s = 2^{1/n-2} \quad (13)$$

The interest points detected along an edge in the image are not well localised because of the ambiguity caused by the similarity of the interest point feature descriptor along the edge. A variant of the Harris corner detector [17] was used to remove interest points detected at image edges. The following

with a value of $r=10$ [15] was used

$$\frac{(D_{xx} + D_{yy})^2}{D_{xx}D_{yy} - D_{xy}D_{yx}} < \frac{(r+1)^2}{r} \quad (14)$$

The derivatives were calculated on the LoG responses of the retina pyramid using the following

$$D_{yy} = \sum \cos\left(\frac{v_{i,y} - v_{e,y}}{v_{i,x} - v_{e,x}} - \frac{\pi}{2}\right) * ((v_i, s') - (v_e, s')) : \forall v_i \in \mathbb{N}(v_e)$$

$$D_{yx} = \sum \cos\left(\frac{v_{i,y} - v_{e,y}}{v_{i,x} - v_{e,x}} + \frac{\pi}{4}\right) * ((v_i, s') - (v_e, s')) : \forall v_i \in \mathbb{N}(v_e)$$

$$D_{xy} = \sum \cos\left(\frac{v_{i,y} - v_{e,y}}{v_{i,x} - v_{e,x}} - \frac{\pi}{4}\right) * ((v_i, s') - (v_e, s')) : \forall v_i \in \mathbb{N}(v_e)$$

$$D_{xx} = \sum \cos\left(\frac{v_{i,y} - v_{e,y}}{v_{i,x} - v_{e,x}}\right) * ((v_i, s') - (v_e, s')) : \forall v_i \in \mathbb{N}(v_e) \quad (15)$$

VII. RESULTS

The following figure contains the responses from the Gaussian retina pyramid. As the responses are internally represented as one-dimensional vectors the authors chose to visualise the responses by reversing the entire sub-sampling process and back-propagating the responses back to retina (image) space.

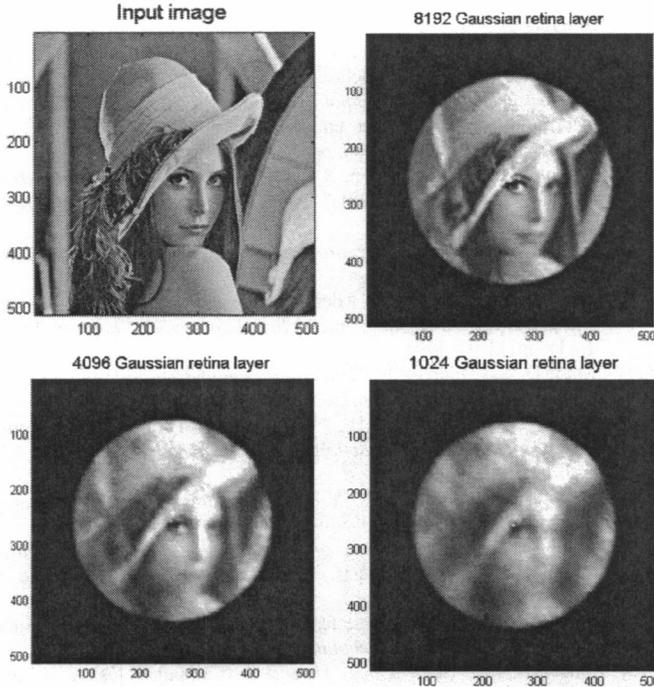


Fig. 6. Input greyscale Lena image and back-propagated responses from the octave separated Gaussian retina pyramid. Even though filter/sub-sample operations have occurred on an irregular sampling tessellation, Gaussian blurring of the visual stimulus can be observed. The space-variant nature of the processing is apparent with the region near Lena's right eye being sampled at a higher spatial frequency than more peripheral regions of the image.

The back-propagation operation resembles the sub-sampling operation over the pseudo-random retina tessellation. If $O(c)$ contains the receptive field response at node v_c in a layer, the

back-propagated response at node v_i in the immediately lower level in the hierarchy is given by the following

$$R(i) = O(c) W_c(v_i) \quad (16)$$

The responses from the LoG retina pyramid were similarly visualised by back-projecting back to the image plane (Fig. 7).

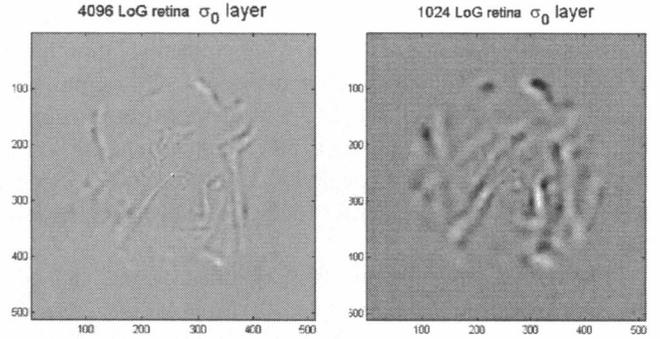


Fig. 7. Back-propagated LoG responses. LoG layer responses were back-propagated through the Gaussian retina pyramid back to the image plane. The band-pass filtering of the LoG layer neurons that extract contrast information can be observed in the back-projected LoG responses. The contrast information extracted near the point of fixation on Lena's right eye has a higher spatial frequency than that extracted from more peripheral regions in the field-of-view.

Table 1
Number of LoG discrete scale-space extrema in the retina pyramid.

	4096 octave LoG layers					1024 octave LoG layers					256 octave LoG layers				
Scaling of support region (Eqn. 9)	s^1	s^2	s^3	s^4	s^5	s^1	s^2	s^3	s^4	s^5	s^1	s^2	s^3	s^4	s^5
No. of $\sigma^2 \nabla^2 G$ scale-space extrema	18	0	0	0	0	0	0	0	0	0	0	0	1	0	0
No. of extrema after normalization	55	52	38	31	34	19	11	13	10	6	1	7	5	4	1

Table 1 contains the number of scale-space extrema detected at LoG layers in the retina pyramid within each octave. In this experiment each octave had seven LoG layers ($n=7$ in Eqn. 9). Not normalising the LoG response results in most extrema being detected in the coarser layers of each octave of the retina pyramid while normalising with the LoG response to random stimuli (\mathcal{D}_2) results in extrema being detected more equitably across scale within the octave.

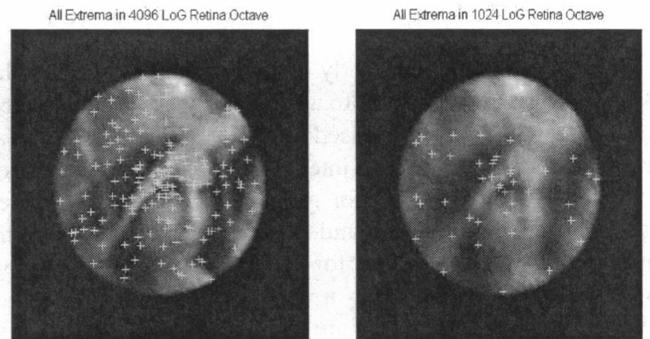


Fig. 8. LoG scale-space extrema (after corner detection). The extrema are displayed on the back-propagated Gaussian layer responses.

VIII. DISCUSSION

This paper showed that multi-resolution hierarchical feature extraction and interest point detection is possible with visual information extracted using a biologically-inspired non-uniform sampling of an artificial retina with pseudo-randomly tessellated receptive fields. While there is previous work on filtering hexagonal grids [18] and space-variant image processing [10], this paper addressed the problem of filtering space-variant, irregular tessellations in multi-resolution feature extraction hierarchies. Gaussian filtering / sub-sampling and Laplacian of Gaussian filtering was demonstrated using an efficient pyramidal approach to multi-resolution visual processing. Discrete LoG scale-space extrema were detected and evaluated with a corner detector to generate interest points.

The retinal tessellation that resulted from the self-organisation in the paper has a locally irregular receptive field distribution. In other signal processing domains [19], an irregular, non-uniform sampling has been used to increase frequency bandwidth and used for aliasing-free digitisation of analogue signals.

By foregoing geometric regularity of the retinal mosaic the approach achieves continuity in sampling density throughout the retina – from the uniform fovea to the extremity of the coarse periphery. The structure of the retina locally resembles a hexagonal lattice with occasional deviations in the hexagonal topology in locations where there is a transition between the dominant influences on the network. These deviations enable the retina tessellation to maintain a sampling density continuum at a macroscopic level with the retina's uniform foveal region seamlessly coalescing into a space-variant periphery. It has been found that a hexagonal lattice is the most efficient way to tile a plane [20] and has useful properties such as equal distance to neighbours and being an approximation to a circular tiling.

Using scale normalised LoG receptive fields was not sufficient to find LoG scale-space extrema that were equitably distributed among the LoG retina layers in each octave. Each LoG receptive field in the retina pyramid had to be normalised by its response to random stimuli for extrema to be equitably distributed among the retina pyramid layers generating interesting points useful for higher level reasoning.

IX. FUTURE WORK

The authors are currently working on extending the described vision front-end into a complete space-variant vision system capable of task-based attention behaviour. Visual information extracted at the interest points in the periphery of the field-of-view are used for generating object hypotheses for an attention mechanism and that from interest points in the high acuity foveal region, for direct task-based reasoning. Space-variant vision using a fixed retina-based sampling mechanism evolved in nature to reason with a dynamic changing visual environment. Animal eyes are bombarded with an ever changing visual environment as the animal navigates in its surroundings or the environment itself

changes. The authors believe that foveated, space-variant vision evolved for reasoning in a dynamic visual environment, i.e. video. The foveated, space-variant vision strategy reduces the combinatorial explosion of visual information while the fixed retina structure is serially targeted at the most salient location in the dynamic scene depending on the task at hand. The authors believe that the work presented in this paper is a useful tool for the future investigation of foveated space-variant vision.

ACKNOWLEDGMENT

This work is supported by the European Union 6th Framework Program project "Integrated Project Research Area CINE" Project ref: IST-2-511316-IP

REFERENCES

- [1] Curcio, C.A., Sloan, K.R., Kalina, R.E., and Hendrickson, A.E., "Human photoreceptor topography", *Journal of Comparative Neurology*, 292: pp. 497-523, 1990.
- [2] Schwartz, E.L., "Spatial mapping in primate sensory projection: Analytic structure and relevance to perception", *Biological Cybernetics*, 25: pp. 181-194, 1977.
- [3] Schwartz, E.L., "Computational Anatomy and functional architecture of the striate cortex", *Vision Research*, 20: pp. 645-669, 1980.
- [4] Johnston, A., "The geometry of the topographic map in striate cortex", *Vision Research*, 29: pp. 1493-1500, 1989.
- [5] Wilson, S.W., "On the retino-cortical mapping", *International Journal of Man-Machine Studies*, 18(4): pp. 361-389, 1983.
- [6] Gomes, H., *Model Learning in Iconic Vision*. PhD Thesis, University of Edinburgh. 2002.
- [7] van der Spiegel, J., Kreider, G., Claeys, C., Debusschere, I., Sandini, G., Dario, P., Fantini, F., Belluti, P., and Soncini, G., *A foveated retina-like sensor using CCD technology*, in *Analog VLSI implementation of neural systems*, C. Mead and M. Ismail, Editors. 1989, Kluwer Academic Publishers: Boston. pp. 189-212.
- [8] Clippingdale, S. and Wilson, R., "Self-similar Neural Networks Based on a Kohonen Learning Rule", *Neural Networks*, 9(5): pp. 747-763, 1996.
- [9] Kohonen, T., *Self-organization and associative memory*, Berlin: Springer-Verlag, 1984.
- [10] Wallace, R.S., Ong, P.W., Bederson, B.B., and Schwartz, E.L., "Space-Variant Image-Processing", *International Journal of Computer Vision*, 13(1): pp. 71-90, 1994.
- [11] Marr, D. and Hildreth, E., "Theory of edge detection", *Proceedings of the Royal Society of London*, B(207): pp. 187-217, 1980.
- [12] Burt, P.J. and Adelson, E.H., "The Laplacian Pyramid as a Compact Image Code", *IEEE Transactions on Communications*, 31(4): pp. 532-540, 1983.
- [13] Schmid, C. and Mohr, R., "Local Grayvalue Invariants for Image Retrieval", *PAMI*, 19(5): pp. 530-535, 1997.
- [14] Mikołajczyk, K., *Detection of local features invariant to affine transformations*, PhD Thesis, Institute National Polytechnique de Grenoble, France, 2002.
- [15] Lowe, D., "Distinctive image features from scale-invariant keypoints", *International Journal of Computer Vision*, 60(2): pp. 91-110, 2004.
- [16] Lindeberg, T., *Scale-Space Theory in Computer Vision*, Kluwer Academic Publishers, 1994.
- [17] Harris, C. and Stephens, M., "A Combined Corner and Edge Detector", *Proceedings of The Fourth Alvey Vision Conference*, 1988.
- [18] Simoncelli, E.P. and Adelson, E.H., "Non-separable extensions of quadrature mirror filters to multiple dimensions", *Proceedings of the IEEE*, 78(4): pp. 652-664, 1990.
- [19] Bilinskis, I. and Mikelsons, A., *Randomized Signal Processing*, Prentice Hall International, 1992.
- [20] Hales, T.C., "The Honeycomb Conjecture", *Discrete Computational Geometry*, 25: pp. 1-22, 2001.