

THREE-DIMENSIONAL FEATURE DETECTION USING OPTIMAL STEERABLE FILTERS

François Aguet, Mathews Jacob*, and Michael Unser

Biomedical Imaging Group
Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
*Beckman Institute, University of Illinois at Urbana-Champaign, USA

ABSTRACT

We present a framework for feature detection in 3-D using steerable filters. These filters can be designed to optimally respond to a particular type of feature by maximizing several Canny-like criteria. The detection process involves the analytical computation of the orientation and corresponding response of the template. A post-processing step consisting of the suppression of non-maximal values followed by thresholding to eliminate insignificant features concludes the detection procedure. We illustrate the approach with the design of feature templates for the detection of surfaces and curves, and demonstrate their efficiency with practical applications.

1. INTRODUCTION

The detection of features in three-dimensions is a computationally intensive task and thus represents an ideal application for the steerable filters described by Freeman et al. [1]. These filters can be very efficiently rotated by taking a linear combination of basis filters. Only very few approaches for 3-D feature detection using steerable filters currently exist, often using empirical approaches for the construction of the filters [2]. Traditionally, feature detection in 3-D is performed by computing the eigenvalues of the Hessian matrix [3]. In this paper, we provide a general framework for feature detection in three-dimensional data using filters that are optimal with respect to several Canny-like criteria [4]. We show how such filters can be designed to optimally respond to a particular type of feature, and describe the process of computing their optimal orientation and the corresponding response. Our approach follows the classical framework used in two-dimensional feature detection, which consists of a filtering step followed by non-maximum suppression and thresholding. Two types of structures frequently encountered in volumetric data are curves and surfaces. We derive optimal templates for the detection of these two types

This work was supported in part by the Swiss National Science Foundation under Grant 200020-101821. The second author would like to thank the Beckman foundation for support.

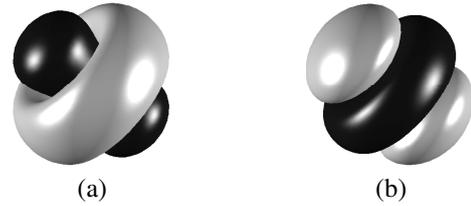


Fig. 1. Isosurface representations of the curve (a) and surface (b) detectors.

of features using the class of steerable filters based on second order Gaussian derivatives. In this context we also introduce a novel non-maximum suppression method for curves.

2. STEERABLE FEATURE DETECTION IN SPACE

The detection of a specific feature at a particular point in 3-D space involves the determination of the optimal orientation

$$(\theta^*(\mathbf{x}), \phi^*(\mathbf{x})) = \arg \max_{\theta, \phi} (f(\mathbf{x}) * h(\mathbf{R}_{\theta, \phi} \mathbf{x})), \quad (1)$$

where $h(\mathbf{R}_{\theta, \phi} \mathbf{x})$ is an appropriate feature template rotated by the Euler angles θ and ϕ . The magnitude of the response is given by

$$r^*(\mathbf{x}) = f(\mathbf{x}) * h(\mathbf{R}_{\theta^*, \phi^*} \mathbf{x}). \quad (2)$$

Since a direct implementation of this detection procedure is computationally intensive, we use the steerable formulation introduced by Freeman et al. [1].

2.1. Family of steerable filters

The framework of steerable filters is ideally suited for our application since the rotated version of a steerable template can be obtained by taking a linear combination of a small number of basis filters. We choose the family of detectors defined by the linear combination of partial derivatives of

an isotropic 3-D Gaussian function $g(x, y, z)$:

$$h(x, y, z) = \sum_{k=1}^M \sum_{i=0}^k \sum_{j=0}^{k-i} \alpha_{k,i,j} \underbrace{\frac{\partial^i}{\partial x^i} \frac{\partial^j}{\partial y^j} \frac{\partial^{k-i-j}}{\partial z^{k-i-j}} g(x, y, z)}_{h_{k,i,j}(\mathbf{x})}. \quad (3)$$

The convolution of the volume $f(\mathbf{x})$ by a rotated version of the steerable filter $h(\mathbf{x})$ then becomes

$$f(\mathbf{x}) * h(\mathbf{R}_{\theta,\phi}\mathbf{x}) = \sum_{k=1}^M \sum_{i=0}^k \sum_{j=0}^{k-i} b_{k,i,j}(\theta, \phi) f * h_{k,i,j}(\mathbf{x}), \quad (4)$$

where the weights $b_{k,i,j}$ are polynomials in $(\cos \theta \sin \phi)$, $(\sin \theta \sin \phi)$, and $\cos \phi$. In the next sections we extend the 2-D formulation established in [5] to derive optimal feature detectors in three dimensions.

2.2. Design of optimal filters

We use the approach specified in [5] to derive the optimal detector. We assume a certain feature $f_0(\mathbf{x})$ with a specified orientation (say along the x axis) and obtain the criterion. The signal energy term is given by the response of the detector at the origin:

$$S = (f_0 * h)(\mathbf{0}) = \int_{\mathbb{R}^3} f_0(\mathbf{x}) h(-\mathbf{x}) d\mathbf{x}. \quad (5)$$

Assuming that the input signal is corrupted by additive white noise of unit variance, the variance of its output is given by the energy of the filter:

$$N = \int_{\mathbb{R}^3} |h(\mathbf{x})|^2 d\mathbf{x}. \quad (6)$$

The extension of the regularization term used in [5] for the suppression of false oscillations is also straightforward; its expression is given by

$$R = \int_{\mathbb{R}^3} (|h_{xx}(\mathbf{x})|^2 + |h_{yy}(\mathbf{x})|^2 + |h_{zz}(\mathbf{x})|^2) d\mathbf{x}. \quad (7)$$

The localization term penalizes the average error in the location estimate of the feature. In 3-D, the features of interest are often elongated along multiple directions (e.g. 0-D features like points, 1-D features like curves and 2-D features like surfaces). In these cases, we have to minimize the localization error along all the co-dimensions. For a m -D feature in \mathbb{R}^3 , the localization error can be quantified as

$$\text{Loc} = - \int_{\mathbb{R}^3} f_0(\mathbf{x}) \Delta_{(3-m)} h(-\mathbf{x}) d\mathbf{x}, \quad (8)$$

where $\Delta_{(3-m)}$ denotes a Laplacian operator oriented along the $3 - m$ -dimensional hyperplane orthogonal to the feature

direction. For example, for a line feature oriented along the x axis, the operator is $\Delta_2 = \partial_{yy} + \partial_{zz}$.

Having obtained the expressions for the individual terms, we follow exactly the same optimization procedure as in [5] to obtain the detectors¹; we maximize

$$\mathcal{C} = S \cdot \text{Loc} - \mu R, \quad (9)$$

subject to the constraint $N = 1$.

We now consider two special cases to illustrate the utility of the approach, focusing on second order ($M = 2$) curve and surface detectors. To derive the curve detector, we set $f_0(\mathbf{x}) = \delta(y, z)$, a Dirac delta function oriented along the x axis. We set² $\mu = 0$ to obtain the optimal curve detector

$$h_{\text{curve}} = \frac{1}{2\sqrt{2\pi}} \left(-\frac{2}{3} g_{xx} + g_{yy} + g_{zz} \right) \quad (10)$$

Similarly, we obtain the surface detector by setting $f_0(\mathbf{x}) = \delta(x)$, which yields

$$h_{\text{surface}} = \frac{1}{8\pi\sqrt{6}} (-4g_{xx} + g_{yy} + g_{zz}) \quad (11)$$

The impulse responses of the curve and surface detectors are illustrated in Fig. 1-a & b, respectively.

2.3. Feature detection

Since the detectors differ only in the weight of the g_{xx} term³, we express the two detectors in a common form:

$$h(\mathbf{x}) = \underbrace{(g_{xx} + g_{yy} + g_{zz})}_{\Delta g} - (\alpha + 1)g_{xx}, \quad (12)$$

where $\alpha = 2/3$ for the curve detector and $\alpha = 4$ for the surface detector. Note that Δg is an isotropic 3-D Laplacian; it is invariant of rotation. $h(\mathbf{x})$ can be rotated to an orientation specified by $\mathbf{v} = (\cos \theta \sin \phi, \sin \theta \sin \phi, \cos \phi)$ as

$$h(\mathbf{R}_{\theta,\phi}\mathbf{x}) = (g_{xx} + g_{yy} + g_{zz}) - (\alpha + 1) \mathbf{v}^T \mathbf{H}_g \mathbf{v}, \quad (13)$$

where \mathbf{H}_g denotes the 3-D Hessian matrix of $g(\mathbf{x})$. We rewrite this expression as

$$h(\mathbf{R}_{\theta,\phi}\mathbf{x}) = \mathbf{v}^T \mathbf{A}_g \mathbf{v} \quad (14)$$

where $\mathbf{A}_g = (g_{xx} + g_{yy} + g_{zz})\mathbf{I} - (\alpha + 1)\mathbf{H}_g$ (\mathbf{I} denotes the identity matrix). The linearity of convolution implies that

$$f(\mathbf{x}) * h(\mathbf{R}_{\theta,\phi}\mathbf{x}) = \mathbf{v}^T \mathbf{A}_{f*g} \mathbf{v}. \quad (15)$$

¹The derivation of the detectors in higher dimensions involves exactly the same approach

²Since second order detectors do not oscillate much, the regularization term is not required. It is essential for higher order detectors.

³We neglect the normalization factor to obtain simpler expressions.

Therefore the response of the filter and its optimal orientation are given by

$$r^* = \lambda_{\max} \quad (16)$$

$$\mathbf{v}^* = \mathbf{e}_{\max} \quad (17)$$

where λ_{\max} and \mathbf{e}_{\max} are the maximum eigenvalue and corresponding eigenvector of \mathbf{A}_{f^*g} .

2.4. Implementation

In order to implement these detectors, one needs to evaluate the 3×3 matrix \mathbf{A}_{f^*g} for each position \mathbf{x} . This is achieved by convolving the input signal with the six separable 3-D basis templates $(g_{xx}, g_{yy}, g_{zz}, g_{xy}, g_{xz}, g_{yz})$, which constitutes the most important part of the computational load. The eigenvalues are determined as the analytical solution of the 3rd degree characteristic equation. When the multiplicity of the principal eigenvalue is equal to one, it implies that the rank of the matrix $\mathbf{M} = \mathbf{A} - \lambda_{\max}\mathbf{I}$ must be 2, and thus that two rows of \mathbf{M} are linearly independent. We also know that the eigenvector corresponding to λ_{\max} is orthogonal to these rows of \mathbf{M} (since $\mathbf{M}\mathbf{v} = \mathbf{0}$). Therefore, the eigenvector can be efficiently obtained by finding two linearly independent rows of \mathbf{M} and computing their cross-product. We also exploit the symmetry of \mathbf{M} to further speed up the computations.

3. NON-MAXIMUM SUPPRESSION

3.1. Surface detector

Non-maximum suppression for the surface detector is trivial and analogous to the procedure for ridge detection in 2D. By construction, the detector's orientation is orthogonal to the feature and hence to the surface. For a point that is detected as being part of a surface, non-maximum suppression is performed by taking the vectors that are orthonormal to the surface in this point, and by interpolating for the value at the corresponding coordinates (see Fig. 2-b). We use linear interpolation, since its precision is sufficient for the application in addition to being computationally simple.

3.2. Curve detector

The procedure is different for the curve detector, since the optimal orientation of the detector is along the direction of the feature. Therefore, to be retained, the value of a point \mathbf{x} needs to be higher than the values that are orthonormal to the feature in this point. Since we are working in a discrete representation, we need to interpolate for the values on the unit circle that is orthogonal to the feature in the point considered (see Fig. 2-a). To simplify this process, we take advantage of the concavity of linear interpolation (i.e., interpolation between two points cannot yield a value that is

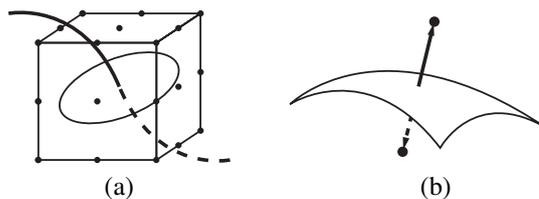


Fig. 2. Representations of the schemes used for non-maximum suppression in the case of curve (a) and surface detection (b). For the curve detector, we consider only the neighborhood points that are needed for the interpolation.

larger than the maximum of the points used in the interpolation). The circle is always contained in the cubic neighborhood of 26 points around \mathbf{x} (Fig. 2-a). The influence of these points in the linear interpolation on the circle is maximal at the position that corresponds to their projection onto the circle. Thus, non-maximum suppression is performed by first computing the projection of the neighborhood points onto the plane that is orthogonal to the feature orientation. Normalization of the resulting vectors yields the locations on the unit circle at which we need to interpolate. If any of the values obtained are superior to the value in \mathbf{x} , the point is suppressed. In some cases, however, the procedure outlined above will not be sufficient. Consider the response of the detector to an impulse. It corresponds to a second derivative of a Gaussian, and for every point within the spot, the orientation of the detector is towards the impulse. It is easy to see that the non-maximum suppression will fail to remove the points in the response. To solve the problem, a second iteration of point suppression is applied, based on the number of direct neighbors in the plane orthogonal to the feature orientation for every point \mathbf{x} . If too many neighbors are detected, the point is discarded. The thresholding step concluding the detection procedure can either be performed manually, or by selecting a criterion such as the percentage of voxels to be retained.

4. EXAMPLES AND APPLICATIONS

4.1. Surface detection

A promising application is the detection of fractures in cataclastic rock from x-ray computer tomograms (Fig. 3-a & b). The tomograms⁴ present a high amount of noise due to the density and thickness of the samples. While the fractures are often invisible from the exterior of the sample, they become visible in the tomogram due to the air they encompass. Fig. 3-c shows the result of the detection after non-maximum suppression and thresholding. Here the detector size and threshold have been set in order to isolate the largest fracture present in the sample.

⁴Courtesy of P. Christe, GEOLEP, EPFL

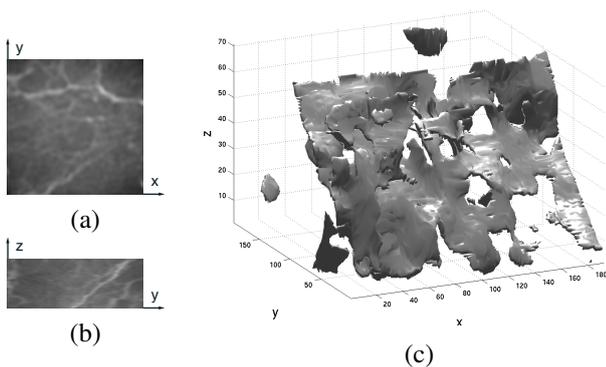


Fig. 3. (a) XY-slice of a volume of cataclastic rock. (b) XZ-slice of the sample. (c) 3-D view of the result of the detection, showing the principal fracture in the sample.

4.2. Curve detection

A problem ideally suited for curve detection is determining the location and orientation of actin filaments in limited angle electron tomography [6]. Among the difficulties with such data are the low signal-to-noise ratio and limited resolution. Additionally, the density of the filaments is high and there are frequent intersections and branchings (see Fig. 4-a & b). In the test data⁵, other structures such as ribosomes are present. These, along with strong noise in feature-less regions, can lead to false detections. Nevertheless, we are able to successfully detect most of the filaments (see Fig. 4-c). Note that in addition to the detected curve segments, we also know their orientation. This information can be extremely useful in determining the branching angle between intersecting filaments, which is of great practical significance for biological interpretation.

The post-processing steps applied here consist of thresholding based on the filter's output to discount small features that are not likely to be filaments. As seen in Fig. 4-c, this can lead to discontinuities in the detections. Future work will consist of developing more sophisticated approaches for post-processing that are beyond simple voxel-wise thresholding (e.g., statistical analysis based on clusters).

5. CONCLUSION

We have proposed a general approach for the design of optimal 3-D steerable feature detectors. Valid for the family of detectors based on Gaussian derivatives of any order, the approach was further specified and illustrated for detectors based on second order derivatives of Gaussians. We then derived two filters suitable for the detection of surfaces and curves, respectively, and showed their efficiency in practical applications.

⁵Courtesy of A.S. Frangakis [6]

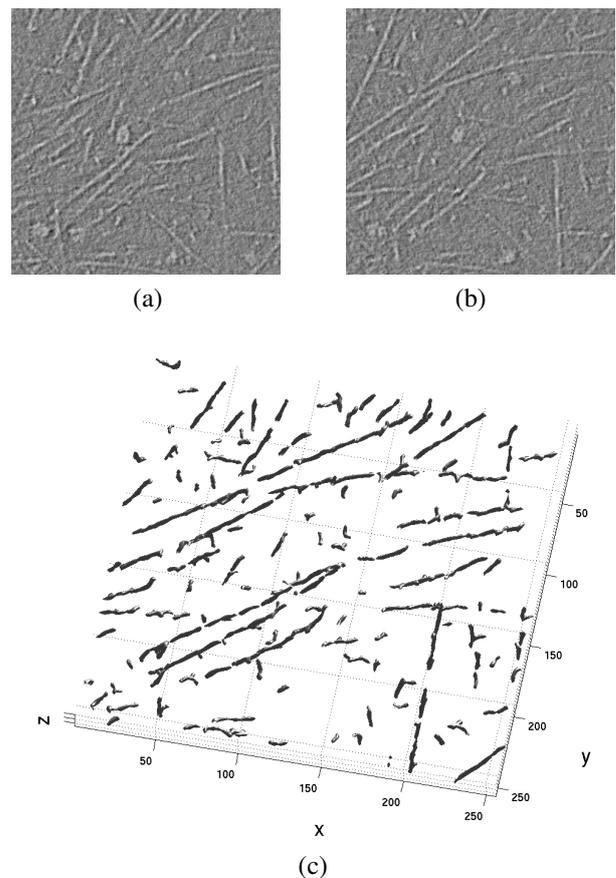


Fig. 4. (a) and (b): two slices of the EM-tomogram in which the filaments shown in (c) were detected. (c) 3-D isosurface representation of the detected actin filaments. In reality the detected filaments are a series of connected points in space. This result was obtained by performing the detection on a section of 20 images (of size 256×256).

6. REFERENCES

- [1] W.T. Freeman and E.H. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, 1991.
- [2] C. Huang and Y. Chen, "Motion estimation method using a 3D steerable filter," *Image and Vision Computing*, vol. 13, pp. 21–32, 1995.
- [3] Frangi A., *Three-Dimensional Model-Based Analysis of Vascular and Cardiac Images*, Ph.D. thesis, Utrecht University, 2001.
- [4] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.
- [5] M. Jacob and M. Unser, "Design of steerable filters for feature detection using canny-like criteria," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 1007–1019, 2004.
- [6] O. Medalia, I. Weber, A.S. Frangakis, Nicastro D., Gerisch G., and Baumeister W., "Macromolecular architecture in eukaryotic cells visualized by cryoelectron tomography," *Science*, vol. 298, pp. 1209–1213, 2002.