

Gauge features for curvilinear target recognition

James M. Coggins

BAE Systems Advanced Information Technologies,
6 New England Executive Park, Burlington, MA 01803

ABSTRACT

Curvilinear targets are common in many imaging modalities. Detection of such targets can be challenging because of their multiscale structure, their frequent obscuration in natural imagery, their turns, intersections, and merges, and the prevalence of false positive detections based on local information. Using a spatial spectroscopy approach, we introduce image analysis methods that use the concept of gauge frames to simplify the identification of curvilinear targets. Fast computational approximation methods are described for gauge fields, and an experiment is described illustrating the power of higher-order derivatives for understanding even relatively simple geometric structures. Methods for extracting coherent curvilinear objects that exploit the larger-scale commonalities of points in the object are described.

Keywords: spatial spectroscopy, differential geometry, scale space, gauge frame, curvilinear objects

1. INTRODUCTION

Curvilinear targets are long, thin structures visible in a variety of image modalities. In aerial photographs or synthetic aperture radar, such targets include roads, trails, rivers, or vehicle tracks. In medical imagery, curvilinear targets include blood vessels or tubes such as the delivery instruments for brachytherapy. In industrial inspection, wires on printed circuit boards or on the topmost layer of an integrated circuit being fabricated or being delayered are curvilinear targets. Such targets are characterized by a consistent appearance along the principal direction of the target with some form of contrast (which may involve contrasting brightness, texture, or height) between the target and surrounding structures. An edge (in intensity, texture, or height) may be visible at times between the curvilinear target and its surround and may contribute to detection of the target. In realistic natural imagery, however, the edge of the target may be indistinct or nonexistent, the target's appearance may be broken by occlusions, shadows, or detection dropouts, and multiple targets may overlap, intersect, or merge. The methods described in this paper have been used successfully in such challenging circumstances.

The approach to curvilinear target detection described here is called spatial spectroscopy [1-4]. **Spatial spectroscopy** is a computer vision methodology based on scale space theory, differential geometry, and statistical pattern recognition, with roots in studies and models from visual psychophysics [1,6,13]. In parallel to electromagnetic spectroscopy, spatial spectroscopy begins with a dispersal of the spatial (Fourier) spectrum into a series of channels, followed by nonlinear recombination of the channels to infer properties of the image content. Section 2 will describe in more detail the aspects of spatial spectroscopy that are relevant for curvilinear target recognition. Approximation methods that enable rapid computation of derivative-of-Gaussian features are also described.

Section 3 introduces the gauge frame and defines features based on the gauge frame that are useful for identifying curvilinear targets.

Copyright 2009 Society of Photo-Optical Instrumentation Engineers. This paper was published in Proc. SPIE Vol. 7335 (Automated Target Recognition XIX), F. Sadjadi and A. Mahalanobis, editors, April 2009 and is made available as an electronic reprint with permission of SPIE. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

Section 4 introduces the relationship between derivatives and symmetries and describes an experiment that shows the essential role of gauge vectors based on higher-order directional derivatives (to order 4) in describing some simple geometries.

Section 5 describes how coherent curvilinear objects can be extracted from the gauge coordinates obtained across scale and derivative orders.

[*james.coggins@baesystems.com](mailto:james.coggins@baesystems.com), 1-781-273-3388 fax 1-781-273-9345

2. SPATIAL SPECTROSCOPY

2.1 Mathematical foundation

$$L(\mathbf{x}+\mathbf{h};\sigma) = \sum_{|\mathbf{n}|=0}^N \left(h_x \frac{d}{dx} + h_y \frac{d}{dy} \right)^{\mathbf{n}} L(\mathbf{x};\sigma)$$

The mathematical foundation of spatial spectroscopy is the truncated 2-dimensional Taylor series expansion, called the N-jet [8-11], defined as

where L is the image function, σ is a scale parameter that defines the size of the (Gaussian) measurement aperture function, $\mathbf{x}=[x \ y]$ is an image location and $\mathbf{h}=[h_x \ h_y]$ is a spatial offset in the image plane. The Taylor Series expansion shows how measurements of scaled derivatives at location \mathbf{x} can provide information about the value of the image at a remote location, $\mathbf{x}+\mathbf{h}$. As the maximum order of differentiation, N , increases, the approximation to the values of the image function at increasingly remote locations improves. As higher-order derivatives are employed, the Taylor Series knits image points into neighborhoods.

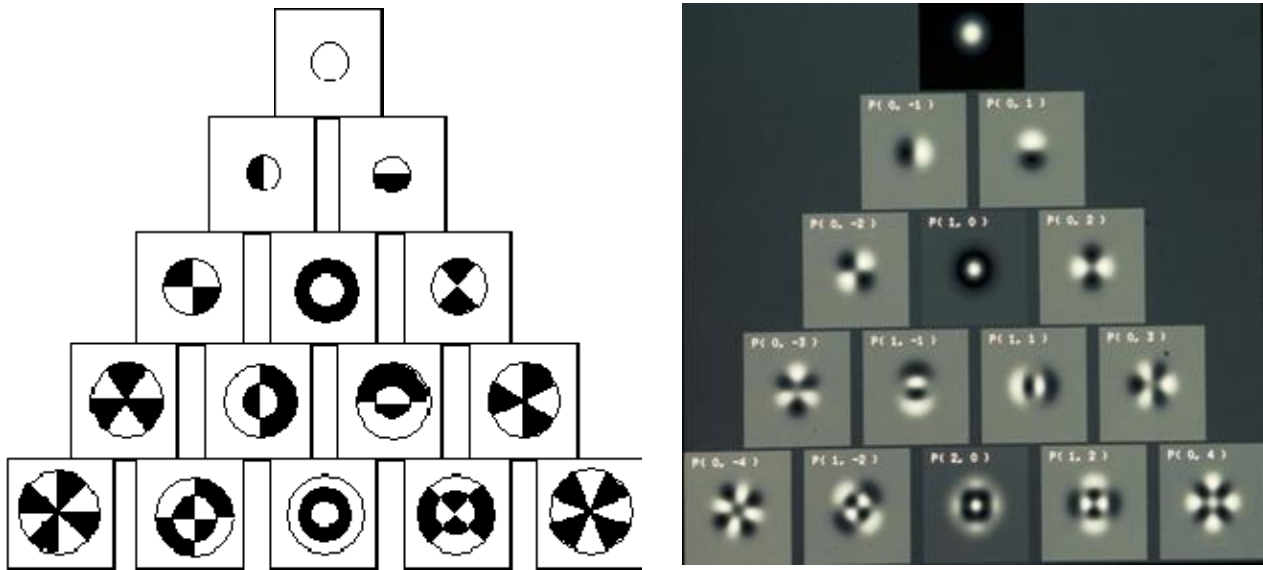


Fig. 1: Derivative-of-Gaussian kernels of order 0 through 4. The derivative-of-Gaussian kernels are complex-valued image functions; the real part is shown in the left image. The diagram on the right shows the signs of the real parts of the derivative kernels. This is the Polar form of the derivative kernels. Kth-order members of other families of these kernels (Cartesian, wavelet) can be obtained by complex-valued linear combinations of the Kth-order members of this series. It is essential to include all members of each derivative order to obtain equivalent information with the corresponding order of the other families.

Unfortunately, only small values of N can be used in practice because of the large number of filters required as N increases. Figure 1 illustrates the 0th through 4th order derivative-of-Gaussian kernels in the Polar family. Fifteen kernels are required to obtain the complete Taylor Series expansion through order 4. This fact reveals the essential value of a multiscale approach to image analysis: Extending the area where the Taylor series approximation is good by increasing N is expensive. It is less costly to increase the scale parameter and recompute a short N -jet at the larger scale. Thus, information about the structure of the neighborhood farther from \mathbf{x} is obtained by expanding the measurement aperture while avoiding an expensive Taylor series expansion into higher orders of differentiation.

2.2 Engineering Foundation

The engineering foundation for spatial spectroscopy is the dispersion of the 2-D Fourier spectrum. The Fourier spectrum consists of complex values, and it is critical for image analysis to disperse both the real and imaginary parts of the spectrum. A well-known experiment [12] demonstrates that most of the information in the Fourier spectrum about form is contained in the phase spectrum, so ignoring the imaginary part of the spectrum dooms any image analysis scheme to weakness, if not failure. Because the Fourier transform of real, symmetric functions is real and symmetric, methods that rely on real, symmetric filters such as filtering by multiscale Gaussians only, or Laplacian (symmetric second derivative) or difference-of-Gaussian pyramids, disperse only the real part of the spectrum. Wavelets can disperse both the real and imaginary parts of the spectrum, but nonlinear filtering methods limit the flexibility for implementation optimizations and complicate the application of nonlinear decision tools after dispersion. Gabor functions can disperse the phase spectrum, but only if both sine and cosine phase components are applied (at a minimum), and even then, coverage of the spectrum is on a different basis from the derivative-of-Gaussian basis. Derivative-of-Gaussian kernels have been shown to be the only admissible basis for an unbiased front-end visual system [14, 15], though optimizations based on knowledge of the image content can result in other kinds of filters being optimal [2, 4]. In spatial spectroscopy, a series of low-pass filters (e.g. 2-D Gaussians) disperse the real part of the spectrum, and derivative operators disperse the imaginary part of the spectrum.

Nonlinear operators recombine the dispersed channels to infer properties of the image content. Since in spatial spectroscopy the dispersal process is all linear, the recombination needs to consist of nonlinear operators. (Otherwise, the linear recombination operators could be combined with the dispersal operators.) The essence of the recombination stage is to make decisions about the image content, and decision operators are inherently nonlinear. The collection of relevant nonlinear operators used to date is surprisingly small. Absolute value is the key nonlinear operator for representing texture (average local energy through scale/orientation channels). In the present work, minimum and maximum operators play an essential role.

2.3 Computational Approximations

In practice, the dispersion of the spectrum can be approximated for rapid computation in several ways. Since the entire spectral dispersal process is linear, one can rearrange components of the derivative-of-Gaussian operation in various ways and apply approximations at multiple points in the process. Figure 2 shows how the derivative-of-Gaussian kernels can be approximated by a few convolutions followed by sampling of the filtered image at offset locations.

The scale dispersal can be performed by binomial filtering. A small (often 3x3) rectangle of constant amplitude is convolved repeatedly with the image. Each iteration increases slightly the footprint of the effective filter, and as the iterations increase, the form of the effective filter approaches a Gaussian. Since the rectangular kernel is very small compared to the image size, a spatial convolution is faster than going through the Fourier transform. The break-even point occurs when the kernel size exceeds \log_2 of the image size, but the breakeven size can be increased through optimized implementation of the spatial convolution.

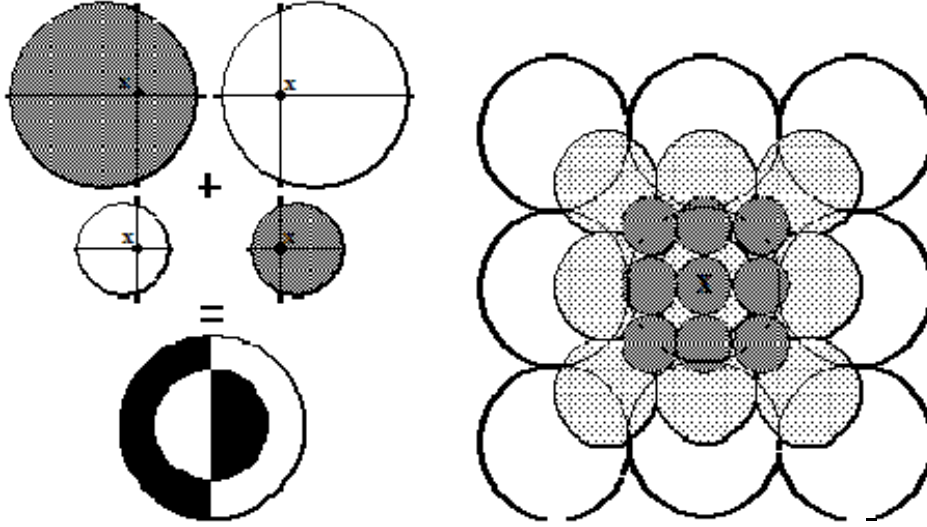


Fig. 2: Computational approximation of N-jet computation. The Gaussian derivatives (not just the Laplacian of Gaussian) can be approximated by linear combinations of multiscale *offset* Gaussians. The description of point \mathbf{x} includes the intensities of filtered images at offset locations. (Here, 3 convolutions yield 27 features describing pixel \mathbf{x} from which approximations to many directional derivatives of many different orders and orientations may be obtained.) To maintain scale invariance, the offsets should scale with σ , and σ should increase geometrically.

Derivatives can be approximated by finite differences between the location of interest and offset sample locations in blurred images. This process of blur and offset sampling is mathematically equivalent (by linearity of the dispersal process) to blurring using difference-of-offset-Gaussian kernels, but the former is much faster to implement. The parameter σ denotes the standard deviation of the Gaussian filter that defines the measurement aperture and δ denotes the distance from a pixel of interest that its neighbors are sampled to capture derivative information. In general, δ is proportional to σ , but the constant of proportionality varies with the application in both value and criticality.

In the present work, directional derivatives are computed by applying a blur at the scale corresponding to the width of the curvilinear target and then sampling the blurred image at surrounding points to estimate a centered 1st derivative or a 1-D oriented 2nd derivative. These derivatives, computed at location \mathbf{x} with offset $\delta\theta$ (distance δ at an orientation defined by unit vector θ) are $\Delta_1=L(\mathbf{x}-\delta\theta)-L(\mathbf{x}+\delta\theta)$; $\Delta_2=2L(\mathbf{x})-L(\mathbf{x}-\delta\theta)-L(\mathbf{x}+\delta\theta)$

3. THE GAUGE FRAME

The coordinate system in which a raster image is measured is usually an arbitrary choice. Even when it is not entirely arbitrary (as in synthetic aperture radar, where the sensor is looking from the left or bottom of the frame, and the sensor look direction has significance for understanding the image content), the $\{\mathbf{x} \ \mathbf{y}\}$ frame has minimal utility for understanding the image content, especially for recognizing curvilinear structures. Curvilinear structures can turn in any direction, so the most descriptive and relevant coordinate frame needs to change constantly along the length of the object. In differential geometry, the concept of a *gauge frame* allows a frame to be defined at every point to optimize some mathematical property of the point's neighborhood structure, and thus, the gauge frame is tied to the image content. Analysis of the image content can be simplified by defining features with respect to the gauge frame [9, 14, 15].

A simple gauge frame is the one defined by first-order derivatives of a (suitably regularized) image. Denoting the image intensity at an arbitrary point by L and directional derivatives by subscripts denoting the vector along which the derivative is computed, the derivatives with respect to the measurement coordinates are denoted L_x and L_y . The gradient is a directional derivative in the direction of maximum

gradient, denoted L_w where vector $w = \langle L_x \ L_y \rangle$ and the gradient magnitude is $|L_w| = (L_x^2 + L_y^2)^{1/2}$. Similarly, the isophote direction, which is orthogonal to the gradient is defined as $v = \langle -L_y \ L_x \rangle$. Note that the derivative in the isophote direction is 0: $L_v = 0$ (by definition of the “isophote”). Thus at every point in the image, one can define the first-order gauge frame $\{v, w\}$ which has the interesting mathematical property that $L_v = 0$, so many terms in the formulas for geometrical quantities involving L_v disappear, greatly simplifying the expressions for such quantities.




Fig. 3: Gauge coordinates simplify the description of image structure. Points on the edge of a disk in image (a) are mapped to a 1st-order derivative feature space (b). Each image point maps to a different point in the feature space. In the 1st-order gauge feature space, the edge points map to a single point in the 1st-order gauge feature space. Studies of similar mappings of geometric objects have proven fruitful in building insight into the nature of the scaled derivative space [5].

The simplest demonstration of the power of gauge coordinates involves the 1-jet expansion of a black disk on a white background (Figure 3).

The use of gauge frames extends to higher-order derivatives [9]. The second order derivatives in the measurement frame are denoted L_{xx} , L_{yy} , and L_{xy} ($=L_{yx}$). The $\{p, q\}$ gauge frame is defined by the directions of principal curvature, with p being the direction of maximal curvature and q the direction of minimum curvature. The vectors p and q are the eigenvectors of the Hessian matrix of second derivatives,

$$H(L) = \begin{bmatrix} L_{xx} & L_{xy} \\ L_{xy} & L_{yy} \end{bmatrix}.$$

The $\{p, q\}$ gauge figures prominently in the method for detecting curvilinear targets. Since the second-order gauge points in the direction of maximum second derivative, that gauge direction is always orthogonal to the principal direction of a curvilinear target, whether the target is lighter or darker than the background. Figure 4 illustrates the use of the 2nd order gauge for detecting vehicle tracks that appear as parallel dark lines. In this case, the 2nd-order gauge at slightly different scales can be used to find the wheel tracks as well as the space between them. The magnitude of the 2nd order gauge vector is maximal when the pixel is located in the center of a track of matching scale.

This example also illustrates that higher-order derivatives can provide more reliable results than lower-order derivatives [15]. Higher-order derivatives impose more constraints on the local neighborhood structure and result in fewer false positive detections than are obtained with lower-order derivatives. For example, the study in Figure 4 uses second-order derivatives which require the existence of a pair of edges at the appropriate separation before signaling a potential detection. First order derivatives are confused by

data dropouts and extraneous (noise) structure in the image. While this may not be surprising, the notion that computation of derivatives is an ill-posed problem has become deeply embedded in the computer vision community, causing practitioners to opt for complex, heuristic edge-combination strategies when a higher-order derivative would go much further toward solving the problem than any 1st-order, edge-based strategy.




Fig. 4: Use of 2nd-order gauge to detect vehicle tracks. The second-order gauge is easier to use in this case than first-order (edge) detectors because the 2nd order derivative spans the lines(or the gap between them), requiring the existence of a pair of edges the right distance apart before declaring a possible detection.

4. DERIVATIVES AND SYMMETRIES

The computational approximation to the directional derivatives involves convolving the image with a Gaussian at the scale of the width of the curvilinear target, sampling the blurred image at offset locations, recombining those measurements to approximate the directional derivatives, and computing the correct gauge vector for analysis of the neighborhoods of each point.

Gauge vectors can be computed for derivative orders 1-4 by sampling a scaled image at a set of points in a circle about the point of interest and forming the relevant gauge vectors by suitable linear combinations of the intensities at the sample points. In order to sample compatible distances and angles at all scales and orders, one needs to sample a multiple of 12 points around the circle. (In the experiment, we sample 96 points.) The first-order derivative is computed by the difference between the center point and a point on the circle. The second order derivatives use the formula for Δ^2 . The third order derivatives map the points into a feature space with angles tripled. Three points at the same distance but 120 degrees apart contribute to each 3rd derivative vector. The fourth order derivatives use points at the same distance but 90 degrees apart. In each case, the kth order derivative wraps the orientation samples k times around a circle in feature space. Each sample defines a vector whose length is the difference between the intensity at the center point and the intensity at the sample. Summing the vectors at each orientation in feature space and dividing the phase of the resultant by the derivative order yields the principal gauge vector for that scale and order. This computation is illustrated in Figure 5.




Fig. 5: Angular folding for computing directional derivatives of higher orders. Left: 12 points on a circle around the point of interest are sampled. The difference between point x and point k measures the first derivatives at various orientations. The resultant vector defines the angle of the 1st order gauge. Center: In this feature space, angles are doubled to create a two-fold symmetry for measuring second order derivatives. Vector magnitudes are the intensity differences of the sampled pixels from the center pixel. The resultant is the 2nd order gauge. Right: Angles are tripled to compute the third-order gauge.

This computation clearly reveals the relationship between directional derivatives and symmetries. In particular, second-order derivatives merge opposite pairs of samples, third order derivatives merge samples at 120 degree intervals, and so on. A consequence of this property is that when the second derivative gauge has maximum length over a neighborhood, the point is on a medial axis of an object. The medial axis is of particular importance for curvilinear structures since it constitutes a skeleton of the curvilinear target. A problem that frequently occurs in following curvilinear objects is that the track of the object fades out at intersections. Points in an intersection have no preferred second-order orientation.

The question now is whether we can discern which of the gauge vectors for derivative orders 1-4 provides the best characterization of the neighborhood of pixel x . An image of intersecting lines was created and the gauge vectors for orders 1-4 were computed as described above at each pixel. The lengths of the gauge vectors were compared, and each pixel was identified by the derivative order that produced the longest gauge vector. The original image and the resulting label image are shown in Figure 6.

What is striking about this result is that all four derivative orders play clear, identifiable roles in describing the pixels in this simple image. First order derivatives respond to edges, of course. Second order derivatives respond strongly to the centers of lines. Third order derivatives respond to T junctions. Fourth-order derivatives respond to crossing junctions. This result points the way to more reliable tracking of curvilinear targets with intersections and junctions: use higher-




Fig. 6. Dominant derivative orders for pixels around curvilinear targets. Left: original image for this experiment. The white lines are 16 pixels wide. Right: Label image indicating the dominant derivative order at each pixel. Parameters used to produce this image are $\sigma=4$, $\delta=12$. Note the presence of 4th order dominance at crossing junctions, 3rd order dominance at T junctions, 2nd-order dominance at centers of lines and 1st order dominance at edges.

order derivatives to demand more constraints on local neighborhoods before declaring a detection and explicitly detect more complex structures when first and second derivatives faded out due to additional symmetries in the local structure that are captured better by higher-order derivatives.

Also worth noting is that the derivative orders play the same roles in background regions as they do on the lines. The background rectangle on the right side enclosed by white lines is detected strongly by the second-derivative operator because its scale is in range of the filter used in this example.

5. GROWING CURVILINEAR OBJECTS

Gauge vectors provide an effective filter for curvilinear objects based on local structure, but a final step is required to detect the curvilinear structures coherently and to eliminate false positive target detections that can occur based on local measurements. The problem is that local measurements do not account for the large-scale coherence that occurs along the curvilinear structure. Three strategies have been used to extract coherent curvilinear objects based on the previous observations about multiscale gauge coordinates.

The first strategy was to reblur the gauge vector image. After computing gauge vectors at each pixel, one can pass through the array of gauge vectors computing for each pixel the resultant from a vector sum of the gauge vectors in the local neighborhood, weighted toward the direction of the gauge vector at the pixel of interest. In background areas, this resultant has a small magnitude because the neighboring vectors have random phases. On the other hand, along curvilinear objects the resultant is huge because the gauge vectors at pixels along the principal direction of the target reinforce each other. This method turned out to be an extremely sensitive change detector. In noisy images, it is rather too sensitive.

The second strategy was to blur the image in the principal direction of the curvilinear object until the blurred intensity began to change because the kernel started overlapping a contrasting region of the image. The point was to take advantage of the larger-scale nature of the consistent appearance along the target. This approach was difficult to control, particularly at pixels near the ends of a target, where highly asymmetric filtering would have been most appropriate. In addition, as the scale of this blur grew, determining a cutoff point was difficult since sometimes the change signaling the end of the target was small relative to the kernel footprint already being integrated.

The third strategy, which has proven the most effective in practice, is to explicitly grow the curvilinear objects by following (perpendiculars to) the gauge vectors starting from seed points. The growth process moves along the target structure as a wave. The centroid of the front of the wave can be recorded periodically. The combination of growth plus the centroid of the wave provides both a coherent raster mask

of the object and a graphical representation of its skeleton in one pass through the image data. It is possible for the wavefront to split when curvilinear objects merge or split. Such events can be detected by noting when the wavefront becomes disconnected. Then one connected component of the wavefront can be pursued and the other component deferred by adding the component to the list of seeds.

Since curvilinear object growth is based on following similarly-oriented gauge vectors, object detections often break at intersections. If one object's gauge dominates an intersection (perhaps because its width is a better match to the scale of the detection filters) the dominant object can be grown through the intersection but other (smaller) objects joining it at the intersection will be broken there. Combining our findings on the role of higher-order derivatives with our results on growing coherent structures, we find that explicit detection of intersections based on higher-order derivatives in advance of the growth stage can permit an adjustment to the growth rules to allow connection of structure detections across intersections.

6. CONCLUSION

The need to recognize curvilinear structures occurs frequently in many imaging modalities. While curvilinear structures encounter the usual image disturbances such as noise and occlusion, they have multiscale, geometric properties that can be exploited to enhance the robustness of their detection. Gauge vectors are a mechanism for simplifying the representation of the geometric structure of curvilinear objects. Gauge vectors are based on directional derivatives and can be computed at any order of differentiation. The combination of differentiation and scale is a powerful means for dispersing the spatial spectrum, permitting simple nonlinear operators to perform decision tasks that infer aspects of the image content. Spatial spectroscopy is a unifying methodology that combines multiscale analysis, differential geometry, and statistical pattern recognition in a powerful suite of practical engineering tools for solving image analysis problems.

REFERENCES

- [1] J. M. Coggins, *A Framework for Texture Analysis Based on Spatial Filtering*, University Microfilms: Ann Arbor, 1982, 168 pages.
- [2] J. Coggins and C. Huang, "Defining Optimal Feature Sets for Segmentation by Statistical Pattern Recognition" *Proc. Mathematical Methods in Medical Imaging II*, D. C. Wilson, ed. Proc. SPIE 2035, 1992, pp. 80-88.
- [3] J. Coggins, "Statistical Approaches to Multiscale Medial Vision" *Mathematical Methods in Medical Imaging I*, D. C. Wilson and J. N. Wilson, editors, Proc. SPIE 1768, 1992, pp. 14-24.
- [4] J. Coggins, "Statistical Investigations of Multiscale Image Structure" in *Proc. Visualization in Biomedical Computing*, Richard Robb, editor, Proc. SPIE 1808, 1992, pp. 145-159.
- [5] J. M. Coggins, "Non-linear Feature Space Transformations", *Digest of IEE Symposium on Applied Statistical Pattern Recognition*, Birmingham, U.K. April 20, 1999.
- [6] J. M. Coggins and A. K. Jain, "A Spatial Filtering Approach to Texture Analysis," *Pattern Recognition Letters*, vol. 3, 1985, pp. 195-203.
- [7] J. M. Coggins and A. K. Jain, "Surface Orientation from Texture," Proceedings of the 1986 International Conference on Systems, Man, and Cybernetics, Atlanta, GA, October 14-17, 1986, pp. 1617-1620.
- [8] J. J. Koenderink and A. J. van Doorn, "Generic Neighborhood Operators", *IEEE PAMI* 14(6), June 1992, pp. 597-605.
- [9] Jan J. Koenderink, *Solid Shape*, MIT Press 1990.
- [10] J. Koenderink, "Operational Significance of Receptive Field Assemblies," *Biological Cybernetics*, 1988.
- [11] J. Koenderink & A. J. van Doorn, "Representation of Local Geometry in the Visual System," *Biological Cybernetics*, 55:367-375, 1987.
- [12] A. V. Oppenheim and J. S. Lim, "The importance of phase in signals," Proc. IEEE 69, 529-541 (1981).
- [13] S. M. Pizer, J. M. Coggins, C. A. Burbeck, "Formation of Image Objects in Human Vision," *Computer Assisted Radiology (Proc. CAR '91)*, Springer-Verlag, Berlin, 1991, pp. 535-542.

- ^[14] B. M. terHaar Romeny, *Front End Vision and Multi-Scale Image Analysis*, Springer, 2003.
- ^[15] B. M. terHaar Romeny, L. M. J. Floack, A. H. Salden, and M. A. Viergever, "Higher-order differential structure of images", *Image and Vision Computing*, v12, pp 317-325, July/August 1994.