ELSEVIER

Signal Processing 82 (2002) 1269–1282

SIGNAL PROCESSING

www.elsevier.com/locate/sigpro

Detection and characterisation of planar fractures using a 3D Hough transform

Augusto Sarti *, Stefano Tubaro

Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy

Received 27 July 2001

Abstract

In this article we propose an algorithmic approach – the detection and the characterization of planar fractures based on the analysis of 3D data relative to rock samples (coming from X-ray/NMR tomography). Data analysis is based on a particular implementation of the Hough Transform for the detection of planes in a 3D space. One of the original aspects of our approach is the pattern detection strategy. In fact, it works in an iterative fashion and consists of the progressive removal of the layers that constitute the cumulative array in the Hough space. At each step, we first determine the leading fracture from the analysis of the Hough cumulative arrays, and then we remove the corresponding cumulative layer in order to enable the detection of the next important fracture. We also show the results of some tests, which prove the proposed method effective even with very complex data sets.

Keywords: Volumetric analysis; Plane detection; Hough transform

1. Introduction

Rock masses generally develop fractures at different scales, from microcracks to faults, due to emplacement and cooling of rock bodies, crustal movements, and regional stress fields at different geologic stages. The spatial distribution of such fractures, in fact, carries a great deal of information for various fields of geoscience, including hydrogeology for fracture-affected flow channels, resource exploration for vein-type mineral deposits and fluids in fractured reservoirs (e.g. [1,6,19]). More recently, the general interest in a statistical characterisation of fractures in rock masses has increased for environmental reasons, for example to determine whether fractured reservoirs could be used to temporarily store large quantities of methane that, otherwise, would have to be released in the atmosphere.

A wide variety of methods has been developed for studying the geometry of rock fractures. Some of them are based on the analysis of the data acquired by sensor arrays that collect signals related to artificially generated pressure waves/impulses that propagate through the earth [19,4,26], therefore they are able to collect data on large-scale fractures and faults.

When dealing with fractures of smaller scale, however, the only data we can rely comes from the analysis of rock samples extracted from boreholes. For this reason, a general view of fracture distribution in a study area must be based on stochastic methods. In the past decade, a variety of methods have been developed with the goal of clarifying "scaling laws" of

* Corresponding author. Tel.: +39-2-2399-3647; fax: +39-2-2399-3413.

E-mail address: augusto.sarti@polimi.it (A. Sarti).

0165-1684/02/$ - see front matter
PII: S0165-1684(02)00249-9

fractures using fractal theory (e.g. [25,7,8,16]). Such methods are usually based on the analysis of fracture length distribution over wide scales. However, in order to reveal spatial correlation structures of fracture attributes, their location needs to be considered as well ([18,5,13]). Indeed, one has to bear in mind that this type of analysis cannot be pushed too far, as the amount of fracture data is small and strongly biased in a study area. As a consequence, fracture attributes that are applicable to distribution modelling are limited to length, appearance pattern, and azimuth (orientation) [17].

In order to be able to conduct a statistical analysis based on such features, we need tools that are able to efficiently and effectively extract fracture data in a non-destructive fashion. One way to do so consists of analysing a "flattened" view of the external side walls of boreholes, which can be obtained through appropriate imaging devices [23] or contact sensors [12], depending on the nature of the rock mass. As these images only show the intersection between fractures and external surface of the borehole, they enable the estimation of the geometry of the fractures of reasonable size in an automatic or a semi-automatic fashion. Using data coming from a set of boreholes extracted from the same area, together with data of some other origin (e.g. the results of the analysis of the propagation of pressure wave impulses) it is then possible to extrapolate a general outline of the fracture distribution and orientation ("up-scaling" [24]).

In order to extract additional information (for example on smaller size fractures such as microcracks) it is also possible to analyse the internal structure of rock samples. This can be done in a non-destructive way using X-ray tomography or nuclear magnetic resonance (NMR) tomography, after a proper treatment of the sample (e.g. water saturation in the case of NMR analysis) [21,3,10]. In general, such methods of analysis generate volumetric data (voxsets) [11], although the magnitude associated to each voxel is of some interest only in the case of X-ray tomography. In fact, while the greyscale values in X-ray data can be related to the material's density, this is no longer true with NMR data. Anyway, the intensity value associated to volumetric data can be quite easily binarised so that the value "1" can be interpreted as "empty" (fracture) while "0" can be treated as "full" (rock mass).

Even after binarisation, locating the most significant

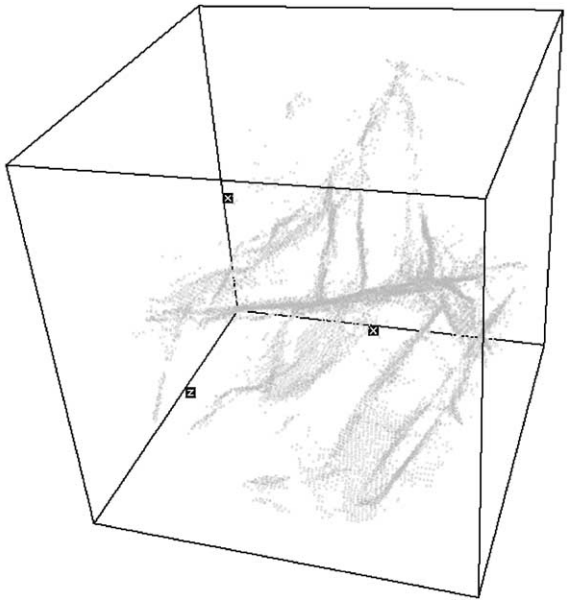

Fig. 1. A perspective view of a binary data set obtained through X-ray-tomography of a rock carrot. The clouds of gray points represent the empty areas (fractures) of the carrots.

fractures of a rock sample and understanding their geometry can be quite a formidable task when left to visual inspection. In fact, aside from visualisation problems related to the volumetric nature of the data, rock fractures are generally discontinuous (made of non-connected regions in the volume of interest) and they often exhibit a variable thickness. In addition, such fractures are usually "fogged" by highly porous material and measurement noise. It is also important to mention that visual inspection can only be conducted by a specialised operator, and can be a time-consuming and expensive procedure [23].

Considering the type of fracture attributes that are applicable to distribution modelling [17] it is appropriate and convenient to model the fractures as planar surfaces, characterised by position, orientation, extension and thickness profile. Fig. 1 shows a perspective view of a mildly fractured (binarised) data set, which gives an idea of the nature of the data that we are dealing with, and shows what kind of difficulties that visual inspection would encounter. This justifies the need of an automatic analysis strategy.

The Hough transform HT is a powerful technique that is widely used for the detection of lines in 2-D

signals (typically images) [9,14]. Extensions of this technique exist for the identification of other parametric shapes, such as sinusoids [14], circles and ellipses [27,22]. In this work we propose an extension of the basic HT to a 3-D domain, which enables us to identify planar fractures through the analysis of volumetric data sets (such as those produced by tomographic analysis).

In particular, we propose a novel method for accurately detecting fractures in inverse order of relevance. The method, in fact, exhibits a strong resilience against the inevitable mutual "masking" that occurs between fractures in the Hough domain. In order to do so, it works an iterative fashion by first detecting the leading fracture and then removing its impact onto the Hough domain. The next fracture in order of relevance will thus be the leading fracture of the residual Hough map. The method we propose also incorporates a technique for refining the localisation of the estimated planes through "local" linear regression, and it allows the user to perform thickness measurements and, more generally, conduct a simple morphologic analysis of the fracture.

It is important to emphasise the fact that our approach to planar fracture detection and localisation works on binary data (which is obtained from volumetric data of tomographic origin through simple thresholding [1]). The reasons why we do not directly use the available greyscale values are various:

- The samples that we analyse came from either X-ray or NMR tomography. While in the first case greyscale values give an indication of the volumetric density, in NMR data these greyscale values are no longer meaningful.
- Even when greyscale values can be directly related to densitometric information (X-ray case), the use of a greyscale approach does not seem to be of much help as:
  1. While the density of rock masses may change a great deal in space, fractures are characterised by a very specific and constant density.
  2. It is not possible to directly apply a greyscale approach to Hough transform such as that of O' Gorman and Clowes [20], as the gradient of the volumetric data would reveal the (irregular) boundaries of the fractures instead of their volume, which could introduce additional bias in the estimation. The availability of gradient information is, in fact, what should speed up the computation of the Hough transform. Without such information, a greyscale method becomes more expensive than a binary one, both in terms of computational load and memory requirements.
- NMR data are very noisy and a simple binarisation improves the situation a great deal.

In the next Section we will provide the basics on the Hough transform and its extension to the 3D domain (3D-HT). Such concepts will then be applied to the analysis of rock carrots. The proposed solutions will be tested on both simulated and real data.

2. Basics on the Hough transform

The Hough transform (HT) was first introduced as a method of detecting patterns of points in image data [27]. With reference to a specific pattern model, the Hough transform examines each point and finds all possible model parameters that agree with it. By collecting all such parameters in a properly defined parameter space (Hough space) we can determine data patterns that comply with the reference model through cluster identification in the Hough space. In practice, the HT converts a complex pattern detection problem in the image space into a more manageable peak detection problem in the parameter space.

The key idea of the method can be better illustrated with reference to the specific problem of identifying sets of collinear points (straight lines) in the 2D case. Let us consider a binary image, and assume that its pixels are "significant" when their value is 1. A set of significant image points (x, y) are collinear (see Fig. 2(A)) when they satisfy a relationship of the form

$$y - \hat{m}x - \hat{c} = 0, \tag{1}$$

where $\hat{m}$ and $\hat{c}$ being the parameters that define slope and intercept (with the y-axis) of the considered line. This equation can be seen as a one-to-many mapping from the image plane to the space of possible parameters. In fact, given a point $(\hat{x}, \hat{y})$ on the image plane, we can identify all the parameters (m, c) that

[1] Notice that fractures are characterised by a known and constant density (before tomographic acquisition, rock samples are soaked in water). This makes thresholding a fairly simple and safe method for binarising data.

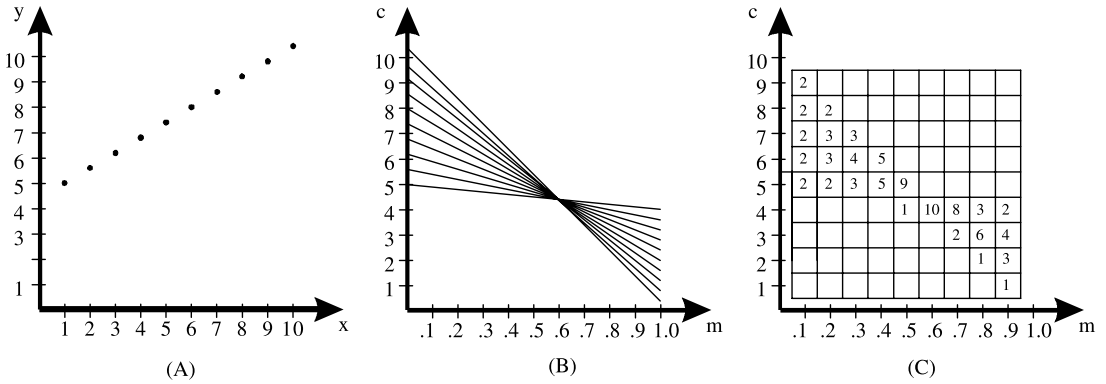

Fig. 2. Basic idea behind a HT for line detection. (A) A set of non-zero collinear points in the data space. (B) Corresponding lines in the (m, c) parameter domain (Hough space). The point of intersection of such lines identifies the parameters of the line of collinearity. (C) Corresponding Hough cumulation array.

"agree" with that point as those that correspond to all the lines that pass through $(\hat{x}, \hat{y})$. Such lines satisfy the equation

$$\hat{y} - \hat{x}m - c = 0. \tag{2}$$

Each image point (x, y) will thus correspond to a line in the (m, c) domain (see Fig. 2(B)).

If we consider a set of collinear image points, their corresponding lines in the (m, c) space intersect in one point, whose coordinates are the parameters of the line that the image points lie on (see Fig. 2(B)).

In principle, the HT can be viewed as sort of an "evidence gathering" procedure. Each image point "votes" for all parameter combinations that describe straight lines to which it could belong. The votes are counted in an accumulation array (Hough matrix). The sum of votes accumulated in one cell indicate the relative likelihood of lines described by parameters within the corresponding parameter cell (see Fig. 2(C)).

Indeed, the above definition of the HT based suffers from the disadvantage of not having a compact support. In fact, in order to be able to construct and analyse the Hough space associated to a specific data set, we need its support to be limited. This is the reason why the classical HT for line detection is based on a different parametrisation. In fact, in order to satisfy the compact support requirements, the two parameters used to describe a line in the Hough space are: the distance of a line from the origin of the reference frame, whose range is easy to limit; and its orientation in radians, which is limited to one period [9,14].

3. 3D Hough transform for plane detection

The definition of a 3D HT (3D-HT) is the natural extension of the 2D HT summarised in the previous Section. This time, however, we will follow the above line of thoughts with reference to the specific concept of null space.

A plane in the data space can be described by an equation of the form

$$\hat{a}x + \hat{b}y + \hat{c}z + 1 = 0, \tag{3}$$

which corresponds to a point of coordinates $(\hat{a}, \hat{b}, \hat{c})$ in the parameter space, and represents the null space of plane (3). In order to define the 3D-HT, however, we should consider the null space of individual points. Given a point **p** of coordinates $(\hat{x}, \hat{y}, \hat{z})$, the set of the parameters of all the planes that pass through **p** constitute its null space. These planes are characterised by all parameters (a, b, c) such that

$$a\hat{x} + b\hat{y} + c\hat{z} + 1 = 0. \tag{4}$$

Therefore, quite obviously, the null space of a point in the data space is a plane in the parameter space. This duality between the data space (x, y, z) and the above-defined parameter space (a, b, c) becomes even

more apparent when looking at the equation of the plane in vector form [27]

$$[\,a\ b\ c\ 1\,]\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = 0. \tag{5}$$

Another interesting case is that of collinear points in the data space, whose null spaces (planes in the parameter space) intersect in just one line that, in turn, is the null space of the line of collinearity in the data space. Notice that two parallel (perpendicular) lines in the data space correspond to two parallel (perpendicular) lines in the parameter space [27].

Given a binary point-set in the discrete 3D data space, its HT is defined as an integer data set in the discrete 3D parameter space, whose points (plane parameters) describe the number of significant data voxels that lie on the considered plane. Voxels are classified as significant when their binary value is 1 (in our specific case the voxels of a fracture are significant).

The above-considered plane parametrisation exhibits several interesting properties, but also suffers some a severe drawback. In fact, even if the data space (x, y, z) is limited

$$\begin{aligned} &x_{\min} \leqslant x \leqslant x_{\max}, \\ &y_{\min} \leqslant y \leqslant y_{\max}, \\ &z_{\min} \leqslant z \leqslant z_{\max} \end{aligned} \tag{6}$$

the parameter space turns out to be unbounded. In fact, the parameter a, b, c can range anywhere from $-\infty$ to $+\infty$. This is a particularly undesirable property when determining the HT as a cumulatation array.

In order to overcome this difficulty, alternative parameterisations can be considered. A reasonable one can be directly derived from the 2D-HT, and refers to the following representation of a plane:

$$x\cos\alpha + y\cos\beta + z\cos\gamma = d, \tag{7}$$

where α, β, γ are the angles between the plane and the x-, y-, z-axis, respectively.

Such angles are bound to satisfy the constraint

$$\cos^2\alpha + \cos^2\beta + \cos^2\gamma = 1, \tag{8}$$

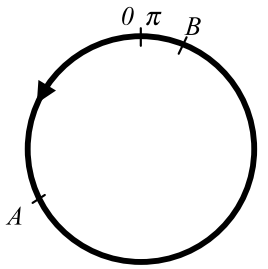

Fig. 3. Computation of the distance between two points A and B on a cyclic coordinate axis with a 0–π range (see Eq. (10)).

therefore only 3 out of 4 parameters turn out to be independent. Keeping this fact in mind, our plane parameterisation becomes (α, β, d). Such parameters can be used to evaluate the plane coefficients as

$$\begin{aligned} a &= -\frac{\cos\alpha}{d}, \\ b &= -\frac{\cos\beta}{d}, \\ c &= -\frac{\sqrt{1-\cos^2\alpha-\cos^2\beta}}{d}. \end{aligned} \tag{9}$$

Eqs. (9) are strongly non-linear, therefore the 3D-HT induced by this parametrisation turns out to be more difficult to characterise. On the other hand, this 3D-HT results in a bounded parameter space whenever the data space is bounded. In fact, as we can see in Eq. (7), d is the distance of the plane from the origin, therefore it is bounded when the sample is of limited size. Since angles are periodic, their range is intrinsically bounded. Actually, since in our analysis it is not important to assign a specific normal direction to a plane (we have no notion of "inside" or "outside"), the range of α and β can be limited to $0 - \pi$.

The fact that α and β are cyclic coordinates has another important consequence on a plane detection strategy. In fact, with reference to Fig. 3, the distance between two points (A and B) along one of such coordinate axes needs to be defined as follows:

$$\begin{aligned} &0 \leqslant A < \pi, \quad 0 \leqslant B < \pi, \\ &\operatorname{dist}(A,B) = \operatorname{dist}(B,A) \\ &\qquad = \min\{|A-B|, (\pi - |A-B|)\}. \end{aligned} \tag{10}$$

3.1. Some remarks on the numerical implementation of a 3D-HT

A numerical implementation of the above-defined 3D-HT is based, like in the 2D case, on the definition of a proper 3D cumulation array. Its construction is done by scanning all significant voxels in the data space (those that correspond to fractures). Given a significant data voxel, we increment by one unit all those cells of the cumulation array that correspond to the parameters of a plane that passes through that data voxel. The determination of such cells can be easily done in an analytical fashion. If $(\hat{x}, \hat{y}, \hat{z})$ are the coordinates of the centre of the data voxel, the parameters of all planes that pass through it are bound to satisfy the equation

$$\hat{x}\cos\alpha + \hat{y}\cos\beta + \hat{z}\cos\gamma = d, \tag{11}$$

where γ can be written as a function of α and β using the relationship $\cos^2\alpha + \cos^2\beta + \cos^2\gamma = 1$. We can thus proceed by considering all the possible pairs of angles α and β, and for each pair we can compute the corresponding value of d. Indeed, as the parameters belong to a discrete set, there will only be a finite number of combinations of angles, and the corresponding distances will have to be quantised to the nearest available value. Once determined the corresponding cell of the cumulation array, we will increment its value by one unit.

Once constructed the 3D-HT, the dominant planes in the data set can be determined through a peak search process on the Hough cumulation array. The searching strategy is, in fact, crucial for a correct identification of the fracture planes. In the next Section we will describe an effective search strategy for the specific application considered for this work.

In order to better understand the above plane detection strategy, let us consider the HT of four coplanar points shown in Fig. 4. As we can expect, there exists a point of cooccurrence for the four surfaces (clouds of points), whose coordinates provides the parameters of the plane of coplanarity of the four data points.

Indeed, the above ideal situation of perfect coplanarity never actually occurs. In fact, in typical real situations peaks are not so pronounced and isolated because

- fractures have non-zero thickness ("loose fit" of planes),
- fractures have non-perfectly planar shape (irregular scattering of points),
- data space is discrete (quantised data coordinates), and
- parameter space is discrete (quantised plane parameters).

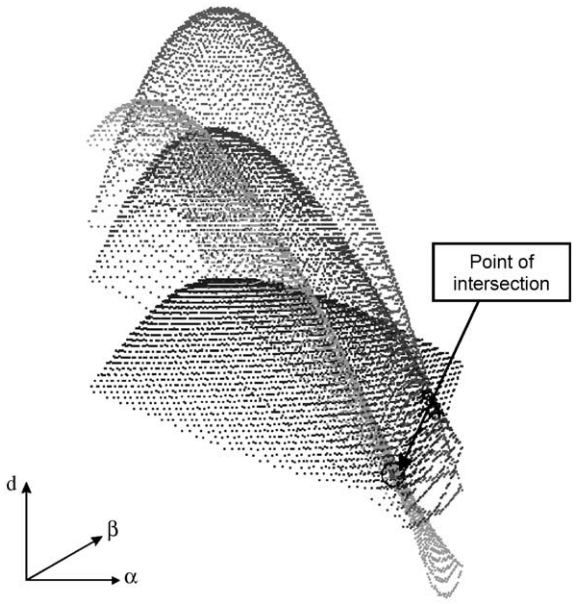

Fig. 4. Hough transform of four coplanar points. A different colour is used for the cloud of points associated to each data point. The distance between adjacent points is related to the quantisation step of α, β and d. All surfaces intersect in just one point which identifies the parameters of the plane that the data points lie on.

As a consequence, there is a region around each peak in the Hough space where the cumulation array takes on significant values, and the shape of this region is strongly influenced by the shape of the HT of a single data point. In fact, Fig. 4 suggests that the shape of this region (*cumulation region*) will be that of a thickened portion of a 3D sinusoid (often similar to a thick bowl). In order to visualise this fact, we simulated a number of $100 \times 100 \times 100$ voxel specimens, whose points of fracture lie on single planes. In our simulations the number of fracture voxels varies from about 5000 to about $15,000$. Fig. 5 shows a perspective view of a typical cumulation region in the Hough space. In this image, grey levels evoke the magnitude associated to each cell (the colour scale ranges from red to black going through orange, yellow and blue),

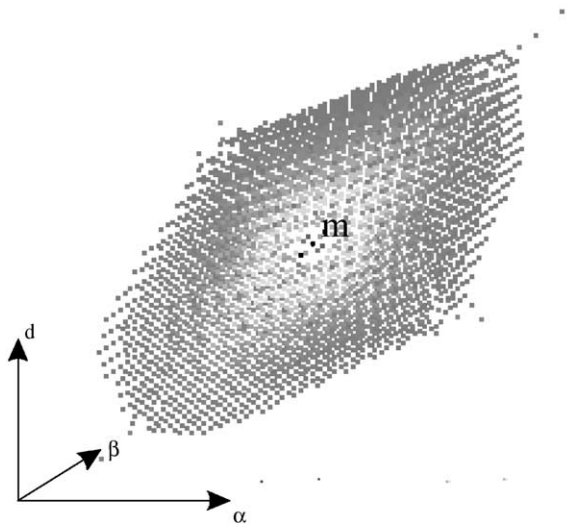

Fig. 5. A perspective view of the cumulation array (Hough space) associated to a set of coplanar voxels (data space). The grey scale is used to give an idea of the value associated to each cell. The cell with the highest value is the black one marked with *m*. In order to help visualise the data, only the cells with value greater than 100 are shown.

and only the cells whose value is greater than 100 are displayed.

When more than one plane of fracture are present, the Hough space exhibits several cumulation regions that are similar to the one shown in Fig. 5. Considering that fractures may differ significantly from one another in the number of significant voxels, the fact that each local peak is surrounded by other peaks of comparable value makes the identification of planes of fracture quite complex. In fact, a simple peak search strategy would not produce reasonable results.

In order to be able to detect not just the leading fracture but also minor ones, we could proceed by first identifying cumulation regions corresponding to the various significant planes of fracture, and then selecting the local maxima (one per region). This approach, however, requires the cumulation regions to be separated from each other. Unfortunately, this does not usually happen as they often overlap one another in the Hough space. This problem, together with several others, led us to develop a novel search strategy, which we will introduce in the next Section.

4. An effective approach for planar fracture detection

When looking at a data set such as the one shown in Fig. 1, the detection of planar structures (fractures) appears to be very complex. In fact, fractures are not continuous, their thickness is often significantly larger than one voxel, and it may vary a great deal along the fracture's extension.

In order to better illustrate the various problems that arise when dealing with real data, we will first consider the characterisation of the leading fracture, and from there we will describe our global search strategy.

4.1. Identification and characterisation of the leading fracture

Considering the definition of the Hough space given in the previous Section, we should expect the leading fracture to correspond to the global maximum in the Hough space. Normally fracture's thickness exceeds one voxel, therefore a measurement of its total volume is significant for its characterisation. In order to measure its volume we can count the significant voxels that lie within a "thick slice" centred on the considered plane. In order to decide a reasonable slice thickness, we can take into account the analysis of different real data sets and the information provided by geophysicists. In the case of an X-Ray CT with a 1 mm cubic voxel, we can use a typical thickness of 5 voxel (fracture's thickness is normally smaller than that). It is also possible to compute the "extension" (area on the plane) of the fracture. This measurement is obtained by computing the sum of the projections on the estimated plane of all the fracture voxels. Normally, data sets are very large (to the order of $100\times100\times100$ voxels), therefore in order to ensure a rapid evaluation of the fracture volume the slice height is oriented along the coordinate axis that is perpendicular the most to the fracture plane (see the 2D case shown in Fig. 6).

When considering a fracture of significant thickness, a maximum in the Hough space may not correspond to the plane that best represents a fracture. This can be more easily understood when considering a simple two-dimensional data set corresponding to a single fracture of constant thickness (see Fig. 7). Since the maximum of the HT corresponds to the plane of "maximum consensus", i.e. the plane (dashed line of

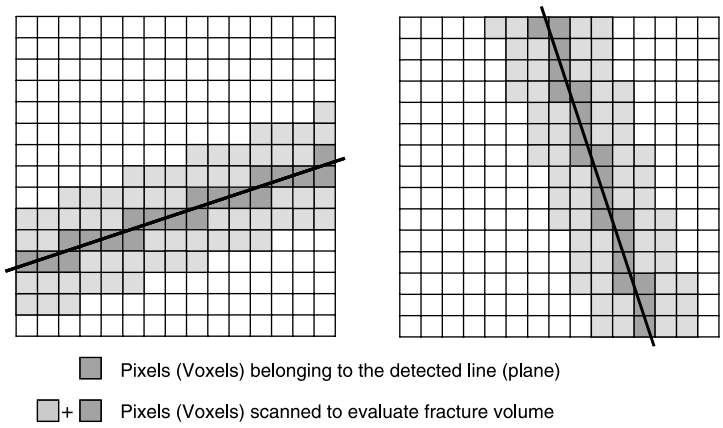

Fig. 6. Thick slice used to evaluate the fracture volume, as measured with reference to the estimated plane. For the sake of clarity we consider a 2D case, therefore voxels here correspond to pixels and planes correspond to lines.

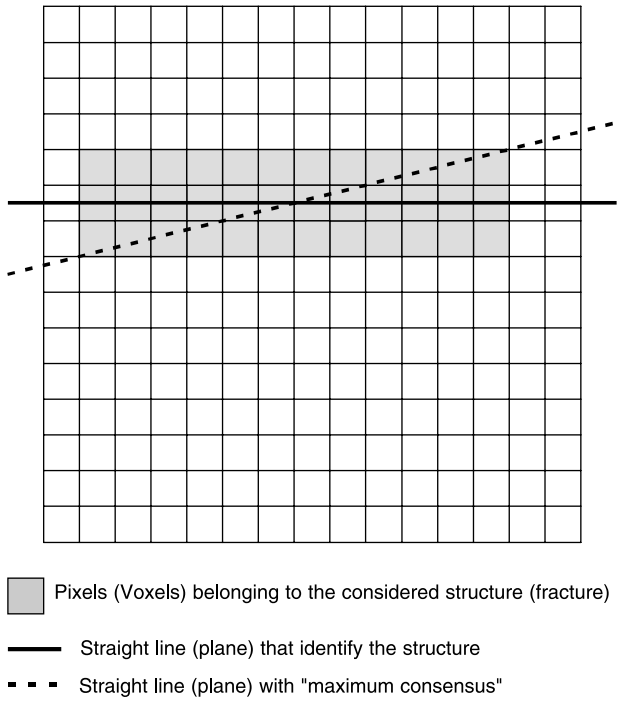

Fig. 7. Difference between the plane that correctly identifies a fracture and the one actually detected through the analysis of the Hough transform. Like in Fig. 6, the situation is depicted just for the 2D case.

Fig. 7) that fits the largest number of fracture voxels (pixels), its orientation does not correspond to that of the fracture as it ends up laying along the fracture's diagonal.

In many situations, the difference between the plane of maximum consensus and the one that best represents the fracture can be very modest or even negligible compared to the HT's quantisation step. Anyway,

this error can be easily compensated for through a regularisation process based on the minimisation of the mean square distance

$$L = \sum_{i=1}^{N} l_i^2, \tag{12}$$

where N is the total number of fracture voxels and l_i is the distance of the centre of the ith voxel with from the fracture plane.

The whole estimation procedure for the dominant plane can be summarised as follows:

1. Computation of the HT of the given data set;
2. Determination of the global maximum of the HT (plane of maximum consensus);
3. Determination of the fracture voxels;
4. Computation of the best-fitting plane through linear regression (minimisation of the mean square distance (12)). This process can be carried out through a full-search procedure over a set of planes whose parameters are centred around that of maximum consensus.

Once determined the best-fitting plane, we can estimated the fracture's volume, its extension, and the thickness distribution (see Fig. 8).

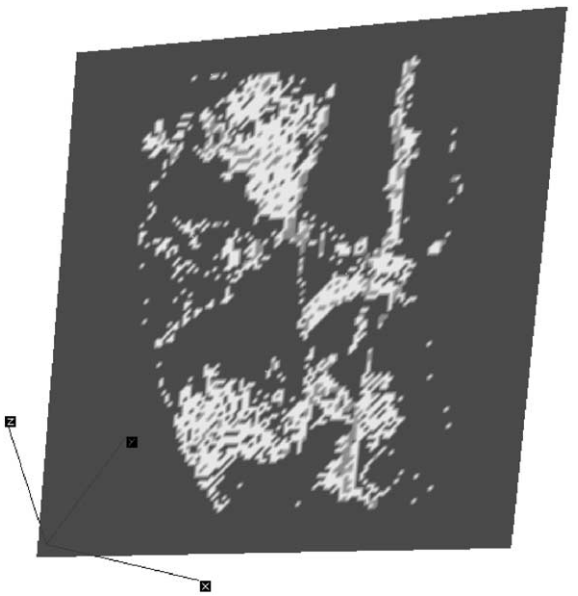

Fig. 8. In a perspective view both the position/orientation and the local thickness of the fracture, relative to a real data set, are shown. The black areas identify the regions of the plane where the fracture is, in fact, not present. The local thickness increases from yellow to green, blue and white.

4.2. Identification of the other fractures

In the previous section we described the procedure adopted for the localisation and characterisation of the leading fracture in the data set. The localisation of the other fractures can be very critical. A quasi-planar point-set generates in the Hough space a cumulation region (a cloud of points with significantly large cumulation values) such as the one shown in Fig. 5. When the data exhibit several planar point-sets, a corresponding number of cumulation regions will appear in the Hough space. Such "clouds" will often intersect one another quite heavily (depending on the relative position and orientation of the fracture planes) with the result that:

- the geometry of the overlapped regions changes dramatically (see Fig. 9);
- "weaker" cumulation regions are often "masked" by stronger overlapping regions.

Because of the above reasons, cluster detection [15,2] followed by local peak search in the Hough space cannot guarantee reliable results.

In order to overcome such difficulties, we developed a novel approach to the identification of the planar data structures, based on a combined analysis of the Hough space and the data space. Here are the main steps of this approach:

1. Computation of the HT;
2. Detection of the dominant plane through the previously-described approach (Hough space analysis);
3. Identification of the fracture voxels (data space analysis) that pertains to the detected fracture;
4. Estimation of the HT associated to the dominant fracture and subtraction of this transform from the HT of the whole data set; the residual HT is now purged of the contribution of the dominant plane, therefore we can proceed with the search of the second to the second most dominant plane by repeating steps 2–4.
5. The process is iterated until all the significant fractures are identified. One reasonable stopping criterion consists of checking whether the volume of

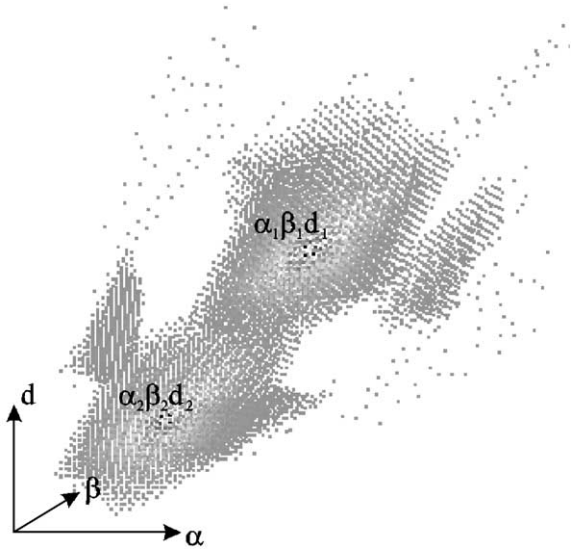

Fig. 9. A perspective view of the cumulation array (Hough space) associated to a data set containing two planar structures plus a number of significant voxels randomly scattered in the data space. A grey scale is used to visualise the values associated to the cells (see also Fig. 5). In order to provide a clearer view, only cells whose value is > 100 are shown.

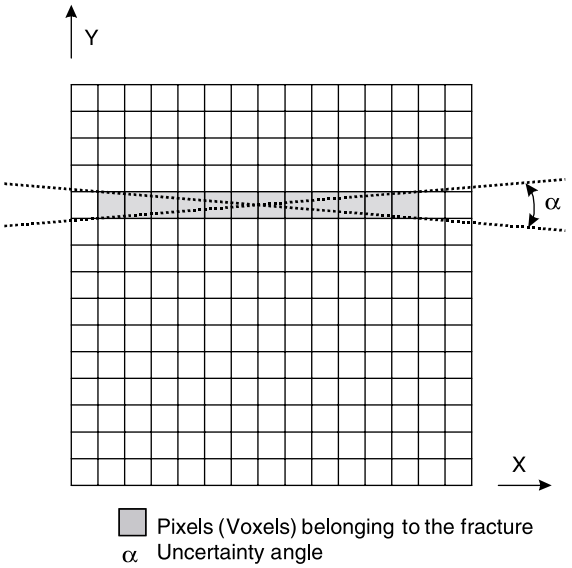

Fig. 10. Fracture localisation uncertainty.

the last detected fracture is negligible compared to the dominant fracture.

As we can see, the point of strength of this plane-search procedure is in the fact that at every step we always search for the dominant plane among what is left from the previous compensation.

It is interesting to notice that each fracture voxel contributes to the Hough map only once. Similarly, the removal of the contribution of the (current) dominant fracture must be done in such a way to count the contribution of each fracture voxel only once as well. This way the residual of the Hough map will never be negative.

5. Experimental results

The proposed approach for fracture detection has been tested on a variety of synthetic and natural data sets. The size of the considered voxsets are $50\times50\times50$ and $100 \times 100 \times 100$ voxels. As for the quantisation of the Hough space (the size of the elementary cell in the cumulation array) we chose for the two angles a step of $2^o \approx 0.035$ rad, while for the distance d (see Eq. (7)) the step size was chosen to be equal to the voxel side (typically 1 linear unit). The choice of a 2^o stepsize for the angular coordinates is related to the fact that too dense a quantisation would only increase the computational load without significantly improving the ability to localise fracture planes. This can be attributed to the fact that a finite voxel size is responsible for a certain degree of ambiguity in the exact definition of the fracture's orientation. For example, a fracture having a square extension of 60 voxels per side, and a thickness of 1 voxel will carry an intrinsic localisation uncertainty of about 2^o (see Fig. 10).

As for the performed tests, we first considered simple synthetic data sets that included a single fracture with various locations, orientations, extensions and thickness (fixed or variable) and a number of significant voxels randomly scattered in the data volume. In such cases the structure detected by the algorithm corresponded to the fracture. We then considered more complex synthetic data sets such as the one of Fig. 11.

This data set includes six fractures of different shapes and dimensions. For all the fractures the thickness is one voxel. The data set also includes a large number of scattered significant voxels. The results

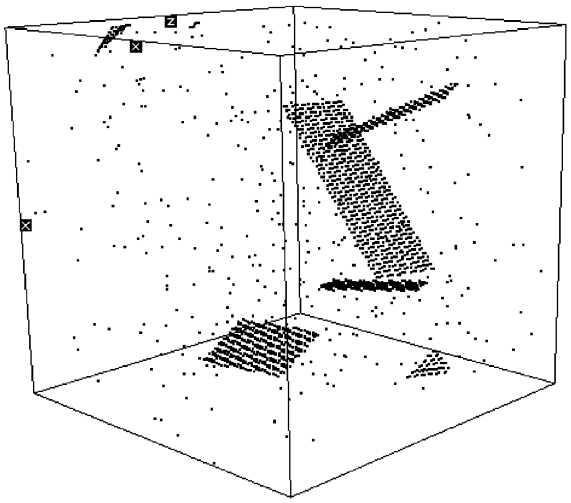

Fig. 11. A perspective view of a synthetic data set that includes six fractures with different shapes and sizes. All the fractures are one voxel thick. There is also a large number of "significant" voxels randomly scattered in the volume.

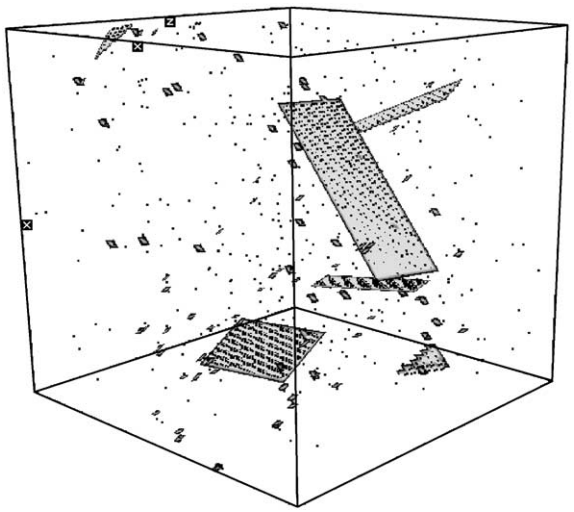

Fig. 12. A perspective view of the results obtained from the analysis of the data set shown in Fig. 11. The six fracture actually present on the data are perfectly identified (see how the grey areas perfectly overlap the black ones). Moreover, some other very small fractures are detected. This is due to the amount of scattered "significant" voxels.

of the analysis are shown in Fig. 12. The six fractures that are actually present in the data are perfectly identified (see how the grey surface patches perfectly

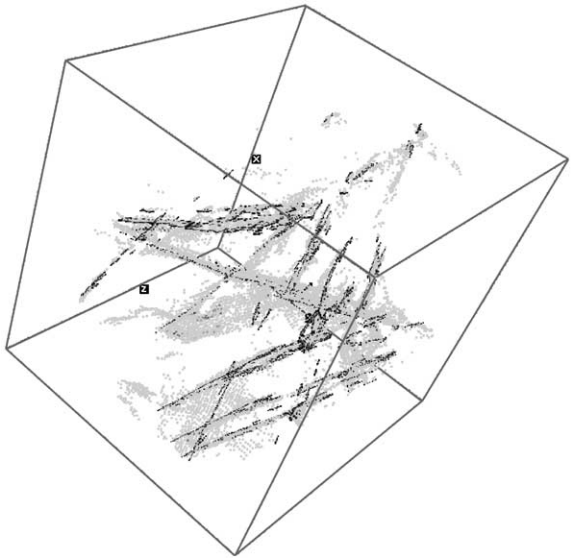

Fig. 13. Another perspective view of the data set of Fig. 1. The fracture models (planes whose transparency depends on fracture thickness) are superimposed.

overlap the black ones), with a very strong resilience against the randomly scattered significant voxels. In fact, additional other very small fractures are detected by the method. Such fractures are, in fact, generated by the random coplanarity of significant voxels within the dense clouds of scattered data. Such additional planes can be easily removed after checking the volume, the extension, and the distribution of the detected fractures. As far as the accuracy of the localisation is concerned, it always turned out to be below the predicted level of intrinsic uncertainty.

This experiment also allowed us to test the criterion we adopted to determine when to end the iterations. The volume of the data set was 10^3 cm^3, while the surface of the leading fracture was 173 cm^2 and that of the other five fractures was always above 0.35 cm^2, which corresponds to 0.2% of 173 cm^2. Knowing this allows us to set the threshold that will decide when to end the iterations. In fact, we can decide not to consider fractures whose area is no greater than 0.2% of the leading one.

We also carried out some tests on real data sets. The results relative to the data set of Fig. 1 are presented in Figs. 13–15. The voxel size is determined considering the size and the accuracy of the data set. Since the

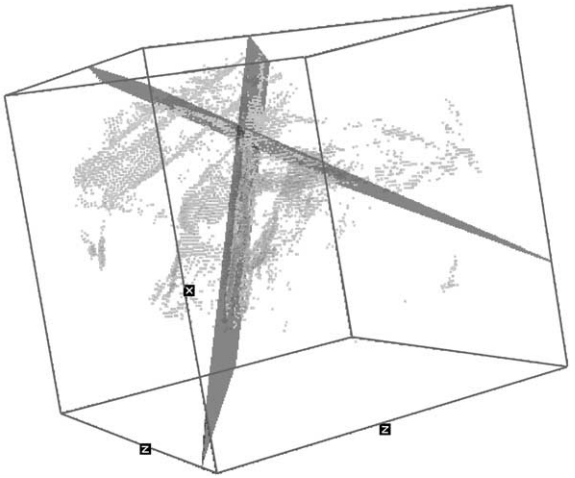

Fig. 14. A different perspective view of the data set of Fig. 1. The two leading fracture planes are superimposed.

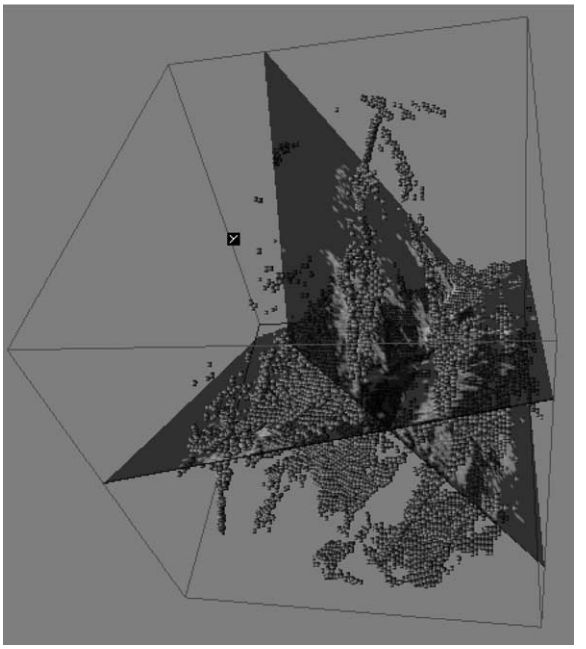

Fig. 15. Perspective view of the data set of Fig. 1, with the two leading fracture planes. The fracture thickness appear as an image on the fracture plane. Fracture planes are partially transparent for visualisation purposes.

spatial resolution of the tomographer we used to extract volumetric data is about 1 mm and the rock samples we analysed are cylinders of about 8–10 cm in diameter, we used a voxset size of about $100 \times 100 \times 500$ voxels. A perspective view of both the original data set and the estimated fracture plane representation are shown in Fig. 13. In Figs. 14 and 15 we can see the planes relative to the two leading fractures. Using the same thresholding method adopted with synthetic data, and ruling out fractures whose area is no greater than 2% of the area of the leading fracture, we end up with 18 fractures, in order of relevance, all accurately positioned in 3D space (and all with a measurement of the relative volume and area), which is more than a human operator would ever be able to achieve. In fact, both the original data set and the obtained results were accurately inspected by an expert analyst using a volume/surface rendering system. This manual inspection process confirmed the quality of proposed technique.

It is important to notice that this method returns the parameters of the fracture planes and the related data in a reverse order of importance. Occasionally, usually with data sets that include a large number of fractures that are close to each other, the order in which the fractures are detected may not be exactly the one related to their extension. This is due to the fact that the cumulation regions of two close fractures (see Figs. 5 and 7) are so heavily overlapped that the minor peak ends up increasing above the level of the major peak. This, however, does not prevent all fractures from being correctly estimated. All that could be required in that case is a post-sorting process based on the fracture extension/volume.

6. Some software implementation details

The software of our algorithm has been implemented in C++. The current version runs on Windows, Linux and IRIX platforms. In particular the processing of a $128 \times 128 \times 128$ voxset similar to the one of Fig. 1, running on a PC equipped with a PIII 700 Mhz and a 256 MB RAM, takes about 5 min.

In order to develop and validate the algorithm we also had to overcome significant visualisation problems for both data and fracture planes. In order to guarantee portability and to simplify

future SW expansions, we developed a volume/surface visualisation SW based on JAVA 3D. All the images of data sets and results have been obtained with this visualisation tool. It is important to underline that a standard PC with a modern accelerated video board enables a very comfortable and natural visualisation and manipulation of the considered data sets and of the relative results.

7. Conclusions

In this article we proposed an algorithmic approach to planar fracture detection and characterisation based on a particular implementation of the HT. One of the original aspects of our approach is in the pattern detection strategy. Our approach, in fact, is iterative, and it consists of the progressive removal of the cumulative layers in the Hough space. At every step, we determine the leading fracture among the residual ones after the removal of the previous cumulative layer.

The proposed solution proved effective even with the most complex data sets, always with better accuracy than what predicted.

It is important to emphasise the fact that our approach to fracture detection is general enough to accommodate a variety of applications of pattern detection. In particular, we can change reference model in order to detect different types of surfaces. Examples of surface models that could be detected are spherical, parabolic or hyperbolic, although the Hough space would have to have more dimensions than three.

Future developments of this technique concern the possibility to use greyscale information to improve plane detection. Considering the discussion given in the Introduction about using greyscale values for this purpose, we could devise a strategy that uses gradient information to quickly obtain a rough estimate of the planes in 3D space, and then uses a greyscale version (without gradient information) to refine the extracted data. Indeed, this second refinement phase, would need a preliminary pre-processing phase to limit data to fracture regions. Having a rough estimate of the planes would also help reducing the search space for the second phase.

Acknowledgements

We wish to thank ENI-AGIP (national Italian oil company) for financially supporting this research activity and for giving us permission to publish this article with their data sets. In particular, we are grateful to Drs. Daniela Mattiello and Massimo Balzarini for their help during the research activities, for collecting and preparing the volumetric data and for the fruitful discussions on the study of fractures in rock samples.

References

[1] P.M. Adler, J.F. Thovert, Fractures and Fracture Networks, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1999.
[2] J.C. Bezdek, Some new indexes of cluster validity, IEEE Trans. Systems Man Cybernet. 28 (3) (June 1998) 301–315.
[3] M. Bomans, K.H. Hohne, G. Laub, A. Pommert, U. Tiede, Improvement of 3D acquisition and visualisation in MRI, Magn. Res. Imaging 9 (1991) 597–609.
[4] A.R. Brown, Interpretation of Three-Dimensional Seismic Data, 5th Edition, Society of Exploration Geophysicists, 1995.
[5] J.P. Chilès, Fractal and geostatistical methods for modelling of a fracture network, Math. Geol. 20 (6) (1988) 631–654.
[6] M.P. Coward, T.S. Daltahan, H. Johnson, et al., Structural Geology in Reservoir Characterisation (Special Publications), Geological Society, London, 1998.
[7] P. Davy, A. Somette, D. Sornette, Experimental discovery of scaling laws relating fractal dimensions and the length distribution exponent of fault systems, Geophys. Res. Lett. 19 (4) (1992) 361–363.
[8] N.H. Dawers, M.H. Anders, C.H. Scholz, Growth of normal faults: Displacement-length scaling Geology 21 (1993) 1107–1110.
[9] R.O. Duda, P.E. Hart, Use of the Hough transformation to detect lines and curves in picture, Commun. ACM 15 (1) (January 1972) 11–15.
[10] W.A. Edelstein, Oil core NMR imaging/spectroscopy instrumentation, Magn. Res. Imaging 9 (1991) 67–865.
[11] J. Foley, A. van Dam, S. Feiner, J. Hughes, Computer Graphics: Principles and Practice, Addison-Wesley, Reading, MA, 1995.
[12] M. van Ginkel, L.J. van Vliet, P.W. Verbeek, M.A. Kraaijveld, E.P. Reding, H.J. Lammers, Robust curve detection using a radon transform in orientation space applied to fracture detection in borehole images, in: R.L. Lagendijk et al. (Eds.), ASCI'01, Proceedings of the Seventh Annual Conference of the Advanced School for Computing and Imaging, Heijen, The Netherlands, May 30–June 1, ASCI, Delft, 2000, pp. 299–306.

[13] E. Gringarten, 3-D geometric description of fractured reservoirs, Math. Geol. 28 (7) (1996) 881–893.
[14] J. Illingworth, J. Kittler, A survey of the Hough transform, Comput. Vision, Graphics Image Process. 44 (1988) 87–116.
[15] A. Jain, R. Dubes, Algorithms for Clustering Data, Prentice-Hall, Englewood Cliffs, NJ, 1988.
[16] K. Koike, K. Kaneko, Characterisation and modelling of fracture distribution in roch mass using fractal theory, Geothermal Sci. Technol. 6 (1999) 43–62.
[17] K. Koike, K. Komorida, Y. Ichikawa, Fracture-distribution modelling in rock mass using borehole data and geostatistical simulation, Annual Conference of the International Association for Mathematical Geology, IAMG 2001, Cancún, Mexico, September 6–12, 2001.
[18] J.C.S. Long, D.M. Billaux, From field data to fracture network modelling: an example incorporating spatial structure Water Resources Res. 23 (2) (1987) 1201–1216.
[19] National Research Council, Rock Fractures and Fluid Flow, Contemporary Understanding and Applications, National Academic Press, Washington, DC, 1996.
[20] F. O'Gorman, M.B. Clowes, Finding picture edges through collinearity of feature points, IEEE Trans. Comput. 25 (4) (1976) 449–456.
[21] L.J. Pyrak-Nolte, C.D. Montemagno, D.D. Nolte, Volumetric imaging of aperture distributions in connected fracture networks, Geophys. Res. Lett. 24 (18) (September 1997) 2343–2346.
[22] J. Sklansky, On the Hough technique for curve detection, IEEE Trans. Comput. C-27 (10) (October 1978) 923–926.
[23] B.B. Thapa, P. Hughett, K. Karasaki, Semi-automatic analysis of rock fracture orientations from borehole wall images, Geophysics 62 (1) (January–February 1997) 129–137.
[24] F.M. Verga, G. Giglio, F. Masserano, L. Ruvo, Calibration of fractured reservoirs with dynamic data, SPE Reservoir Simulation Symposium, February 11–14, 2001, Houston, TX, USA.
[25] M. Vignes-Adler, A. La Page, P.M. Adler, Fractal analysis of fracturing in two African regions, from satellite imagery to ground scale, Technophysics 196 (1992) 69–85.
[26] P. Weimer, T.L. Davis (Eds.), Application of 3-D Seismic Data to Exploration and Production, Society of Exploration Geophysicists, 1996.
[27] R.K.K. Yip, P.K.S. Tam, D.N.K. Leung, Modification of Hough transform for circles and ellipsis detection using a two-dimensional array, Pattern Recognition 25 (9) (1992) 1007–1022.