Tilburg University

Automatic extraction of brushstroke orientation from paintings

Berezhnoy, I.J.; Postma, E.O.; van den Herik, H.J.

Published in:
Machine Vision and Applications

Publication date:
2009

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Berezhnoy, I. J., Postma, E. O., & van den Herik, H. J. (2009). Automatic extraction of brushstroke orientation from paintings. *Machine Vision and Applications*, 20(1), 1-8.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Automatic extraction of brushstroke orientation from paintings

POET: prevailing orientation extraction technique

Igor E. Berezhnoy · Eric O. Postma ·
H. Jaap van den Herik

Received: 5 July 2006 / Accepted: 5 July 2007 / Published online: 27 March 2008
© Springer-Verlag 2008

Abstract Spatial characteristics play a major role in the human analysis of paintings. One of the main spatial characteristics is the pattern of brushstrokes. The orientation, shape, and distribution of brushstrokes are important clues for analysis. This paper focuses on the automatic extraction of the orientation of brushstrokes from digital reproductions of paintings. We present a novel technique called the (prevailing orientation extraction technique (POET)). The technique is based on a straightforward circular filter and a dedicated orientation extraction phase; it performs at a level that is undistinguishable from that of humans. From our experimental results we may conclude that POET supports the automatic extraction of the spatial distribution of oriented brushstrokes. Such an automatic extraction will aid art experts in their analysis of paintings.

Keywords Prevailing orientation extraction technique · Texture · Orientation extraction

1 Introduction

When an art expert analyses a painting, a variety of techniques are available. They range from chemical analysis of pigments via analysis of visual observations to dendrochronological analysis of the panel supports. In this paper we

focus on digital visual analysis techniques for oil-on-canvas paintings. In the analysis of impressionist oil-on-canvas paintings, brushstrokes play a key role. In particular, the spatial characteristics of brushstrokes, such as their orientation, shape, and distribution are important [13]. Art experts believe that these characteristics hold a unique signature of the painter and, therefore, may aid judgements about the authenticity of a painting. Up till now the analysis of spatial characteristics was performed manually by skilled art experts. However, manual extraction of spatial characteristics is a difficult, and time-consuming task [13].

Recent advances in artificial intelligence (in particular in image processing) allow the art expert to be supported by digital techniques. Quantitative and objective analysis may facilitate the quality and consistency of the visual assessment. In this paper we present a new technique, called the prevailing orientation extraction technique (POET), designed for the automatic extraction of brushstrokes texture orientations. Our focus on brushstrokes texture orientation, rather than on brushstrokes themselves (as in [10]), is motivated by the difficulty of segmenting individual brushstrokes in Van Gogh's works. Van Gogh employed a painting style that gave rise to many overlapping brushstrokes, making the identification of individual strokes very hard [13].

In order to be actually beneficial to the art expert, a computerized technique must perform at a level comparable to that of humans. Therefore, in our study the POET is evaluated by comparing its orientation judgements to those produced by human subjects. We show that the POET performs at such a level.

The outline of the paper is as follows. In Sect. 2, we present the data used for the empirical evaluation of the POET. Section 3 describes the POET and the procedure for obtaining orientation judgements of human subjects. Section 4 gives the experimental setup and Sect. 5 presents the results. Section 6

I. E. Berezhnoy (✉) · E. O. Postma · H. Jaap van den Herik
Maastricht University, MICC, Maastricht, The Netherlands
e-mail: i.berezhnoy@micc.unimaas.nl

E. O. Postma
e-mail: postma@micc.unimaas.nl

H. Jaap van den Herik
e-mail: herik@micc.unimaas.nl

Fig. 1 Nine examples of patches extracted from paintings by Vincent van Gogh

discusses the results and Sect. 7 provides conclusions and pointers to future work.

2 The Data

In order to make the reader acquainted with the main spatial characteristics of paintings we start defining a data set of patches. They contain oriented brushstrokes as a test bench for our technique. The data were obtained from an image set of 169 digitized reproductions of paintings by Vincent van Gogh. The reproductions were obtained by scanning ektachromes at a resolution of 2,000 dpi with 48 bits colour depth. To normalize the spatial scale of the paintings, all images were resized to match the resolution of the lowest-resolution painting: 196.3 dpi. Normalization is necessary for adequate comparison.

From the normalized image set, we randomly selected 200 grayscale patches of 100×100 pixels each. Figure 1 shows nine examples of such patches. It should be noted that in most of these examples, a clear prevailing brush-stroke orientation is difficult to determine. Often, more than one oriented pattern is present within a single patch.

3 Automatic orientation extraction

The design of the prevailing orientation extraction technique (POET) is inspired by observation and analysis of human performance on the task and motivated by a failure and/or

computational inefficiency of traditional approaches such as based on edge detectors [5,9]. The POET is inspired by the observation that the brightest and largest oriented contours within a patch of brushstrokes determine the perceived prevailing orientation. The design of the POET is motivated by a failure to extract the prevailing orientation from the Fourier spectrum directly [2] as well as by the desire to improve the performances of well-known techniques, such as, orientation-estimation by means of smooth derivatives filters [4,7] and orientation estimation using multi-scale principal component analysis [6].

The POET consists of two stages: a filtering stage and an orientation-extraction stage. In the filtering stage the image is convolved with a circular filter (CF) yielding a convolved image in which oriented parallel contours are enhanced. In the orientation-extraction stage, the convolved image is processed to extract the prevailing orientation. Figure 2 illustrates the successive stages of an original image patch that is processed by the POET.

Figures 2a, b illustrate the filtering stage. Figure 2a shows the original image patch containing oriented brushstrokes. The patch is convolved with a circular filter yielding the convolved image shown in Fig. 2b. Figures 2c–e show the three steps constituting the orientation-extraction stage. This stage creates a binary representation of the convolved image containing oriented objects (the white shapes in figures 2c–e). In the following two subsections we describe both stages of the POET in more detail.

3.1 The filtering stage

The filtering stage is based on the application of a circular filter. The circular filter is a filter deliberately designed to satisfy two criteria: (1) orientation invariance, i.e., the filter should enhance oriented contours in all directions, and (2) a band-pass response, i.e., the filter should enhance contours within a specific range of spatial frequencies. Using the frequency sampling method [8], the design of a filter meeting both criteria is straightforward. The method for designing the filter is illustrated in Fig. 3. In meeting the orientation-invariance criterion a circular band of frequencies is selected in the frequency domain. Two parameters define the orientation-invariant pass band of the filter, the pass-band width W and the centre frequency R . The upper and lower cut-off frequencies of the filter are defined as $R - W/2$ and $R + W/2$, respectively. The pass-band criterion is met by selecting appropriate values for R and W . Figure 3a shows the idealized frequency domain response of the filter ($R = W = 0.5$). Figure 3b shows the actual response. The shape of the filter in the spatial domain is shown in Fig. 3c.

We now return to Fig. 2. After setting up the parameters of the filter, an image patch (Fig. 2a) is convolved with the filter. The resulting convolution image can be seen in Fig. 2b.

Fig. 2 Illustration of successive steps in the POET consisting of the filtering stage (a–b) and orientation–extraction stage (c–e)

Fig. 3 Illustration of the frequency sampling method. **a** Desired frequency response. **b** Actual frequency response. **c** The corresponding filter in the spatial domain

Fig. 4 Thresholding the convolved image (shown in Fig. 2c) for threshold values $1/K, 2/K, \dots, K/K$. $K = 33$

3.2 The orientation–extraction stage

The orientation–extraction stage transforms the convolved (filtered) image into a set of binary oriented objects which correspond to the bright brushstrokes present in the patch. The properties of the objects will be used to extract the principal orientation of the brushstrokes.

To obtain a binary image, we apply a simple and efficient multilevel thresholding (see [1] for a recent overview). After mapping the convolution values onto the unit interval (such that the minimum value equals 0 and the maximum value 1), all convolution values smaller than k/K are set to zero, and all others to 1, where K represents the number of threshold levels and k is an index for the threshold level. Given a sufficiently large value of K ($K > 20$), an appropriate value of k has to be found so that the binary image contains oriented objects that correspond to segments of the brushstrokes.

Figure 4 illustrates the binary maps obtained for the convolved image shown in Fig. 2b for $k = 1, 2, \dots, K$ with $K = 33$. Our objective is, given K levels, to select a threshold value k/K that maximizes the number of oriented objects in the binary image because these objects contribute to the perceived prevailing orientation. An oriented object is defined as a connected cluster of at least eight non-zero pixel values of which the enclosing ellipse has a major axis length that exceeds 70% of its minor axis length.

Our procedure to estimate the prevailing orientation of objects in the binary maps is based on two observations. The first observation is that for small and large values of k , the total number of objects is smaller than for intermediate values of k . The reason for this is obvious. For small k , many objects are connected forming larger objects (see the top row of Fig. 4), whereas for large k , only a few objects exceed the threshold (see the bottom row of Fig. 4). For intermediate

Fig. 5 The number of objects (*upper graph*) and oriented objects (*lower graph*) as a function of k ($K = 33$) for a single image patch

values of k , there is a value of k for which the number of oriented objects is maximal (k still to be determined). The second observation, based on preliminary experiments, is that in general the value of k for which the number of oriented objects is maximal corresponds to the value of k for which the number of objects is maximal. The two graphs in Fig. 5 illustrate this for a single convolved image. The graphs show the number of objects (*upper graph*) and oriented objects (*lower graph*) as a function of k ($K = 33$). For this particular image, the maximum number of objects and oriented objects is reached for $k = 10$.

These two observations allow us to reduce the computational costs for determining the appropriate value of k . We simply select the value of k for which the number of objects is maximal.

Having established the appropriate value of k , we sort the oriented objects according to the major axis length of their enclosing ellipse. From the two front-ranked objects, the one with the largest eccentricity of its enclosing ellipse is selected. The orientation of this object is defined as the prevailing orientation of the image patch.

In our experiment, human orientation judgements serve as a reference for evaluating the POET's performance. The procedure for acquiring the human judgements is described in Sect. 4.2.

4 Experimental setup

This section consists of three parts. In the first part (4.1) we describe the two main parameters of the circular filter and their constraints. The values of both parameters were optimized in order to achieve an optimal performance of the

POET. In the second part (4.2) the human subject orientation judgement procedure is explained. The third part (4.3) describes the techniques applied in order to evaluate POET's performance.

4.1 The POET's parameter settings

To optimize POET's performance, we performed an exhaustive search in the parameter space. We systematically varied the two main parameters, the pass-band width W and the centre frequency R of the circular filters' frequency response. Both parameters were varied in steps of 0.01 on the unit interval where 0 corresponds to the minimum spatial frequency (zero or DC point) and 1 to the maximum spatial frequency. Both parameters were constrained to valid values, i.e., the following three conditions should all be fulfilled: $0 < R < 1$ and $R - W/2 \geq 0$ and $R + W/2 \leq 1$. POET's performance is evaluated by comparing its orientation judgements with those of the human judges. The human orientation judgements were obtained in a separate experiment. The set-up of the experiment is described in the next subsection.

4.2 Set-up of the human orientation judgement experiment

In the human orientation judgement experiment, 15 subjects took part. The procedure of the experiment was as follows. Individual patches were presented on a computer screen. The size of the patch on the screen was 0.2×0.2 meters and the subjects had a viewing distance of about 0.5 m. Each subject had to specify two points: a beginning and an end of the line representing the prevailing brushstroke orientation of the patch. When the two points were specified, the line connecting the two points was superimposed on the patch to allow subjects to inspect and revise their orientation judgement visually. Subjects could revise their judgement as often as they desired. All subjects were instructed to complete the task at their own pace. In total, each subject judged the orientation of 200 patches. On average, they completed the task within 10–15 min. Figure 6 shows a typical example of a line superimposed on an image patch.

4.3 Evaluating POET's performance

The evaluation of the POET is based on a comparison of its estimates of the prevailing orientation of image patches with those obtained by human judges. Our interest is in the smallest difference between the orientation judgements of the POET and a subject of the same patch. In addition, we compare the judgements of the POET with two other techniques for estimating the orientation of textured patterns.

Fig. 6 A typical example of a patch and the line reflecting a subjects' judgement of the prevailing brushstroke orientation

4.3.1 Computing angular differences

The pair-wise comparisons of the POET and human prevailing orientation judgements is performed as follows. The angular difference is defined as the smallest angle between the two orientations. Representing the orientations to be compared by α_i and α_j (with $\alpha_i > \alpha_j$) we define their angular difference D_α in degrees as:

$$D_\alpha(i, j) = \min\{d_\alpha(i, j), 180 - d_\alpha(i, j)\}, \quad (1)$$

where

$$d_\alpha(i, j) = \min\{(\alpha_i - \alpha_j)_{\text{mod}360}, (180 + \alpha_i - \alpha_j)_{\text{mod}360}\} \quad (2)$$

Using this definition, the angular difference is confined to the interval 0 to 90° .

Fig. 7 Histograms of the orientation judgements for the fifteen human subjects (hs1 to hs14). Each histogram shows the number of times a certain orientation (modulo 180°) was assigned to a patch

4.3.2 Criterion for success

The main criterion for the successful automatic extraction of brushstroke orientation is a performance that is indistinguishable from human performances. To assess whether this criterion is met, determine the differences between the judgements of the POET and those of the human subjects by means of the mean squared angular distance (MSAD).

5 Results

The presentation of results consists of three parts. In the first part, we present the results of the human orientation judgements. In the second part, these judgements are compared to those obtained with the POET. Finally in the third part we compare the performance of the POET to two other orientation estimation techniques.

5.1 The human orientation judgement

Figure 7 displays the fourteen histograms of the human orientation judgements. Each histogram shows how often a patch was assigned a certain orientation (modulo 180°) by a particular subject. The overall variety of the histograms suggests that the 14 subjects were not always consistent in their judgements.

Figure 8 shows the overall histogram of the human judgements, obtained by summing all histograms shown in Fig. 7. The histogram reveals that many patches were assigned a vertical (0°) or horizontal (90°) orientation. Visual examination of the data set revealed that the preference for these orientations stems from the orientational distribution of brushstrokes

Fig. 8 Histogram of the orientation judgements of all subjects

Table 1 Overview of the parameter values R and W that yield the best match with one or more human judgements

Filter	R	W
A	0.18	0.18
B	0.19	0.10
C	0.12	0.10
D	0.21	0.14
E	0.13	0.10
F	0.21	0.15
G	0.16	0.14
H	0.21	0.10

in the data set rather than from a perceptual effect, i.e., the “oblique effect”[11].

5.2 The judgements compared

The parameters R (central spatial frequency) and W (bandwidth) of the POET were optimized to match the human judgements as closely as possible. The optimization was performed for each subject separately yielding in total eight different pairs of optimized parameter values (four of these pairs were matched to more than one human orientation judgement, the remaining four to single human judgements). Table 1 lists the eight unique pairs of optimized parameter values of R and W . Each pair is labelled with a separate label (a letter). The first column of the table indicates the filter labels (A to H). The second and the third columns show the values of the two filter parameters. These values indicate that spatial frequencies centered at 0.1–0.18 cycles pixel yield the best match with human judgements. The corresponding brushstroke contours are about 5–8 pixels in width.

Table 2 Overview of the filters that yield the best match for a particular human subject

Human Subject	Filter	MSAD
hs1	A	697.00
hs2	B	532.00
hs3	C	878.00
hs4	C	861.00
hs5	D	396.00
hs6	D	464.00
hs7	B	785.00
hs8	E	756.00
hs9	B	807.00
hs10	F	675.00
hs11	G	695.00
hs12	E	709.00
hs13	B	584.00
hs14	H	471.00

Table 3 Overview of the filters that yield the smallest average angular distance over all human subjects

Filter	Mean MSAD
A	709.79
B	680.93
C	841.36
D	463.07
E	858.5
F	755.36
G	869.86
H	573.79

Filter D showed the best performance

In Table 2, the first column lists the human subjects; the second column gives the optimal filter for the human subjects. The third column displays the MSAD value for the filter–human subject combination.

Subsequently, we take these six filters whose performances match the human performances most closely and compute an average MSAD per filter over all human subjects. These averages are listed in Table 3. The filter with the smallest average MSAD error is selected as the filter for the POET. From Table 3 we may see that filter D, with $R = 0.21$ and $W = 0.14$ (see Table 1), gives the lowest MSAD, and is therefore selected as the filter for the POET.

We cross-compared the judgements of the POET with those of each human subject (hs). This resulted in the matrix shown in Table 4. Each entry specifies the MSAD, i.e., the MSE of the angular difference between the row and column judgments. We deliberately included entries below the diagonal to facilitate the interpretation of the results. Overall,

Table 4 A cross-comparison of all judgements

	POET	MS-PCA	SF	hs1	hs2	hs3	hs4	hs5	hs6	hs7	hs8	hs9	hs10	hs11	hs12	hs13	hs14
POET	0	424	1,841	584	271	464	490	396	464	559	313	560	426	541	354	616	445
MS-PCA	424	0	1,947	680	463	587	647	504	521	597	439	808	524	538	466	783	451
SF	1,841	1,947	0	1,893	1,891	1,902	1,981	1,922	2,018	2,091	1,907	1,844	1,989	2,051	1,867	1,948	1,985
hs1	584	680	1,893	0	675	486	876	629	630	898	532	846	582	629	660	1,161	811
hs2	271	463	1,891	675	0	578	533	456	483	595	350	678	528	734	459	710	539
hs3	464	587	1,902	486	578	0	864	408	503	809	442	674	510	718	575	961	658
hs4	490	647	1,981	876	533	864	0	583	658	619	616	673	893	818	616	779	593
hs5	396	504	1,922	629	456	408	583	0	437	568	379	550	590	739	315	685	583
hs6	464	521	2,018	630	483	503	658	437	0	509	343	733	493	545	444	911	625
hs7	559	597	2,091	898	595	809	619	568	509	0	598	753	841	750	565	870	624
hs8	313	439	1,907	532	350	442	616	379	343	598	0	600	418	511	361	792	523
hs9	560	808	1,844	846	678	674	673	550	733	753	600	0	741	858	586	843	759
hs10	426	524	1,989	582	528	510	893	590	493	841	418	741	0	685	564	907	521
hs11	541	538	2,051	629	734	718	818	739	545	750	511	858	685	0	605	1,057	667
hs12	354	466	1,867	660	459	575	616	315	444	565	361	586	564	605	0	684	420
hs13	616	783	1,948	1,161	710	961	779	685	911	870	792	843	907	1,057	684	0	866
hs14	445	451	1,985	811	539	658	593	583	625	624	523	759	521	667	420	866	0

Each entry specifies the mean squared angular distance (MSAD) between the row and column judgements

from the values in this table we see that the POET judgements agree fairly well with those of most human subjects.

Figure 9 illustrates the differences between the human orientation judgements (gray dots) and the POET judgements (solid line). The graph shows the angular differences $D_\alpha(i, j)$ with respect to the horizontal axis ($\alpha_j = 0$) against all patches sorted in order of ascending orientation according to the POET. Evidently, POET's judgements agree quite well with most human judgements.

5.3 The POET versus standard techniques

In this subsection we compare the performance of the POET to two orientation-estimation techniques: single-scale steerable filters (SF) [4, 7] and multi-scale principal components (MS-PCA) [6]. The two parameters for the orientation estimation with steerable filters (E2) are the number of orientations and the scale. The prevailing orientation is defined as the orientation with the maximum energy. In the MS-PCA the image patch was subdivided into 12×12 sub-images, for each of which the orientation was determined from the principal components at multiple scales. The prevailing orientation was defined as the most frequently observed orientation (the mode of a 60-bin histogram of the angles).

We optimized the parameters of these two techniques in a similar way as for the POET and obtained the results listed in the second and third rows and columns of Table 4. The Mean Squared Angular Difference (MSAD) is smallest

Fig. 9 Graph of angular differences $D_\alpha(i, j)$ between orientation judgement α_i and the horizontal axis $\alpha_j = 0$ as a function of patch number (sorted in order of ascending orientation according to the POET). The solid line represents POET's judgements, the gray dots are individual human judgements

for the POET. The steerable filters clearly show the worst performance.

Table 5 shows the MSAD for each technique, averaged over all human subjects.

As an additional evaluation, we computed the correlation coefficients of the judgements generated by the three techniques and the fourteen human subjects. In Table 6, the correlations for each technique (columns) are listed per human

Table 5 Overview of the performance of the POET and two alternative techniques: the multi-scale principal components analysis (MS-PCA) and the steerable filters (SF)

Method	Mean MSAD
POET	463
MS-PCA	572
SF	1, 949

Table 6 Agreement between automatic judgements of the three techniques (columns) and the human judgements (rows) expressed in terms of correlation coefficients

	POET	MS-PCA	SF
hs1	0.75	0.69	0.23
hs2	0.86	0.75	0.23
hs3	0.76	0.70	0.23
hs4	0.76	0.65	0.22 NS
hs5	0.80	0.74	0.22 NS
hs6	0.74	0.73	0.16 NS
hs7	0.71	0.66	0.16 NS
hs8	0.84	0.77	0.24
hs9	0.77	0.65	0.23 NS
hs10	0.80	0.75	0.22 NS
hs11	0.76	0.74	0.17 NS
hs12	0.84	0.77	0.27
hs13	0.73	0.60	0.21 NS
hs14	0.80	0.78	0.22 NS
Mean	0.78	0.71	0.22

All correlations are significant ($p < 0.001$), except for those with the indications NS

subject (rows). All correlations are significant ($p < 0.001$), except for a subset of those obtained with the steerable filters (indicated by NS in the table).

The bottom row of Table 6 shows the average correlation coefficients. Again, the best correlation is obtained for the POET.

Figure 10 provides a visual illustration of the correlations by plotting the automatic judgements (horizontal axis) against the human judgements. The three groups of 14 plots display the correlations of the POET (left), the MS-PCA (middle), and the steerable filters (right), with the 14 human judgements. From a global visual examination, the relative performances of the three techniques are clear: the POET yields the best agreement (the points are nearest to the diagonal), multi-scale principal components analysis performs second best, and the steerable filters yield the worst performance.

6 Discussion

In this paper we presented a novel technique called the POET which generates orientation judgements that are indistinguishable from those of human subjects. The success of the POET is likely due to the fact that it identifies the main oriented objects in the image and use these for estimating the prevailing orientation. The performance obtained with MS-PCA is only slightly worse than that obtained with the POET. However, MS-PCA is much more computationally demanding than the POET partly due to the Eigenvalue decomposition required for computing the principal components.

The steerable filters did not show a very good performance. This is likely due to the fact that these filters average over all orientations in the image so that local oriented structures can mask the prevailing orientation of the main objects in the image. Many of the patches in our dataset contain more than one oriented structure. The POET selects the prevailing orientation and ignores the others. In this respect, the POET differs from many alternative techniques such as the

Fig. 10 Plots of the automatic judgements (*horizontal axis*) against the human judgements (*vertical axis*) for the three techniques. From left to right the POET, multi-scale principal components analysis (MS-PCA), and the steerable filters (SF)

steerable filters. The good performance of MS-PCA is likely due to the multi-scale estimation of local orientation in combination with the selection of the most frequently occurring local orientation. Possibly, a combination of the POET with this multi-scale approach may lead to a further improvement of POET's performance.

The filtering technique underlying the POET (and SF and MS-PCA) may be contrasted to statistical techniques that rely on the distribution of grey values within image regions. An interesting statistical technique for determining texture orientation is based on *interaction maps*, i.e., two-dimensional maps that represent the pair-wise variations of pixel values [3]. This technique is capable of extracting perceptual characteristics of a texture image: anisotropy, symmetry, and regularity. More importantly, it may be applied to our task to estimate the prevailing orientation of textured images.

7 Conclusion

From the experiments performed and the results obtained we may conclude that POET has a performance indistinguishable from the performances of human subjects. The POET outperforms two alternative techniques: steerable filters and multi-scale principal components analysis. Although MS-PCA performed quite good it is computationally more demanding than the POET. Therefore, we may conclude that the POET offers an attractive and effective technique for estimating the prevailing orientation of brushstrokes from digital images. In our future research, we intend to extend the POET by incorporating elements of other techniques [3, 12] to further improve the orientation–estimation performance.

References

1. Arora, S., Acharya, J., Verma, A., Panigrahi, P.K.: Multilevel thresholding for image segmentation through a fast statistical recursive algorithm (2006). <http://arxiv.org/pdf/cs/0602044>
2. Bigün, J., Granlund, G.H.: Optimal orientation detection of linear symmetry. In: Proceedings of the IEEE First International Conference on Computer Vision, pp. 433–438. London, Great Britain (1987)
3. Chetverikov, D., Haralick, R.: Texture anisotropy, symmetry, regularity: recovering structure and orientation from interaction maps. In: Proceedings of the British Machine Vision Conference, pp. 57–66 (1995)
4. Farid, H., Simoncelli, E.: Differentiation of multi-dimensional signals. *IEEE Trans. Image Process.* **13**(4), 496–508 (2004)
5. Felsberg, M., Sommer, G.: The monogenic signal. *IEEE Trans. Signal Process.* **12**(49), 3136–3144 (2001)
6. Feng, X., Milanfar, P.: Multiscale principal components analysis for image local orientation estimation. In: Proceedings of the 36th Asilomar Conference on Signals, Systems and Computers, Vol. 1, pp. 478–482. IEEE Press, Pacific Grove (2002)
7. Freeman, W., Adelson, E.: The design and use of steerable filters. *Trans. Pattern Anal. Mach. Intell.* **13**, 891–906 (1991)
8. Gonzalez, R., Woods, R.: *Digital Image Processing*, 2nd edn. Prentice Hall, Englewood Cliffs (2002)
9. Knutsson, H.: Representing local structure using tensors. In: Proceedings of Scandinavian Conference on Image Analysis (1989)
10. Lettner, M., Kammerer, P., Sablatnig, R.: Texture analysis of painted strokes. In: Proceedings 28th Workshop of the Austrian Association for Pattern Recognition (OAGM/AAPR), Schriftenreihe der OCG, Vol. 179, pp. 269–276 (2004)
11. McMahon, M., MacLeod, D.: The origin of the oblique effect examined with pattern adaptation and masking. *J. Vis.* **3**(3), 230–239 (2003)
12. Pouliquen, F.L., Costa, J.D., Germain, C., Baylou, P.: A new adaptive framework for unbiased orientation estimation. *Pattern Recog.* **38**, 2032–2046 (2005)
13. van Dantzig, M.M.: *Pictology*. E.J.Brill, Leiden (1973)