SHARPENING ORIENTATION SELECTIVITY FOR EFFICIENT IMAGE FILTERING

Jeffrey Ng and Anil A. Bharath

Faculty of Engineering, Imperial College London
South Kensington, London, UK SW7 2AZ
{jeffrey.ng, a.bharath}@imperial.ac.uk

ABSTRACT

The steering framework significantly reduces the computational cost of detecting and extracting oriented low-level spatial features in images by linearly combining the responses of a small set of basis filters at fixed orientations to yield filter responses at arbitrary orientations. However, the characteristics of the steered filter, such as orientation and scale selectivity, remain the same as those of the basis filters. We show that linear combination with the appropriate basis filters and weights can achieve more. We present a novel approach for sharpening the orientation selectivity of polar-separable filters by linear combination and well-known trigonometric identities which reduces the need for more basis filters with higher orientation selectivity. We describe the principles for and issues with designing quadrature sharpening filters. We also qualitatively and quantitatively compare 4-basis sharpened filter kernels with an equivalent steerable set of 6-basis filters.

1. INTRODUCTION

The detection and extraction of low-level oriented features in images is used in many tasks such as image analysis, image denoising, contour extraction and image segmentation. Low-level features are also useful as robust image primitives for higher-level tasks such as object detection and recognition [1]. A measure of robustness is provided by the steering framework [2] which allows practical linear filtering under adaptive orientation control, efficient computation of filter responses at multiple orientations, and accurate estimation of dominant orientation. The steering framework provides an efficient means of extracting gradient information in a particular direction with reduced interference from other oriented image structures in the same spatial neighbourhood.

X-Y separable filters have low computational requirement as they only involve two 1D filtering operations in the horizontal and vertical directions respectively. In order to retain completeness of the information contained in the image across orientations, the filters need a broad orientation selectivity characteristic (orientation response curve) close to a cosine function. The broad orientation selectivity increases the potential for intereference in filter responses from oriented image structures close in orientation in the same spatial neighbourhood, e.g. corners and junctions with acute angles. Freeman and Adelson [2] proposed quadrature filters with higher orientation selectivity that can be used in small basis sets at fixed orientations and whose reponses can be linearly combined with analytically derived steering functions to yield filter responses at arbitrary orientations. The higher selectivity of the filters allows the dominant orientation in a spatial neighbourhood to be more accurately estimated. Perona [3] used the Singular Value Decomposition in order to obtain steerable filter masks. Bharath and Ng [4] proposed polar-separable quadrature filters whose angular frequency (orientation) characteristics can be separately specified from their spatial frequency (scale) characteristics. Once angular frequency functions that provide constant response power across orientations have been specified, appropriate steering functions are learned from the orientation response characteristics of the basis filters and the desired orientation response characteristics for each steerable orientation.

The orientation selectivity of the basis filters determines the size of the basis set and thus the number of filtering operations. Higher and tighter selectivity of the basis filters requires more basis filters to cover the whole range of orientations. Each additional filter to the small basis set significantly increases computational requirements. In the previous steering approaches, the orientation and scale selectivity of the steered filters are the same as the characteristics of the basis filters. However, with polar separable selectivity characteristics, the linearly combined orientation selectivity characteristics of the steered filter can be sharpened (increased) for the widely-used basis filter family with squared-cosine characteristics [5, 2]. Therefore, sharpening increases the efficiency of the steering framework. In neuroscience, there is experimental evidence that the neurons in the primary visual cortex can refine their orientation selectivity across the stages of the visual processing pathways [6].

The paper is organised in the following way. We describe in Section 2 the polar-separable design of the basis filters in the Fourier domain and the choice of the orientation and scale selectivity characteristics. In Section 3, we discuss the trigonometric identity that allows us to sharpen the orientation selectivity of our steered filters and we also address some practical issues, namely phase interference from the trigonometric identities, to achieve sharpening. We compare the impulse responses of our *sharpened* filters involving a basis set of 4 filters and orientation-selectivity equivalent *steerable* set of 6 basis filters, which can be important for real-time operation. In Section 5, we provide a summary and a brief discussion on future work.

2. FILTER DESIGN AND STEERING

The characteristics of the filters in the steering framework are chosen to match the requirements of the filtering task, such as accuracy of orientation estimation, and the characteristics of the image structures to be detected, such as scale. The flexibility to address these different requirements regarding both orientation and scale is provided by polar-separable filter kernels specified in the Fourier Domain. In order to achieve sharpening, an extra set of requirements is introduced which relates the orientation selectivity characteristics of the basis filters to that of the sharpened filter.

2.1. Polar-Separable Design

We specify the polar separable filter kernel $G_\alpha^\theta(\omega, \phi)$ in the Fourier domain by the radial frequency (scale selectivity) function $G_\alpha(\omega)$ and angular frequency (orientation selectivity) function $G^\theta(\phi)$ where α and θ are the scale and orientation of the desired filter respectively

$$G_\alpha^\theta(\omega, \phi) = G_\alpha(\omega) G^\theta(\phi) \tag{1}$$

The radial spatial frequency characteristic of the filter controls the distribution of the power spectra over scales. Although an in-depth discussion of the properties of different radial frequency functions is outside the scope of this work, it is worth pointing out that scale-stability of the following function allows the comparison of filter responses across scales in a scale-space analysis framework. More specifically, we adopt Erlang functions [7] of order $n = 7$ and scale $\alpha = 0.5$.

$$G_\alpha(\omega) = \left(\frac{\alpha e}{n}\right)^n \omega^n e^{-\alpha\omega} \tag{2}$$

The angular frequency characteristic of the filter determines the selectivity of its response to a specific range of orientations. In order to provide uniform coverage of filter responses across orientations and to facilitate steering by linear combination, the angular power (sum of squares) of the basis set of oriented filters for covering orientations $[0, \pi]$ needs to be flat. We choose a squared cosine function, clipped by a rectangular function, which provides a flat angular power response in a set of four orientations, i.e. $\theta \in \{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}\}$.

$$G^\theta(\phi) = cos^2(\phi - \theta) rect(\phi - \theta) \tag{3}$$

where *rect* is the unit rectangular function

$$rect(\phi) = \begin{cases} 1 & \text{, if } |\phi| \leq \frac{\pi}{2} \\ 0 & \text{, otherwise} \end{cases} \tag{4}$$

The *rect* function causes the kernel to be zero over half of Fourier space. This yields even-symmetric and odd-symmetric filter masks in the spatial domain that are tuned to lines and edges respectively. In the traditional steering framework, the magnitude of the power response (sum of squares) of the basis filters should be constant across orientations and as close as possible to 1 to reduce the degree of overcompleteness. Overcompleteness distributes the energy of an impulse across many basis filters and reduces localisation to specific basis filters. However, overcompleteness and appropriate interference of trigonometric functions allows the angular frequency response to be sharpened, which we expand in Section 3.

In order to obtain the filter masks in the spatial domain, we use the inverse discrete Fourier transform of the filter kernel $G_\alpha^\theta(\omega, \phi)$ sampled over the range $[-\pi, \pi]$ in both the horizontal and vertical frequency dimensions at increments of $\pi/64$. We obtain filter kernels $F_\alpha^\theta(m, n)$ where the significant coefficients are found at the centre of the kernel. The filter masks are subsequently cropped to a size of 15×15.

2.2. Steering Filter Responses

We first describe the principles and mechanisms behind filter steering. The angular frequency functions (Eqn (3)) of the basis filters are based on cosine functions at multiple orientation phases and they can be linearly combined to yield angular frequency functions at intermediate orientation phases. Owing to the polar separability of the filter kernels in the Fourier domain and the linearity of the inverse discrete Fourier transform, the responses of the basis filters can also be linearly combined to yield filter responses at arbitrary orientations. More specifically, given the complex filter responses $f_k(m, n)$, $k = 0..K-1$ of the K basis filters at location (m, n), the response of a filter oriented at an arbitrary angle θ can be obtained by linearly combining the responses of the basis filter with a steering or weighting function $s_k(\theta)$.

$$f(m, n, \theta) = \sum_{k=0}^{K-1} s_k(\theta) f_k(m, n) \tag{5}$$

Similarly to [4], we use standard matrix algebra, $\mathbf{B} = \mathbf{FW}$ and the Moore-Penrose pseudo-inverse to obtain appropriate steering coefficients (column vector $\mathbf{W}$) from the desired angular frequency functions for each orientation (column vector $\mathbf{B}$), and the angular frequency functions $G^\theta(\phi)$ of the basis filters (matrix of column vectors $\mathbf{F}$). The steering coefficients are then regressed into polynomial-based steering functions $s_k(\theta)$ with θ as the regressor.

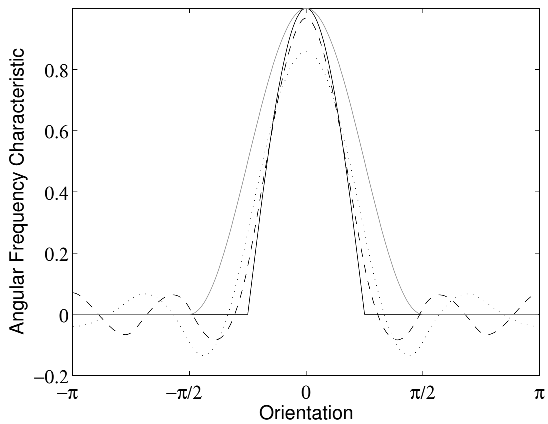

Fig. 1. The angular frequency characteristics of: $cos^2(\theta)$ of the basis filter (gray line), desired $\cos(2\theta)$ of the sharpened filter (black solid line), sharpened filter with $cos^2(\theta)$ basis filters (dotted line) and the improved sharpened filter obtained with $cos^{2.01}(\theta)$ basis filters (dashed line).

3. FILTER SHARPENING

The steering framework is based on the assumption that the orientation-domain phase of the cosine-based angular frequency functions can be arbitrarily shifted by a linear combination of the same cosine-based angular frequency functions evenly spread out at appropriate fixed orientations under the principle of superposition. Polar separability of the filter kernels in the Fourier Domain and the linearity of the inverse discrete Fourier Transform allows the shifting or steering to be applied to the filter responses themselves. In addition to superposition, a well known trigonometric identity allows a linear combination of the fixed angular frequency functions to yield another function with smaller orientation bandwidth, i.e. higher orientation selectivity.

$$\cos(2\theta) = \cos^2(\theta) - \sin^2(\theta) \quad (6)$$

$$= \cos^2(\theta) - \cos^2(\theta + \pi/2) \quad (7)$$

In Fig. 1, we show the $cos^2(\theta)$ angular frequency characteristic of the basis filter in gray and the desired sharpened $\cos(2\theta)$ angular frequency characteristic in solid black. In the traditional steering framework, designing a basis filter set with a $\cos(2\theta)$ angular frequency characteristic requires six basis filters to obtain constant angular power response. With sharpening, the basis set can be reduced to only four basis filters with $cos^2(\theta)$ angular frequency characteristic. There is a signficant computational reduction of filtering operations by one third. However, there is an important issue with a basis filter set with those characteristics. The angular frequency characteristic of the basis filters are set to zero over half of Fourier space to obtain complex quadrature filters in the spatial domain. This introduces a non-linearity that is not compatible with the trigonometric identity in Eqn (7). The same phase interference that sharpens the angular frequency characteristic causes the angular response to rise again past the clipping points of the $rect(.)$ function used in Eqn (3). The Moore-Penrose pseudo inverse attempts to compensate for the non-linearity by reducing the interference but this causes the angular frequency peak of the sharpened filter to be lowered as shown with the dotted line in Fig. 1.

Unfortunately, the angular frequency response of the basis filters at $\pm\pi$ apart have zero response at each other's peak and thus cannot be used in the linear sharpening to increase the peak of the sharpened filter. We slightly relax the constraint of Eqn (7) by introducing a small epsilon factor of 0.01 in the power of the cosine that retains the phase interference, at a slight loss of orientation selectivity, but provides the support from the nearby basis filters to greatly refine the shape of the peak of the sharpened filter, as shown by the dashed line in Fig. 1. The clipping non-linearity of the $rect(.)$ function is approximated by oscillations of relatively low amplitude. Another alternative consists of relaxing the steering constraints and jittering the orientations of the fixed basis filters slightly but this complicates both steering and sharpening.

4. EXPERIMENTS

The sharpened angular frequency characteristic is a good approximation to the desired $\cos(2\theta)$ characteristic. In this section, we investigate how well the sharpened impulse responses of a basis set of four filters approximates the impulse responses of an equivalent steerable basis set of six filters.

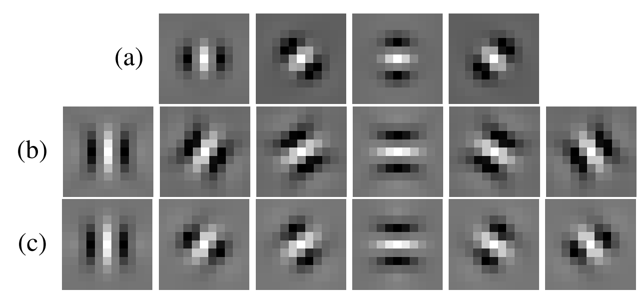

Fig. 2. Rows from top to bottom: The impulse responses of the even-symmetric (a) filters in the basis set with $cos^2(\theta)$ angular frequency characteristic, (b) steerable filters with the $\cos(2\theta)$ angular selectivity at six orientations, and (c) sharpened filters steered at the same orientations as (b).

In the first row of Fig. 2, we show the impulse response of the even-symmetric part of the quadrature filters in the basis set of our sharpening framework. On the next row, we show the desired impulse responses of the sharpened filters, synthesised with $\cos(2\theta)$ angular frequency functions at six orientations. We finally show the impulse response of the result of our sharpening, which includes steering, at the same six orientations as the previous row. We also show the even-symmetric

Orientation	0	$\pi/6$	$\pi/3$	$\pi/2$	$4\pi/6$	$5\pi/6$
MSE even sharpened filter	0.0072	0.0306	0.0306	0.0072	0.0306	0.0306
MSE odd sharpened filter	0.0236	0.0142	0.0142	0.0236	0.0142	0.0142

Table 1. Table of Mean Squared Error (MSE) of equivalent sharpened filters (2-Norm of even and odd-symmetric 6-basis steerable filters are both 0.1146).

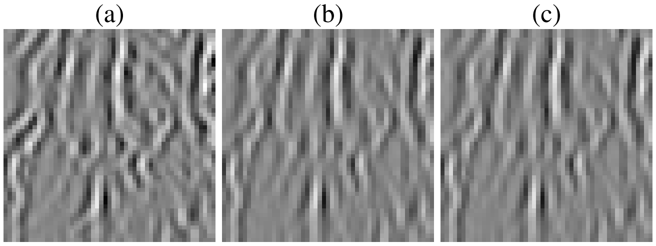

Fig. 3. The zoomed even-symmetric responses (hair of Lenna) of the (a) basis filters with $cos^2(\theta)$ angular frequency characteristic, (b) steerable filters with the $\cos(2\theta)$ angular selectivity, and (c) sharpened filters at orientation 0.

responses for the vertical orientation obtained from the hair region of Lenna in Fig. 3 and point out that the sharpened responses (c) are more sharply tuned to vertical structures similar to the desired $\cos(2\theta)$ response (b) than the basis filter (a). We observe that responses of the sharpened filters (c) visually match the desired responses (b). We also provide Mean Squared Error values of the sharpened kernels in Table I.

5. CONCLUSION

The steering framework greatly improves the efficiency of obtaining the response of oriented filters at arbitrary orientations. The steering principle operates on the basis that the angular frequency characteristics of a small set of basis filters can be linearly combined to yield the angular frequency characteristic of the filter at any arbitrary orientation. The accuracy of estimating dominant orientation in the steering framework and computing the gradient energy with reduced interference from other oriented structures in the neighbourhood is limited by the angular selectivity characteristics of the filters. Increasing the angular selectivity characteristics results in an increase in the number of filters in the basis set in order to appropriately cover the whole range of orientations. Adding basis filters significantly increases the computational burden of the filtering operations in the steering framework.

We have shown that by using the appropriate trigonometric identities for the angular characteristics of the basis filters, we can actually sharpen the orientation selectivity of the linearly combined filter in addition to steering its orientation arbitrarily. In this paper, we have shown that performing sharpening on a basis filter set of four orientations is a good approximation to applying steering on a basis filter set of six orientations. Our approach involving the Moore-Penrose pseudo-inverse to obtain the sharpening coefficients balances the requirements to sharpen the orientation selectivity while retaining flat response non-linearity over half of Fourier space. Although our basis filter set is overcomplete, the small number of basis filters causes the trade-off to favour the latter requirement. We have overcome the problem by adding a small value to the power of the angular frequency functions of the basis filters and we have improved the approximation of the sharpened filter responses to the desired filter characteristics.

Possible avenues of future research include investigation into the reduction of the oscillations that the Moore-Penrose pseudo inverse has chosen to approximate the flat non-linearity in the angular frequency functions. Other classes of angular frequency functions may also yield themselves to sharpening by other trigonometric identities.

6. REFERENCES

[1] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *Proceedings of Computer Vision and Pattern Recognition*, 2001, vol. 1, pp. 511–518.

[2] W.T. Freeman and E.H. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, 1991.

[3] P. Perona, "Deformable kernels for early vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 5, pp. 488–499, 1995.

[4] A.A. Bharath and J. Ng, "A steerable complex wavelet construction and its application to image denoising," *IEEE Transactions on Image Processing*, vol. 14, no. 7, pp. 948–959, 2005.

[5] Gösta Granlund and Hans Knutsson, *Signal Processing for Computer Vision*, Kluwer Academic Publishers, 1994.

[6] D.L. Ringach, R.M. Shapley, and M.J. Hawken, "Orientation selectivity in macaque v1: Diversity and laminar dependence," *Journal of Neuroscience*, vol. 22, no. 13, pp. 5639–5651, 2002.

[7] A. Papoulis, *Probability, Random Variables and Stochastic Processes*, McGraw–Hill, 1984.