The Role of Non-CRF Inhibition in Contour Detection

Cosmin Grigorescu
Univ. of Groningen
P.O. Box 800, 9700 AV
Groningen, The Netherlands
cosmin@cs.rug.nl

Nicolai Petkov
Univ. of Groningen
P.O. Box 800, 9700 AV
Groningen, The Netherlands
petkov@cs.rug.nl

Michel A. Westenberg
Univ. of Groningen
P.O. Box 800, 9700 AV
Groningen, The Netherlands
michel@cs.rug.nl

ABSTRACT

We propose a biologically motivated computational step, called non-classical receptive field (non-CRF) inhibition, to improve the performance of contour detectors. Non-CRF inhibition is exhibited by 80% of the orientation selective neurons in the primary visual cortex of macaque monkeys and has been demonstrated to influence the visual perception of man as well. We introduce an image processing operator, the *bar cell operator*, which consists of a Gabor energy operator augmented with non-CRF inhibition. This operator responds strongly to isolated lines, edges and contours, but exhibits a weaker or no response to edges that make part of texture. We evaluate the contour detection performance of the proposed operator for images of natural scenes with associated ground truth edge maps. The bar cell operator consistently outperforms the Canny edge detector, mostly due to a reduced number of false positives.

Keywords

contour detection, edge detection, Canny operator, Gabor energy, non-classical receptive field inhibition, texture.

1 INTRODUCTION

An important finding in the neurophysiology of the visual system of monkeys and cats, made in the beginning of the 1960s, was that the majority of neurons in the primary visual cortex act as edge detectors. A typical neuron from this cortical area will respond vigorously to an edge or a line of a given orientation and position in the visual field. In the following years, the study of the properties of such neurons has been an active area of research and, in 1981, a Nobel prize for medicine and physiology was awarded to D. Hubel and T. Wiesel who pioneered this work [Hub82]. Computational models for two types of orientation selective cells, called the simple cell [AP79, MTT78b] and the complex cell [MTT78a, SH85], gave the basis for biologically motivated edge detection algorithms in im-

age processing and computer vision. In particular, a family of two-dimensional Gabor functions was proposed as a model of the receptive fields of simple cells [Dau85, JP87].

Later neurophysiological research revealed a considerable functional diversity in the rather broad class of orientation selective cells. Besides the subclasses of simple and complex cells, further subclasses were identified, such as end-stopped cells (originally also called hypercomplex cells [HW68, PvdH91]), contour cells [BvdZP97, PH01], and grating cells [vdHPD91, vdHPD92]. The computational models for all these subclasses assumed that the only condition for a cell to elicit a vigorous response is that the appropriate stimulus be present within a specific region of the visual field. This region is presently referred to as the classical receptive field (CRF).

Detailed studies have unveiled, however, that the behaviour of orientation selective cells is more complex than suggested by these early computational models. In particular, measurements have shown that once a cell is activated by a stimulus in its CRF, another, simultaneously presented stimulus outside that field can have an effect on the cell response. This, mostly inhibitive effect, is referred to as *non-classical receptive*

field (non-CRF) inhibition and is exhibited to a different extent by 80% of the orientation selective cells [KvE92, JGWS01].

For instance, Nothdurft et al. [NGvE99] measured the response of neural cells when different texture surrounds were present outside the CRF. First, they mapped the CRF of a cell by determining the optimal orientation position and size of bar which made the cell elicit a strong response. Then, they placed texture consisting of oriented elements of the same type as the optimal stimulus in the area outside the CRF. Approximately one-third of the orientation selective cells exhibited an inhibition effect caused by the surrounding texture irrespective to the orientation of the surrounding texture elements, see Fig. 1. In general, an orientation selective cell with non-CRF inhibition will respond most strongly to a single bar, line, or edge in its receptive field and will show reduced response with the addition of further bars to the surrounding. The response decreases with the distance from the CRF. Neural cells which show this behaviour were called *bar cells* [vdHPD92] and a computational model was proposed for them in [PK97].
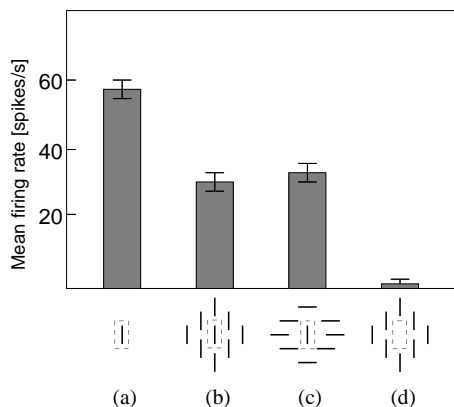
Figure 1: *(a) The response of a neuron to a stimulus composed of a single bar of optimal size and orientation inside the CRF (dotted square). A decreased resonse is recorded when texture consisting of identical bars is present in the area outside the CRF: the surrounding bars have the same orientation (b) and orthogonal orientation (c) relative to the optimal stimulus. (d) In absence of the optimal stimulus, the response is reduced to the level of spontaneous activity (Courtesy of C. Nothdurft and Visual Neuroscience).*

The above mentioned neurophysiological behaviour of bar cells correlates well with the results of various psychophysical experiments, which have shown that the perception of an oriented stimulus, such as a line, can be influenced by the presence of other such stimuli (distractors) in its neighbourhood. This influence can, for example, manifest itself in a decreased saliency of groups of parallel lines [Kan79], Fig. 2.
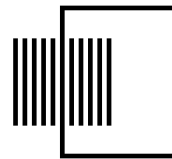
Figure 2: *Reduced perceptual saliency of a line when embedded in other parallel lines: the left section of the rectangle is "lost" in the surrounding grating.*

As non-CRF inhibition seems to be a common property of orientation selective neurons, and proves to play a significant role in our perception of edges and lines, we considered that a more close examination of this mechanism for edge detection in image processing and computer vision is worthwhile. Our main hypothesis is that non-CRF inhibition suppresses edges which make part of texture, while it does not suppress edges that belong to the contours of objects. An edge detection algorithm which employs this inhibition mechanism will thus primarily detect contours of objects, and it will not react to edges which belong to texture regions. The edge maps generated by such an edge detector will be more useful for contour-based object recognition tasks, such as shape comparison [GP02], than traditional edge detectors which do not make a difference between contour and texture edges.

The paper is organized as follows. Section 2 describes the computational model. The simple cell and complex cell models and the related Gabor and Gabor energy filters are briefly discussed, and an operator which models non-CRF inhibition, the *bar cell operator*, is introduced. In Section 3, we evaluate the performance of the bar cell operator. A suitable performance measure is introduced, and experimental results obtained with the bar cell operator and the Canny edge detector are compared. Finally, we summarize the results and draw conclusions in Section 4.

2 COMPUTATIONAL MODEL

2.1 Simple Cells and Gabor Filters

The spatial summation properties of simple cells can be modelled by a family of two-dimensional Gabor functions [Dau85]. We use a modified parameterization to take into account restrictions found in experimental data [PK97]. A receptive field function of such a cell, in engineering terms the impulse response,

$g_{\lambda,\sigma,\theta,\varphi}(x,y)$, $(x,y) \in \Omega \subset \mathbb{R}^2$, which is centered around the origin, is given by:

$$g_{\lambda,\sigma,\theta,\varphi}(x,y) = e^{-\frac{\tilde{x}^2+\gamma^2\tilde{y}^2}{2\sigma^2}} \cos(2\pi\frac{\tilde{x}}{\lambda}+\varphi),$$
$$\tilde{x} = x\cos\theta + y\sin\theta, \qquad \tilde{y} = -x\sin\theta + y\cos\theta,$$
(1)

where $\gamma = 0.5$ is a constant, called the spatial aspect ratio, that determines the ellipticity of the receptive field. The standard deviation σ of the Gaussian factor determines the linear size of the receptive field. The parameter λ is the wavelength and $1/\lambda$ the spatial frequency of the cosine factor. The ratio σ/λ determines the spatial frequency bandwidth, and, therefore, the number of parallel excitatory and inhibitory stripe zones which can be observed in the receptive field. In this paper, we fix the value of the ratio σ/λ to $\sigma/\lambda = 0.56$, which corresponds to a half-response bandwidth of one octave. The angle parameter θ, $\theta \in [0,\pi)$, determines the preferred orientation. The parameter φ, $\varphi \in (-\pi,\pi]$, is a phase offset that determines the symmetry of $g_{\lambda,\sigma,\theta,\varphi}(x,y)$ with respect to the origin: for $\varphi = 0$ and $\varphi = \pi$ it is symmetric (or even), and for $\varphi = -\frac{\pi}{2}$ and $\varphi = \frac{\pi}{2}$ it is antisymmetric (or odd); all other cases are asymmetric mixtures.

The response $r_{\lambda,\sigma,\theta,\varphi}(x,y)$ of a simple cell with a receptive field function $g_{\lambda,\sigma,\theta,\varphi}(x,y)$ to an input image with luminance distribution $f(x,y)$ is computed by convolution:

$$r_{\lambda,\sigma,\theta,\varphi}(x,y) = f(x,y) * g_{\lambda,\sigma,\theta,\varphi}(x,y) \qquad (2)$$

In image processing and computer vision, the filter defined by (1) and (2) is known as the (linear) Gabor filter.

2.2 Complex Cells and Gabor Energy Filters

The Gabor energy is related to a model of complex cells which combines the responses of a pair of simple cells with a phase difference of $\frac{\pi}{2}$. The results of a pair of symmetric and antisymmetric filters are combined, yielding the Gabor energy $E_{\lambda,\sigma,\theta}(x,y)$ as follows:

$$E_{\lambda,\sigma,\theta}(x,y) = \sqrt{r_{\lambda,\sigma,\theta,0}^2(x,y) + r_{\lambda,\sigma,\theta,-\frac{\pi}{2}}^2(x,y)},$$
(3)

where $r_{\lambda,\sigma,\theta,0}(x,y)$ and $r_{\lambda,\sigma,\theta,-\frac{\pi}{2}}(x,y)$ are the outputs of a symmetric and an antisymmetric filter, respectively. It can be shown that the Gabor energy is equal to the square root of the local power spectrum of the image [GPKon]. In the following, we will use Gabor energy maps $E_{\lambda,\sigma,\theta_i}(x,y)$ for a number of N_θ different orientations, with θ_i given by $\theta_i = i\,\pi/N_\theta$, $i = 0,1,\ldots,N_\theta - 1$.

2.3 Non-CRF Inhibition

We extend the Gabor energy operator presented above with an inhibition term to qualitatively reproduce the above mentioned non-CRF inhibition behaviour of most orientation selective cells. For a given point in the image, the inhibition term is computed in a ring-formed area surrounding the CRF centered at the concerned point.

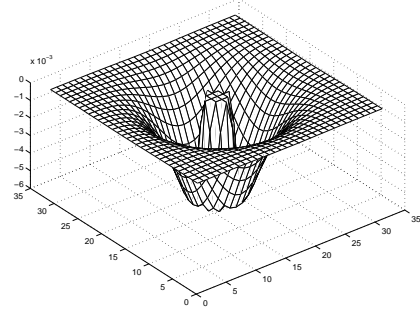Let $DoG_\sigma(x,y)$ be the difference of Gaussians defined by:

$$DoG_\sigma(x,y) = \frac{1}{\sqrt{2\pi(4\sigma)^2}}\,e^{-\frac{x^2+y^2}{2\,(4\sigma)^2}} - \frac{1}{\sqrt{2\pi\sigma^2}}\,e^{-\frac{x^2+y^2}{2\sigma^2}}$$
(4)

We define the weighting function $w_\sigma(x,y)$ as follows:

$$w_\sigma(x,y) = \frac{1}{\|\,H(\,DoG_\sigma\,)\,\|}\,H(\,DoG_\sigma(x,y)\,),$$
$$H(z) = \begin{cases} 0 & z < 0 \\ z & z \geq 0, \end{cases}$$
(5)

where by $\|.\|$ we denote the L_1 norm. The function $H(z)$ ensures that the operator has only positive response. Figure 3 shows the plot of this function.

Figure 3: *Weighting function which models the contribution of the non-CRF surround.*

We model non-CRF inhibition by computing an inhibition term $t_{\lambda,\sigma}(x,y)$. First, we construct an energy map $\widehat{E}_{\lambda,\sigma}(x,y)$ with values of maximum Gabor energy response:

$$\widehat{E}_{\lambda,\sigma}(x,y) = \max\{E_{\lambda,\sigma,\theta_i}(x,y) \mid i = 0,1,\ldots,N_\theta{-}1\},$$
(6)

The inhibition term $t_{\lambda,\sigma}(x,y)$ is computed as a convolution of the maximum energy map $\widehat{E}_{\lambda,\sigma}(x,y)$ with the weighting function $w_\sigma(x,y)$:

$$t_{\lambda,\sigma}(x,y) = \widehat{E}_{\lambda,\sigma}(x,y) * w_\sigma(x,y) \qquad (7)$$

The computation of the above suppression term does not take into account the different orientations for which the maxima of the Gabor energy response

$\widehat{E}_{\lambda,\sigma}(x,y)$ is achieved. Therefore, we refer to this type of inhibition as *isotropic* non-CRF inhibition.

We introduce a new operator $b_{\lambda,\sigma}(x,y)$ which takes as its inputs the maximum energy map $\widehat{E}_{\lambda,\sigma}(x,y)$ and the inhibition term $t_{\lambda,\sigma}(x,y)$:

$$b_{\lambda,\sigma}(x,y) = H(\widehat{E}_{\lambda,\sigma}(x,y) - \alpha t_{\lambda,\sigma}(x,y)), \quad (8)$$

with $H(z)$ defined as in (5). The factor α controls the strength of the inhibition of the surround on the maximum Gabor energy term. If there is no texture in the surrounding of a given point, the response of this operator at that point will be equal to the maximum Gabor energy term. An isolated edge passing through that point will be detected by the introduced operator in the same way as it is detected by the Gabor energy operators. However, if there are other edges in the surrounding, the inhibition term $t_{\lambda,\sigma}(x,y)$ may become so strong that it cancels completely the contribution of the maximum Gabor energy term, resulting in zero response of the operator introduced above. Defined in this way, the concerned operator will respond to isolated lines, edges, and bars, but it will not respond to groups of such stimuli that make part of texture, see Fig. 4(c). We will refer to this operator briefly as the *'bar cell operator'*, in analogy with the function of the type of visual neuron that exhibits a similar behaviour [vdHPD92, PK97].

(a) (b) (c)
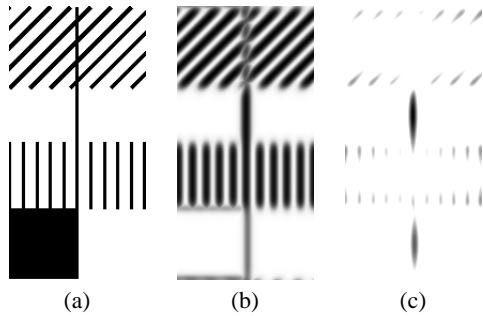
Figure 4: *(a) Synthetic input image. (b) The Gabor energy operator responds to lines and edges independently of the context, i.e., the surrounding in which these lines and edges are embedded. (c) The bar cell operator with isotropic inhibition responds selectively to isolated lines and edges only.*

2.4 Binary Edge Map Construction

We construct binary edge maps from the bar cell response $b_{\lambda,\sigma}(x,y)$ by two post-processing operations called nonmaxima suppression and hysteresis thresholding [Can86, SHB99].

An orientation map $\Theta(x,y)$ with the orientation for which the maximum Gabor energy response $\widehat{E}_{\lambda,\sigma}(x,y)$ is achieved can be computed as follows:

$$\Theta(x,y) = \theta_k, \quad \text{where}$$
$$k = \operatorname{argmax}\{E_{\lambda,\sigma,\theta_i}(x,y) \mid i = 0, 1, \ldots, N_\theta - 1\}. \quad (9)$$

From the orientation map $\Theta(x,y)$ and bar cell response $b_{\lambda,\sigma}(x,y)$, which specify the normal to the local edge direction and the local edge strength, respectively, nonmaxima suppression thins the edges in $b_{\lambda,\sigma}(x,y)$ to one-pixel wide candidate edges. Hysteresis thresholding provides the final binary edge map from the candidate edges by computing two threshold values, t_h and t_l. The first value, t_h, is computed based on percentage p of the candidate edge pixels that should be retained in the final edge map. The low hysteresis threshold value, t_l, is a fraction of t_h; in our experiments, we chose $t_l = 0.5\,t_h$.

We decided to perform the same post-processing operations as in the Canny edge detector [Can86] in order to simplify comparison at a later stage. Prior to post-processing, the Canny edge detection operator computes the gradient magnitude and direction with a scale-dependent differential operator [Td90].

3 PERFORMANCE EVALUATION

Most state-of-the-art methods for evaluation of edge detector performance use natural images with an associated ground truth specified by a subject [BKD01]. Some studies [SGB01] show that the performance of a detector is task-dependent. For a task like object recognition, for example, some detectors may perform better than others despite similar results obtained for synthetic images. The proposed bar cell detector aims explicitly at better detection of object contours in presence of surrounding texture. Therefore, we evaluate its performance for extraction of object contours in natural images rich in textured background. We selected a set of 20 images which depict either man-made objects on textured background or animals in their natural habitat; for each image, an associated ground truth binary edge map was drawn by hand. Figure 5, first and second column, presents a subset of four such images together with their corresponding ground-truth edge maps.

3.1 Performance Measure

Let E_{GT} and B_{GT} be the set of edge pixels and background pixels of the ground truth edge image, respectively, and E_{D} and B_{D} be the set of edge pixels and background pixels of the operator-detected edge image, respectively. The set of correctly detected edge pixels is $E = E_{\mathrm{D}} \cap E_{\mathrm{GT}}$. False negatives, i.e. ground-truth edges missed by the edge detector, are given by

the set $E_{\mathrm{FN}} = E_{\mathrm{GT}} \cap B_{\mathrm{D}}$, while false positives (spurious edges) are given by the set $E_{\mathrm{FP}} = E_{\mathrm{D}} \cap B_{\mathrm{GT}}$.

We define the *performance measure* of an edge detector as:

$$P = \frac{\mathrm{card}(E)}{\mathrm{card}(E) + \mathrm{card}(E_{\mathrm{FP}}) + \mathrm{card}(E_{\mathrm{FN}})}, \quad (10)$$

in which $\mathrm{card}(X)$ denotes the number of elements of set X.

The performance measure P is a scalar taking values in the interval $[0, 1]$. If all true edge pixels are correctly detected and no background pixels are falsely detected as edge pixels, then $P = 1$. For all other cases, the performance measure takes values smaller than 1, being closer to zero as more edge pixels are falsely detected and/or missed by the edge detector operator.

Since edges cannot always be detected at exact integer image coordinates, we consider that an edge pixel is correctly detected if a corresponding ground truth edge pixel is present in a 5×5 square neighborhood centered at the respective pixel coordinates. The false negatives and false positives are determined by eliminating those pixels which are correctly detected from the ground truth edges and detected edges, respectively.

3.2 Experimental results

In the following, we assess the performance of the bar cell contour detector, and compare it with the performance of the Canny edge detector. The Canny edge detector has two parameters: σ, the standard deviation of a Gaussian smoothing kernel, and p, the percentage of candidate edge pixels which are retained in the final edge map. The bar cell contour detector has an additional parameter, α, which is the texture inhibition factor. For the Canny edge detector, we used 8 scales, $\sigma = \{1.0, 1.2, 1.4, 1.6, 1.8, 2.0, 2.2, 2.4\}$. For the bar cell contour detector we used 4 scales covering the same domain, $\sigma \in \{1.2, 1.6, 2.0, 2.4\}$ and 2 texture attenuation factors, $\alpha \in \{1.0, 1.2\}$. For both methods, we used 5 values of the percentage of candidate edge pixels, $p = \{50\%, 40\%, 30\%, 20\%, 10\%\}$. This results in 40 parameter combinations for each of the methods. The number of orientations used by the bar cell model was fixed to $N_\theta = 12$.

For comparable values of σ and p, the edge maps delivered by the bar cell operator had better performance than Canny's edge maps. Figure 5 shows the best performance edge maps obtained for four of our test images for both Canny's edge detector (third column) and bar cell contour detector (fourth column). For a better illustration, we also computed the percentage of false positives e_{fp} as the number of false positives di-

vided by the number of correctly detected edge pixels ($e_{\mathrm{fp}} = \mathrm{card}(E_{\mathrm{FP}}) / \mathrm{card}(E)$), and the percentage of false negatives e_{fn} as the number of false negatives divided by the number of ground truth edge pixels ($e_{\mathrm{fn}} = \mathrm{card}(E_{\mathrm{FN}}) / \mathrm{card}(GT)$). The performance measures, parameters, and percentages of false positives and false negatives are displayed below each image.
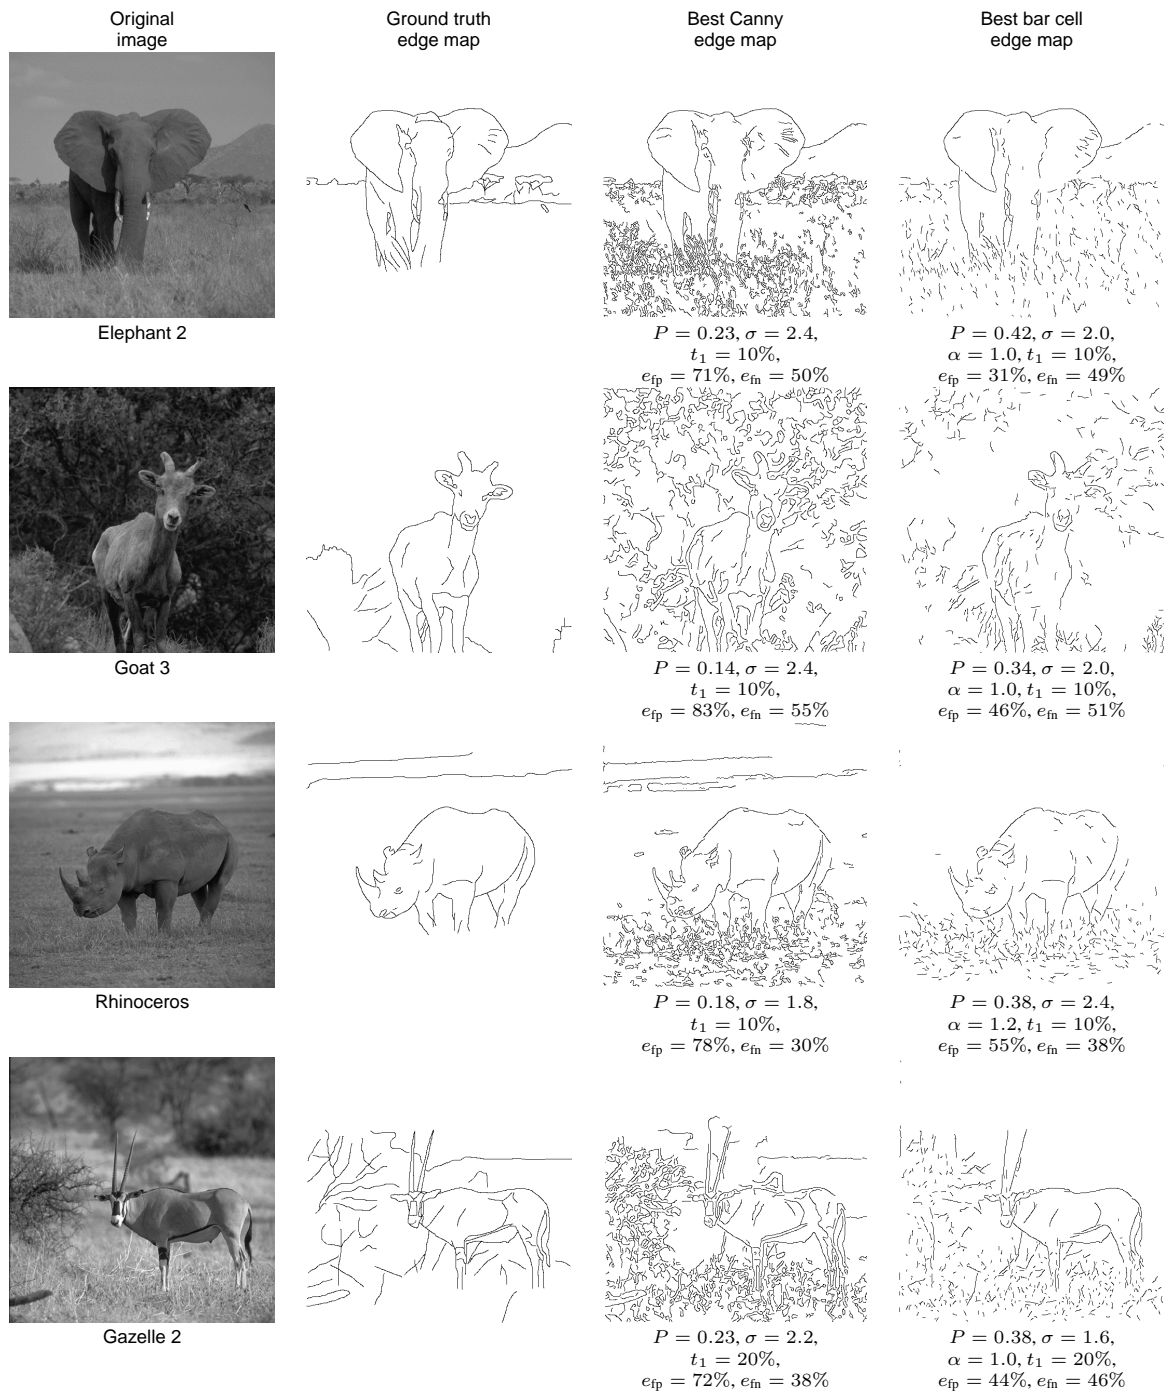
These results show that, indeed, the bar cell contour detector suppresses edges in the presence of surrounding texture. In some cases, such as the image on the second row ("Goat 3"), the Canny edge map contains so many spurious edges that it is hard to distinguish between the contours of the object and the other edges. In contrast, the object contours in the bar cell edge map can be easily recognized. For all twenty images used in our experiment the best performance measure is consistently higher for the bar cell contour detector, and this is mostly due to a reduced percentage of false positives.

4 SUMMARY AND CONCLUSIONS

The non-CRF inhibition algorithm presented in this paper treats classes of edges and lines in two different ways: single contour lines and edges, on one hand, being considered as non-texture features, are not effected by the inhibition, while groups of lines and edges, on the other hand, viewed as texture features, are suppressed. As already noted in the introduction, this different treatment correlates well with our visual perception.

The model of non-CRF inhibition we use in this study is simple and straightforward. The response of an orientation and scale specific operator in a given position is suppressed by the responses of the same operator in other neighboring positions. Our model makes use of a single parameter, α, the coefficient with which the weighted summation inhibition term is taken into account. The value of this parameter can be determined in an optimization problem derived from a specific goal, e.g. maximization of the performance of the operator for a certain set of images.

Inhibition mechanisms have been applied previously to biologically motivated edge detectors in order to improve certain aspects of their function. A symmetric Gabor filter, will, for instance, respond not only along a line but also alongside the line at a certain distance from it. Similarly, the largest response of an antisymmetric Gabor filter to a line will be displaced from the line. In [Hei95, PKL93], various inhibition mechanisms have been proposed to remove these flanking responses. These works differ from the current work in two major aspects. First, the inhibition mechanisms act within the CRF. Second, the purpose of the inhibition is quite different: it deals with the removal of

Original image	Ground truth edge map	Best Canny edge map	Best bar cell edge map
Elephant 2		$P = 0.23, \sigma = 2.4,$ $t_1 = 10\%,$ $e_{\text{fp}} = 71\%, e_{\text{fn}} = 50\%$	$P = 0.42, \sigma = 2.0,$ $\alpha = 1.0, t_1 = 10\%,$ $e_{\text{fp}} = 31\%, e_{\text{fn}} = 49\%$
Goat 3		$P = 0.14, \sigma = 2.4,$ $t_1 = 10\%,$ $e_{\text{fp}} = 83\%, e_{\text{fn}} = 55\%$	$P = 0.34, \sigma = 2.0,$ $\alpha = 1.0, t_1 = 10\%,$ $e_{\text{fp}} = 46\%, e_{\text{fn}} = 51\%$
Rhinoceros		$P = 0.18, \sigma = 1.8,$ $t_1 = 10\%,$ $e_{\text{fp}} = 78\%, e_{\text{fn}} = 30\%$	$P = 0.38, \sigma = 2.4,$ $\alpha = 1.2, t_1 = 10\%,$ $e_{\text{fp}} = 55\%, e_{\text{fn}} = 38\%$
Gazelle 2		$P = 0.23, \sigma = 2.2,$ $t_1 = 20\%,$ $e_{\text{fp}} = 72\%, e_{\text{fn}} = 38\%$	$P = 0.38, \sigma = 1.6,$ $\alpha = 1.0, t_1 = 20\%,$ $e_{\text{fp}} = 44\%, e_{\text{fn}} = 46\%$

Figure 5: *Natural scenes with objects on textured background (first column), their corresponding ground truth edge maps (second column), the best edge maps obtained with the Canny edge detector (third column), and the best edge maps obtained with the bar cell contour detector (last column).*

flanking responses, rather than with the suppression of texture edges.

In this paper we model only the isotropic inhibitory behaviour of bar cells. However, there is physiological evidence that anisotropic inhibition, i.e. edges of the same orientation as the main stimulus have stronger suppression effect than edges of different orientations, is equally important: approximately one third of the cells with non-CRF modulation exhibit isotropic inhibition while another third are accounted for anisotropic inhibition [NGvE99]. Elsewhere [GPW02] we studied the effect of anisotropic inhibition as well. For natural images like those shown in Figure 5 we obtain slightly better results using the isotropic inhibition.

The surround inhibition can be incorporated as an additional processing step in most edge detection operators. More specifically, it can be added to the Canny edge detector as an intermediate step between the gradient computation and edge thinning and binarization. This inhibition step may be expected to improve contour detection performance in images that contain objects of interest on a cluttered or textured background.

5   REFERENCES

[AP79] B. W. Andrews and D. A. Pollen. Relationship between spatial frequency selectivity and receptive field profile of simple cells. *J. Physiology (London)*, 287:163–176, 1979.

[BKD01] K. Bowyer, C. Kranenburg, and A. Dougherty. Edge detector evaluation using empirical ROC curves. *Computer Vision and Image Understanding*, 84:77–103, 2001.

[BvdZP97] R. Baumann, R. van der Zwan, and E. Peterhans. Figure-ground segregation at contours. *European J. Neuroscience*, 9:1290–1303, 1997.

[Can86] J. F. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.

[Dau85] J. G. Daugman. Uncertainty relations for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *J. Optical Society of America A*, 2:1160–1169, 1985.

[GP02] C. Grigorescu and N. Petkov. Distance sets for shape filters and shape recognition. Technical Report IWI-2002-9-02, Institute of Mathematics and Computing Science, University of Groningen, 2002.

[GPKon] S. E. Grigorescu, N. Petkov, and P. Kruizinga. Comparison of texture features based on Gabor filters. *IEEE Trans. Image Processing*, accepted for publication.

[GPW02] C. Grigorescu, N. Petkov, and M. A. Westenberg. Improved contour detection by non-classical receptive field inhibition. *2-nd Workshop on Biologically Motivated Computer Vision (BMCV'2002), Tüebingen, Germany, to appear in Lecture Notes in Computer Science LNCS-2525*, 2002.

[Hei95] F. Heitger. Feature detection using suppression and enhancement. Technical Report TR-163, Communication Technology Laboratory, Swiss Federal Institute of Technology, 1995.

[Hub82] D. H. Hubel. Explorations of the primary visual cortex, 1955-1978. *Nature*, 299:515–524, 1982. (1981 Nobel Prize lecture).

[HW68] D. H. Hubel and T. N. Wiesel. Receptive fields and functional architecture of monkey striate cortex. *J. Physiology (London)*, 195:215–243, 1968.

[JGWS01] H. E. Jones, K. L. Grieve, W. Wang, and A. M. Sillito. Surround suppression in primate V1. *J. Neurophysiology*, 86(10):2011–2028, 2001.

[JP87] J. P. Jones and L. A. Palmer. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiology*, 58:1233–1258, 1987.

[Kan79] G. Kanizsa. *Organization in Vision, Essays on Gestalt Perception*. Praeger, New York, 1979.

[KvE92] J. J. Knierim and D. C. van Essen. Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiology*, 67:961–980, 1992.

[MTT78a] J. A. Movshon, I. D. Thompson, and D. J. Tolhurst. Receptive field organisation of complex cells in the cat's striate cortex. *J. Physiology (London)*, 283:79–99, 1978.

[MTT78b] J. A. Movshon, I. D. Thompson, and D. J. Tolhurst. Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J. Physiology (London)*, 283:53–77, 1978.

[NGvE99] H.-C. Nothdurft, J. Gallant, and D. van Essen. Response modulation by texture surround in primate area V1: Correlates of "popout" under unesthesia. *Visual Neuroscience*, 16(1):15–34, 1999.

[PH01] E. Peterhans and F. Heitger. Simulation of neural responses defining depth order and contrast polarity at illusory contours in monkey area V2. *J. Computational Neuroscience*, 10(2):195–211, 2001.

[PK97] N. Petkov and P. Kruizinga. Computational models of visual neurons specialised in the detection of periodic and aperiodic oriented visual stimuli: bar and grating cells. *Biological Cybernetics*, 76(2):83–96, 1997.

[PKL93] N. Petkov, P. Kruizinga, and T. Lourens. Orientation competition in cortical filters – an application to face recognition. In *Computing Science in The Netherlands 1993*, pages 285–296, Amsterdam, 1993. Stichting Mathematisch Centrum.

[PvdH91] E. Peterhans and R. von der Heydt. Elements of form perception in monkey prestriate cortex. In A. Gorea, Y. Fregnac, Z. Kapoulis, and J. Findlay, editors, *Representations of Vision. Trends and Tacit Assumptions*, pages 111–124. Cambridge Univ. Press, Cambridge, 1991.

[SGB01] M. C. Shin, D. B. Goldgof, and K. W. Bowyer. Comparison of edge detectors using an object recognition task. *Computer Vision and Image Understanding*, 84(1):160–178, 2001.

[SH85] H. Spitzer and S. Hochstein. A complex cell receptive field model. *J. Neurophysiology*, 53:1266–1286, 1985.

[SHB99] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis, and Machine Vision*. Brooks/Cole Publishing Company, Pacific Grove, CA, USA, 1999.

[Td90] H. Tagare and R. deFigueiredo. On the localization performance measure and optimal edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(12):1186–1190, 1990.

[vdHPD91] R. von der Heydt, E. Peterhans, and M. R. Dürsteler. Grating cells in monkey visual cortex: Coding texture. In B. Blum, editor, *Channels in the Visual Nervous System: Neurophysiology, Psychophysics and Models*, pages 53–73. Freund Publ. House Ltd., London, 1991.

[vdHPD92] R. von der Heydt, E. Peterhans, and M. R. Dürsteler. Periodic-pattern-selective cells in monkey visual cortex. *J. Neuroscience*, 12:1416–1434, 1992.