

# Applications of Image Analysis in Orientation Space

M. van Ginkel, L.J. van Vliet, P.W. Verbeek  
Pattern Recognition Group  
Department of Applied Physics  
Delft University of Technology  
e-mail: {michael,lucas,piet}@ph.tn.tudelft.nl

## 1 Introduction

At the heart of many problem-solving techniques lies the reduction of the problem into smaller sub-problems. Image analysis is no exception, and it is therefore not surprising that a lot of attention has been given to the detection of the fundamental building blocks of images: lines and edges. These building blocks, or primitives, can be characterised by a fairly small number of features: contrast, scale, orientation and phase.

These building blocks and their features form the basis for the description of slightly more complex structures and descriptors. Examples are corners, T-junctions, wavelets<sup>1</sup> (structures) and curvature (feature). The distinction between primitives and complex structures is not always completely clear and may depend on the techniques used to extract them. The technique described in this paper for example, is based on linear filtering theory. Within that framework it is convenient to consider wavelets as primitives, alongside lines and edges.

We will show a number of problems which we can solve on the basis of these primitives and, in particular, making use of one of their features: orientation. We will argue below that to cope with moderately complex structures, a special representation for orientation is required: orientation space. But first we must establish what exactly we mean by orientation:

### *Definition*

A local orientation is a direction in which the intensity variations are larger than the immediately surrounding directions. The direction with the maximal variation is the dominant orientation.

This definition is not exact: to be useful, maximal variation must be a local notion, introducing the freedom to choose the type and size of analysis window.

Now we come to the important point: why do we need a special orientation representation? Most line and edge detectors have been developed with the assumption that the line or edge to be detected is isolated in the image. This is in contradiction with the idea that these primitives are used to describe more complex structures. It is therefore necessary to explicitly include that possibility that multiple edges or lines are present at a given location. We will deal with this and the more general problem of representing multiple measurement results at one spatial location by introducing the concept of a measurement space.

A conventional measurement has the following structure

$$I(\vec{x}) \xrightarrow{\text{measurement}} m(\vec{x}) \quad (1)$$

with  $I(\vec{x})$  the  $N$ -dimensional input image, and  $m(\vec{x})$  the result of the measurement at each point  $\vec{x}$ . The idea behind a measurement space is to add a new dimension, i.e. a new independent

<sup>1</sup>meant literally: a small localised wave packet, not in the strict sense as it is used in the literature on wavelet transformations

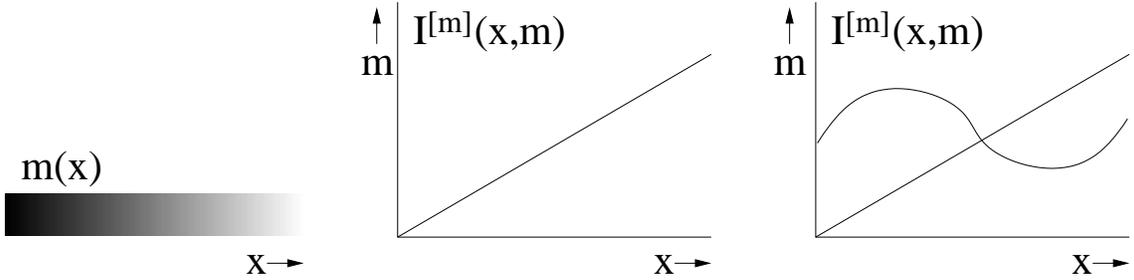


Figure 1: Left: conventional representation. The dependent variable is depicted as grey-values as a function of the independent variable. The ramp is of course one-dimensional, but has been stretched along the  $y$ -axis, because it cannot be seen otherwise. In the middle we see a measurement space. Both variables are now treated as independent variables. Note the sparseness of this representation. On the right there is a second measurement in addition to the one displayed in the middle picture. Representing the second measurement alongside the first one is no problem in measurement space, but impossible in the representation on the left.

variable, corresponding to the measurement. This gives us

$$I(\vec{x}) \xrightarrow{\text{measurement}} I^{[m]}(\vec{x}, m) \quad (2)$$

A simple example of the measurement space representation is given in figure 1.

We will call a measurement space for orientation an orientation space, denoted by  $I^{[\phi]}(x, y, \phi)$ . If two lines, with orientations  $\phi_1$  and  $\phi_2$  intersect at the point  $(x, y)$ , then the fact that both orientations are present can be represented in orientation space:

$$I^{[\phi]}(x, y, \phi) = \delta(\phi - \phi_1) + \delta(\phi - \phi_2). \quad (3)$$

In the following section we will give the details of constructing an orientation space, especially for sampled images. For now it suffices to say that it is based on applying an elongated, orientation selective, filter to the image. The filter is applied under a equi-distant set of angles. One important restriction introduced in the next section will be that the orientation selectivity must be finite. The consequence is that the delta functions in the equation above, must be replaced by some other distribution (a Gaussian for example) in practice.

How can we construct an orientation space using a filter bank? Consider a set of edges or lines, each with an orientation that is sufficiently different from that of the others. The elongated filter will only respond to a given line or edge if it is properly aligned with it. Each of the lines or edges is therefore reflected by its own distinct response. This is the basis of the techniques in this paper.

Before given a brief overview of the paper and the applications that we consider, let us examine the orientation space representation of a simple object, in order to develop more feeling for the approach. The simple object we have chosen is a circle. Its representation in orientation space is given in figure 2. At any given point of the  $(r \cos \phi, r \sin \phi)$  the local orientation, i.e. the orientation of the normal, is  $\phi \bmod \pi$ . The orientation space representation of the circle is therefore  $(r \cos \phi, r \sin \phi, \phi \bmod \pi)$ .

The structure of the paper is as follows. In the next section we look at the technical details of our orientation space implementation. The remainder of the paper is devoted to demonstrating the power of the approach using the following examples:

- Segmentation of overlapping semi-transparent or line drawing-like objects (section 3).
- Detection of boundaries between orientation fields (section 5).
- Texture characterisation, specifically line bundle like textures (section 6).

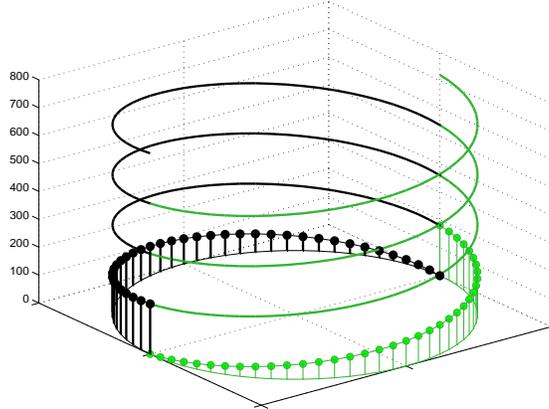


Figure 2: Orientation space representation of a circle. Several periods of the curve along the  $\phi$ -axis are shown. The first period is displayed as a stem plot. Notice that the circle is transformed into a double helix.

The last two applications are based on the results of a section on extracting orientation information from orientation space. The applications we discuss in this paper are not only interesting by themselves, but also provide an opportunity to obtain some insight into:

- How to tune the orientation space parameters for a specific problem
- The types of problems for which orientation space is useful
- The limitations of the orientation space approach

One final note. Some of the results are best displayed in colour, but in the printed version of the paper we use grey-value only. The colour versions of the results can be viewed on the web [19].

## 2 Orientation space

Orientation space was first introduced by Chen and Hsu [3] and later, independently, by van Vliet and Verbeek [18]. Chen uses the orientation of the longest line in sight to construct his orientation space. In [18] a local Hough transform is proposed to construct orientation space. The approach we use here, based on a linear filter bank, was proposed in [8]. A similar approach using Gabor filters was recently used in [2]. An interesting approach for adaptive filtering was introduced by [14] and [11]. These authors propose an invertible orientation space transform, so that oriented, line like, structures can be enhanced.

A linear orientation space is defined as follows:

$$I^{[\phi]}(x, y, \phi) = I(x, y) * h(x, y; \phi) \quad (4)$$

The orientation selective template filter  $h(x, y)$  is rotated over an angle  $\phi$  to obtain  $h(x, y, \phi)$ . Convolution is denoted by “\*”.

The choice of  $h(x, y)$  is largely free. One constraint on  $h(x, y)$  follows from the observation that we should deal separately with orientation and scale. This is achieved by imposing that the Fourier transform of  $h(x, y)$  has the following structure:

$$\mathcal{F}\{h\}(f, \theta) = \Phi_{[f]}(f)\Phi_{[\theta]}(\theta), \quad (5)$$

with  $f$  and  $\theta$  the polar coordinates in Fourier space,  $\Phi_{[f]}$  the radial component of the filter and  $\Phi_{[\theta]}$  the angular part.

There is one major constraint on  $h$ , which follows naturally when we investigate sampled orientation spaces. The Nyquist theorem [16] states that in order to sample a signal correctly, it must be bandlimited. In order to compute the convolution both the input image  $I(x, y)$  and the filter  $h(x, y)$  must be bandlimited. This is the first minor constraint on  $h$ . The convolution result, i.e.  $I^{[\phi]}(x, y, \phi)$  for a fixed value of  $\phi$  is bandlimited. Sampling along the  $x$  and  $y$  axes is therefore allowed, but what about sampling along the  $\phi$  axis?

In [8] we derived the conditions under which  $I^{[\phi]}$  may be sampled along the  $\phi$  axis. These conditions are in the form of constraints on the filter  $h(x, y)$ . These constraints are, in fact, the same as those imposed on a steerable filter [6]. A steerable filter is a filter for which the filter response under an arbitrary angle can be synthesized from the filter response under a fixed set of orientations. Although the viewpoint of orientation space and steerable filters differs, the former emphasizing the representation, the latter the synthesis of a filter under an arbitrary orientation, the mathematical structure is the same. If we sample along the  $\phi$  axis, we are in fact computing the filter response for a finite number of orientations. If we sample properly we can compute  $I^{[\phi]}$  for arbitrary  $\phi$  from the sampled version, or, equivalently, for an arbitrary orientation of the filter, which is exactly what a steerable filter is.

The filter we use has the same radial component as the one in [8], but we use a slightly different angular filter. The angular filter proposed in [8] is the filter with the optimal orientation selectivity for a given number of samples, but the filter response suffers from some secondary responses. Here we use a Gaussian profile with approximately the same orientation selectivity. This filter is not truly bandlimited, but the aliasing effects are negligible. The angular filter is as follows:

$$\Phi_{[\theta]}(\theta; N) = \exp\left(-\frac{(N\theta)^2}{2\pi^2}\right), \quad (6)$$

where  $N$  is the amount of filters (i.e. the number of samples along the  $\phi$  axis), and is related to the orientation selectivity  $s_a$  (=the standard deviation of the Gaussian kernel) as  $s_a = \pi/N$ . Note that there is trade-off between selectivity and localisation. As  $s_a$  becomes smaller, smaller orientation differences can be detected, but the filter becomes more and more elongated in the spatial domain, thus losing localisation.

The radial component of the filter is given by:

$$\Phi_{[f]}(f; f_c, b_f) = \left(\frac{|f|}{f_c}\right)^{\frac{f_c^2}{b_f^2}} \exp\left(-\frac{f^2 - f_c^2}{2b_f^2}\right) \quad (7)$$

which is a bandpass filter with an approximately Gaussian profile with standard deviation  $b_f$ . It attains its maximum for  $f_c$ :  $\Phi_{[f]}(f_c; f_c, b_f) = 1$  and goes to zero as  $f$  goes to zero.

The resulting filter  $\mathcal{F}\{h\} = \Phi_{[f]}(f; f_c, b_f)\Phi_{[\theta]}(\theta; N)$  is a quadrature filter, because  $\Phi_{[\theta]}(\theta; N)$  guarantees that  $\mathcal{F}\{h\}(f, \theta)$  is effectively zero for  $\theta$  outside the interval  $(-\pi/2; \pi/2)$ . A filter with this structure is complex valued in the spatial domain, with the property that the real and imaginary parts of the filter are each other's Hilbert transform. A filter with such a structure is called a quadrature filter. Using a quadrature filter ensures that we will detect both even (lines) and odd (edges) structures.

The orientation space defined by this filter can be conveniently given in the following form in the Fourier domain:

$$\mathcal{F}_{x,y}\{I^{[\phi]}\}(f, \theta, \phi; f_c, b_f, N) = \mathcal{F}\{I\}(f, \theta)\Phi_{[f]}(f; f_c, b_f)\Phi_{[\theta]}(\theta - \phi; N) \quad (8)$$

The subscript  $x, y$  on the Fourier transform operator indicates that it operates only along the  $x$  and  $y$  dimensions, *not* along the  $\phi$  axis.

Summarizing, orientation space has a number of advantages over conventional orientation representations. The complete list is as follows:

- Representation of overlapping or intersecting structures is possible (multi-valued representation).

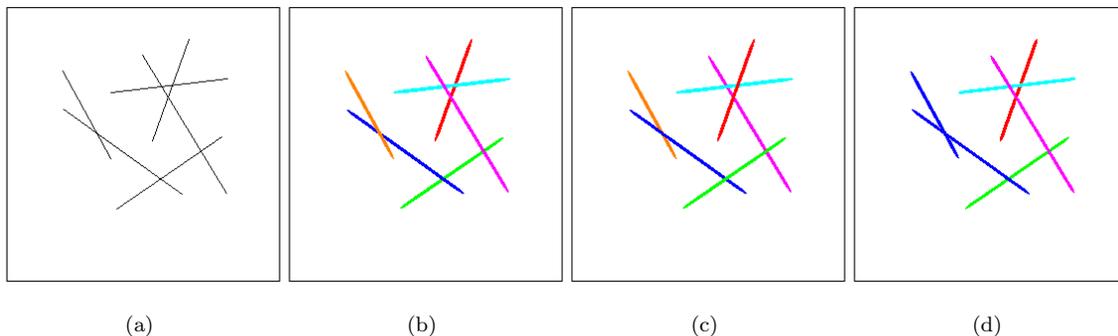


Figure 3: a) A number of overlapping lines. b) The segmentation result for  $N = 33$  (2D visualisation). The labels are displayed in different grey-values c) As (b), but using a different visualisation. The differences show that at the intersection points the lines are still assigned distinct labels. d) Segmentation result for  $N = 17$ . Note that the left most two lines can no longer be distinguished.

- Control over the orientation selectivity (this is important because of the trade-off between selectivity and localisation).
- Improved SNR. The orientation space transform spreads noise over the whole orientation space, whereas the signal is mapped to very specific locations.

### 3 Segmentation of overlapping objects

Segmentation of *semitransparent* or *line-drawing-like* overlapping objects is conceptually the simplest application of orientation space. In this section we will show a few simple examples demonstrating the principle. The basic assumption of a segmentation procedure is spatial connectivity. The pixels belonging to an object are connected, and pixels belonging to another object are not connected to the pixels of the first object. When objects overlap, this assumption is violated. By using the principle of connectivity we are actually using spatial position as a feature measurement.

In order to distinguish between two overlapping objects we need some feature or measurement, which *does* differ for the two objects at the locations they share. It is unlikely that the objects also have the same orientation at these overlap locations. Local orientation, a simple and robust feature, is therefore a good candidate. Touching objects form a notable exception. They cannot be distinguished using this method, since these have the same orientation at the overlap/intersect points.

It is important to realise that in order to distinguish between the two objects we must have a measurement for *each* of them. This is a very important constraint, since it excludes the situation where one object occludes another. This is, in fact, the reason why we explicitly restricted ourselves to semitransparent and line-drawing-like objects in the first place. An important class of images that satisfies this constraint consists of those obtained by projection like imaging techniques, for example X-ray imaging.

Another important constraint follows from the finite orientation selectivity of the orientation space representation, and the trade-off between this selectivity and localisation. The consequence is that the edges or lines defining the object boundaries may or may not be distinguishable depending on the orientation difference between them and the orientation selectivity of the orientation space. Especially when the objects are strongly curved, localisation of the filter response is important, thus limiting the attainable orientation selectivity.

We test the approach on two synthetic images. The first, shown in figure 3a, consists of a number of overlapping lines. The segmentation procedure is as follows. First, we compute the

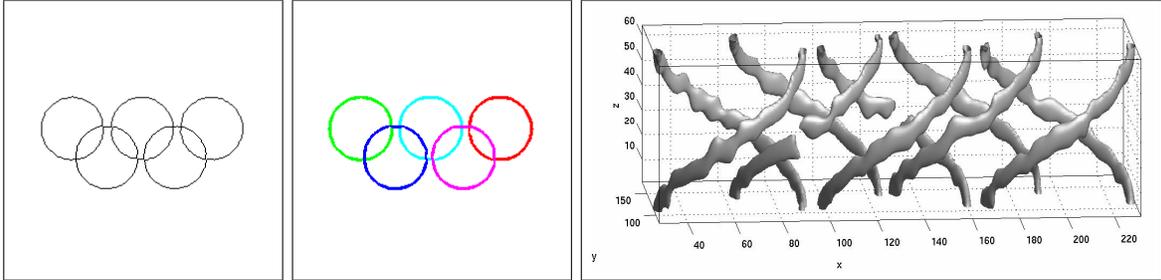


Figure 4: a) A number of overlapping circles in a familiar pattern. b) Result of the segmentation and labeling. c) Iso-surface plot of the orientation space representation. The iso-surface shown corresponds to the threshold used in the segmentation.

modulus of the orientation space representation. All lines have the same grey value and width. The magnitude of the filter response is therefore identical for all the lines. We can therefore use a very simple segmentation procedure. We have thresholded the orientation space image at half the maximum value (i.e. the maximum filter response) in that image. The resulting three-dimensional image was then labeled (distinct objects are assigned distinct labels) using a connected components algorithm.

To visualise the result, we use a simple projection method. In figure 3b&d a two-dimensional labeled image is shown that is obtained from the three-dimensional labeled image by taking the maximum label along the  $\phi$  axis, i.e:

$$l^{[2D]}(x, y) = \max_{\phi} l^{[3D]}(x, y, \phi) \quad (9)$$

In figure 3c the same procedure is applied, only using a minimum operator rather than a maximum operator. The non-object pixels do not participate in the minimum operation. If at a given  $(x, y)$  location two objects are present in the three-dimensional labeled result, we will see one of them in the maximum projection result, and the other in the minimum projection. This demonstrates that both objects are correctly assigned a unique label. The parameters used to obtain these results are given in the table below.

parameter	figure 3b,c	figure 3d
centre frequency $f_c$	0.2	0.2
bandwidth $b_f$	$0.8f_c$	$0.8f_c$
number of slices $N$	33	17
selectivity $s_a$	0.095	0.18
threshold $T$	$\frac{1}{2} \max  I^{[\phi]}(x, y, \phi) $	$\frac{1}{2} \max  I^{[\phi]}(x, y, \phi) $

Note that the bandwidth parameter  $b_f$  is large with respect to the filter's centre frequency. A large frequency bandwidth implies a narrow filter in the spatial domain the filter. With this setting the filter is tuned to detect lines (or edges).

The experiment demonstrates the fact that objects which do not sufficiently differ in their orientation cannot be separated. The two lines at the left have a fairly small difference in orientation. In figure 3b,c the orientation selectivity is sufficient to distinguish between the two, whereas in figure 3d it is not, and the two lines are assigned the same label.

The second test image contains a number of intersecting circles. Here the objects have a nice smoothly changing contour, and the segmentation procedure should therefore be able to assign each object just one label. The circles do intersect, but do not touch each other, so the basic conditions for the segmentation procedure are satisfied. In this experiment we have used a threshold based

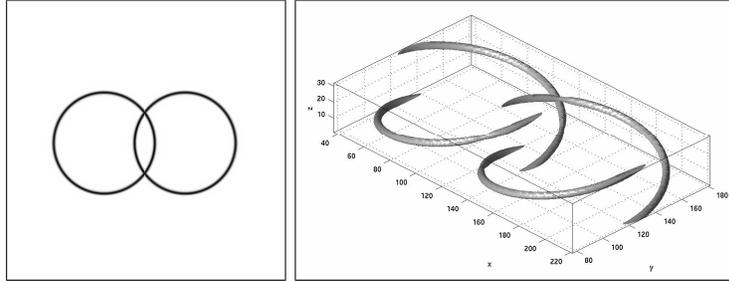


Figure 5: a) Two properly sampled circles and b) an iso-surface plot of their orientation space representation. Contrary to the results in figure 4, the orientation space response is smooth.

on a 2-means clustering method [4]. The segmentation is successful, as the visualisation of the labeled result in figure 4 shows. The parameters that were used are as follows:

parameter	figure 4
centre frequency $f_c$	0.2
bandwidth $b_f$	$0.8f_c$
number of slices $N$	17
selectivity $s_a$	0.18
threshold $T$	2-means (=isodata)

There are two things that can be learned from this example. First, this test image demonstrates some of the inherent limitations of orientation space, and, in particular, the trade-off between orientation selectivity and localisation. The image is just on the border of what the segmentation procedure can deal with. Some of the circles nearly touch. And even worse: at those locations a third circle intersects the other two. In these close quarters the trade-off between orientation selectivity and localisation becomes critical. This can be seen in an iso-surface rendering of the orientation space, figure 4, which shows that there are holes in the structures near the critical areas. The fact that the rings are still correctly labeled is actually more or less an accident: a ring needs to be cut in two places in order to be split into two distinct objects, and both affected rings are cut in just one place.

Second, the orientation space representation is not smooth, but rather irregular. This is due to the fact that the circles have been drawn using a standard line drawing routine and the resulting circles are not good representatives of a true circle. We need properly sampled and bandlimited circles. The fact that we do not, has two effects: the slightly erratic shape of the orientation space, and "leakage" of orientation energy along the whole orientation axis. This demonstrates the importance of proper sampling. We find confirmation of this conclusion in figure 5, which shows the orientation space representation of two overlapping, properly sampled, bandlimited circles. One can clearly see the smooth response.

## 4 Intermezzo: extraction of orientation information from orientation space

The presence of one or more oriented structures in a neighbourhood is reflected by the presence of an equal amount of peaks in orientation space. The relative strength of these peaks reflects the relative strength of the structures. The position of the peaks reflects the orientation of the structures. The applications in the following two sections are based on the relative strength of the peaks. The topic of this section is to establish what these peaks represent and how we extract them from orientation space.

## 4.1 A simple model for multi-orientation neighbourhoods

We will use the following simple model:

$$I_c(x, y) = \sum_i g_i(x \cos \chi_i + y \sin \chi_i) \quad (10)$$

to describe the neighbourhood  $I_c(x, y)$  centered at the point  $(x_c, y_c)$  in the image  $I(x, y)$ , i.e.:

$$I_c(x, y) = I(x_c + x, y_c + y) \quad (11)$$

The interpretation of this model is as follows: at each location  $(x_c, y_c)$  in the image  $I(x, y)$  we describe the neighbourhood  $I_c(x, y)$  as a linear superposition of patterns. Each pattern  $i$  varies along just one orientation, given by  $\chi_i$ . The intensity variations themselves are given by the one-dimensional functions  $g_i$ .

To establish the response in orientation space to this model, we need to examine the Fourier transform of the superposition:

$$\mathcal{F}\left\{\sum_i g_i(x \cos \chi_i + y \sin \chi_i)\right\}(f, \theta) = \sum_i \mathcal{F}\{g_i\}(f) \frac{1}{f} \delta(\theta - \chi_i) \quad (12)$$

The term  $\delta(\theta - \chi_i)/f$  is nothing more than the polar representation of a line under an angle  $\chi_i$  (the factor  $1/f$  is due to the Jacobian of the coordinate transformation). If we combine this result with the definition of orientation space as given in equation 8, the following result is obtained:

$$\mathcal{F}_{x,y}\{I_c^{[\phi]}\}(f, \theta, \phi) = \sum_i \Phi_{[f]}(f) \mathcal{F}\{g_i\}(f) \frac{1}{f} \Phi_{[\theta;G]}(\chi_i - \phi) \delta(\theta - \chi_i) \quad (13)$$

The final step is to compute the inverse Fourier transform of the last result. Due to the linearity of the Fourier transform each term of the sum can be transformed independently of the others. Each term consists of a part depending on the parameters  $f$  and  $\theta$ , and an other depending on  $\phi$ . The parameters  $f$  and  $\theta$  are the polar representation of the spatial frequency components, and the term depending on these is therefore subject to the inverse Fourier transform. The inverse Fourier transform has no effect on the second term since it does not depend on  $f$  and  $\theta$ . The orientation space for the model can therefore be written in the following form:

$$I_c^{[\phi]}(x, y, \phi) = \sum_i B_i(x, y) \Phi_{[\theta;G]}(\chi_i - \phi) \quad (14)$$

This is the response in orientation space to the model. We are not so much interested in the model and the response to it, as in characterising the local neighbourhood. We find a good descriptor for the neighbourhood in  $I^{[\phi]}(x_c, y_c, \phi)$ , which we will write as:

$$I_c^{[\phi]}(0, 0, \phi) = \sum_i A_i \Phi_{[\theta;G]}(\chi_i - \phi) \quad (15)$$

This multi-valued point descriptor for the neighbourhood does not describe the neighbour exactly, but it does capture its main properties. Along the  $\phi$  axis we find a number of peaks, centered at the points  $\chi_i$ , with amplitudes  $A_i$ . The locations of these peaks correspond directly to the orientations of the underlying structure. The amplitude represents the strength of the underlying structure. This strength does not depend on the signal alone, because it is of course subject to the filtering operations. The measured strength therefore depends on the filter parameters  $f_c$  and  $b_f$ .

In this way we can interpret orientation space  $I^{[\phi]}(x, y, \phi)$  for a given point  $(x, y)$  as a descriptor for the structure at the point  $(x, y)$ . Orientation space as a whole is then the set of descriptors for each neighbourhood in the image:

$$I^{[\phi]}(x, y, \phi) = \sum_{i=1}^{P(x,y)} A_i(x, y) \Phi_{[\theta;G]}(\chi_i(x, y) - \phi) \quad (16)$$

with  $P(x, y)$  the number of peaks at  $(x, y)$ . In practice we will often assume  $P(x, y) = 2$ , describing any remaining peaks and noise by a residue term:

$$I^{[\phi]}(x, y, \phi) = \left[ \sum_{i=1}^2 A_i(x, y) \Phi_{[\theta:G]}(\chi_i(x, y) - \phi) \right] + R(x, y, \phi) \quad (17)$$

We use a simple method to extract these peaks from orientation space: find the point along the  $\phi$  axis with the largest amplitude. Now select three points: the point itself, and its two neighbours at a distance  $s_a$  (rounded to grid points). Fit a parabola through this point. The top of the parabola is our estimate for the amplitude of the peak, the position of the top is our estimate for the corresponding orientation. In this way we can find the location of the peak with subpixel accuracy. To find the second strongest peak, we repeat this procedure after removing the first peak. This is done by subtracting a Gaussian with  $\sigma = s_a$  and its amplitude chosen equal to the measured one.

It is important to stress that the method described in this section is not exact, but it has the advantage of being simple and fast. We also tried more advanced methods, such as Levenberg-Marquard [1] and the EM algorithm [1, 10], but these did not seem to be very robust. Also, these methods require a reasonable initial estimate, which would typically be obtained by a method similar to the one we are using. A related point is that the orientation space typically has a good SNR, alleviating the need for more sophisticated approaches. In the next two sections, two applications are shown, which use the peak finding method described above. The results seem to indicate that the method works well enough in practice.

## 5 Orientation analysis at the boundary between two orientation fields

Many image processing operators are designed to detect a particular isolated structure. The structure tensor [13, 12], for instance, assumes a one-dimensional structure. If this assumption is not satisfied the result of such an operator is, in general, ill-defined. Typically there are regions where the proper conditions for correct operation do indeed exist, but problems arise at the interface of these regions.

Consider two touching orientation fields. Within each of the fields a conventional orientation estimator will yield the correct orientation estimate. Any estimator requires a certain analysis window and when this window lies at the interface between the two fields, the conventional orientation estimators will yield some kind of average orientation. This situation can sometimes be detected. The method in [13, 12], for instance, includes a measure that indicates whether the assumptions of the estimator are satisfied (a single orientation). But even if it can be detected, the situation cannot be resolved. To do this a truly multi-valued approach, i.e. the orientation space representation, is required.

The approach for estimating the orientation is simple. In the vicinity of the boundary there will be two responses in orientation space. Using the technique described in the previous section the amplitude of each response can be retrieved as well as the corresponding orientation. If the largest part of the analysis window lies in the first field, the response in orientation space corresponding to the orientation of the first field will be the strongest, since the evidence for this orientation is accumulated over a larger area. Only when half the analysis window lies in one field and the other half in the second, will the two responses be approximately equal.

This latter property actually allows us to detect the position of the boundary. The quotient  $q_o$  defined as

$$q_o = \frac{A_2}{A_1}, \quad (18)$$

i.e. the strength of the second strongest peak divided by that of the strongest, goes to zero as the distance to the boundary increases. At the boundary itself, where  $A_1 = A_2$ , the quotient becomes one.

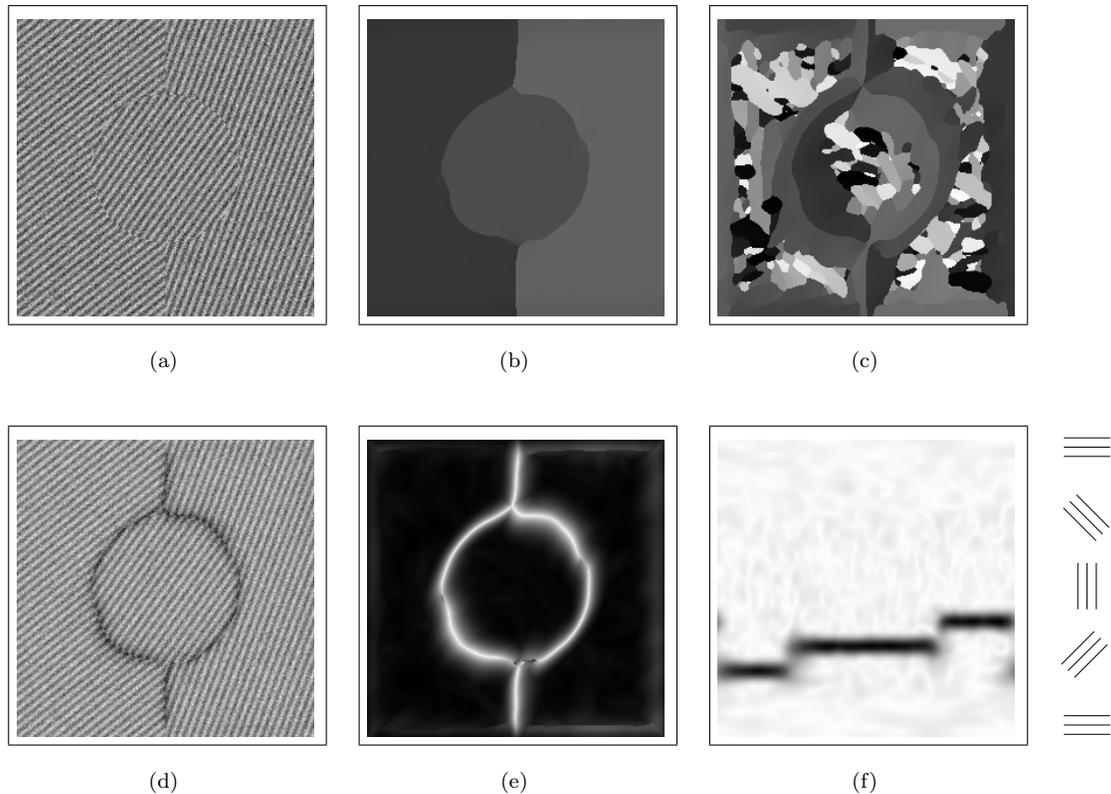


Figure 6: a) Synthetic image of three orientation fields. b) Orientation with the strongest orientation response. c) Orientation with the second strongest orientation response. d) The detected boundaries between the orientation fields overlaid on top of the test image. e) The detected boundary. f) An  $x - z$  slice in orientation space halfway the image. The markers next to the image are a visual aid to see to which orientation the responses correspond.

This idea is demonstrated in two examples. In the first we have generated a synthetic image, consisting of three orientation fields with identical characteristics except for their orientation, see figure 6a. The orientations difference between the left and centre, and the right and centre regions is 15 degrees. The orientation difference between the left and right region is 30 degrees.

In the same figure the orientation corresponding to the strongest response and that corresponding to the second strongest response are shown in (b) and (c) respectively. It is clear that the orientation in (b) gives a consistent orientation estimate in each of the three regions, over the whole of those regions, even near their interfaces. It should be noted that the results in figure (c) do not and should not contain real information except near the interfaces, where it shows the orientation estimate for the orientation at the other side of the interface. These interface regions are indicated by a high value of the quotient  $A_2/A_1$ . The responses due to the border of the image should be ignored.

Figure 6e shows the quotient  $q_o$  for this image. Note the sharp response at the boundaries between the regions. The quotient is displayed on top of the original image (d), showing that the localisation of the detected boundary is very good, although it depends on the alignment of the boundary with the pattern. If the boundary is aligned with the pattern, the localisation is better than when it is perpendicular to the boundary.

The true boundary of the centre region is a circle, but the detected edge is not, although nearly so. To the human visual system the result in (e) is as good as the true edge: try to find points in the overlay image (d) where the detected edge does not seem to correctly indicate the boundary

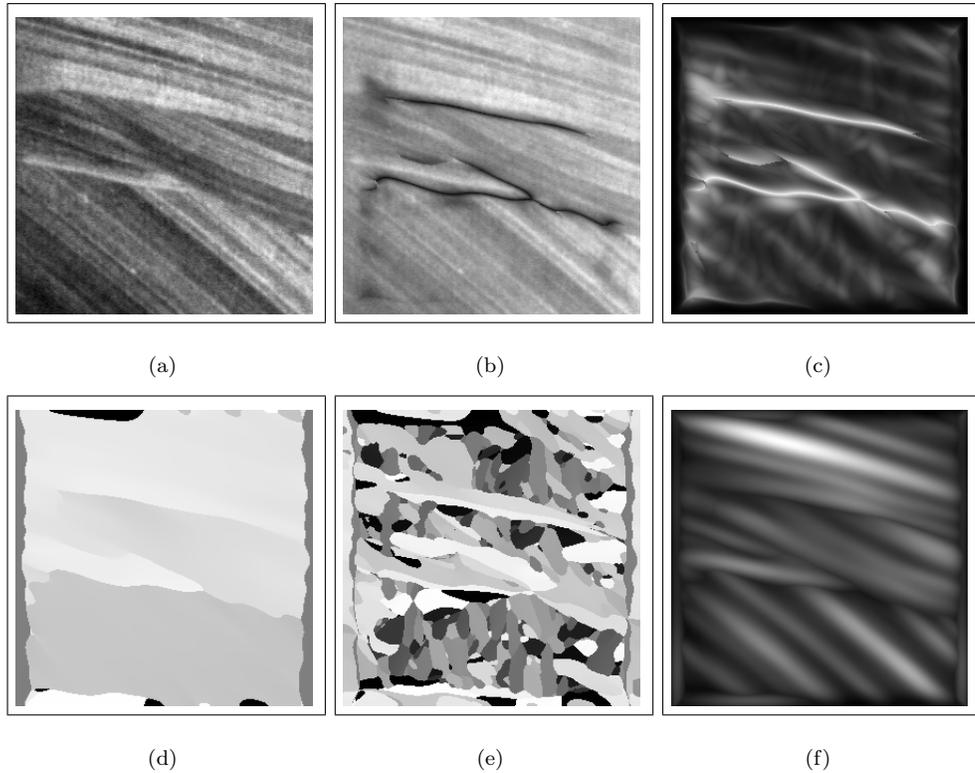


Figure 7: a) A borehole image. b) The detected boundaries overlaid on top of the borehole image. c) The quotient  $A_2/A_1$ . d) The orientation of the strongest orientation event. e) Orientation of the second strongest orientation event (only valid near the boundaries). f) Amplitude of the strongest orientation event. Note the substantial variability of the amplitude. This is in stark contrast with orientation estimate, which *is* very reliable.

between the orientation fields.

The parameters used to obtain these results are given in the following table:

parameter	setting
centre frequency $f_c$	0.15
bandwidth $b_f$	$0.5f_c$
number of slices $N$	49
selectivity $s_a$	0.064

The second example is an image of geological origin. Several orientation fields can be discerned. Again the task is to extract the boundaries between the fields. The results in figure 7 show that the approach is able to detect these boundaries. Near the boundary of the image, the response of boundary detector has been set to zero to prevent false responses to the image boundary. We can make a number of observations.

The top region has a fairly constant orientation, but in the region below it the orientation changes from left to right. On the left its orientation is significantly different from that of the top region, but on the right the two orientations are almost the same. This is correctly reflected in the absence of a response by the boundary detector at the right.

The bifurcation in the lower boundary correctly captures a subtle difference in orientation in the middle region.

The most disappointing result is that although the detector does correctly detect the boundaries, the *localisation* is very poor in places. The reason for this is an implicit assumption in the methodology that is flawed. If the amplitudes of the two orientation fields differ, the point where  $A_1/A_2 = 1/2$  does *not* mark the position of the boundary. The apparent location of the boundary moves into the region with the weaker response.

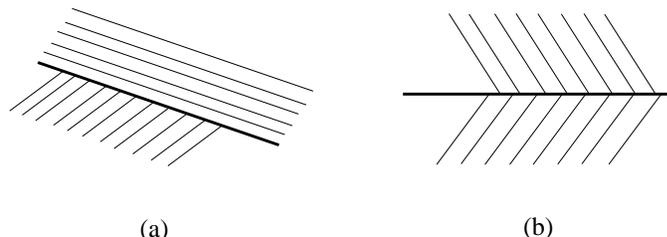


Figure 8: a) Situation where an edge detector *can* be used to detect the boundary between the orientation fields. b) Situation where an edge detector *cannot* be used. The black bar shows the position of the boundary between the orientation fields; it does *not* reflect actual image data.

The obvious question is: can we improve the localisation? One way is to consider only one response. The boundary should then lie at the point where the response is at half its maximum. In this case we assume that the boundary separates two plateaus where the response is flat. Figure 7f shows that this assumption is severely flawed as well. The amplitude of the orientation space representation does *not* consist of large areas of nearly constant response.

An other obvious choice is the use of an edge detector, but this is only possible when there is a physical edge (or line) separating the two fields. Note the subtle difference with the previous alternative where the filter bank was tuned to the orientation fields rather than to the transient (edge) separating them. Using an edge detector is possible in situations such as the one in figure 8a, but not when the orientation of the boundary does not coincide with the orientation of one of the two orientation fields, as in figure 8b.

A more general approach is to use a line ending detector, but this option has not been investigated.

## 6 Texture characterisation using orientation space

The technique described in this section is a slight variation on the one in the previous section. Here we will apply the technique to characterise textures, instead of using it to detect the boundary between two "simple" orientation fields. The types of structures we will try to characterise are small patches with a deviating orientation in an otherwise regular orientation field, for example bifurcations.

The idea is simple: the relative strengths of the strongest orientation event, the second strongest orientation event and the residue reflect the underlying structures. We will not actually compute the relative strengths, but instead visualise them directly. This can be done in two ways. The first is the most attractive option but requires the use of colour. The green channel is used for the amplitude of the first peak, red for the second, and blue for the residue. The colour results are available through the web [19] as an alternative to the visualisation used here: displaying the individual channels in grey-value. The images show:

- in green** ( $A_1$ ) the regions with the highest regularity, characterised by a single orientation,
- in red** ( $A_2$ ) regions where something special happens: touching orientation fields, deformations, bifurcations, and
- in blue** ( $R$ ) chaotic regions, where the orientation energy is distributed over a wide range of orientations.

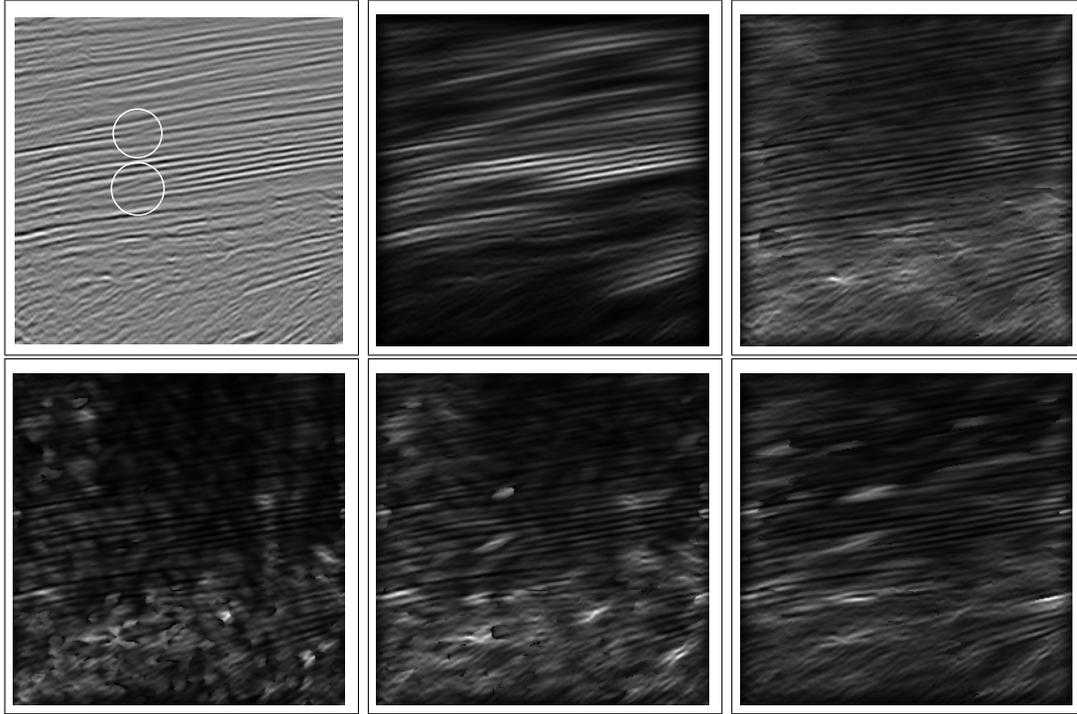


Figure 9: a) A seismic image. b) Amplitude of strongest orientation response overlaid on top of the seismic image by multiplying the two. c) Orientation residue. d,e,f) Amplitude of the second strongest orientation response for  $N = 17$ ,  $N = 33$  and  $N = 65$  respectively.

The first example is a seismic image. Seismic imaging is an important tool for obtaining large scale information about the subsurface. Acoustical waves are sent into the subsurface using a seismic vibrator. Traveling through the subsurface these waves are reflected at interfaces between layers of different acoustical impedance. These reflections are recorded at the surface using multiple detectors. The reflectivity and the position the subsurface structures can be inferred from the differences in arrival time of the reflections using a process called seismic inversion.

Stacked layers of materials with different acoustical impedances result in a very regular wavelet-like structures in the seismic image. It is the deviating geological structures, such as unconformities, faults and channels for instance, that give rise to irregularities in the seismic image. Detecting these deviating structures in a seismic image is therefore of importance for the analysis of the subsurface.

The frequency content of these images is confined to a relatively small interval. Tuning the frequency selective part of the orientation space filter is therefore simple, and the appropriate parameter settings, as given below, have been used in the experiments.

parameter	setting
centre frequency $f_c$	0.15
bandwidth $b_f$	$0.5f_c$

Contrary to the segmentation application, where we were trying to detect lines, here we are interested in wavelet like structures. This is reflected in the relatively narrow bandwidth  $b_f$ , and the corresponding wide spatial filter, containing multiple filter lobes. Two filter parameters have now been determined, and just one degree of freedom, the orientation selectivity, remains. To investigate the influence of this parameter, the approach has been tested using the following settings:  $s_a = 0.18(N = 17)$ ,  $s_a = 0.095(N = 33)$  and  $s_a = 0.049(N = 65)$ . These cover the range

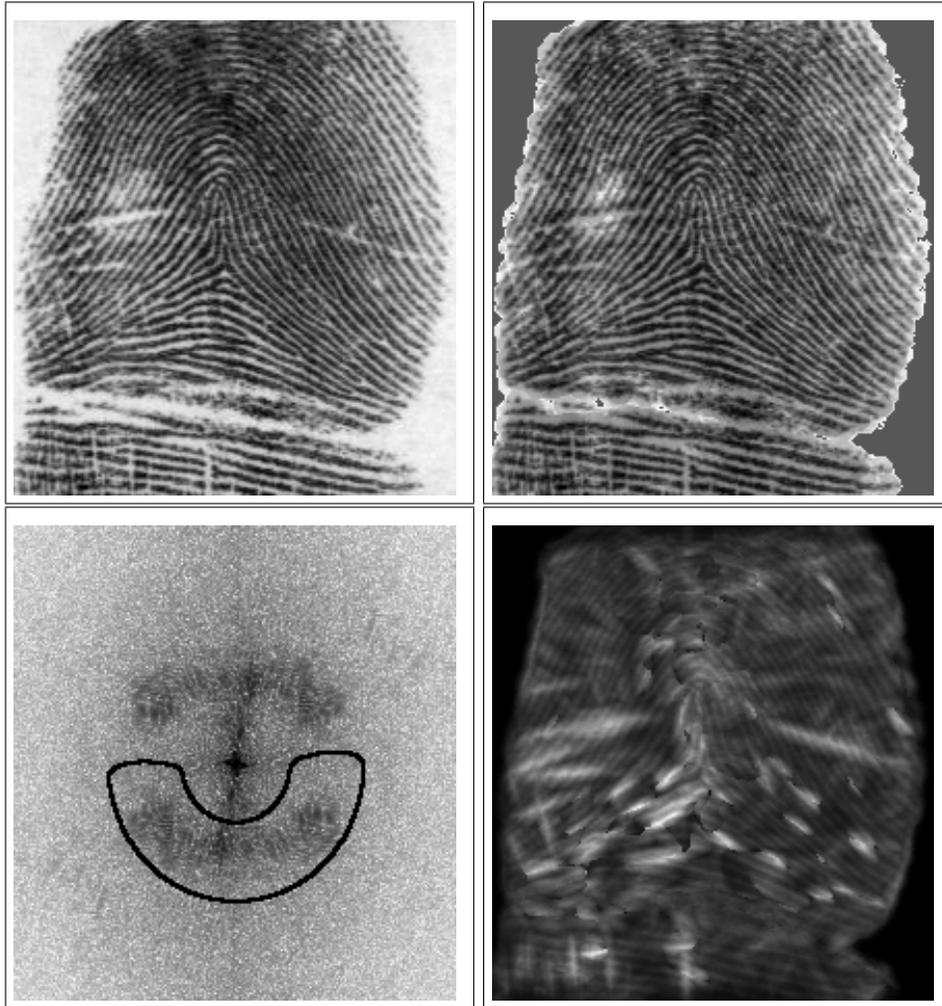


Figure 10: a) A fingerprint image. b) Corrected fingerprint image using a local contrast stretch. c) Fourier transform of the (uncorrected) fingerprint image. The region indicated by the black curve shows where the filter's or one of its rotated copies' transfer value is larger than 0.5. d) The amplitude of the second strongest orientation event. Bifurcations and scratches are highlighted.

from a relatively low to a very high orientation selectivity.

The results are shown in figure 9. The seismic image itself is given in (a), where we have circled two interesting patches that deviate in orientation from their immediate surroundings. The amplitude  $A_1$  and the residue  $R$  for  $N = 65$  are given in (b) and (c) respectively. These are not shown for the other settings of the selectivity (they do not depend on the selectivity too much). More interesting is the dependency of the amplitude of the second strongest peak on the orientation selectivity. This is shown in (d), (e) and (f) for  $N = 17$ ,  $N = 33$  and  $N = 65$  respectively.

The two patches that are indicated in (a) can be detected quite well if the orientation selectivity is sufficiently large. With  $N$  set to 17, the two features mentioned are not detected. The difference between their orientation and that of the surrounding structure is too small. With  $N$  set to 33 they are detected and well localised. As we increase the selectivity to  $N = 65$ , the response becomes stronger, but at the same time less well localised.

Bifurcations, T-junctions and similar special points are often used as landmarks. These landmarks

are commonly used to find matching points in images representing the same physical data but acquired using different imaging modalities, for instance CT and MRI. Finding the correspondence between such images is called registration. Another application can be found in identification by fingerprints. This is commonly done using fiducial marks, i.e. the points that provide important topological information such as bifurcations. The experiment with the seismic image gives a strong indication that the texture characterisation approach may also be suitable to detect bifurcations.

The fingerprint image in figure 10a contains a fairly large number of bifurcations. The fingerprint has been damaged by a number of scratches. These too, distort the main orientation field, and we may therefore expect that they will be highlighted as well. There is one complication: the contrast varies over the image. This affects the measured absolute orientation strength. The ratios between the orientation strengths at a given location will not be affected however. Because we display absolute orientation strength in the visualisation, we will get bad results where the contrast is low. To overcome this problem a local contrast stretch has been applied to the image, resulting in figure 10b.

To get correct results the frequency selective part of the filter must be properly tuned. To determine the parameters of the filter the Fourier spectrum of the image was computed and inspected. The filter parameters were then set in order to cover most of the information in the spectrum. The spectrum and the the region where the transfer function is larger than 0.5 are shown in (c).

The results in figure 10 were obtained using the following parameter settings:

parameter	setting
centre frequency $f_c$	0.2
bandwidth $b_f$	$0.5f_c$
number of slices $N$	33
selectivity $s_a$	0.095

The results in (d) clearly show the successful detection of the bifurcations as well as the scratches.

## 7 Discussion

We have shown a number of applications demonstrating the potential of the orientation space approach. We have discussed a range of possible applications, as well as advantages and limitations of this approach. We believe the orientation space approach to be a valuable addition to the set of generic image processing tools.

In the section on orientation space we state that the orientation space approach offers a better SNR. This is due to the fact that the orientation space filter bank acts like an adaptive filter: by selecting the filter that gives the best response, we are adapting the orientation parameter of the filter (hence the name steerable filter). The filter bank is therefore a candidate in situations where an adaptive edge/line detector is required.

Other applications of orientation space not discussed in this paper are curvature estimation [7] and parameterised curve detection [9].

Some of the results suggest avenues for further investigation. The most important category is formed by the extraction of bifurcations, T-junctions, line endings and similar structures. Orientation space like approaches to these problems have been investigated by [5, 15, 17]. Another interesting problem is to distinguish between scratches and bifurcations in the fingerprint example.

## Acknowledgements

This work was partially supported by the Rolling Grants program 94RG12 of the Netherlands Organization for Fundamental Research of Matter (FOM).

## References

- [1] C.M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.
- [2] J. Chen, Y. Sato, and S. Tamura. Orientation space filtering for multiple orientation line segmentation. *Pattern Analysis and Machine Intelligence*, 22(5):417–429, May 2000.
- [3] Y.S. Chen and W.H. Hsu. An interpretive model of line continuation in human visual perception. *Pattern Recognition*, 22(5):619–639, 1989.
- [4] R.O. Duda, P.E. Hart, and D.G. Stork. *Pattern Classification*. John Wiley & Sons, second edition, 2001.
- [5] W.T. Freeman. *Steerable Filters and Local Analysis of Image Structure*. PhD thesis, Massachusetts Institute of Technology, 1992.
- [6] W.T. Freeman and E.H. Adelson. The design and use of steerable filters. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 13(9):891–906, September 1991.
- [7] M. van Ginkel, P.W. Verbeek, and L.J. van Vliet. Curvature estimation for overlapping curved patterns using orientation space. In B. ter Haar Romeny, D.H.J. Epema, and J.F.M. Tonino, editors, *Proceedings of the third annual conference of the Advanced School for Computing and Imaging (Lommel, Belgium)*, pages 173–178, June 9-11 1998.
- [8] M. van Ginkel, P.W. Verbeek, and L.J. van Vliet. Orientation selectivity for orientation estimation. In *Proceedings of the 10th Scandinavian Conference on Image Analysis (Lappeenranta, Finland)*, volume I, pages 533–537, June 9-11 1997.
- [9] M. van Ginkel, L.J. van Vliet, P.W. Verbeek, M.A. Kraaijveld, E.P. Reding, and H.J. Lammers. Robust curve detection using a radon transform in orientation space applied to fracture detection in borehole images. In *ICCV'2001 [submitted]*, 2001.
- [10] S. Haykin. *Neural Networks, a comprehensive foundation*. Prentice-Hall, second edition, 1999.
- [11] S.N. Kalitzin, B.M. ter Haar Romeny, and M.A. Viergever. Invertible orientation bundles on 2d scalar images. In *Proceedings of the First International Conference on Scale-Space Theory in Computer Vision*, pages 77–88. Lecture Notes in Computer Science, Springer-Verlag, July 1997.
- [12] M. Kass and A. Witkin. Analyzing oriented patterns. *Computer Vision, Graphics and Image Processing*, 37:362–385, 1987.
- [13] H. Knutsson. Representing local structure using tensors. In *The 6th Scandinavian Conference in Image Analysis*, pages 244–251, June 19-22 1989.
- [14] J.B. Martens. Local orientation analysis in images by means of the hermite transform. *IEEE transactions on Image Processing*, 6(8):1103–1116, August 1997.
- [15] M. Michaelis and G. Sommer. Junction classification by multiple orientation detection. In Jan-Olof Eklundh, editor, *ECCV'94, Third European Conference on Computer Vision*, pages 101–108, Stockholm, Sweden, May 1994.
- [16] A.V. Oppenheim, A.S. Willsky, and I.T. Young. *Signals and Systems*. Prentice-Hall, 1983.
- [17] E.P. Simoncelli and H. Farid. Steerable wedge filters for local orientation analysis. *IEEE transactions on Image Processing*, 5(9):1377–1382, 1996.
- [18] L.J. van Vliet and P.W. Verbeek. Segmentation of overlapping objects. in L.J. van Vliet and I.T. Young (eds), *Abstracts of the ASCI Imaging Workshop*, Venray, The Netherlands, 5-6, October 1995.
- [19] <http://www.ph.tn.tudelft.nl/~michael/qqnvpbvb.html>.