J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*

IMAGE TEXTURE

An image texture may be defined as a collection of elements or patterns wherein the individual elements themselves may or may not have a well-defined structure. Textures associated with most of the man-made objects have some regularity associated with them, in addition to a well-defined pattern structure. Examples of such textures include pictures of barbed wire, a brick wall, or a marble floor with periodically repeating tiles. On the other hand, a coastal line in an aerial photograph showing sand and water does not have any structure. Texture appears in natural pictures very frequently. The surface of a polished wooden table is a good example. A Landsat satellite picture showing the vegetation in the Amazonian area or the floating ice in the Antarctic are some other examples of such natural textures. A popular data set for researchers in this field is the digitized images from the Brodatz album (1). Some examples of textured images are shown in Fig. 1.

Image texture analysis in the past two decades has primarily focused on texture classification, texture segmentation, and texture synthesis. In addition, texture mapping has been studied extensively in computer graphics for generating realistic images for visual simulations, computer animation, and 3-D rendering of elevation maps. In texture classification the objective is to assign a unique label to each homogeneous region. For example, regions in a satellite picture may be classified into ice, water, forest, agricultural areas, and so on. In medical image analysis, texture is used in applications such as segmenting magnetic resonance (*MR*) images of brain into gray and white matter, or detecting cysts in the X-ray computed tomography (*CT*) images of the kidneys. While image classification results in segmentation, there is also considerable interest in achieving segmentation without prior knowledge of the textures.

Texture is also useful in recovering 3-D shape. A homogeneous 3-D texture under perspective projection assumption will induce distortions in the projected image. Figure 2 illustrates this. Variations in the image, such as the density and shape changes of the texture 'blobs', provide information about the 3-D shape of the object (2). Much of the previous work is concerned with synthetic image data where the focus was on shape recovery. However, as mentioned earlier, identifying the basic texture primitive itself is a major research problem. Alternative approaches include using spatial frequency information in recovering the shape (3,4,5).

Perhaps one of the few successful applications of textures is in content based image search of large image and video databases (6). Texture information can be used to annotate and search images such as aerial photographs or color pictures of nature. Figures 3 and 4 illustrate some examples taken from an aerial photograph database. In the applications section we discuss more about the construction of a texture dictionary to efficiently navigate large pictorial databases. Combined with color and shape, texture can be used to select a wide variety of patterns even when semantic level information is absent.

Texture Classification and Segmentation

Texture classification refers to the problem of assigning a particular class label to a given textured region. If the images are preprocessed to extract homogeneous textured regions, then the pixel data within these regions can be used for estimating the class labels. Here standard pattern classification techniques may be applied

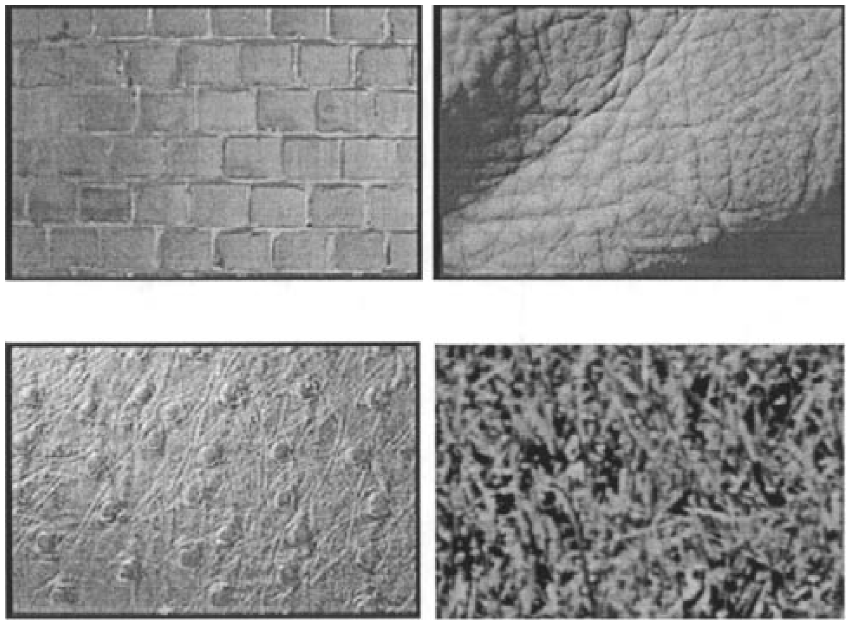

Fig. 1. Brick wall, elephant skin, ostrich skin leather, grass from Brodatz album.

assuming that there is only one texture in the region. Some of the initial work on texture analysis considered using spatial image statistics for classification purposes. These include image correlation (7), energy features and their extensions (8), features from co-occurrence matrices (9), run-length statistics (10), and difference statistics. As texture analysis research matured, two distinct approaches have emerged: in one, researchers seek to understand the process of texture generation and this led to generative texture models which could be used for classification as well as texture creation. This emphasis can be seen in much of the work on random field models for texture representation such as the 2-D nonsymmetric half plane models (11) and noncausal Gauss Markov random field models and their variations (12,13,14,15). A review of some of the recent work can be found in Ref. 16. Once the appropriate model features are computed, the problem of texture classification can be addressed using techniques from traditional pattern classification. Although significant progress has been made using these methods, several problems remain as the methods are sensitive to illumination and resolution changes, and transformations such as rotation.

The second emphasis has its roots in modeling human texture perception, particularly that of pre-attentive texture discrimination. Psychologists have studied texture for the purposes of understanding human visual perception for many decades now. Figure 5 shows an example of a texture mosaic where the boundary between the L's and +'s can be easily identified by humans without requiring detailed analysis of the micropatterns. Pre-attentive texture segmentation refers to this ability of humans to distinguish between textures in an image without any detailed scene analysis. Central to solving this problem are issues related to defining these texture features and their computation. Some of the early work in this field can be attributed to Julesz (17) for his theory of textons as basic textural elements. Spatial filtering approach has been used by many researchers for detecting texture boundaries not clearly explained by the texton theory (18,19). Texture discrimination is generally modeled as a sequence of nonlinear filtering operations without any prior assumptions about the texture generation process. Independent of the human psychovisual considerations, spatial filtering for texture analysis is now a mature area. Some of the recent work involves multiresolution filtering for both classification and segmentation (20,21,22). In the following we briefly describe the current research on both model based and

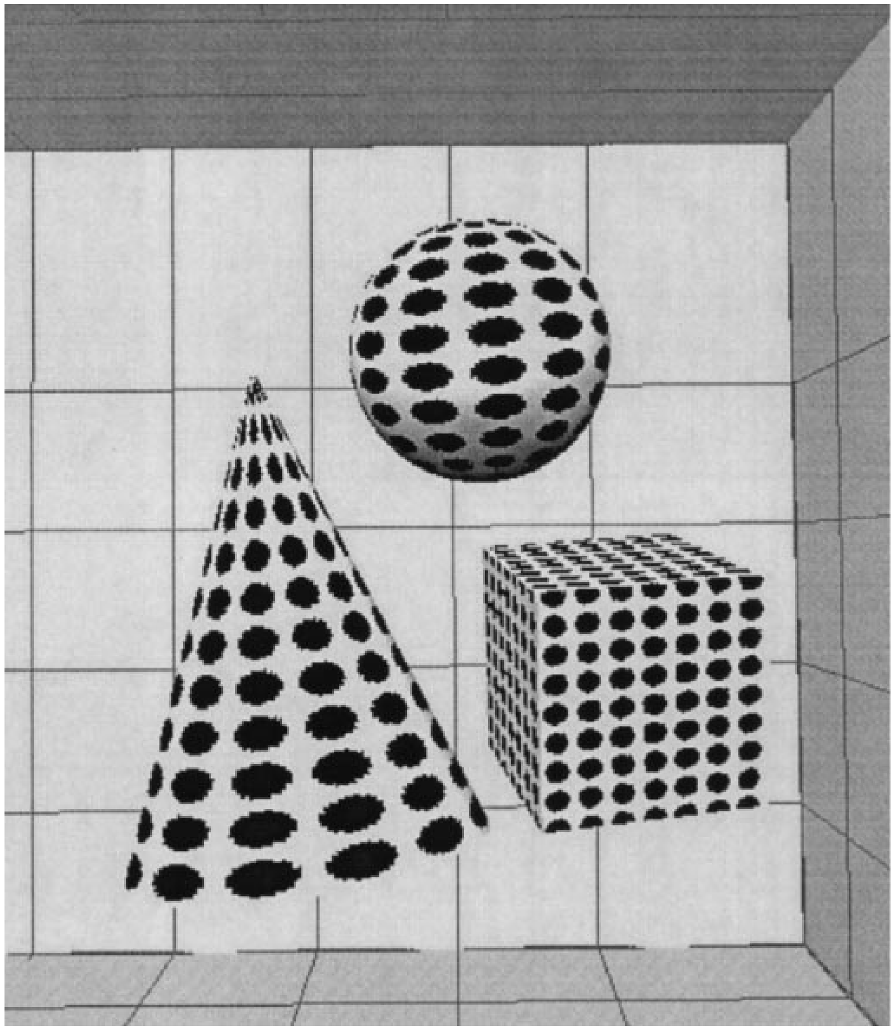

Fig. 2. Circular blobs mapped on to different 3-D shapes. The perspective distortion of the texture in the images provide depth cues.

spatial filtering approaches to texture analysis. Our aim is to expose the reader to the breadth of literature on these topics. Selected references are provided at the end for additional reading.

Gray Level Image Statistics and CO-Occurrence Matrices

A digitized image is an array of intensity values. For a textured image, these intensity values, in general, are distributed in a random manner. However, the statistics underlying these distributions are helpful in characterizing these textures. A picture of a beach has significantly different variations in gray values compared to a cloud image. The gray level co-occurrence matrix characterizes second order gray level relationships. Consider any two pixels separated by (D, θ) in the image. Here D is the distance between the two pixels, the

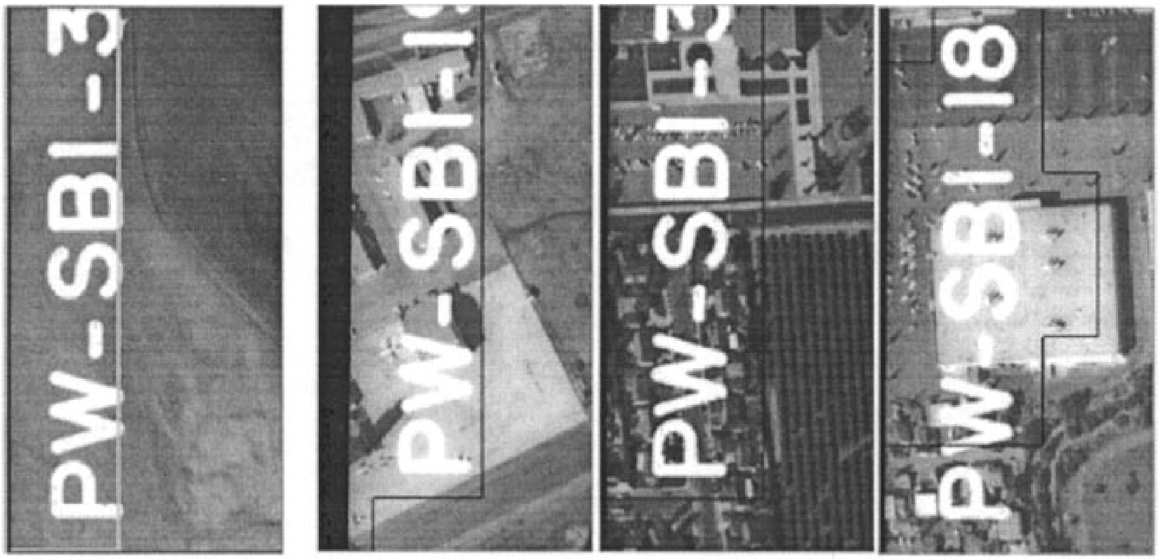

Fig. 3. Search and retrieval using texture. "Texture" in aerial photographs can help in search and retrieval of similar image patterns. Identification text superimposed on an airphoto is selected as a texture pattern in this example, and regions containing similar identification marks are retrieved.

line connecting the two pixels makes an angle θ with the X-axis. Let one of the pixels have an intensity value I and the other pixel a value J. We can then count all occurrences of pixel pairs in the image which are separated by (D, θ). Let this be $f(I, J)$. One can thus construct a matrix $f(m, n)$ wherein the elements represent frequency count for the particular pair of intensities and for a specific distance D. If no directionality distinction is made between the two pixels, that is, $f(I, J) = f(J, I)$, we get a symmetric co-occurrence matrix.

A different matrix is constructed for each (D, θ). Further abstraction of information is necessary for computational reasons. Typical features that are computed from $f(I, J)$ include the energy $(\Sigma_{I,J} f(I, J))$ and entropy $(-\Sigma_{I,J} f(I, J) \log f(I, J))$. Detailed description of these features can be found in Ref. 23.

While the co-occurrence matrix provides an intuitive mechanism to capture spatial relationships, computationally it is expensive. Note that the (D, θ) space must be sampled, and even for a coarse quantization such as 5 distances and 6 orientations, one needs to compute 30 co-occurrence matrices. Since the derived statistics directly depend on the gray values, this measure is sensitive to gray scale distortions. Further, this method only captures intensity relationships in a fine grain texture and may not be well suited for textures whose primitives are spatially large as in, for example, a picture of a brick wall.

Random Field Models

A typical image is represented over a rectangular array and the statistical distribution of the pixel intensities can be modeled as a random field. A simple model is to represent a pixel intensity at location s, $y(s)$, as a linear combination of pixel values $\{y(s + r), r \in N\}$ within a small neighborhood N and additive noise (24). The neighborhood N is typically a block of pixels surrounding s thus leading to noncausal models. The Gaussian Markov random field (*GMRF*) models is a specific class of such two-dimensional noncausal models that has been quite popular and widely studied within the texture analysis literature. Let Ω denote a set of grid points

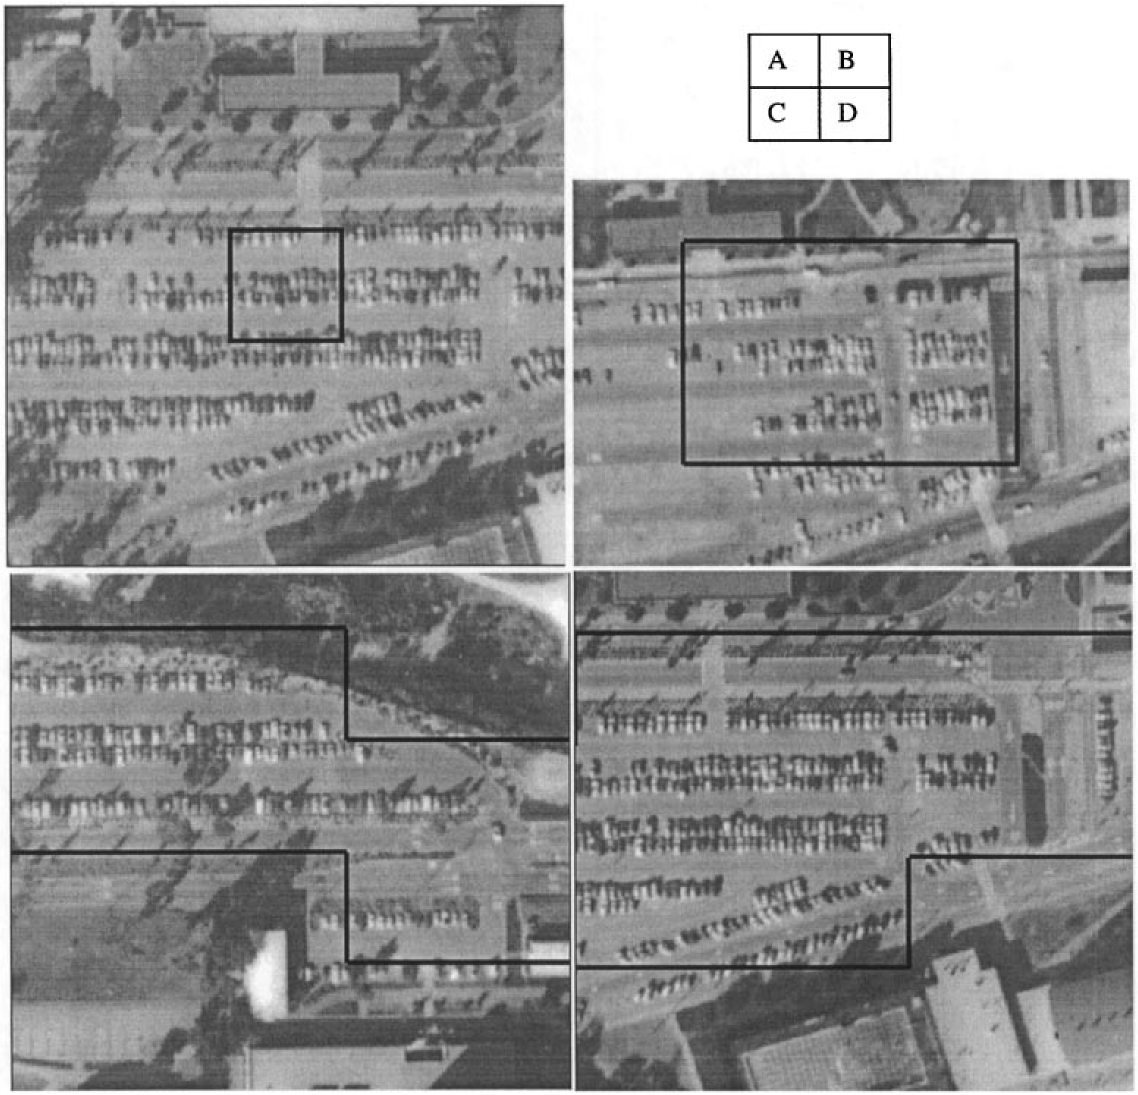

Fig. 4. Another example of using texture for pattern retrieval. In this, a texture descriptor from a parking lot is used to search for other parking lots in a aerial image database. The top three matches contain similar patterns. (A): query image, B,C,D: top three retrievals based on texture in (A).

on a two-dimensional lattice of size $M \times M$. Then a random process $Y(s)$ is said to be Markov if

$$\Pr(y(s)|\text{all } y(r), r \neq s) = \Pr(y(s)|y(s+r), r \in N)$$

The neighborhood N of s can be arbitrarily defined. However, in most image processing applications it is natural to consider neighbors which are also spatially closer. For the case of GMRF, the neighborhood set of pixels is symmetric. For instance, the four nearest neighbors (North, South, East, and West) of a pixel form the simplest neighborhood set which is referred to as the first order GMRF model. Adding the diagonal neighbors gives the

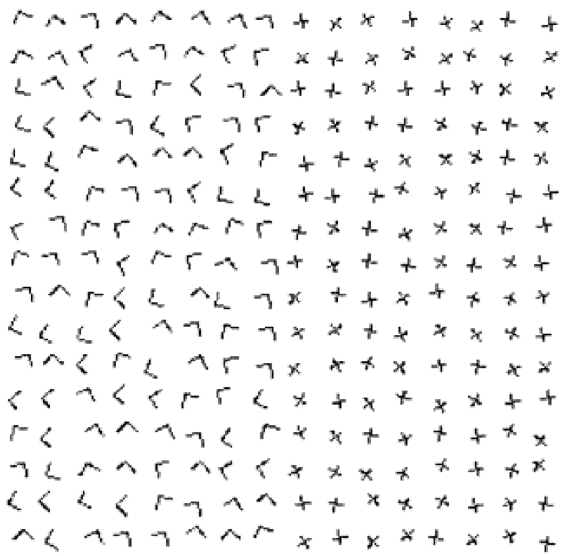

Fig. 5. A texture consisting of randomly oriented L and + micro patterns. The line segments are of equal length.

second order GMRF model and so on. Cross and Jain (25) provide a detailed discussion on the applications of Markov random fields (*MRF*) in modeling textured images. A nontrivial problem with any MRF model is the estimation of model parameters. Usually the structure of the model (such as first or second order GMRF) is assumed to be known even though estimating the structure itself is a challenging issue. A comparison of different GMRF parameter estimation schemes can be found in Ref. 24.

While some of the initial motivation for the use of MRF models was for texture synthesis, it has also been used for texture classification and segmentation purposes. Given an image consisting of an unknown number of texture classes, the problem of segmentation involves both parameter estimation and classification. However, for parameter estimation one needs homogeneous image regions, which can only be obtained after the image is segmented! Both parameter estimation and segmentation are computationally expensive operations. If the texture class and hence the class parameter information is known, image segmentation can be formulated as a maximum a posteriori (*MAP*) estimation problem. In modeling images with more than one texture, in addition to the MRF model describing a texture patch, an additional random process to characterize the distribution of textures in the image, is introduced. This texture label process is usually modeled using discrete Markov models with a single parameter measuring the amount of clustering between neighboring pixels. Let Y be the image intensity, modeled as a GMRF conditioned on the class label, and L be the class label process. Then the posterior distribution of texture labels for the entire image given the intensity array is

$$\Pr(L|Y) = \frac{\Pr(Y|L)\Pr(L)}{\Pr(Y)}$$

where $\Pr(Y|L)$ is the conditional probability of observing the given intensity array given the label distribution, $\Pr(L)$ is the probability of a label distribution L. Maximizing the right-hand size gives the MAP estimate. In general, finding an optimal solution is not feasible because of the nonconvex nature of the MAP cost function. Geman and Geman (26) employ a stochastic relaxation method for finding a solution. This approach can be shown in theory to yield the global optimum but requires impractical annealing schedules. Most of its software implementations have a fixed number of iterations and the results are usually good.

Gaussian Markov Random Field Parameter Estimation. While there are many methods proposed for estimating GMRF parameters, it is difficult to guarantee both consistency and stability of the estimates. Consistency refers to the property that the estimates converge to the true values of the parameters and stability refers to the positive definiteness of the covariance matrix in the expression for the joint probability density of the MRF. Parameter estimation can be formulated as an optimization problem. However, as mentioned earlier, parameter estimation requires segmented image regions whereas segmentation requires good estimates of the parameters. Lakshmanan and Derin (27) propose an optimization framework for simultaneous parameter estimation and segmentation. They compute the maximum likelihood estimates of the parameters and a MAP solution for segmentation.

Multiresolution Analysis And Markov Random Fields. Much of the computations involving MRF models are computationally expensive: the cost functions for parameter estimation and segmentation are nonconvex and hence involve iterative algorithms. One possibility is to use a multiresolution approach. Coarser resolution sample fields are obtained by subsampling. In general GMRF lose their Markov property under subsampling (the second-order GMRF with separable autocovariance matrix is an exception). An excellent discussion on multiresolution GMRF models can be found in Ref. 28. Krishnamachari and Chellappa (29) present techniques for GMRF parameter estimation at coarser resolution from finer resolution parameters. They use the coarse resolution parameters to segment the coarse resolution image and the segmentation results are propagated to finer resolutions. Besides improving the speed of computations, segmentation at coarser resolutions provides good initial conditions for following finer resolution images, thus improving the segmentation results compared to working at the original resolution.

Concluding the discussion on GMRF models for texture analysis, we observe that they provide an elegant mathematical framework for describing texture and for deriving algorithms for texture classification and segmentation. Multiresolution models are particularly interesting from a computational point of view. Recently, GMRF models have been applied to unsupervised segmentation of color texture images wherein both spatial and interband interactions are modeled using random fields (30).

Spatial Filtering

A serious drawback of random field models that characterize intensity distributions in modeling textures is that they are sensitive to gray level distortions induced by changes in imaging conditions such as lighting. In contrast, typical spatial filtering methods use the variations in the intensity as opposed to the absolute values of the intensity itself for texture classification and segmentation purposes. Laws' work on texture energy features (8) is one of the early attempts to apply spatial filtering followed by some nonlinearities (such as computing the energy) for texture discrimination. In this formulation, images are processed by a number of filters, each designed to extract a certain type of image feature. Typical features of interest include edges and lines. While the initial design of these filters by Laws was quite ad-hoc, he obtained significantly better results compared to co-occurrence based methods. In recent years, spatial filtering methods have been extensively studied for texture classification/segmentation tasks.

Malik and Perona (19) proposed a computational framework for modeling pre-attentive texture discrimination. Images are filtered using even symmetric kernels (such as Gaussians and Gaussian second derivatives) followed by half-wave rectification. Weak responses are suppressed using local negative feedback among competing feature locations. Texture gradient is then computed and boundaries are identified. The authors argue against the use of energy type features as well as the use of odd-symmetric filters in the preprocessing stage. Their experimental results indicate that the proposed algorithm yields results comparable to human discrimination on a set of textures frequently used in the literature.

Directional filters, such as those based on Gabor functions, have been used by many researchers for texture analysis (31,32,33,34,35). Gabor functions are modulated Gaussians and the general form of a two-dimensional Gabor function can be written as:

$$g(x, y: u_0, v_0) = \exp\left(-[x^2/2\sigma_x^2 + y^2/2\sigma_y^2] + 2\pi i[u_0x + v_0y]\right) \tag{1}$$

where σ_x and σ_y define the widths of the Gaussian in the spatial domain and (u_0, v_0) is the frequency of the complex sinusoid. Some of the early work on using Gabor functions in image processing and vision can be credited to Daugman (36). Daugman suggests that the 2-D receptive field profiles of simple cells in the mammalian visual cortex can be well modeled by Gabor functions. These functions also have some nice theoretical properties such as minimizing the joint uncertainty in space and frequency which may have some implications in coding and recognition applications.

In Ref. 33, the authors present a multiresolution framework for boundary detection. Features of interest in an image are generally present in various spatial resolutions, and a multiscale representation facilitates feature extraction and analysis. The image is first filtered with a bank of Gabor filters tuned to different orientations and scale. This is followed by local competitive and co-operative feature interactions to suppress weak features while reinforcing the stronger ones. An interesting aspect in Ref. 33 is the use of interscale interactions where features at neighboring scales interact, and the resulting output is sensitive to curvature changes and line endings in the image. A grouping stage combines outputs tuned to similar orientations. Finally, texture and intensity gradients are computed to detect boundaries in the image. The paper reports results on a wide variety of images including some interesting examples of illusory boundaries.

Continuing on the use of Gabor filters and a multiresolution representation, Manjunath and Ma (21) propose a filter design that generates a bank of Gabor filters given the lower and upper cut off frequencies. This self-similar wavelet dictionary is obtained by dilations and rotations of the kernel $g(x, y)$.

$$g(x, y) = \left(\frac{1}{2\pi\sigma_x\sigma_y}\right)\exp\left[-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2}\right) + 2\pi jWx\right] \tag{2}$$

$$G(u, v) = \exp\left\{-\frac{1}{2}\left(\frac{(u-W)^2}{\sigma_u^2} + \frac{v^2}{\sigma_v^2}\right)\right\} \tag{3}$$

where $\sigma_u = \frac{1}{2}\pi\sigma_x$ and $\sigma_v = \frac{1}{2}\pi\sigma_y$

The filter bank can then be computed as:

$$\begin{aligned} g_{mn}(x, y) &= a^{-m}G(x', y'), \qquad & a > 1, m, n = \text{integer} \\ x' &= a^{-m}(x\cos\theta + y\sin\theta), \qquad & y' = a^{-m}(-x\sin\theta + y\cos\theta) \end{aligned}$$

where $\theta = n\pi/K$ and K is the total number of discrete orientations. Let U_l and U_h denote the lower and upper center frequencies of interest and S be the number of scales in the multiresolution decomposition. Then the following equations can be used to compute the filter parameters:

$$a = \left(\frac{U_h}{U_l}\right)^{1/(S-1)}, \sigma_u = \frac{(a-1)U_h}{(a+1)\sqrt{2\ln 2}}$$

$$\sigma_v = \tan\left(\frac{\pi}{2K}\right)\left[U_h - 2\ln 2\left(\frac{\sigma_u^2}{U_h}\right)\right]\left[2\ln 2 - \frac{(2\ln 2)^2\sigma_u^2}{U_h^2}\right]^{-1/2}$$

where $W = U_h$ and $m = 0, 1, \ldots, S-1$. The image is then convolved with each of the filters in the dictionary and the mean and standard deviation of the resulting outputs are used in constructing a texture feature vector. A weighted Euclidean distance metric is used to compare two feature vectors. Extensive experiments on the texture in the Brodatz album (1) indicate that this feature vector does quite well in characterizing a wide variety of textures. A detailed comparison and experimental validation of different texture feature representations is made in (21). Later on we present an image database application where similar patterns are retrieved based on texture.

Rotation Invariant Texture Classification

The general approach to developing rotation-invariant techniques has been to modify successful non-rotation-invariant techniques. Since standard MRF models are inherently dependent on rotation, several methods have been introduced to obtain rotation invariance. Kashyap and Khotanzad (37) developed the circular autoregressive model with parameters that are invariant to image rotation. Choe and Kashyap (38) introduced an autoregressive fractional difference model that has rotation (as well as tilt and slant) invariant parameters. Cohen, Fan, and Patel (39) extended a likelihood function to incorporate rotation (and scale) parameters. To classify a sample, an estimate of its rotation (and scale) is required.

For feature-based approaches, rotation-invariance is achieved by using anisotropic features. Porat and Zeevi (40) use first and second order statistics based upon three spatially localized features, two of which (dominant spatial frequency and orientation of dominant spatial frequency) are derived from a Gabor-filtered image. Leung and Peterson (41) present two approaches, one that transforms a Gabor-filtered image into rotation invariant features and the other of which rotates the image before filtering; however, neither utilizes the spatial resolving capabilities of the Gabor filter. You and Cohen (42) use filters that are tuned over a training set to provide high discrimination among its constituent textures. Greenspan, et al., (43) use rotation-invariant structural features obtained via multiresolution Gabor filtering. Rotation invariance is achieved by using the magnitude of a discrete Fourier transform (*DFT*) in the rotation dimension.

Haley and Manjunath (22) have investigated applications of Gabor features for rotation invariant classification. A polar analytic form of a two-dimensional Gabor wavelet and a much more detailed set of microfeatures is computed. From these microfeatures, a micromodel which characterizes spatially localized amplitude, frequency, and directional behavior of the texture, is formed. The essential characteristics of a texture sample, its macrofeatures, are derived from the estimated selected parameters of the micromodel. Classification of texture samples is based on the macromodel derived from a rotation invariant subset of macrofeatures. In experiments using the Brodatz album, comparatively high classification rates are obtained. A detailed feature parametric analysis and feature quality analysis is provided in Ref. 44.

Summarizing the discussion on spatial filtering methods for texture analysis, significant progress has been made in the use of band pass filters for extracting texture features since the early work of Laws (8). As in the case of random field models, scale and rotation invariant analysis remain as challenging issues.

Applications

In recent years image texture has emerged as an important primitive to search and browse through large collections for similar looking patterns. An image can be considered as a mosaic of textures and texture features associated with the regions can be used to index the image data. For instance, a user browsing an aerial image database may want to identify all parking lots in the image collection. A parking lot with cars parked at regular intervals is an excellent example of a textured pattern. Similarly, agricultural areas and vegetation patches are other examples of textures commonly found in aerial imagery and satellite photographs. An example of a typical query that can be asked of such a content based retrieval system could be “retrieve all Landsat images of Santa Barbara which have less than 20% cloud cover” or “Find a vegetation patch that looks like this region.” In the Alexandria digital library (*ADL*) (45) project at the University of California at Santa Barbara, researchers are developing a prototype geographic information system that will have some of the image search features described above.

Manjunath and Ma (21) investigate the role of textures in annotating image collections and report on the performance of several state-of-the-art texture analysis algorithms with performance in image retrieval being the objective criterion. Their texture analysis scheme based on a Gabor wavelet decomposition described earlier in this article performed quite well in this application compared to methods such as those using random field models and orthogonal wavelet filters. Ma and Manjunath (46) provide a detailed description of a system that searches aerial photographs based on texture content. They demonstrate that texture could be used to select a large number of geographically salient features including vegetation patterns, parking lots, and building developments. Using texture primitives as visual features, one can query the database to retrieve similar image patterns. Much of the results presented are with airphotos although a similar analysis can be applied to Landsat and Spot satellite images. This is currently being integrated into the ADL project (45), whose goal is to establish an electronic library of spatially indexed data, providing internet access to a wide collection of geographic information. A significant part of this collection includes maps, satellite images, and airphotos. For example, the Maps and Imagery Library at the UCSB contains over 2 million of historically valuable aerial photographs. A typical air photo can take over 25 MB of disk space, and providing access to such data raises several important issues, such as multiresolution browsing and selecting images based on content. Figure 4 and Fig. 5 show some examples of texture based retrieval in an airphoto database.

What distinguishes image search for database related applications from traditional texture classification methods is the fact that there is a human in the loop (the user), and there is a need to retrieve more than just the best match. In typical applications a number of top matches with rank-ordered similarities to the query pattern will be retrieved. Comparison in the texture feature space should preserve visual similarities between patterns. This is an important but difficult problem in content-based image retrieval. Toward this objective, a hybrid neural network algorithm to learn the pattern similarity in the texture feature space is proposed (47). This approach uses training data containing the pattern similarity information (provided by human indexers) to partition the feature space into many visually similar clusters. A performance evaluation of this approach using the Brodatz texture indicate that a significantly better retrieval performance can be achieved. In addition to retrieving perceptually more relevant data, an additional advantage of this approach is that it also provides an efficient indexing tree to narrow down the search space.

An interesting component of the system described in Ref. (46) is the texture thesaurus for similarity search. A texture thesaurus can be visualized as an image counterpart of the traditional thesaurus for text search. It creates the information links among the stored image data based on a collection of code words and sample patterns obtained from a training texture pattern set. Similar to parsing text documents using a dictionary or thesaurus, the information within images can be classified and indexed via the use of a texture thesaurus. The design of the texture thesaurus has two stages. The first stage uses a learning similarity algorithm to combine the human perceptual similarity with the low level feature vector information, and the second stage utilizes a hierarchical vector quantization technique to construct the code words. The texture thesaurus so constructed is domain-dependent and can be designed to meet the particular need of a specific image data type by exploring the training data. Further, the thesaurus model provides an efficient indexing tree while maintaining or even improving the retrieval performance in terms of human perception. The visual code word representation in the thesaurus can be used as information samples to help users browse through the database.

Summary

Image texture research has seen much progress during the last two decades. Texture based image classification has found applications in satellite and medical image analysis and in industrial vision systems for applications such as defect detection. Texture mapping for visualization and computer animations is now a well established area in computer graphics. Texture appears to be a promising image feature for search and indexing of large image and video databases. In both model based and spatial filtering approaches, current research is on deriving scale and rotation invariant texture features. Recent work on rotation invariant texture computations appear quite encouraging whereas scale invariance remains elusive.

Acknowledgments

The airphoto examples shown in Figs. 4 and 5 are from a demonstration program written by Dr. Wei-Ying Ma.

BIBLIOGRAPHY

1. P. Brodatz *Textures: A Photographic Album for Artists and Designers*, New York: Dover, 1966.
2. D. Blostein N. Ahuja Shape from texture: integrating texture element extraction and surface estimation, *IEEE Trans. Pattern Anal. Mach. Intell.*, **11**: 1233–50, 1989.
3. R. Bajcsy L. Lieberman Texture gradient as a depth cue. *Comput. Graphics Image Process.*, **5**: 52–67, 1976.
4. B. J. Super A. C. Bovik Shape from texture using local spectral moments, *IEEE Trans. Pattern Anal. Mach. Intell.*, **17**: April 1995. p. 333–343.
5. J. Krumm S. A. Shafer A characterizable shape-from-texture algorithm using the spectrogram, *Proc. IEEE-SP Int. Symp. Time-Freq. Time-Scale Anal.*, Oct. 25–28 1994, pp. 322–325.
6. H. Voorhees T. Poggio Computing texture boundaries from images, *Nature*, **333**: 364–367, 1988.
7. P. C. Chen T. Pavlidis Segmentation by texture using correlation, *IEEE Trans. Pattern Anal. Mach. Intell.*, **5**: 64–69, 1983.
8. K. Laws *Textured image segmentation*, Ph.D. thesis, University of Southern California, 1978.
9. R. Haralick R. Bosley Texture features for image classification, *3rd ERTS Symp., NASA* **SP-351**, 1219–1228, 1973.
10. M. M. Galloway Texture analysis using gray level run lengths, *Comput. Graphics Image Process.*, **4**: 172–179, 1975.
11. C. W. Therrien An estimation theoretic approach to terrain image segmentation, *Comput. Vision, Graphics Image Process.*, **22**: 313–326, 1983.
12. F. S. Cohen D. B. Cooper Simple parallel hierarchical and relaxation algorithms for segmenting noncausal Markovian fields, *IEEE Trans. Pattern Anal. Mach. Intell.*, **9**: 195–219, 1987.
13. S. Chatterjee R. Chellappa Maximum likelihood texture segmentation using Gaussian Markov random field models, *Proc. IEEE Conf. Comput. Vision Pattern Recognition*, June 1985.
14. H. Derin H. Elliott Modeling and segmentation of noisy and textured images using Gibbs random fields, *IEEE Trans. Pattern Anal. Mach. Intell.*, **9**: 39–55, 1987.
15. B. S. Manjunath R. Chellappa A note on unsupervised texture segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.*, **13**: 472–483, 1991.

16. R. Chellappa R. L. Kashyap B. S. Manjunath Model-based texture segmentation and classification, in *Handbook of Pattern Recognition and Computer Vision*, C. H. Chen, L. F. Pau, and P. S. P. Wang (eds.), Singapore: World Scientific, 1993, pp. 279–310.
17. B. Julesz Textons, the elements of texture perception, and their interactions, *Nature*, **290**: 12, March 1981.
18. J. R. Bergen E. H. Adelson Early vision and texture perception, *Nature*, **333**: 363–364, 1988.
19. J. Malik P. Perona Preattentive texture discrimination with early vision mechanisms, *J. Opt. Soc. Amer.*, **A7**: 923–932, 1990.
20. T. Chang C.-C. J. Kuo Texture analysis and classification with tree structured wavelet transform, *IEEE Trans. Image Process.*, **2**: 429–441, 1993.
21. B. S. Manjunath W. Y. Ma Texture features for browsing and retrieval of image data, *IEEE Trans. Pattern Anal. Mach. Intell.*, **18**: 1996, 837–842.
22. G. Haley B. S. Manjunath Rotation invariant texture classification using a complete space-frequency model, *IEEE Trans. Image Process.*, 1998, in press.
23. R. M. Haralick L. Shapiro *Computer and Robot Vision*, Addison-Wesley, 1992, vol. 1, chap. 9.
24. R. Chellappa Two-dimensional discrete Gaussian Markov random field models for image processing, in L. N. Kanal and A. Rosenfeld (eds.), *Progress in Pattern Recognition 2*, Amsterdam: Elsevier, 1985.
25. G. R. Cross A. K. Jain Markov random field texture models, *IEEE Trans. Pattern Anal. Mach. Intell.*, **5**: 25–39, 1983.
26. S. Geman D. Geman Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intell.*, **6**: 721–741, 1984.
27. S. Lakshmanan H. Derin Simultaneous parameter estimation and segmentation of Gibbs random fields using simulated annealing, *IEEE Trans. Pattern Anal. Mach. Intell.*, **11**: 799–813, August 1989.
28. S. Lakshmanan H. Derin Gaussian Markov random fields at multiple resolutions, in R. Chellappa (ed.), *Markov Random Fields: Theory and Applications*, New York: Academic Press, 1993, pp. 131–157.
29. S. Krishnamachari R. Chellappa Multiresolution Gauss-Markov random field models for texture segmentation, *IEEE Trans. Image Process.*, **6**: 1997.
30. D. K. Panjwani G. Healey Markov random field models for unsupervised segmentation of textured color images, *IEEE Trans. Pattern Anal. Machine Intell.*, **17**: 939–954, 1995.
31. A. Bovik M. Clark W. Geisler Multichannel texture analysis using localized spatial filters, *IEEE Trans. Pattern Anal. Mach. Intell.*, **12**: 55–73, 1990.
32. A. K. Jain F. Farrokhnia Unsupervised texture segmentation using Gabor filters, *Pattern Recognition*, **23**: 1167–1186, 1991.
33. B. S. Manjunath R. Chellappa A Unified approach to boundary perception: edges, textures and illusory contours, *IEEE Trans. Neural Networks*, **4**: 96–108, Jan 1993.
34. D. Dunn W. E. Higgins J. Wakeley Texture segmentation using 2-D Gabor elementary functions, *IEEE Trans. Pattern Anal. Mach. Intell.*, **16**: 1994.
35. H. Greenspan R. Goodman R. Chellappa C. H. Anderson Learning texture discrimination rules in a multiresolution system, *IEEE Trans. Pattern Anal. Machine Intell.*, **16**: 894–901, 1994.
36. J. G. Daugman Uncertainty relation for resolution in space, spatial frequency and orientation optimized by two-dimensional visual cortical filters, *J. Opt. Soc. Amer.*, **2**: 1160–1169, 1985.
37. R. L. Kashyap A. Khotanzad A Model-based method for rotation invariant texture classification, *IEEE Trans. Pattern Anal. Mach. Intell.*, **8**: 472–481, 1986.
38. Y. Choe R. L. Kashyap 3-D shape from a shaded and textural surface image, *IEEE Trans. Pattern Anal. Mach. Intell.*, **13**: 907–918, 1991.
39. F. S. Cohen Z. Fan M. A. Patel Classification of rotated and scaled textured image using Gaussian Markov random field models, *IEEE Trans. Pattern Anal. Mach. Intell.*, **13**: 192–202, 1991.
40. M. Porat Y. Y. Zeevi Localized texture processing in vision: analysis and synthesis in the Gaborian space, *IEEE Trans. Biomed. Eng.*, **36**: 115–129, 1989.
41. M. M. Leung A. M. Peterson Scale and rotation invariant texture classification, *Proc. 26th Asilomar Conf. Signals, Syst. Comput.* Pacific Grove, CA, October 1992.
42. J. You H. A. Cohen Classification and segmentation of rotated and scaled textured images using tuned masks, *Pattern Recognition*, **26** (2): 245–258, 1993.

43. H. Greenspan *et al.* Rotation invariant texture recognition using a steerable pyramid, *Proc. IEEE Int. Conf. Image Process.*, Jerusalem, Israel, October 1994.
44. G. M. Haley *Rotation invariant texture classification using a complete space-frequency model*, M.S. Thesis, Elect. Comput. Eng. Depart., Santa Barbara, CA: Univ. California, June 1996.
45. T. Smith *et al.* A Digital library for geographically referenced materials, *IEEE Computer*, 54–60, May 1996.
46. W. Y. Ma B. S. Manjunath A texture thesaurus for browsing large aerial photographs, *J. Amer. Soc. Inf. Sci.*, **49** (7): 633–648, 1997.
47. W. Y. Ma B. S. Manjunath Texture features and learning similarity, *Proc. IEEE Int. Conf. Comput. Vision Pattern Recognition*, San Francisco, CA, June, 1996, pp. 425–430.

B. S. MANJUNATH
University of California
R. CHELLAPPA
University of Maryland