2014/04/29 v0.4.3 IPOL article class

IPOL
IMAGE PROCESSING ON LINE

Published in Image Processing On Line on YYYY–MM–DD.
Submitted on YYYY–MM–DD, accepted on YYYY–MM–DD.
ISSN 2105–1232 © YYYY IPOL & the authors CC–BY–NC–SA
This article is available online with supplementary materials,
software, datasets and online demo at
http://dx.doi.org/10.5201/ipol

Hough Spaces for the Detection of Alignments

PREPRINT July 6, 2015

Abstract

We review some different versions the Hough transform for the detection of straight lines in images. These choices concern the interpretation of the values of the input image (intensities or gradients), the voting strategy for each point (a single point or a whole curve in the transform space), and the coordinates in the transform space (polar or projective). We study the effect of filtering the original or the transform images on the quality of the most significant detection. We recall the relationship with the Radon and Fourier transforms. We explain how the same ideas apply to other parametric clustering problems, like the detection of strata in borehole imaging. The important issue of detection thresholds is left untreated.

Keywords: alignments, Hough transform, sinusoids, borehole imaging

1 Introduction

Overview. The Hough transform is a venerable method for the detection of alignments in images. Its main advantages are that it produces very precise detections and is very robust to noise and occlusions. Its main disadvantage is that, by itself, it lacks a satisfying decision theory. The Hough transform says: if there are straight lines on this image, they must be exactly here and there. However, it does not answer the question of whether there are any straight lines after all. In this report, we conveniently ignore this difficult decision theory and we recall the geometric details of the Hough transform and its variants.

History of the problem. The method was invented by Paul Hough in 1959 [7] to detect tracks in bubble chamber photographs (Figure 1). It can be described vividly as a voting procedure. Each *candidate* track is represented by a pair of numbers (a,b), corresponding to the straight line of equation $y = ax+b$. Then, a bubble observed at the position (x,y) *votes* for all the possible candidate lines passing through this point, namely the straight line of equation $b = -xa+y$ in the ab-plane. The number of votes that a point (a,b) receives is equal to the number of bubbles traversed by the track $y = ax+b$. Thus, the tracks are recovered as the candidates (a,b) that receive the most votes.

This method was patented by Hough in 1962 [12]. The image processing community got hold of it after it was popularized by Duda and Hart [4], where it gained the name "Hough transform" [5]. It was generalized to detect lines in non-binary images and to use the polar representation of lines, instead of the ab-plane. The same ideas were also generalized [1] to the problems of finding circles, ellipses and even arbitrary curves [8].

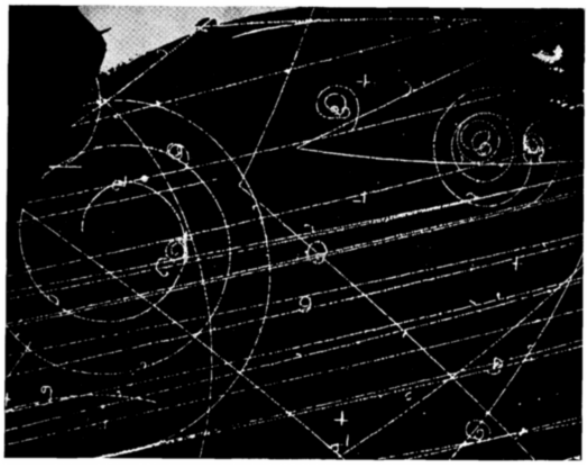

Figure 1: The bubble chamber image used by Hough to illustrate his method [7]. The image shows thousands of tiny white bubbles against a black background. Most of the bubbles are aligned on straight or spiral tracks. The goal is to detect automatically the position of each straight track.

Context: transformation, detection and decision. The Hough transform method for the detection of alignments is a very particular, low-dimensional, example of clustering by accumulation, like mean-shift or K-means []. All the variants have the same structure. Given an image $I : \Omega \to \mathbf{R}$ defined on a rectangle $\Omega \subseteq \mathbf{R}^2$, its transform is a function $I^H : \Theta \to \mathbf{R}$ defined on a certain set of parameters Θ. Each $\theta \in \Theta$ represents a straight line (or whatever feature is to be detected). Finally, the result of the detection is given by the local maxima of the function I^H. The transform I^H is computed by the following algorithm:

1. Set $I^H(\theta) = 0 \quad \forall \theta \in \Theta$
2. For each point $\mathbf{x} \in \Omega$ of the image domain:
   2.1. Let $\ell_{\mathbf{x}} \subseteq \Theta$ be the set of parameters compatible with $\mathbf{x}$
   2.2. For each $\theta \in \ell_{\mathbf{x}}$, increase the value of $I^H(\theta)$ by 1

In practice, the votes in step 2.2. are weighted by $\|\nabla I(\mathbf{x})\|$ so as to give preference to regions of high gradient. The many variants of the Hough transform that can be found in the literature correspond to answers of the following questions:

Q1: How are straight lines parametrized as points $\theta \in \Theta$?
Q2: Given a point $\mathbf{x} \in \Omega$, how to define the set of compatible lines $\ell_{\mathbf{x}}$?
Q3: How to discretize the space Ω?
Q4: How to discretize the space Θ?
Q5: How to pre-process the image $I : \Omega \to \mathbf{R}$?
Q6: How to post-process the transform $I^H : \Theta \to \mathbf{R}$?
Q7: Which local maxima of I^H to choose?

This last question is very difficult. Notice that the function I^H will always have at least one local maximum, even when there are no alignments. In this report we try to answer all these questions but the last one.

Summary of the parametrizations. Let us summarize the most common answers to question Q1 above. The set of all straight lines in the plane is a manifold of dimension 2 (called the projective plane). Thus, any local chart of this manifold is a possible choice for $\Theta \subseteq \mathbf{R}^2$. Since the projective plane is not isomorphic to any subset of the plane, the local chart Θ is necessarily incomplete, and some compromises must be accepted. There are two desirable properties for the parametrization:

P1: The set of all straight lines intersecting Ω is a bounded subset of Θ.
P2: The set of all straight lines passing through a given point $\mathbf{x} \in \Omega$ is a straight line in Θ.

The first property is desirable because it allows to define the function I^H on a bounded domain, without leaving any possible alignment out. The second property is desirable because it simplifies the computation of the transform: vote counting is performed along straight lines on the space Θ. Unfortunately, these two conditions are incompatible:

Proposition 1 (no nice parametrization)**.**
There does not exist any parametrization Θ of straight lines on the plane that satisfies conditions P1 and P2 at the same time.

Proof. A straight line is not bounded. □

Definition 1 (bounded and rectilinear parametrizations)**.**
A parametrization Θ is bounded *if it satisfies condition P1.*
A parametrization Θ is rectilinear *if it satisfies condition P2.*

Definition 2 (standard parametrizations)**.**
Let us consider the straight line ℓ of equation $px+qy+r=0$. The affine parametrization *of ℓ is the rectilinear parametrization given by $(a,b) = (-p/q, -r/q) \in \mathbf{R}^2$. The* oriented polar parametrization *of ℓ is the bounded parametrization given by $(\rho,\theta) = (r/\sqrt{p^2+q^2}, \arctan(q,p)) \in \mathbf{R} \times [0, 2\pi)$. The* polar parametrization *of ℓ is given by $(\rho,\theta) = (r/\sqrt{p^2+q^2}, \arctan(q/p)) \in \mathbf{R} \times [0, \pi)$.*

Remark. *The affine parametrization cannot represent vertical lines. Any other rectilinear parametrization has a similar problem: there is always a slope that can not be represented, and a sequence alignments converging to that slope are parametrized by an unbounded sequence of vectors θ.*

Remark. *Both polar parametrizations allow negative values of ρ, this is necessary for being a smooth chart around the straight lines that cross through the origin of coordinates. The oriented polar parametrizations are a double covering of the chart: each straight line is represented twice, corresponding to the two possible orientations of its normal vector.*

The standard parametrizations that we have defined above are in some sense canonical:

Proposition 2 (equivalence of rectilinear parametrizations)**.**
Any two rectilinear parametrizations are related by an homography of the plane Θ.

Proposition 3 (invariance of polar parametrizations)**.**
The differential form $\mathrm{d}\rho \wedge \mathrm{d}\theta$ is the only measure on Θ which is invariant under rigid motions of Ω.

Proof. See the first chapter of the book by Santaló [10]. □

Due to these results, we will only consider the affine and polar parametrizations of the set of straight lines. However, in the image processing literature there is a plethora of other names for parametrizations: cascaded Hough transform, parallel coordinates, Forman, Muff, Fan-Beam, and the Circle transform. See the booklet by Herout-Dubská-Havel [6] and the references therein for a very comprehensive overview. Some of these parametrizations are just projective dualities with a different convention. Others correspond to atlases of the projective plane formed by two or three bounded charts. The rest are bounded curvilinear representations that have sometimes nice properties but they cannot be invariant to rigid motions on Ω.

Summary of voting strategies. Let us summarize the most common answers to question Q2 above. Given a point $\mathbf{x} \in \Omega$, for what set of lines $\ell_{\mathbf{x}} \in \Theta$ shall it vote? In the case of bubble chamber images, the answer is clear: if $I(\mathbf{x})$ is black there is no bubble and there is no vote; if $I(\mathbf{x})$ is white, there is a bubble and it votes for all the straight lines passing through it. However, for general gray-scale images it is not clear what to do. We consider the following choices for the candidates:

C1: (binary voting) If $I(\mathbf{x}) > 0$, the point $\mathbf{x}$ votes for all the straight lines through $\mathbf{x}$.
C2: (intensity voting) The point $\mathbf{x}$ votes with weight $I(\mathbf{x})$ for all the straight lines through $\mathbf{x}$.
C3: (edge voting) Run method C1 for the Canny pixels of I.
C4: (gradient voting) Run method C2 for the image $\|\nabla I\|$.
C5: (oriented voting) The point $\mathbf{x}$ votes for the straight line through $\mathbf{x}$ perpendicular to $\nabla I(\mathbf{x})$

Note that in C1–C4, each point $\mathbf{x}$ votes for a whole curve in Θ, but in C5 each point votes for a single curve. Thus, method C5 is much more efficient. The method C2 does not seem very natural: it is supposed to find alignments of points of high intensity, not of low intensity. However, this method has interesting mathematical properties; most notably, it is a linear transform of the input image. Due to the linearity, it can be used to find alignments of bright and dark points alike. Alignments of bright points in I will appear as local maxima of I^H, and alignments of dark points of I will appear as local minima of I^H. See Figure 2 for examples of techniques C2–C5 over the same image.

Summary of discretization strategies. Let us summarize the most common answers to questions Q3 and Q4 above. The image domain Ω is always discretized by a regular grid. Once coordinates Θ have been chosen, the space of parameters is also discretized by a regular grid. However, there are two very different ways to fill-in the values of the discrete array $I^H : \Theta \to \mathbf{R}$. In the continuous setting the two ways are equivalent due to Fubini's theorem, but in the discrete case they give very different results.

The traditional way to compute the Hough transform is called *the writing paradigm*. This is the algorithm described above. In the language of voters and candidates: each voter $\mathbf{x} \in \Omega$, finds his list of candidates $\ell_{\mathbf{x}} \subseteq \Theta$ and then casts one vote for each of the candidates $\theta \in \ell_{\mathbf{x}}$. The alternate way to compute the Hough transform is *the reading paradigm*: each candidate $\theta \in \Theta$ finds the list of voters $\ell_\theta \subseteq \Omega$ that are going to vote for him, and collects all their votes.

Both the writing and the reading paradigms require interpolating the position of points in a regular grid, and this is typically done by nearest-neighbor interpolation. This interpolation is nonlinear (it is not commutative with the sums), so the two methods give very different results, with different artifacts each. These artifacts can be softened by filtering the input and output images. An exact sampling of these transforms is a difficult problem [2], out of the scope of this report.

Importance of smoothing the images. An appropriate smoothing of the input and output images is critical, especially in the case "C5" of orientation voting. It is well-known that the gradient direction of a quantized image is heavily quantized [3]. The effect of this orientation quantization is that most vertical columns of the image I^H are left empty, making the detection of maxima impossible. This problem is solved by de-quantizing the input images: they are always smoothed always by a Gaussian kernel of small standard deviation ($\sigma = 0.8$, for example), prior to computing the gradient direction using finite differences.

Then, it is necessary to smooth the transformed image. This smoothing eases the localization of maxima in the domain Θ. For the examples of this report we have always used a Cauchy kernel of scale 1 pixel. We found that this fat-tailed kernel is very good at joining many local maxima into a single one, even if the initial maxima are much farther than the scale of the kernel 5.

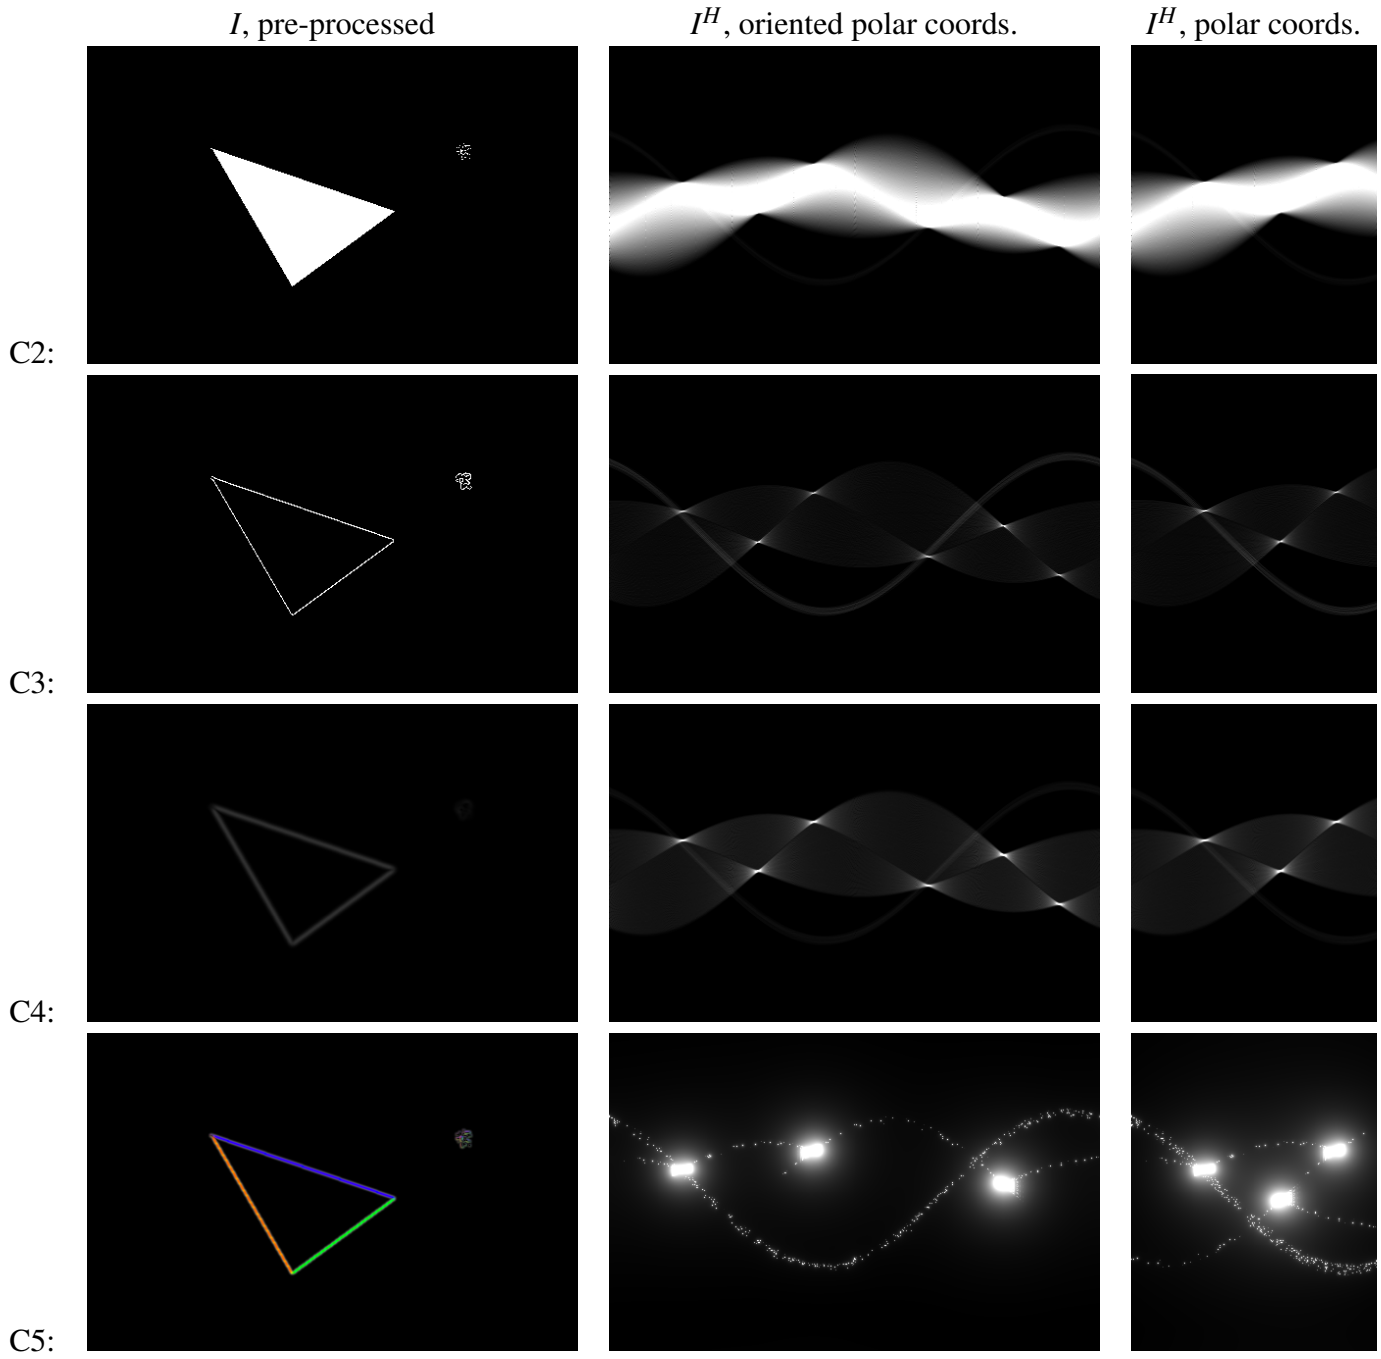

Figure 2: The Hough transform of a synthetic image for the voting methods C2–C5, using polar and oriented polar coordinates. Notice that for methods C2–C5, the oriented polar coordinates are redundant and each side of the triangle appears two times (corresponding to the two possible orientations of the side). For the voting method C5, each point on the boundary only votes for the correct orientation, so the oriented polar coordinates give a different information.

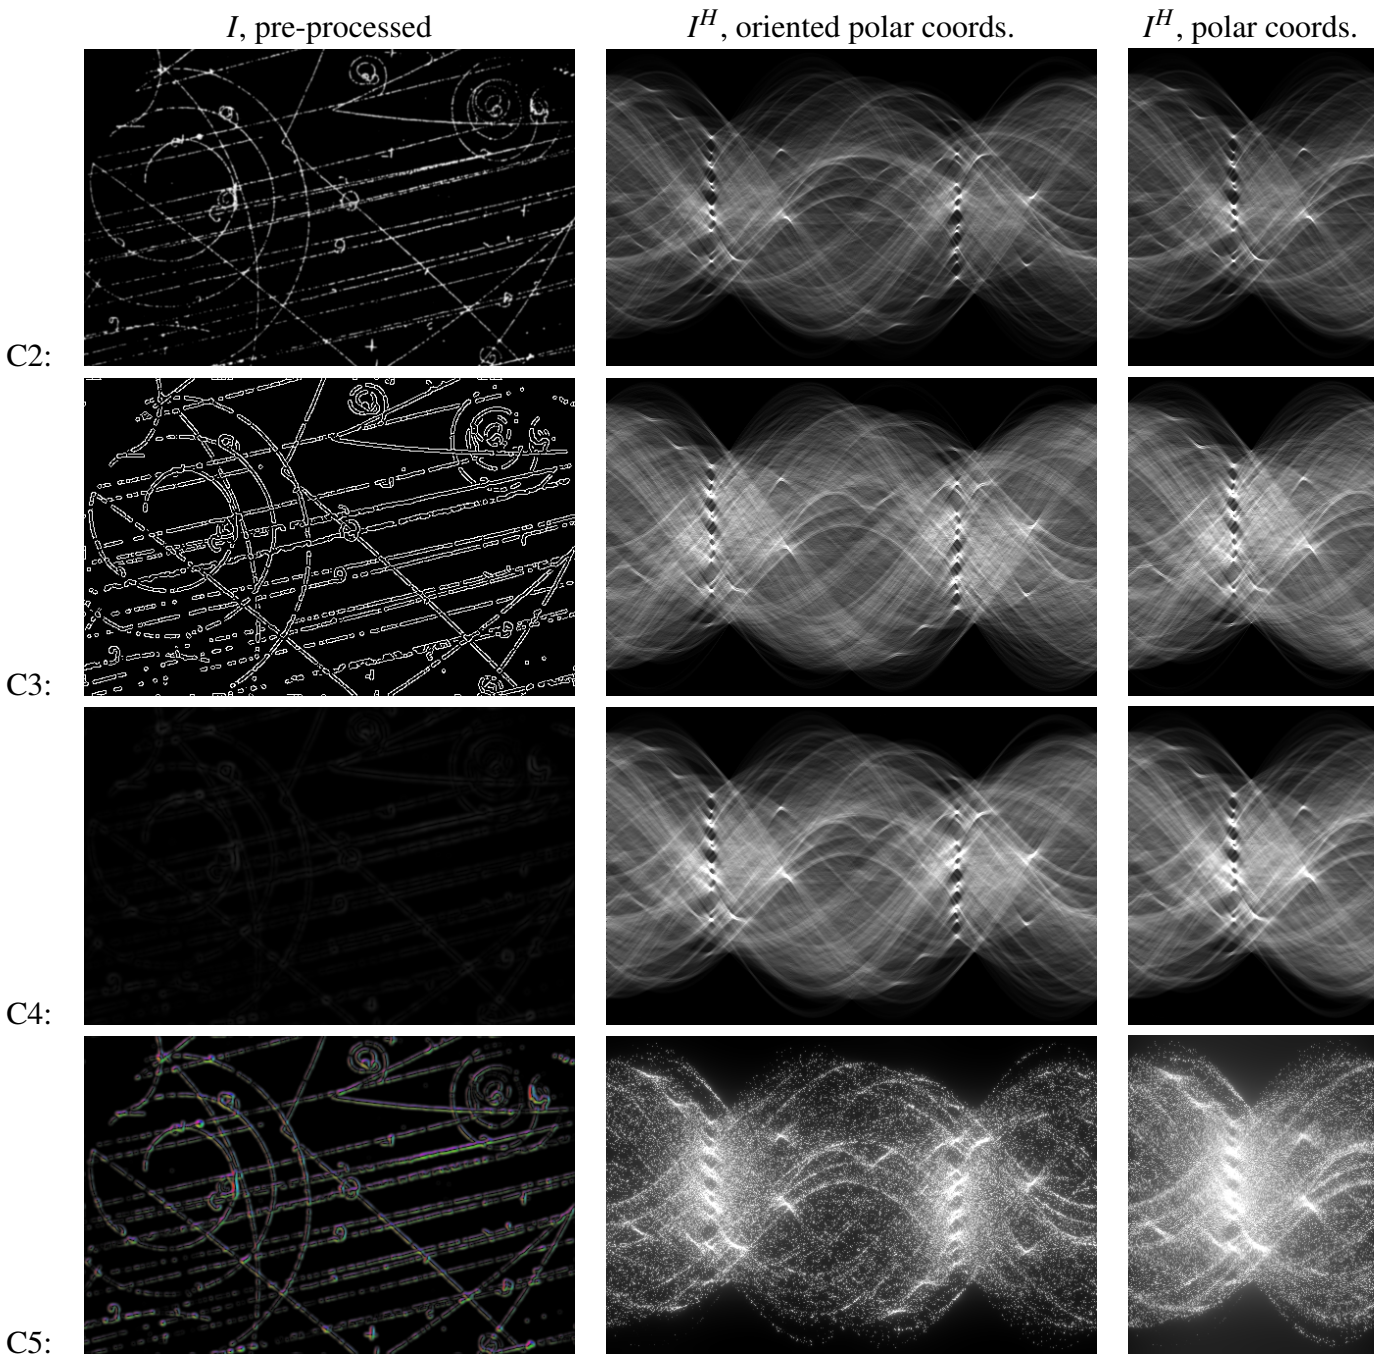

Figure 3: The Hough transform of a bubble chamber image for the voting methods C2–C5, using polar and oriented polar coordinates. Here the only experiment that makes sense is C2. Observe that bundles of parallel lines appear in the transform as vertically aligned local maxima.

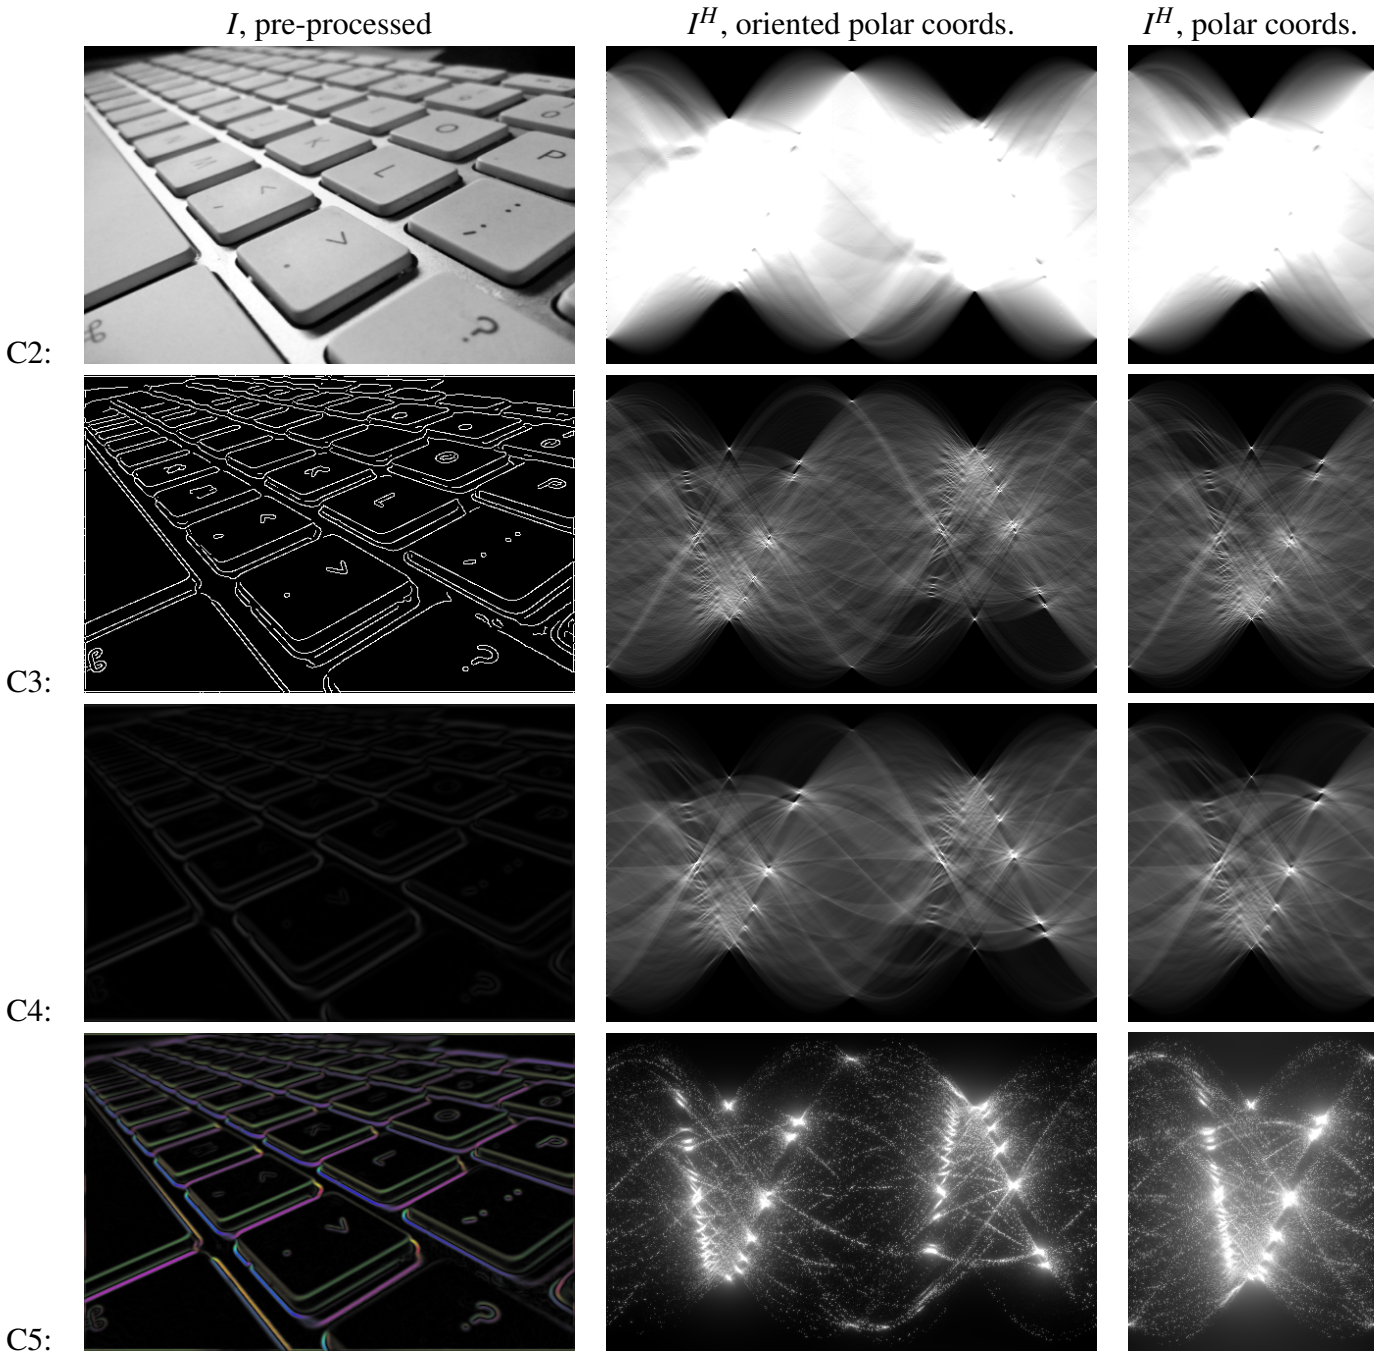

Figure 4: The Hough transform of a real gray-scale photograph, for methods C2–C5, using polar and oriented polar coordinates. For gray-scale images the best localization of the maxima is clearly obtained by C5.

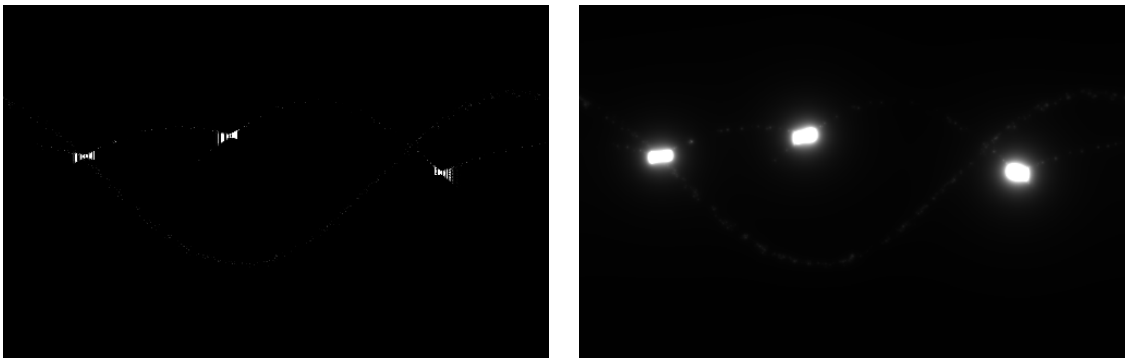

(a) I^H (b) Cauchy$(1) * I^H$

Figure 5: Raw-accumulation of orientation votes vs. Smoothed-accumulation by a Cauchy kernel.

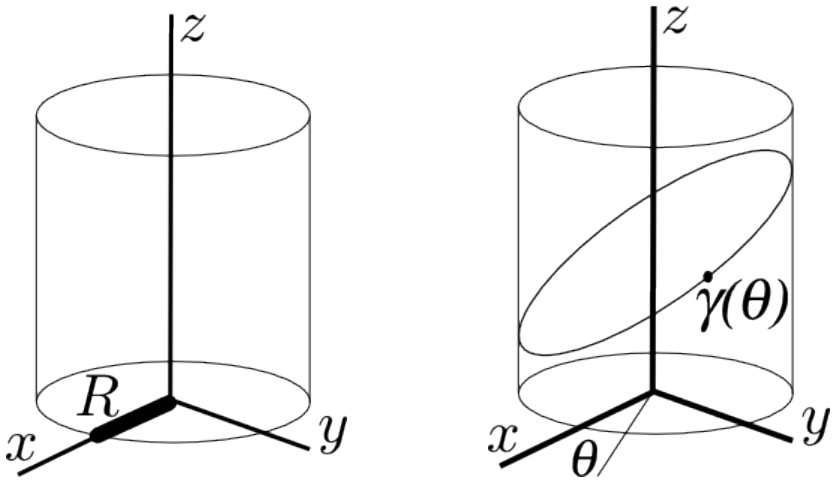

(a) Coordinates around a cylinder (b) Intersection of cylinder and plane

Figure 6:

Interest in borehole imaging. The borehole imaging technology has been developed in the oil industry to analyze the structure associated with sedimentary layers. The resulting images provide useful information concerning reservoir evaluation. For the purposes of this report, a borehole is a cylindrical vertical hole of infinite depth. A borehole imaging device travels inside this cylindrical hole and measures a physical quantity along the boundary, producing a cylindrical image that is unfolded into a rectangular domain (Figure 7). These rectangular images typically have a width of 360 pixels, corresponding to the 360 degrees of the cylinder, and an arbitrarily large height.

The Hough technique is naturally adapted to the detection of planar structures intersected by the borehole. These structures appear as sinusoids in the developed rectangular image. All the considerations that we explained for the case of alignments are transported to the detection of sinusoids, we only have to choose A) a voting strategy for each pixel in the image domain and B) a parametrization of the planes. The choice of A) is easy because the images are gray-scale, thus a local orientation can be computed (method "C5"). The choice of B) is somewhat arbitrary; we propose the representation of a plane by the equation $z = ax + by + c$, and then by the pair (a,b). This representation is rectilinear: given a gradient orientation at a point $\mathbf{x}$, there is a one-parameter family of planes compatible with this gradient, and they form a straight line in the ab-plane. See Figures 8 and 9.

The dip picker transform Let us consider a vertical cylinder of radius R (Figure 6) and a plane of equation $z = ax + by + c$. The intersection of the plane and the cylinder is an ellipse that has the following parametric equation:

$$\gamma(\theta) = \begin{pmatrix} aR\cos\theta \\ bR\sin\theta \\ aR\cos\theta + bR\sin\theta + c \end{pmatrix}$$

When developed on the (θ, z) plane, it is a sinusoid. These sinusoids have three degrees of freedom, corresponding to the three parameters of the plane. Now, if we observe a gradient on the developed image, this gradient must be perpendicular to the tangent $\gamma'(\theta)$ of the sinusoid through. This condition fixes two of the degrees of freedom, leaving a one-dimensional space of possibilities. More precisely, if we observe an image gradient (u,v) at point (θ,z), the coefficients (a,b,c) of the planes that are compatible with this observation satisfy the equation

$$-a\sin\theta + b\cos\theta + \frac{u}{v} = 0 \tag{1}$$

The set of solutions of this equation is a straight line on the ab-plane.

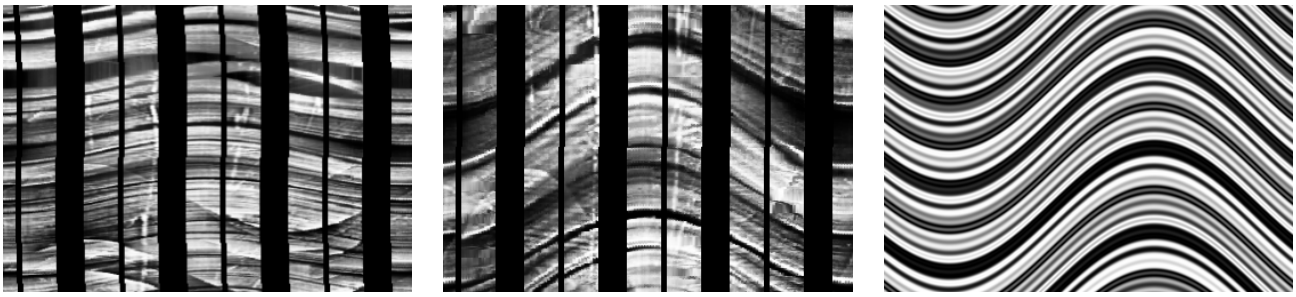

(a) Part of a borehole image (b) Central third of (a) (c) Synthetic ideal case

Figure 7: Borehole images have width 360 and infinite height. The height interval and vertical scale are adjusted interactively for visualization. Here the images (a) and (b) show two different intervals around the same central height. Image (c) is a simulated ideal case of borehole with perfectly parallel strata. The black vertical bands on the real images are due to blind spots in the imaging device.

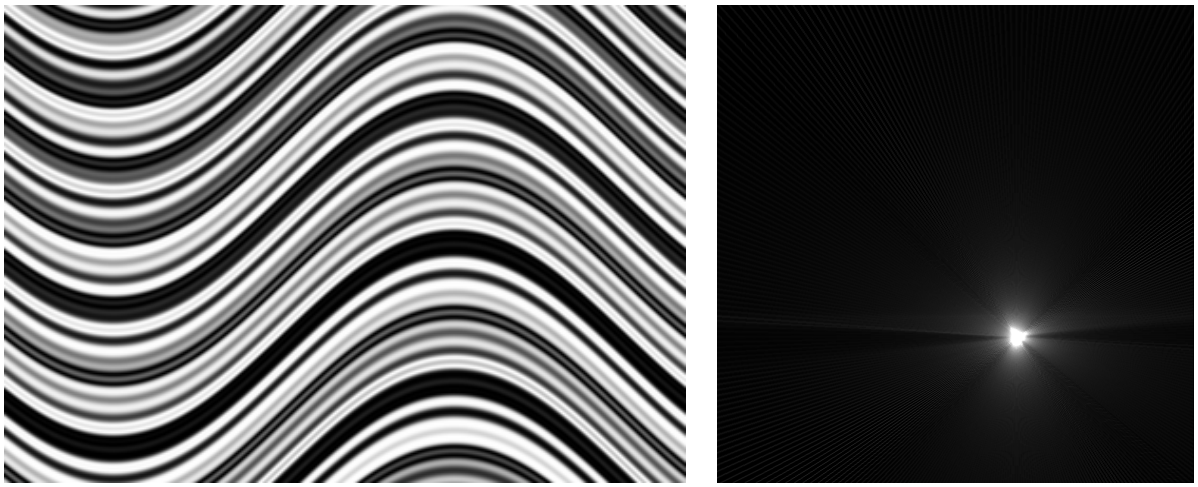

(a) Synthetic borehole image I (b) Dip picker transform I^H

Figure 8: Ideal case of the dip picker transform. The transform image is clearly unimodal, indicating a single bundle of parallel strata.

Definition 3. *The* dip picker transform *is the voting procedure obtained by using equation (1) to obtain a set of candidates for the Hough transform algorithm.*

The Hough transform for dip picking is not a recent idea. It is described by a Halliburton patent [11], using the intensity transform; and later by a Schlumberger patent [9], using the better-suited technique of voting according to the gradient direction.

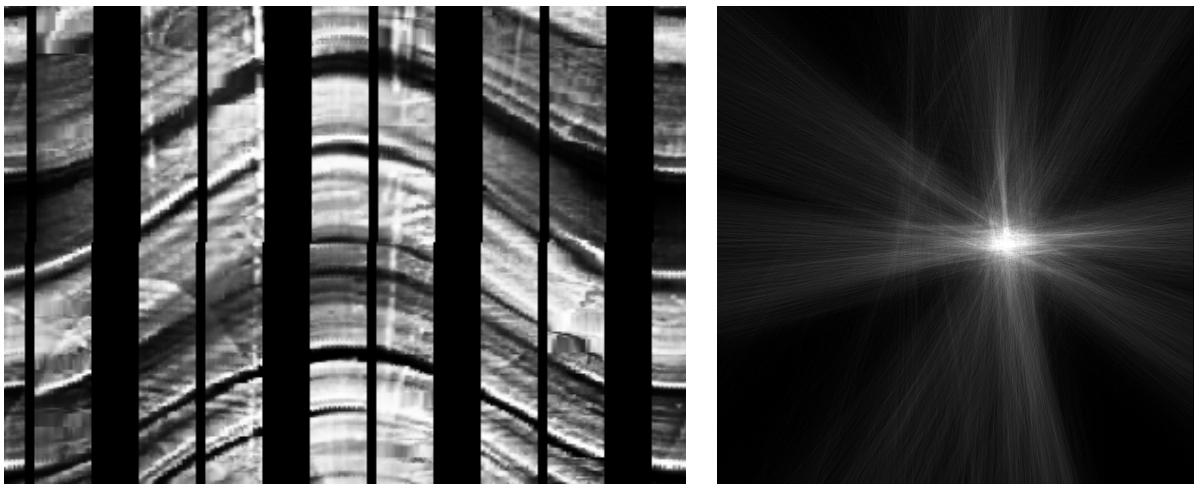

(a) Real borehole image I (b) Dip picker transform I^H

Figure 9: Real case of the dip picker transform. The transform image has one large mode, indicating a bundle of parallel strata, but there are also other structures. The vertical structures are invisible in this part of the domain.

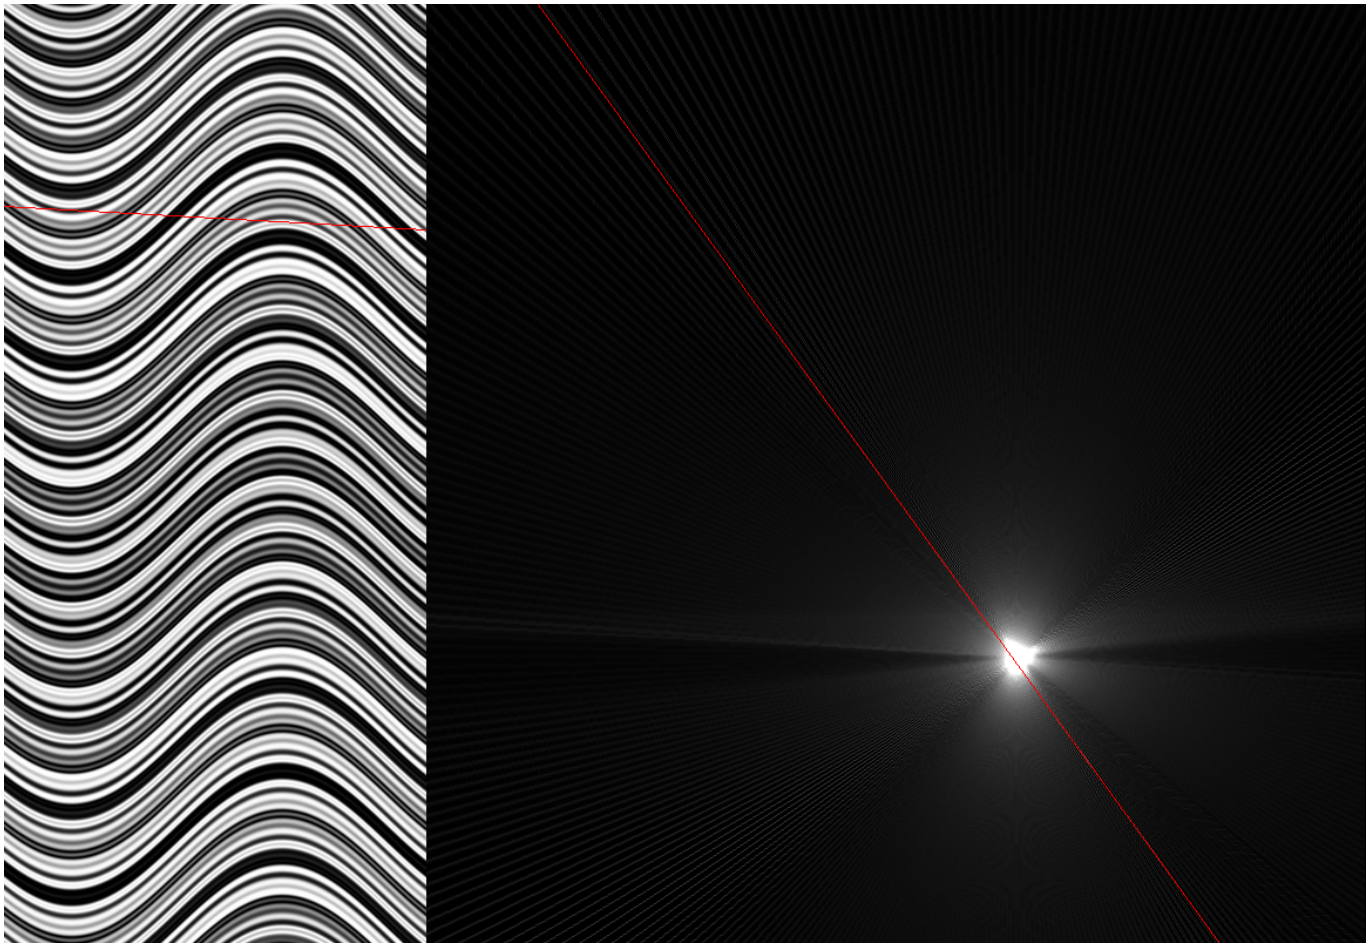

Figure 10: Illustration of the dip picker transform. Right: a tangent to a sinusoid. Left: the straight line on the *ab*-plane corresponding to this tangent, according to equation (1).

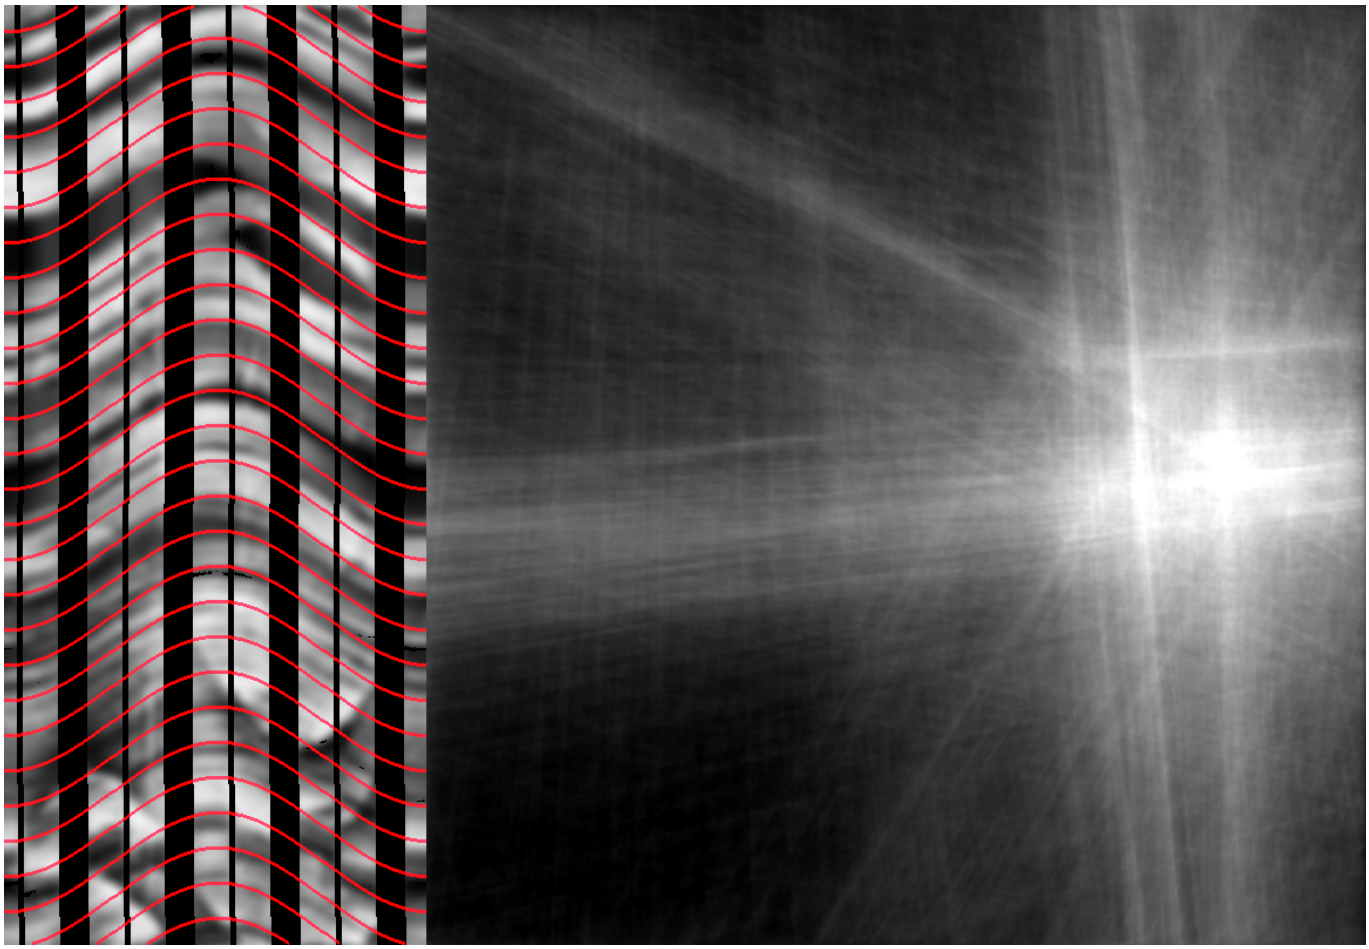

Figure 11: The main bundle of sinusoids detected as the global maximum of the dip picker transform.

Conclusion For gray-scale images it seems to be much better to vote by orientation [9] than by intensity [11]; the method is more general, and the results are more stable, more precise and faster to compute. Some care must be taken to de-quantize the orientation field computed from the images. In the future we would like to try other orientation fields such as the structure tensor or contour stencils (instead of the smoothed gradient).

To detect bundles of parallel or nearly parallel strata, it is very convenient to use a two-dimensional parameter space that discards the z-depth of each plane. The orientation-Hough transform for these images can be computed in real-time for small portions of the cylinder. When the inclination of the strata is assumed to be bounded (e.g., no vertical or near-vertical planes), then the best parametrization is the pair (a,b) to represent the plane of equation $z = ax + by + c$.

References

[1] DANA H BALLARD, *Generalizing the hough transform to detect arbitrary shapes*, Pattern recognition, 13 (1981), pp. 111–122.

[2] GREGORY BEYLKIN, *Discrete radon transform*, Acoustics, Speech and Signal Processing, IEEE Transactions on, 35 (1987), pp. 162–172.

[3] AGNÈS DESOLNEUX, SAÏD LADJAL, LIONEL MOISAN, AND JEAN-MICHEL MOREL, *Dequantizing image orientation*, Image Processing, IEEE Transactions on, 11 (2002), pp. 1129–1140.

[4] RICHARD O DUDA AND PETER E HART, *Use of the hough transformation to detect lines and curves in pictures*, Communications of the ACM, 15 (1972), pp. 11–15.

[5] PETER E HART, *How the hough transform was invented*, Signal Processing Magazine, IEEE, 26 (2009), pp. 18–22.

[6] A. HEROUT, M. DUBSKÁ, AND J. HAVEL, *Real-Time Detection of Lines and Grids: By PClines and Other Approaches*, SpringerBriefs in Computer Science, Springer London, 2012.

[7] PAUL VC HOUGH, *Machine analysis of bubble chamber pictures*, in International Conference on High Energy Accelerators and Instrumentation, vol. 73, 1959.

[8] JOHN ILLINGWORTH AND JOSEF KITTLER, *A survey of the hough transform*, Computer vision, graphics, and image processing, 44 (1988), pp. 87–116.

[9] NAOKI SAITO, NICHOLAS N BENNETT, AND ROBERT BURRIDGE, *Method of determining dips and azimuths of fractures from borehole images*, Sept. 28 1999. US Patent 5,960,371.

[10] L.A. SANTALÓ, *Introduction to integral geometry*, Actualités scientifiques et industrielles, Hermann, 1953.

[11] DAVID O TORRES, *Method for determining dip and strike angles in borehole ultrasonic scanning tool data*, Nov. 10 1992. US Patent 5,162,994.

[12] HOUGH PAUL VC, *Method and means for recognizing complex patterns*, Dec. 18 1962. US Patent 3,069,654.