MODULATION DOMAIN TEXTURE RETRIEVAL FOR CBIR IN DIGITAL LIBRARIES

Joseph P. Havlicek [1], *Jinshan Tang* [2], *Scott T. Acton* [2], *Rob Antonucci* [3], *and Fabrice N. Ouandji* [1]

[1]School of Electrical and Computer Engineering, University of Oklahoma, Norman, OK
[2]Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA
[3]Science Systems and Applications Inc., Lanham, MD

ABSTRACT

We present a new texture retrieval algorithm that, for the first time, performs content-based image retrieval (CBIR) in the modulation domain by computing powerful low-level texture features based on computed AM-FM image models. Performance of the new algorithm is analyzed with respect to competing methods where texture features are computed from Gabor filter magnitude responses. Our experimental results show that the new algorithm achieves a significant performance advantage. We describe how the new algorithm will be used for the texture component of a novel CBIR service called DIRECT, which is designed to provide image based searches in distributed digital libraries without the need for manual entry of annotations and image metadata.

1. INTRODUCTION

In large, distributed, and heterogeneous libraries that may contain tens of thousands of images or more, search engines based on text annotations or other manually extracted metadata are infeasible because of the need for librarians to laboriously catalog and index the images and compile the annotations. Thus, there is a critical need for content-based image retrieval (CBIR) services capable of searching for images based on automatically extracted visual information without the requirement for manual intervention. Although a relatively small number of systems have made progress toward bridging the *semantic gap* discussed in [1] between low-level features and interpretive descriptions of image meaning, the vast majority of practical, currently deployable CBIR systems search for images that are similar in terms of low-level color, shape, and texture features.

In this paper we present a new texture retrieval algorithm that, for the first time, uses computed AM-FM models for performing CBIR in the modulation domain. This algorithm is being developed as the texture component of a decentralized CBIR service for distributed digital libraries. We test the performance of the new algorithm against a database of 17,408 textured image tiles and find that the modulation domain approach offers a significant performance gain relative to competing techniques.

2. BACKGROUND

In order to obtain a meaningful performance comparison between various texture retrieval techniques, we restrict our attention to a well-defined retrieval task that has been investigated by others [2–6]. We begin with a set of N_I homogeneous texture images referred to as the *original* images. Each original image is partitioned into N_T smaller disjoint tiles to generate an image database containing $N_I \times N_T$ images of the smaller size. Typically, $N_T = 16$. A query image is then selected from this database. The task is to retrieve the N_R images from the database that are most similar to the query image, where $N_R \geq N_T$.

Retrieved images that came from the same original image as the query image are considered to be *correct* matches, whereas those that came from a different original image are considered to be *incorrect* matches. Let N_C denote the number of correct matches retrieved. The retrieval rate is then defined by $R_C = N_C/N_T$. Since the query image and the other $N_T - 1$ images that came from the same original as the query image are always present in the database, we have for *any* texture retrieval algorithm that $\lim_{N_R \to N_I \times N_T} R_C = 100\%$. Intuitively, the retrieval rate R_C may be interpreted as the answer to the following question: the database contains N_T images that came from the same original as the query image (including the query image itself); what percentage of these were among the top N_R matches returned by the algorithm?

It is well-known that texture features computed from the magnitude responses of a bank of Gabor filters provide excellent performance in this retrieval task, but do so at a computational cost that is generally greater than the cost of other methods [2, 3, 5–7]. For original images that are Brodatz or Brodatz-like textures, the retrieval technique proposed in [3] has rarely been beaten. We refer to this technique as the *Manjunath-Ma algorithm*.

In the Manjunath-Ma algorithm, each image in the database is filtered with a bank of M complex-valued Gabor filters in a polar spectral tessellation, where any two adjacent filters intersect at their half-peak frequencies. For an image A, let $y_m(\mathbf{x})$ be the complex-valued response of the m^{th} filter, where $\mathbf{x} \in \mathbb{R}^2$ and $m \in [0, M-1]$. With $\mu_m^{(A)}$ and $\sigma_m^{(A)}$ the mean and standard deviation of $|y_m(\mathbf{x})|$, the texture feature vector describing the image A is

$$\mathbf{f}^{(A)} = \left[\mu_0^{(A)}\ \sigma_0^{(A)}\ \mu_1^{(A)}\ \cdots\ \mu_{M-1}^{(A)}\ \sigma_{M-1}^{(A)}\right]^T. \tag{1}$$

Let s_m^{μ} be the standard deviation of the feature $\mu_m^{(X)}$ computed over all images X in the database and let s_m^{σ} be the standard deviation of the feature $\sigma_m^{(X)}$, also computed over all images X in the database. Then the distance between two images A and B with feature vectors $\mathbf{f}^{(A)}$ and $\mathbf{f}^{(B)}$ is given by

$$d\left(\mathbf{f}^{(A)}, \mathbf{f}^{(B)}\right) = \sum_{m=0}^{M-1} \frac{\left|\mu_m^{(A)} - \mu_m^{(B)}\right|}{s_m^{\mu}} + \frac{\left|\sigma_m^{(A)} - \sigma_m^{(B)}\right|}{s_m^{\sigma}}. \tag{2}$$

This work was supported in part by the U.S. National Science Foundation under grant DUE-01211596

In the image retrieval task, the algorithm returns the N_R images from the database that are closest to the query image in the sense of the metric (2). Thus, the features used in the Manjunath-Ma algorithm are the first moments and second central moments of the Gabor filter magnitude responses; these are often referred to as *Gabor texture features*. In [2], these features were found to perform better than similar moments computed from orthogonal and biorthogonal wavelet transform coefficients of both the standard and tree-structured types. Similarly, they were found to significantly outperform Tamura features and edge histograms in [4], as well as tree-structured and pyramid-structured wavelet transform features in [3,4].

Only in a few instances have competing methods delivered performance equal to or better than these Gabor texture features. In [3], features computed on MRSAR model parameters [8] in conjunction with the Mahalanobis distance achieved a retrieval rate within 1.5% of that delivered by the Gabor texture features. In [4], the MRSAR features actually performed slightly better than Manjunath-Ma against a database of images from Corel photo galleries. However, Manjunath-Ma performed slightly better on a database of Brodatz textures. For a set of images from the MIT VisTex database, certain complex wavelets used in conjunction with a novel Bayesian distance metric achieved a retrieval rate approximately 3% better than the Gabor texture features [5]. These data make it clear that for the texture retrieval problem introduced in this section, the Gabor texture features of the Manjunath-Ma algorithm consistently achieve a retrieval rate that is better than or within a very few percent of the best known competing methods. Moreover, a series of psychophysical experiments in [7] found that texture similarity as measured by (2) agreed well with human texture perception.

3. MODULATION DOMAIN ALGORITHM

For a real-valued image $s : \mathbb{R}^2 \to \mathbb{R}$, a single-component AM-FM image model is given by $s(\mathbf{x}) = a(\mathbf{x})\cos[\varphi(\mathbf{x})]$ [9–11]. This AM-FM model provides a rich description of the local texture structure. The AM function $a(\mathbf{x})$ captures the local texture contrast and intensity. The FM function is the vector field $\nabla\varphi(\mathbf{x})$. It's orientation and magnitude describe, respectively, the local texture orientation and coarseness (or granularity). For a real image $s(\mathbf{x})$, the AM and FM functions are not unique. In analogy to the 1-D *analytic signal* popularized in communication theory, we disambiguate this situation by applying the directional 2-D Hilbert transform $\mathcal{H}$ described in [12] to construct the *analytic image* $t(\mathbf{x}) = s(\mathbf{x}) + j\mathcal{H}[s(\mathbf{x})]$. As is true for any complex-valued signal, the AM and FM functions of $t(\mathbf{x})$ are unique.

One computes an AM-FM image model by performing *demodulation* to estimate the unknown modulating functions $a(\mathbf{x})$ and $\nabla\varphi(\mathbf{x})$. All practical image demodulation algorithms involve approximations that may suffer from large errors if the modulating functions are not locally smooth in a certain sense described as *local coherency* [9, 11]. Therefore, while single-component AM-FM models exist for all images in theory, it is generally impossible to compute them. In practice, we circumvent this difficulty by considering the image to be a sum of AM-FM components wherein each component admits locally coherent modulating functions. This results in the multi-component AM-FM image model $s(\mathbf{x}) = \sum_{i=1}^{M} s_i(\mathbf{x})$ and associated analytic image

$$t(\mathbf{x}) = s(\mathbf{x}) + j\mathcal{H}[s(\mathbf{x})] = \sum_{i=1}^{M} s_i(\mathbf{x}) + j\mathcal{H}[s_i(\mathbf{x})] = \sum_{i=1}^{M} a_i(\mathbf{x})\exp[j\varphi_i(\mathbf{x})]. \quad (3)$$

Direct access to the AM-FM components $t_i(\mathbf{x}) = a_i(\mathbf{x})\exp[j\varphi_i(\mathbf{x})]$ in (3) is generally unavailable. Estimation of the modulating functions $a_i(\mathbf{x})$ and $\nabla\varphi_i(\mathbf{x})$ may be accomplished by analyzing $t(\mathbf{x})$ with a multiband filterbank to isolate the components from one another on a jointly localized basis in space and spatial frequency. Since joint localization is clearly required, filterbanks comprising sampled complex-valued Gabor functions have been used often for this purpose. The demodulation problem then becomes one of estimating the component modulating functions from the filterbank channel responses. At some particular $\mathbf{x} \in \mathbb{R}^2$, suppose that component $t_i(\mathbf{x})$ dominates the response $y_m(\mathbf{x})$ of filterbank channel m and let G_m and g_m be the channel frequency response and impulse response, so that $y_m(\mathbf{x}) = t(\mathbf{x}) * g_m(\mathbf{x}) \approx t_i(\mathbf{x}) * g_m(\mathbf{x})$. Then the AM and FM functions of $t_i(\mathbf{x})$ are estimated using the spatially local demodulation algorithms

$$\nabla\varphi_i(\mathbf{x}) \approx \mathrm{Re}\left[\frac{\nabla y_m(\mathbf{x})}{j y_m(\mathbf{x})}\right], \quad (4)$$

$$a_i(\mathbf{x}) \approx \left|\frac{y_m(\mathbf{x})}{G_m[\nabla\varphi_i(\mathbf{x})]}\right|. \quad (5)$$

Detailed bounds on the approximation errors inherent in these algorithms and also in their discrete versions were developed in [11].

The M pairs of modulating functions $a_i(\mathbf{x})$ and $\nabla\varphi_i(\mathbf{x})$ in (3) estimated via (4) and (5) provide a dense and powerful description of the local texture structure at a fundamental level that forms the basis for a plausible theory of texture perception in biological vision systems [9, 13–15]. Now consider the relationship between the AM-FM model described in this section and the Gabor texture features utilized by the Manjunath-Ma algorithm described in Section 2. The Gabor magnitude responses may be regarded as course approximations of the true image amplitude modulations $a_i(\mathbf{x})$. The frequency modulations $\nabla\varphi_i(\mathbf{x})$ are not explicitly considered in the Manjunath-Ma algorithm. However, a large magnitude response in a given channel indicates the presence of a strong component in the image with orientation and granularity that are coarsely approximated by the channel center frequency. Hence, we argue that the Gabor texture features of the Manjunath-Ma algorithm are coarse approximations of the image modulations and that this, in fact, explains the success that they have achieved. It then stands to reason that significant performance gains should be realizable by incorporating explicit estimates of the AM and FM functions into the technique. We express the FM functions in terms of their magnitudes $R_i(\mathbf{x}) = |\nabla\varphi_i(\mathbf{x})|$ and orientations $\theta_i(\mathbf{x}) = \arg\nabla\varphi_i(\mathbf{x})$. For each filterbank channel, we then compute the first and second central moments of the modulating functions to obtain a six-element vector

$$f_m = [\mu_a\ \sigma_a\ \mu_R\ \sigma_R\ \mu_\theta\ \sigma_\theta]^T, \quad (6)$$

where μ_a and σ_a are the mean and standard deviation of the AM function $a_i(\mathbf{x})$ in (5) and the moments of the FM function are defined analogously. The modulating function moments in the vector

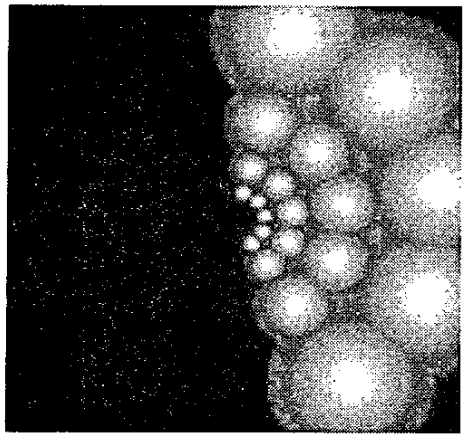

Fig. 1. Spectral depiction of the filterbank used for the proposed modulation domain CBIR algorithm.

f_m characterize the image component $t_i(\mathbf{x})$ in (3) that dominates the response of filterbank channel m. The vectors f_m are then concatenated to obtain an overall modulation domain feature vector

$$\mathbf{f} = [f_1^T \mid f_2^T \mid \ldots \mid f_{M-1}^T]^T \tag{7}$$

that describes the texture content of the image $s(\mathbf{x})$. The distance between two such feature vectors is defined in the same spirit as (2): it is an ℓ^1-norm where each component of the difference is normalized by its standard deviation computed over the entire database.

The Gabor filterbank used for the modulation domain algorithm described in this section is depicted in the spatial frequency plane in Fig. 1, where the frequency origin is located at the center, the positive horizontal frequency axis points right, and the positive vertical frequency axis points down. This figure shows the log-magnitude spectra of all channels. Special nonlinear scaling has been applied for display to accentuate the intersections between filters. Design of biologically motivated filterbanks of this type is described in great detail in [16]. The filterbank of Fig. 1 comprises twenty channels at five orientations. At each orientation, the filter radial center frequencies begin at 8.5 cycles per image and follow a geometric progression with common ratio 1.9. All of the filters in the bank have a half-peak bandwidth of 1.5 octaves. There are three significant differences between this filterbank and the one used by the Manjunath-Ma algorithm. First, the filters in Fig. 1 are all circularly symmetric. Second, with the modulation domain algorithm proposed here, the input images are always floating point images normalized to have zero mean and extremes in the range $[-1, +1]$. Thus, there is no need to zero the DC response of the channel filters. Finally, the filterbank of Fig. 1 is much denser than the one used in the Manjunath-Ma algorithm. In the filterbank of [3], the half-peak contours of any two adjacent filters intersect at a single frequency. For the filterbank of Fig. 1, the half-peak contours of any *four* adjacent filters intersect at a single frequency. While this complicates the filterbank design as described in [16], it provides a tangibly improved coverage of the frequency plane.

4. EXPERIMENTAL RESULTS

The experimental framework was described in some detail in Section 2 so that the relative performance of several existing texture retrieval techniques could be discussed with regards to a specific, well-defined, and well-studied retrieval task. In particular, we pointed out that, for Brodatz and Brodatz-like texture images, the Manjunath-Ma algorithm given in [3] consistently achieves a retrieval rate that is either the best known or is within 5% of the best known performance. Therefore, in this section we construct a specific set of experiments to test performance of the modulation domain algorithm described in Section 3 against the Manjunath-Ma algorithm. Both algorithms were implemented in ANSII C under the GNU g++ compiler.

The original images were a set of $N_I = 1,088$ homogeneous grayscale texture images of size 512×512. These included images from the MIT VisTex database, Brodatz texture images, and other Brodatz-like images organized into the general categories brick, flowers, water, buildings, food, leaves, bark, fabric, and grass. Each original image was partitioned into $N_T = 16$ disjoint tiles to obtain a retrieval database of 17,408 images of size 128×128. The number of these images falling into each category is given in the second column of Table 1.

For each query image, the top $N_R = 16$ matches from the entire database were returned by each algorithm, where N_C was set equal to the number of these 16 that came from the same 512×512 original image as the query image. The retrieval rate for each algorithm was then calculated according to $R_C = N_C/16 \times 100\%$. Every one of the 17,408 images in the database was used as the query image and retrieval rates were averaged within each image category. Results for the proposed modulation domain algorithm are given in column three of Table 1, while those for the Manjunath-Ma algorithm are given in column four.

As can be seen from Table 1, the performance advantage obtained by using the full AM-FM model was significant and ranged from 9% to 30%. This performance gain is largely due to the explicit consideration of the FM magnitude and orientation in the proposed algorithm. With regards to all of the competing techniques considered in Section 2, it should be kept in mind that, for this retrieval task, only the MRSAR features of [8] and the complex wavelet features of [5] have ever been found to outperform the Manjunath-Ma Gabor texture features, and both of these did so by margins of less than 5%. Thus, from the data in Table 1, we may infer that, for this task, the modulation domain texture retrieval algorithm proposed here may be expected to outperform *all* of the techniques discussed in Section 2 by margins of *at least* 5% - 25%. Results for a typical run of the experiment are depicted in Fig. 2.

In [2–6], the retrieval rate R_C was also studied as a function of the number of returns N_R. For the modulation domain texture algorithm of Section 3, some limited, preliminary results of this type are given in Table 2. The experiment was repeated for a reduced set of only $N_I = 75$ original 512×512 images from the VisTex database. Again with $N_T = 16$, this provided a database of 1,200 images of size 128×128. Twenty-five queries were run against this database for $N_R = 16$, 20, and 23. Retrieval rates R_C for the modulation domain algorithm and the Manjunath-Ma algorithm are given in columns two and three of Table 2, respectively, where, despite the reduced size of the experimental data set, the results are seen to be quite consistent with those in Table 1. Finally, it should be noted that the results we have obtained for the Manjunath-Ma

Category	No. Images	Retrieval Rate (R_C)	
		Proposed Alg.	Competing Alg.
Bark	3072	64.5%	51.7%
Brick	1536	75.4%	60.8%
Buildings	256	87.9%	57.4%
Fabric	2816	94.2%	85.2%
Flowers	1792	75.9%	57.9%
Food	2560	75.4%	62.2%
Leaves	3584	65.7%	49.8%
Water	1792	68.2%	55.9%

Table 1. Experimental retrieval rates for the proposed modulation domain algorithm and the competing algorithm of [3].

No. Returns (N_R)	Retrieval Rate (R_C)	
	Proposed Alg.	Competing Alg.
16	84.0%	70.5%
20	89.5%	74.8%
23	91.0%	77.3%

Table 2. Preliminary data giving the retrieval rate as a function of the number of returns N_R for the proposed algorithm and the competing algorithm given in [3].

algorithm in Tables 1 and 2 are in *remarkably* close agreement with the performance figures originally published in [3].

5. DEPLOYMENT IN THE DIRECT CBIR SERVICE

The main features of DIRECT, the *Decentralized Image Retrieval for Education* system, are summarized in this section. DIRECT is a service that provides search by query CBIR functionality for digital libraries using existing interfaces and without imposing new standards, protocols, or behaviors on content providers. Thus, DIRECT may be configured for interoperability with, e.g., the NASA JOINed Digital Library (JDL), the Digital Library for Earth System Education (DLESE), and the National Science Digital Library (NSDL). The DIRECT concept of operations provides a global mode where retrieval is performed on whole images using color and texture features and also a local mode where it is performed on image segments. For the local mode, segmentation is performed automatically when the images are acquired and indexed, and Fourier shape descriptors are used as an additional feature set.

A DIRECT prototype was launched in the JDL in December, 2002, where only the global mode was implemented. For each image, DIRECT computes a quantized sixteen-element histogram for each color in the RGB color space, where the quantization is performed using a table that exploits the color discrimination of the human visual system [17]. This results in a 48-element color feature vector for each image. The distance between a pair of color feature vectors computed from two images is defined using the intersection distance. The texture algorithm currently deployed in the prototype is wavelet-based. The RMS values of the wavelet coefficients at each scale are concatenated to construct texture feature vectors. The distance between texture vectors is computed using the usual Euclidean metric. This algorithm will be replaced by the modulation domain approach introduced in this paper. The similarity between images is quantified by a weighted sum of appropriate distance metrics on the color and texture feature vectors. Users have the capability to specify the weights at query time.

In a heterogeneous library, DIRECT runs autonomously from the individual collections and provides a network of services by creating a brokered peer-to-peer system where portals function as clients, collections function as service providers, and the DIRECT index and CBIR algorithm suite function as both clients and service providers. In this sense, DIRECT is analogous to and combines the strengths of file sharing systems such as Gnutella, the Freenet project, and Napster. Structurally, the DIRECT index looks very similar to the Napster centralized index and exhibits similar performance. However, DIRECT is a peer service as in Gnutella and Freenet, and must gather the information it needs through mechanisms already provided by the collections. Portals that use DIRECT (such as DLESE's graphical user interface) can dictate the mode in which the service is administered.

When a new collection is added to the library or when an existing collection adds new resources, DIRECT searches the collection and acquires metadata for the new resources using the same interface that would be used by a portal. These metadata are used to identify image and web page resources and to obtain their URL's. Image resources are acquired and processed immediately. For web page resources, DIRECT crawls the pages to gather any images that were not independently cataloged. Feature vectors for performing CBIR are computed from all acquired images.

6. CONCLUSION

We presented a new, high-performance modulation domain texture retrieval algorithm that is significant because it is the first to use computed AM-FM image models for performing CBIR. In addition to estimating the local texture amplitude, this new algorithm also explicitly estimates the image frequency modulations, thereby obtaining a powerful characterization of the local texture orientation and granularity. This provides a substantial performance gain compared to existing techniques where the texture features are generally based on Gabor magnitude responses alone.

Indeed, many established image processing techniques that utilize Gabor filter magnitude responses and their moments may be interpreted as coarsely approximating computed AM and FM functions. We have found that a performance advantage can often be achieved by estimating the modulations explicitly. Another advantage of this approach is that the demodulation algorithm (4), (5) estimates properties of the input signal and is essentially independent of the particular filterbank structure. In theory, this implies that modulations from images or image segments of different sizes and shapes can be directly and meaningfully compared despite the fact that their computation will generally require differing filterbank implementations, thus overcoming one of the long standing limitations characteristic of current techniques based on filterbanks, wavelets, and related transforms.

7. REFERENCES

[1] A.W.M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, "Content-based image retrieval at the end of the early years," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 22, no. 12, pp. 1349–1380, Dec. 2000.

[2] W.Y. Ma and B.S. Manjunath, "A comparison of wavelet transform features for texture image annotation," in *Proc.*

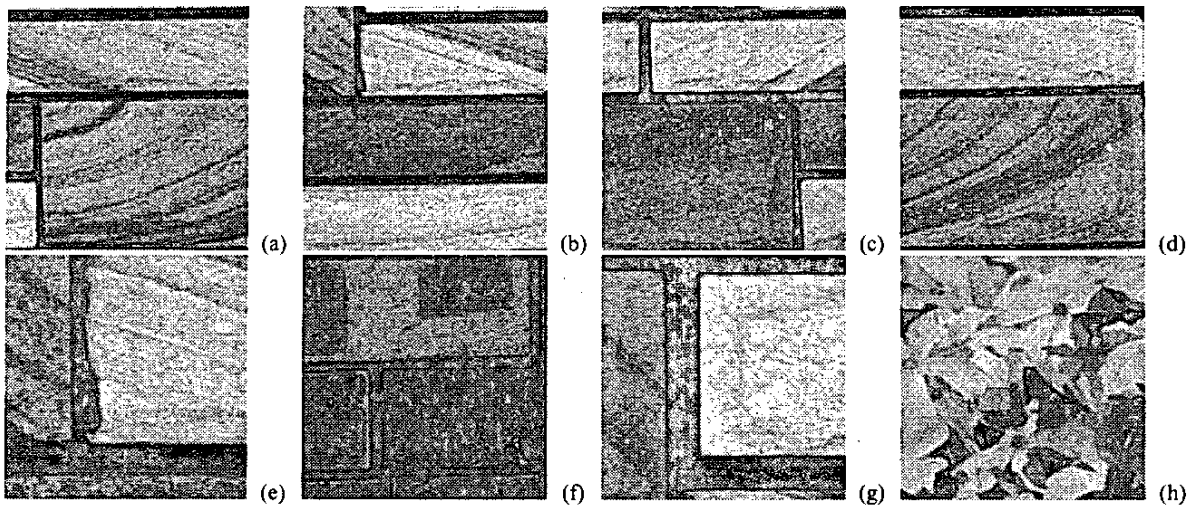

Fig. 2. Retrieval results from a typical trial of the experiment. (a) Query image; it was returned by both algorithms. (b)-(c) Two images from the same original as the query image. Both algorithms returned these two. (d) Another image from the same original as the query image. This is an example of a correct match that was returned by the proposed algorithm, but not by the algorithm of [3]. (e)-(f) Out of the top $N_R = 20$ returns from the proposed algorithm, these are the "best" and "worst" that did *not* come from the same original as the query image. (g)-(h) Out of the top $N_R = 20$ returns from the algorithm of [3], these are the "best" and "worst" that did *not* come from the same original as the query image.

IEEE Int'l. Conf. Image Proc., Washington, DC, Oct. 23-26, 1995, vol. 2, pp. 256–259.

[3] B.S. Manjunath and W.Y. Ma, "Texture features for browsing and retrieval of image data," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 16, no. 8, pp. 837–842, Aug. 1996.

[4] W.-Y. Ma and H. J. Zhang, "Benchmarking of image features for content-based retrieval," in *Proc. 32nd IEEE Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, November 1-4 1998, pp. 253–257.

[5] P. de Rivaz and N. Kingsbury, "Complex wavelet features for fast texture image retrieval," in *Proc. IEEE Int'l. Conf. Image Proc.*, Kobe, Japan, October 24-28 1999, vol. 1, pp. 109–113.

[6] C. Wolf, J.-M Jolion, W. Kropatsch, and H. Bischof, "Content based image retrieval using interest points and texture features," in *Proc. Int'l. Conf. Pattern Recog.*, Barcelona, Spain, September 3-7 2000, vol. 4, pp. 234–237.

[7] B. Zhu, M. Ramsey, and H. Chen, "Creating a large-scale content-based airphoto image digital library," *IEEE Trans. Image Proc.*, vol. 9, no. 1, pp. 163–167, Jan. 2000.

[8] J. Mao and A. Jain, "Texture classification and segmentation using multiresolution simultaneous autoregressive models," *Pattern Recog.*, vol. 25, no. 2, pp. 173–188, 1992.

[9] A. C. Bovik, N. Gopal, T. Emmoth, and A. Restrepo, "Localized measurement of emergent image frequencies by Gabor wavelets," *IEEE Trans. Info. Theory*, vol. 38, no. 2, pp. 691–712, Mar. 1992.

[10] P. Maragos and A. C. Bovik, "Image demodulation using multidimensional energy separation," *J. Opt. Soc. Amer. A*, vol. 12, no. 9, pp. 1867–1876, Sep. 1995.

[11] J. P. Havlicek, D. S. Harding, and A. C. Bovik, "Multidimensional quasi-eigenfunction approximations and multicomponent AM-FM models," *IEEE Trans. Image Proc.*, vol. 9, no. 2, pp. 227–242, Feb. 2000.

[12] J. P. Havlicek, J. W. Havlicek, and A. C. Bovik, "The analytic image," in *Proc. IEEE Int'l. Conf. Image Proc.*, Santa Barbara, CA, Oct. 26-29, 1997.

[13] J. P. Jones and L. A. Palmer, "An evaluation of the two-dimensional Gabor model of simple receptive fields in cat striate cortex," *J. Neurophysiol.*, vol. 58, no. 6, pp. 1233–1258, 1987.

[14] J. G. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *J. Opt. Soc. Am. A*, vol. 2, no. 7, pp. 1160–1169, Jul. 1985.

[15] A. C. Bovik, M. Clark, and W. S. Geisler, "Multichannel texture analysis using localized spatial filters," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 1, pp. 55–73, Jan. 1990.

[16] J. P. Havlicek, A. C. Bovik, and D. Chen, "AM-FM image modeling and Gabor analysis," in *Visual Information Representation, Communication, and Image Processing*, C. W. Chen and Y. Zhang, Eds., pp. 343–385. Marcel Dekker, New York, 1999.

[17] W.Y. Ma and B.S. Manjunath, "NeTra: A toolbox for navigating large image databases," in *Proc. IEEE Int'l. Conf. Image Proc.*, Santa Barbara, CA, Oct. 26-29, 1997, pp. 568–571.