

A Separable Filter for Directional Smoothing

V Lakshmanan

Abstract— **Anisotropic and directional filters can smooth noisy images while preserving object boundaries. Data from remote sensing instruments often have missing pixels due to geometric or power limitations. In such cases, these non-isotropic filters are very inefficient, because transform methods cannot be used when there is missing data or when logical operations need to be performed.**

A directional filter is introduced in this paper that retains the ability to handle missing data and is separable, making it computationally efficient. We demonstrate the directional filter on weather radar data where it can be used to smooth along fronts. Since the filter introduced here can be parameterized for scale, orientation and aspect ratio, this filter can be used in any directional filtering application where transform methods cannot be used, but computational efficiency is desired.

I. INTRODUCTION

In automated remote sensing applications that involve image processing or pattern recognition, image smoothing is a common preprocessing step, to remove information or noise at a scale more detailed than required by further processing stages. Most conventional image processing filters, such as the Gaussian filter, the median filter or a local average [1], are isotropic.

A directional averaging filter [1], [2] can be used to prevent edges from blurring while smoothing by computing a spatial average in several directions and using all the averages to determine the new image values. [1] suggests using the spatial average that is closest to the original pixel value. [3] suggests using the maximum spatial average. Other possibilities include using the direction with the least variance inside the kernel. Directional filters can be used in this sort of a filter bank for pattern recognition [4] and texture segmentation [5]. In anisotropic diffusion [6], the new image value at pixel is determined by the original pixel value and a weighted average of the gradients in several directions.

A common use of directional filters is as part of a linear combination of deformed basis functions [7], [8]. The deformation is "steered" by an adaptable parameter. The region of support of directional basis functions in steerable applications is typically large, so a computationally efficient filter would be useful in such cases as well.

A. Weather Images

In weather images, it would be useful to use a directional filter that will preserve the storm front boundaries by smoothing "along" the front (an elongated storm region) but not "across" it. Ellipses have been found to be a good approximation to large scale storm envelopes. The shapes of automatically identified storms were approximated by ellipses found through a principal components analysis in

[9]. Filtering techniques that compute the windowed cross-correlation of storms with elliptically-shaped filters have been found to be successful. A bank of elliptical filters was used in [3] to smooth the large scale envelope of a storm in weather radar imagery.

A study of different aspect ratios of the ellipses found that ellipses with minor axes of about 15km and major axes of about 64km performed best in extracting the envelope of a typical storm [3]. Weather radars commonly used in the United States provide resolution of about 1km per pixel radially and a range of more than 250km. A weather radar makes a new volume scan every 300 seconds on average. Thus, we need to filter a radar image that will be 500×500 pixels using a filter that is 15×64 pixels in less than 300 seconds. In this paper, we will derive a directional filter that achieves the necessary speed by having a separable form (so that the two-dimensional filtering can be done in two 1-D steps).

The directional filter bank could be speeded up by employing a Fast Fourier Transform to compute the spatial averages. Filtering using this modified technique takes 15 seconds [10], well under the update interval of radar scans and fast enough for real-time implementation. However, this comes at a cost: it is necessary to assign a numerical value to every pixel in the weather radar image including those pixels for which data is missing. So, in a Fourier Transform based method, pixels for which data are missing are also included in the cross-correlation or filter convolution computations. This can result in poor performance in regions neighboring the missing data. If missing data are to be recognized and not used in computing the cross-correlation, the filtering process will become shift-variant [10]. The correspondence of transform multiplication to spatial cross-correlation holds only for linear, shift-invariant filters.

These, then, are the requirements of the filter required to smooth weather radar imagery:

1. The filter should be non-isotropic, with the scale and aspect ratio parameterizable.
2. The angular orientation of the directional filter should be locally adaptable so that the filter "follows" the shape.
3. The filter should be fast, capable of processing radar images in real-time.
4. The filter should not use missing data in its calculations (should not require that missing data be set to some default value).

In this paper, we introduce a separable directional filter. We use this filter in a filter bank [3], [10] and demonstrate the results. This filter is useful beyond radar weather imagery; it can be parameterized to extract features at any arbitrary orientation, scale or aspect ratio at high speeds.

The author is with the University of Oklahoma & National Severe Storms Laboratory, Norman OK, USA, and can be reached at lakshman@ou.edu

II. SEPARABLE, DIRECTIONAL FILTER

A. Development

A filter that can be used to match elliptically-shaped regions that are aligned with the co-ordinate axes is given by: $f(x, y) = 1$ if $x^2/a^2 + y^2/b^2 \leq 1$ and 0 otherwise. where a and b are the axes lengths. The values of a and b define the scale of the filter and their ratio defines the aspect ratio or the elongatedness of the ellipses. To match elongated regions at an arbitrary orientation θ with the x-axis, we can use the matrix of rotation to transform each of the points inside the ellipse. This yields

$$f(x, y) = \begin{cases} 1 & \text{if } \frac{(x \cos \theta - y \sin \theta)^2}{a^2} + \frac{(x \sin \theta + y \cos \theta)^2}{b^2} \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The impulse response of this filter for an aspect ratio of 4:1 and at an orientation of 30° is shown in Figure 1a.

It is easy enough to see that the filter described in Equation 1 cannot be written in the form $f_1(x)f_2(y)$ and is therefore not separable. Consequently, the filtering cannot be divided into two stages of row-wise filtering and column-wise filtering. With large values of a and b , the filtering problem becomes computationally very expensive. In Section II-A.1, we will prove that no oriented low-pass filter can be separable. In Section II-A.2, it is shown that there do exist filters that are both oriented and separable, if we extend the domain of functions that we choose as filters.

A.1 Low-pass Oriented Filter

From Figure 1a (and even without the assumption of an elliptical shape), any oriented smoothing filter, f , is such that:

$$\begin{aligned} f(x, y) &\neq f(-x, y) \\ f(x, y) &\neq f(x, -y) \\ f(x, y) &= f(-x, -y) \end{aligned} \quad (2)$$

The two inequalities are a result of the meaning of ‘‘oriented’’, since the orientation we mean is with respect to the coordinate axes. The symmetry constraint is imposed by the low-pass nature of the filters we need.

Suppose a low-pass filter, $f(x, y)$, exists that is both oriented and separable. Then, we can write: $f(x, y) = g_1(x) \times g_2(y)$ This, along with the inequalities of Equation 2 yields:

$$\begin{aligned} g_1(x) \times g_2(y) &\neq g_1(x) \times g_2(-y) \\ g_1(x) \times g_2(y) &\neq g_1(-x) \times g_2(y) \end{aligned} \quad (3)$$

For any value of (x, y) where $f(x, y) \neq 0$ (so that neither $g_1(x)$ nor $g_2(y)$ is zero), we then have:

$$g_1(x) \neq g_1(-x) \quad g_2(y) \neq g_2(-y) \quad (4)$$

This implies that within the region of support of the filter, we can define functions $h_1(x)$ and $h_2(y)$ such that:

$$\begin{aligned} g_1(x) &= h_1(x) \times g_1(-x) \\ g_2(y) &= h_2(y) \times g_2(-y) \end{aligned} \quad (5)$$

This, along with the symmetry constraint of Equation 2 yields $h_1(x) = h_2(y)^{-1}$ which would lead us to the absurd conclusion that the filter values in the x and y directions are dependent, unless the functions $h_1(x)$ and $h_2(y)$ are constant. Let $h_1(x) = \alpha$. Then, we have: $g_1(x) = \alpha g_1(-x)$ and $g_2(y) = g_2(-y)/\alpha$ At the origin, $f(x, y) \neq 0$, since the filter is a low-pass filter and so, the above equations hold. They imply that $g_1(0) = \alpha g_1(0)$, thus yielding the value of α to be 1. This contradicts Equation 4. Hence, there exists no low-pass filter, $f(x, y)$, that is both oriented and separable.

A.2 Domain Change

We have, however, made an assumption that can be relaxed. It is not necessary for the filter to be of the form given $g_1(x) \times g_2(y)$. In particular, the filter could be of the form:

$$f(x, y) = f_1(x) \times f_2(y) + f_3(x) \times f_4(y) \quad (6)$$

which is the sum of two separable filters. Although not separable in the mathematical sense, the sum of two separable filters can be implemented so as to be computationally efficient. In this section, we will find a filter $f(x, y)$ that is oriented and can be expressed as the sum of two separable filters.

One class of separable functions is of the form: $f(x, y) = \exp(-(x^2/a^2 + y^2/b^2))$ where the a and b signify how fast the response falls to zero. Thus, a and b still determine the scale of the response and their ratio still determines the elongation of the regions that will be matched although they are not the lengths of the major and minor axes.

As before, using the matrix of rotation to transform each of the points in the filter, we obtain the oriented filter:

$$f(x, y) = \exp\left(-\left(\frac{(x \cos \theta - y \sin \theta)^2}{a^2} + \frac{(x \sin \theta + y \cos \theta)^2}{b^2}\right)\right) \quad (7)$$

which can be re-written as: $f(x, y) = f_1(x) \times f_2(y) \times f_3(x, y)$ where

$$\begin{aligned} f_1(x) &= \exp(-x^2(\frac{\cos^2 \theta}{a^2} + \frac{\sin^2 \theta}{b^2})) \\ f_2(y) &= \exp(-y^2(\frac{\cos^2 \theta}{b^2} + \frac{\sin^2 \theta}{a^2})) \\ f_3(x, y) &= \exp(-xy(\frac{1}{b^2} - \frac{1}{a^2})\sin(2\theta)) \end{aligned} \quad (8)$$

The impulse response of this filter for an aspect ratio of 4:1 and at an orientation of 30° is shown in Figure 1b.

The presence of $f_3(x, y)$ keeps this function from being separable. The exponent in the cross-term will be small when x, y is close to the center of the filter, when the orientation, θ , is small or when the aspect ratio, b/a , is close to 1. Under any of these conditions, we can approximate f_3 by the truncated Taylor series: $(1 - xy(1/b^2 - 1/a^2)\sin(2\theta))$ We can then define two new filters, $f_4(x)$ and $f_5(y)$ so that:

$$f_4(x) = x f_1(x) \quad f_5(y) = y f_2(y) \quad (9)$$

Then, the filter $f(x, y)$ can be expressed as the weighted difference of two separable filters:

$$f(x, y) = f_1(x) \times f_2(y) - \alpha f_4(x) \times f_5(y) \quad (10)$$




Fig. 1. Filters oriented at $\theta = 30^\circ$. (a) Ellipse at aspect ratio of 4:1 – not separable, so computationally expensive (b) Localized filter, with higher weights to the center value (c) Separable approximation to (b) where b is chosen to be 1.5 times the filter window size (d) A separable filter of the form given in Equation 12. This filter has an uneven weight distribution, with excessive weights given to the extremes of the ellipse.

where α is given by: $(a^2 - b^2)\sin(2\theta)/(a^2b^2)$.

The filtering process to obtain $f_1(x) \times f_2(y)$ is to first filter the rows of the image with $f_1(x)$. The columns of the result are filtered by $f_2(y)$. A similar process is repeated with the image using $f_4(x)$ and $f_5(y)$. The weighted difference between the two results is the filter response. Incorporating logic to include only valid pixels in the cross-correlation, the result will be scaled by the sum of the weights of $f_1(x)$ and $f_2(y)$ over valid pixels.

It can be shown that the error in truncating the Taylor series expansion of $\exp(-x)$ at the n th term when $0 < x < 1$ is bounded by the $(n+1)$ th term of the expansion. If we define the truncation error to be the extent to which we underestimate the true value of $f(x, y)$, then the truncation error will be positive and bounded by:

$$err(x, y) = x^2 f_1(x) \times y^2 f_2(y) \times 0.5(\sin(2\theta)\left(\frac{1}{a^2} - \frac{1}{b^2}\right))^2 \quad (11)$$

Since the functions $f_1(x)$ and $f_2(y)$ fall rapidly to zero for increasing values of x and y , the error in the approximation should be small regardless of the aspect ratio or orientation. The impulse response of the filter in Equation 10 is shown in Figure 1c. The value of b has been chosen to be 1.5 times the filter window size and the value of a has been chosen so that the aspect ratio, b/a is 4.

What is important about the developed filter is not that it is approximately elliptical but that it is clearly directional, assigning higher weights to values along one axis, and lower weights to values orthogonal to that axis. The elliptical shape is used, only as a target, to arrive at an approximation that provides separable, directional smoothing.

B. Computational Efficiency

Filtering the rows of an image with $f_1(x)$ or $f_4(x)$ takes $(2p + 1)N^2$ computations where $(2p + 1)$ is the size of the filter window (e.g: 64 in the case of a 15×64 window) and N the size of the image. The time to filter the columns of an image is the same. We need to perform this filtering four times on each pixel of an image and compute the weighted difference once. Hence, the computational need of this technique is of the order of $(8p+5)N^2$. For comparison, the computational requirement of the inseparable filter given in Equation 1 is $(2p + 1)^2N^2$, on the order of p , less efficient.

If we employ these filters in a filter-bank of q filters, the total computation cost is of the order of $(8p + 5)qN^2$ computations. The computational cost of a transform method under these conditions is $N^2 + (2q + 3)N^2 \log N$ [10]. Thus, the two techniques are of comparable computational efficiency if the filter size, p is small. The transform method will be more efficient for large values of p . In practice, the computational efficiency of the filter introduced here will depend on the number of valid pixels in the image and on the number of logical operations performed.

III. RESULTS AND DISCUSSION

One method of smoothing is to construct a filter bank [2] of several directional filters and choose the result of one of the directional filters at each point [1]. In [3], [10], a filter bank of 18 filters with θ increments of 10° was used. The radar image is filtered with each of the filters by computing the cross-correlation between the image and the filter window and computing the average of only those pixels in the window that correspond to valid data.

We used the filter described in [3] (i.e. taking the maximum of the responses to filters at different orientations) and given in Equation 10 and defined the filter response as the weighted average of those pixels in the window that correspond to valid data. We chose the values of a and b so that the aspect ratio was 4 and b was 1.5 times the filter window size. The radar reflectivity image in Figure 2a was filtered by a filter bank with filters of the form of Equation 10. The filtering shown here is performed directly in the logarithmic (dBZ) space, since the end-results will be used for clustering values [11]. If the filtered values are to be used for computing moisture content or precipitation, the filtering may have to be performed in linear (Z) space. The result is shown in Figure 2b,d. For comparison, the same image is shown filtered using an isotropic Gaussian filter in Figure 2c. Just west of the radar (bottom of the image) is an isolated storm cell with a strongly directional structure, with higher reflectivities to the south. An isotropic smoothing filter, or a filter that requires a default value to be set for missing data, will not retain this strong directionality, as can be seen in Figure 2c. The directional filter of this paper, which can be implemented efficiently because it is separable, does retain the directionality.

Further study needs to be done with the separable filter to find the aspect ratio, filter size and ratio of b to the filter

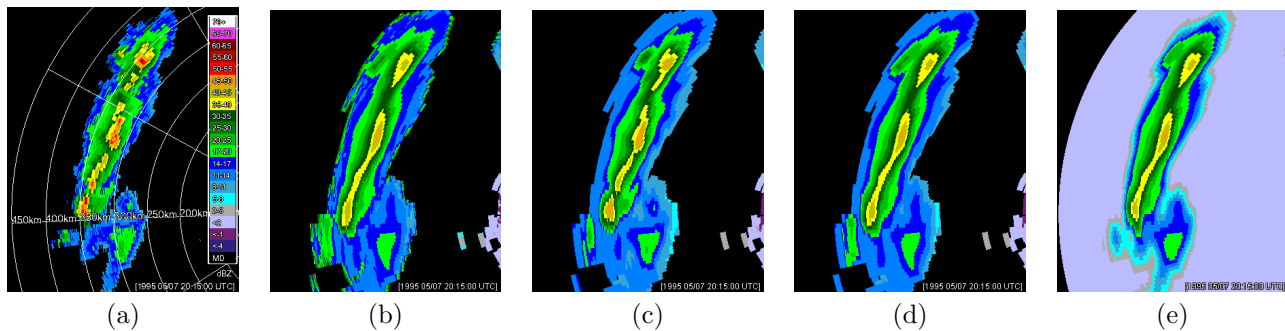


Fig. 2. (a) Part of a weather radar reflectivity image overlaid with 50km-spaced range rings, upto 450km. (b) Smoothed using the filter bank described in Section III where the directional filter used is the separable filter introduced in this paper and given in Equation 10. (c) The same image smoothed using an isotropic Gaussian filter of $\sigma=3$. The smoothing has removed the directionality of some isolated storm cells and leaves behind many local maxima. The directional filter image, (b), does not show those artifacts. (d) Smoothed using Equation 7, which is not separable. The result is clearly directional and similar to that in (b), a much faster implementation. (e) The same filter as in (d), but treating all missing data values as zero, which is what would be done in transform methods. The envelope of the smoothed storm is much larger than the true dimension. In b,c,d, the weighted average was computed based only on pixels that had non-missing data.

No.	5×21 filters			15×64 filters		
	Sep.	Insep.	FFT	Sep.	Insep.	FFT
1	22	44	22	62	492	22
2	22	45	15	64	506	14
3	23	46	14	65	520	14
4	24	47	14	67	529	15
5	23	47	15	68	532	14
6	24	48	14	69	541	14
7	25	49	14	70	548	14

TABLE I

SECONDS TO FILTER A RADAR IMAGE USING THE SEPARABLE FILTER INTRODUCED HERE, AN INSEPARABLE FORM [3] AND A TRANSFORM METHOD [10].

size that best matches the application. The choice of 4:1, 5×21 and 1.5 that were chosen in producing the image in Figure 2d is an initial one; no claim is made that it is optimal for radar weather imagery.

Looking at the form of Equation 10, one might wonder if any other smoothing filter (such as the rectangular window) might be substituted for $f_1(x)$ and $f_2(y)$. Defining the two functions as:

$$\begin{aligned} f_1(x) &= 1 \quad \text{if} \quad x^2 \left(\frac{\cos^2 \theta}{a^2} + \frac{\sin^2 \theta}{b^2} \right) \leq 1 \\ f_2(y) &= 1 \quad \text{if} \quad y^2 \left(\frac{\cos^2 \theta}{b^2} + \frac{\sin^2 \theta}{a^2} \right) \leq 1 \end{aligned} \quad (12)$$

and zero elsewhere, we get the filter shown in Figure 1d. The resulting filter has higher weights assigned to points on the boundary of the ellipse than to the center pixel and so, it is clearly not suitable for smoothing. Indeed, the reader can verify that replacing x^2 with $|x|$ in the definitions of $f_1(x)$ and $f_2(y)$ also results in a badly formed filter. In this paper, we have arrived at a filter that happens to satisfy the criteria of being low-pass, oriented and computationally efficient. Whether there exists a family of filters that meet these criteria is a topic for further work.

The time taken to filter a sequence of images using the filter introduced in this paper is shown in Table I and com-

pared with the time taken when using the filter in Equation 1. These tests were performed on a Sun Sparc Ultra 5. The times taken to filter the same sequence of images using the transform technique of [10] are repeated for comparison. While this filter bank cannot perform the filtering task as efficiently as the transform method, it retains the ability to perform logical operations in the filtering process. In this, it is significantly faster than using the inseparable form of Equation 1. Thus, we have a computationally efficient filter that retains the ability to handle missing data.

Acknowledgements Funding for this research was provided under NOAA-OU Cooperative Agreement NA17RJ1227 and the National Science Foundation Grants 9982299 and 0205628.

REFERENCES

- [1] A. Jain, *Fundamentals of Digital Image Processing*. Englewood Cliffs, New Jersey: Prentice Hall, 1989.
- [2] R. Bamberg and M. Smith, "A filter bank for the directional decomposition of images," *IEEE Trans. Signal Proc.*, vol. 40, no. 4, pp. 882–893, 1992.
- [3] M. Wolfson, B. Forman, R. Hallowell, and M. Moore, "The growth and decay storm tracker," in *8th Conference on Aviation*, (Dallas, TX), pp. 58–62, Amer. Meteor. Soc., 1999.
- [4] S. Park, M. Smith, and R. Mersereau, "Automatic recognition of SAR targets using directional filter banks and higher-order neural networks," in *SPIE Aerosense*, vol. 2, pp. 1286–1290, 1999.
- [5] J. Rosiles and M. Smith, "Texture segmentation using a biorthogonal directional decomposition," in *Sys. Cyber. Inf.*, 2000.
- [6] J. Malik and P. Perona, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. on Patt. Anal. Mach. Intell.*, vol. 12, no. 7, pp. 629–639, 1990.
- [7] W. Freeman and E. Adelson, "The design and use of steerable filters," *IEEE Trans. Patt. Anal. and Mach. Intell.*, vol. 13, pp. 891–906, 1991.
- [8] E. P. Simoncelli and H. Farid, "Steerable wedge filters for local orientation analysis," *IEEE Trans. Image Proc.*, vol. 5, no. 9, pp. 1377–1382, 1996.
- [9] M. Dixon, *Automated Storm Identification, Tracking and Forecasting – A Radar-Based Method*. PhD thesis, University of Colorado and National Center for Atmospheric Research, 1994.
- [10] V. Lakshmanan, "Speeding up a large scale filter," *J. of Oc. and Atm. Tech.*, vol. 17, pp. 468–473, April 2000.
- [11] V. Lakshmanan, R. Rabin, and V. DeBrunner, "Multiscale storm identification and forecast," *J. Atm. Res.*, pp. 367–380, July 2003.