

Speed v. Accuracy for High Resolution Colour Texture Classification

A. Monadjemi, B. T. Thomas, and M. Mirmehdi
Department of Computer Science, University of Bristol,
Bristol, BS8 1UB, England
{monadjem,majid,barry}@cs.bris.ac.uk

Abstract

Methods for extracting features and classifying textures in high resolution colour images are presented. The proposed features are directional texture features obtained from the convolution of the Walsh-Hadamard transform with different orientations of texture patches from high resolution images, as well as simple chromatic features that correspond to hue and saturation in the HLS colour space. We compare the performance of these new features against Gabor transform features combined with HLS and *Lab* colour space features. Multiple classifiers are employed to combine both textural and chromatic features for better classification performance. We demonstrate a considerable reduction in computational costs, whilst maintaining close accuracy.

1 Introduction

Scene classification using texture analysis is a prime example of a computationally expensive process where there is usually a need for trade-off between speed and accuracy. Examples of such bartering are sometimes inherent in the nature of past works on texture analysis or less often the subject of explicit analysis [1, 2, 3, 4, 5]. The problem is exacerbated as the size of the image under analysis increases, involving more and more computations. Indeed, in [5] we presented a frequency space analysis of very high resolution images (4032×2688 pixels) aimed only at increasing the accuracy of texture segmentation. The trade-off issue and the penalties expended by the computational costs were not considered important, and it was found that a 6% increase in accuracy could be achieved albeit at some considerable computational expense. Here, we show that by using a faster approach, similar levels of accuracy can be maintained.

Past works in greylevel texture analysis have increasingly found success in the classification of texture features derived from Gabor filters [6, 7, 1, 4] or Wavelets [2, 3, 8] (to name but a few). The discriminatory power of such features have recently been further strengthened through the use of chromatic features based on such colour spaces as RGB, *Lab*, or HLS [9, 10]. All such works have considered images of a size typically of the order of 64×64 , 128×128 or 256×256 pixels.

The aim of this work is to combine texture features with colour features and apply them to the classification of 128×128 patches *from very high resolution* 4032×2688 pixel colour images of outdoor scenes containing roads, cars, pavements and trees. We examine different feature sets and show a trade-off between accuracy and the speed of computation. Our images embody the fine and coarse resolution characteristics of various

textures, for example images of roads and pavements are particularly rich in high frequency information content. Some researchers disregard the higher frequencies since the power spectra of natural images show an exponential decay with frequency and in most cases the image acquisition is made through a conventional optical system that filters out the very high frequencies [11]. Nevertheless, in [5] we showed that higher frequency information extracted from high resolution images can improve the classification performance. It must be emphasised that in the high resolution images we have used for our experiments, in [5] and in this paper, the *lower frequencies* with which we are working are close to maximum frequencies present in normal resolution images in other works, whilst the *higher frequencies* are well beyond them [5].

Here, we examine the performance of Gabor and *Lab* features as a highly accurate set for colour texture classification against a slightly less accurate but much faster set of novel features. These new proposed features are a combination of directed textural features extracted using the application of Walsh-Hadamard transforms (WHT) to oriented images as well as chromatic features that correspond to, but are more easily computed than, hue and saturation in the HLS or HSI spaces. Gabor features allow us to conveniently capture the low to high frequencies present in the Fourier space of high resolution images. In the same way, we are proposing a method of capturing these frequencies through the rotation of the image by varying angles before applying the WHT. Unser [12] used Hadamard matrices along with other local transforms such as DCT and KLT in texture measurement. He applied different small size filters and a filter sliding scheme in the spatial domain to evaluate the effectiveness of these filters in texture analysis. Kim and Cho [13] also have implemented Walsh functions in a texture segmentation task using 16 different 4×4 Walsh matrices as the textural feature extractors.

To compare the performance of the Gabor and WHT, along with combined chromatic features, we perform texture classification using neural networks. Single classifiers can lead to inefficiently trained and/or complex classifiers. Multiple classifiers are gaining popularity and can be designed such that extra classification steps can be carried out only when necessary. Multiple classifiers can be simply defined as when using more than one classifier, or feature set or both, and combining their outputs by employing a *Combination Rule* to obtain the final answer. The most important motivations for using multiple classifiers are the advantages gained in simplifying the complexities of classification problems and amplification of the strengths of each individual classifier or feature set in the global classification procedure. Cappelli et al. have used multiple classifiers in fingerprint classification [14], and Wan and Fraser have implemented them for different remote sensing classification problems [15]. We perform our experiments using a single classifier at first and then illustrate the improvements gained by using two different arrangements of a multiple classifier system. In one arrangement the final decision is based on a weighted average of two separate classifiers while in the other, a secondary classifier is only used when we do not have enough confidence in the decision of the primary classifier. We refer to these different combinations of classifiers as W_AVG and PRI_SEC. Wang et al. [16] have shown that considerable improvement can be achieved when the feature set can be divided into separate subsets depending on their classification power and characteristics. This maps very well in terms of our use of textural and chromatic features.

We describe our dataset in Section 2. In Section 3, the Gabor feature set is outlined along with the new proposed WHT features and the new chromatic features. Experimental results are presented in Section 4 and the paper is concluded in Section 5.

2 High Resolution Data Set

We have a dataset consisting of 724 colour image patches of 128×128 pixels extracted from high resolution 4032×2688 images of outdoor scenes. In recent years, our research group has developed a neural network based system for classifying images of typical outdoor scenes to an area accuracy of approximately 90% [17]. Texture information is represented in this system using Gabor filters. A common problem is that many regions in typical outdoor scenes are too small to allow a significant range of spatial frequency to be included in the feature set. In [5], we presented a pilot study designed to establish if high resolution images would provide a sufficient increase in texture information to justify the extra computational complexity. We found that a 6% increase in accuracy could be achieved at some considerable computational expense. Here, we show that by using a faster approach, i.e. through our proposed oriented WHT and chromatic features, similar levels of accuracy can be attained.




Figure 1: Sixteen samples from four classes, from top-left: tree16, tree17, tree18, tree19, road15, road12, road8, road14, pave8, pave9, pave10, pave7, car5, car2, car1, and car8.

The images come from four categories: trees, pavements, cars, and roads (typical examples in Figure 1). These patches of high resolution images contain shiny, fairly smooth bodies of cars (but including wheels, door-handles, lights etc., rough and coarse surfaces of pavements, fine resolution granularity of road surfaces, and fine and coarse structures within trees and bushes. These provide a wide range of characteristics and frequencies in the data set. Many such frequencies are diminished or lost in lower resolution images.

3 The Features

In this section, the Gabor transform features, the proposed directional Walsh-Hadamard features and the proposed simple chromatic features are presented and discussed.

3.1 Gabor Features

Gabor filters are widely used for multi-frequency, multi-directional analysis in image processing. Specifically, they have shown high performance as feature extractors for texture discrimination and unsupervised texture classification [6, 7, 1, 4, 11, 9]. The important strength of Gabor filters is that they facilitate oriented or directional band pass filtering of the input texture. This allows the filter to extract notable textural features which are directional and, in a frequency sense, band-limited. A basic practical disadvantage in Gabor

filtering is their high computational costs. The Gabor filter in the frequency domain is:

$$G(u, v) = e^{-\pi\left(\frac{u_p^2}{\sigma_x^2} + \frac{v_p^2}{\sigma_y^2}\right)} \cdot e^{-2\pi j(x_0 u + y_0 v)} \quad (1)$$

where $u_p = (u - \omega_x) * \cos(\theta) + (v - \omega_y) * \sin(\theta)$ and $v_p = -(u - \omega_x) * \sin(\theta) + (v - \omega_y) * \cos(\theta)$ are the rotated/displaced coordinates in the frequency plan, ω_x and ω_y are filter central frequencies (modulation factors) in x and y directions, θ is filter orientation parameter, σ_x and σ_y are filter standard deviations in x and y directions, and x_0 and y_0 are horizontal and vertical displacements in the spatial domain. We keep $x_0 = 0$, $y_0 = 0$, and set $\omega_x = \omega_y$, and $\sigma_x = \sigma_y$ in all experiments. Our Gabor filter bank contained twelve filters arranged in 3 frequency and 4 orientation bands as,

$$G_{ij}(\omega_i, \theta_j), \quad \omega_i = [16, 32, 64], \quad \theta_j = [0, 45, 90, 135] \quad (2)$$

3.2 WHT (Hadamard) Features

In the family of orthogonal linear transforms of time/spatial domain signals, which mostly employ sinusoidal-based kernel functions (e.g. Fourier or cosine transforms), the Walsh transform is defined as:

$$W(u, v) = \frac{1}{N} \sum_{x=0}^{N-1} \sum_{y=0}^{N-1} I(x, y) [(-1)^{\psi(u, v, x, y)}] \quad (3)$$

where N is the image size, I is the image and ψ determines the transform's parametric kernel function $-1^{\psi(\cdot)}$. The Walsh transform is one of the exceptions in this family that implies sequency-based kernel functions and decomposes the input signal into rectangular wave primitives in the transform domain [18]. The kernel function can be selected from a diverse set of possibles. For instance, in the Hadamard natural transform (a member of the Walsh family) ψ is:

$$\psi(u, v, x, y) = \sum_{i=0}^{m-1} [b_i(x)b_i(u) + b_i(y)b_i(v)] \quad (4)$$

where $b_i(z)$ is the i^{th} bit of z in binary representation. The Walsh-Hadamard form of "digital frequency" or sequency, which is the number of zero-crossings or sign changes of the signal, is analogous to frequency in the Fourier transform. However, sequency is twice the size of the frequency of a signal.

As shown in Figure 2, the orthogonal set of rectangular waveforms that generate the WHT kernel can only approximate a sinusoidal waveform by weighted summation of their square wave elements. Hence, we expect our WHT features to be weaker representations of the texture in comparison to those of the Gabor. Nevertheless, the WHT has important computational advantages. For instance, it is a real transform, it only needs addition and subtraction operations, and if the input signal is a set of integer-valued data (as in the

case of digital images), we need only use integer operations. Furthermore, there is a fast algorithm for Walsh transforms by simple substitution of the exponential kernel of the Fast Fourier transform with the $-1^{\Psi(\cdot)}$ kernel of Walsh. The transform matrix, usually referred to as Hadamard, can also be saved in binary format resulting in a decrease in memory requirements.




Figure 2: Synthesising an example sine wave: (Left) Original signal $x(t) = 63\cos(0.16t)$, (Right) Walsh-Hadamard approximation $f(t) = 41h(2,t) + 17h(4,t) + 8h(8,t)$.

We used a *sequency-ordered Hadamard* matrix [18] where the rows (and columns) are ordered according to their sequency, i.e. in the first row there are no sign changes, and in the n^{th} row there are $n-1$, e.g. see Figure 3(left). The 2D Walsh-Hadamard transform can be defined as $WHT(I) = H.I.H'$ where I is the image and H and H' are the Hadamard matrix and its transpose. Note that for a Hadamard matrix $H = H'$. Figure 3(right) shows the sequency space of the sequency-ordered transform and its bands.




Figure 3: (Left) Sequency-ordered 8×8 Hadamard, (Right) Sequencies in WHT domain

We generate *oriented* Hadamard based features to represent the directionality of texture in the same manner of Gabor features. In this scheme, the Hadamard matrix remains constant but the image function is *rotated* by $\alpha = 0^\circ, 45^\circ, 90^\circ, 135^\circ$. The rotation is applied to each element in the top row of the image matrix. At border pixels corresponding elements are used from a repeated imaginary version of the same image matrix. For example, in the simple 4×4 image matrix below, a 45° rotation at position b gives $\{b, g, l, m\}$ (also see A_{45° in equation (6)) and a 135° rotation at same position b gives $\{b, e, l, o\}$ (similarly see A_{135° in (6)):

$$\begin{bmatrix} a & \mathbf{b} & c & d \\ e & f & \mathbf{g} & h \\ i & j & k & \mathbf{l} \\ m & n & o & p \end{bmatrix} \begin{bmatrix} a & \mathbf{b} & c & d \\ \mathbf{e} & f & g & h \\ i & j & k & l \\ \mathbf{m} & n & o & p \end{bmatrix} \quad (5)$$

The full rotation set where $\alpha = 0^\circ, 45^\circ, 90^\circ, 135^\circ$ can be defined for a simple 4×4 image matrix as follows:

$$A_{0^\circ} = \begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \\ m & n & o & p \end{bmatrix} A_{45^\circ} = \begin{bmatrix} a & f & k & p \\ b & g & l & m \\ c & h & i & n \\ d & e & j & o \end{bmatrix} A_{90^\circ} = \begin{bmatrix} a & e & i & m \\ b & f & j & n \\ c & g & k & o \\ d & h & l & p \end{bmatrix} A_{135^\circ} = \begin{bmatrix} a & h & k & n \\ b & e & l & o \\ c & f & i & p \\ d & g & j & m \end{bmatrix} \quad (6)$$

Note that this is not an ordinary geometrical rotation. For example, we create the rows of the A_{45° image by considering the pixels that sit in a 45° direction in image A_{0° and so on. This means that the resulting horizontal rows capture the information at the specified angles.

Of course we extend this concept to 128×128 images in this work. The artifacts of repetition at the borders are much less significant with larger images than those in the examples above. The manner of our *rotations* means that after the WHT transformation we need only extract features for *row sequency information only* corresponding to the directions used. This reduces the $WHT(A_\alpha) = H \cdot A_\alpha \cdot H'$, to a more simplified transform where $WHT(A_\alpha) = A_\alpha \cdot H'$ and we can ignore, without loss, the column sequency information. Indeed, this takes out the redundancy created by the fact that $A_{0^\circ} = A'_{90^\circ}$ and $H = H'$. From each transformed matrix, the maximum value, the mean and the standard deviation were computed resulting in a set of 12 WHT features.

The relative arrangement of pixels is of importance in texture analysis, for example in cooccurrence matrices. Similarly, sequency based features which represent the number of zero-crossings of pixels in a particular direction can convey a notable amount of textural information. We can measure the WHT energy in $WHT(A_\alpha) = A_\alpha \cdot H'$ as the absolute value of the WHT output along each column. For example, as Figure 4 depicts, the sequency representation of a typical fine resolution texture will show more energy in higher sequency bands compared to a coarse resolution texture. The rightmost graph in Figure 4 also illustrates the lack of response of the WHT transform to the coarse vertical texture when it is rotated (i.e A_{90°). One main advantage of the proposed Hadamard based feature extraction scheme is that by using 128×128 Hadamard matrices, rather than the usual 3×3 in [12] or 4×4 in [13], we can extract higher frequency/sequency information. This is easily facilitated by our use of higher resolution images.




Figure 4: From Left - Example average energies for fine resolution texture, coarse resolution texture, and coarse resolution texture at 90° rotation.

3.3 Chromatic Feature Extractors

In [10], the authors applied Gabor filtering to each R, G, and B band to obtain *colour* Gabor features. Here we perform our Gabor or Hadamard transforms on greylevel images and add separate chromatic features for the classification stage. This is appropriate since each image patch will be classified as whole. We define a pool of colour features obtained from *Lab* and HLS spaces and use them in our experiments to decide which colour space performs best given the speed v. accuracy tradeoff issue we are concerned with here.

For the *Lab* space, we use the mean μ and standard deviation σ of the chromaticity bands *a* and *b* in each image patch. In the HLS space, the mean and standard deviation of the hue and saturation bands are used. Additionally, we define two new pseudo-hue and pseudo-saturation features H_p and S_p which are much faster to compute:

$$H_p = \begin{cases} \frac{R}{3} & \text{if } R \text{ is Maximum} \\ \frac{G+255}{3} & \text{if } G \text{ is Maximum} \\ \frac{B+(2 \times 255)}{3} & \text{if } B \text{ is Maximum} \end{cases} \quad S_p = \text{Max}(R, G, B) - \frac{R+G+B}{3} \quad (7)$$

H_p is a mapping from RGB to a hue-like value, where $\max(R,G,B)$ will be mapped to an appropriate $\frac{1}{3}$ division of the possible range [0,255]. S_p measures the difference between the maximum value and average of a colour. The μ and σ of H_p and S_p in an image patch provide an adequate indication of its main colour and extent of variation.

4 Classification Experiments

Our experiments consist of classifying our images using Gabor only, Gabor with colour features, Hadamard only, and Hadamard with colour features. The initial tests were based on a single classifier, and then we used multiple classifiers and investigated the influence of the texture features against the colour features in the classification task.

We divided our 724 images into a training set of 396, a test set of 160 and a validation set of 168 patches. For classification, we employed a back-propagation neural network with one hidden layer, optimised for the best number of nodes. To evaluate classification performance, we used the Mean Square Error, MSE, as the difference between the ground truth G (i.e. the *expected* outputs of classifiers), and the network classification C (i.e. the *actual* outputs of the classifiers) across N classes (here $N = 4$). A second metric, the *Classification Accuracy*, CA, was evaluated as the percentage of correct class assignments across the complete labelled test set. MSE and CA are therefore defined as:

$$MSE = \frac{\sum_{i=1}^N (G_i - C_i)^2}{N} \quad CA = \frac{\text{No. of correct class assignments}}{\text{Total no. of samples (i.e. 160)}} \times 100 \quad (8)$$

In the first experiment, the Gabor features performed better than the WHT features as expected with MSE errors 0.181 and 0.234 respectively (see rows 1 and 2 in Table 1). However, the performance of the WHT features were still very respectable in comparison. Next, we classified using only the features from our different colour spaces to determine their power of discrimination. The results were good (rows 3 to 5 in Table 1) since our classes of objects are fairly distinguished anyway. However, pavements and roads, and cars and trees (in various seasons), can have similar colours and colour features alone is not enough. The *ab* features performed best while the new simpler $H_p S_p$ features were

better than the normal HS features. We then combined our texture features with the colour features, the results for which are shown in rows 6 to 11 in Table 1. These were found to be all very close to each other with the best result in each texture category being Gabor+ H_pS_p and WHT+ H_pS_p features.

Test No.	Test Features	No. of Features	Performance		
			MSE	NI	CA%
1	Gabor	12	0.181	18	88.7%
2	WHT	12	0.234	26	83.7%
3	ab	4	0.119	13	91.9%
4	HS	4	0.172	17	89.4%
5	H_pS_p	4	0.143	15	90.7%
6	Gabor + ab	16	0.097	11	93.1%
7	Gabor + HS	16	0.103	12	92.5%
8	Gabor + H_pS_p	16	0.083	10	93.7%
9	WHT + ab	16	0.105	12	92.5%
10	WHT + HS	16	0.117	12	92.5%
11	WHT + H_pS_p	16	0.094	11	93.1%

Table 1: Single classifier results: for all colour spaces above, the features were μ and σ of each colour band used. NI is the number of incorrect classifications out of 160 images.

Feature Extraction	Gabor	WHT	ab	HS	H_pS_p
Time in msec	2.75	0.24	0.30	0.73	0.05

Table 2: Average execution times for the various feature extraction methods.

Table 2 shows the average time in milliseconds, on a Pentium III-700 MHz PC, required to compute each set of features. Clearly, WHT features are more than 10 times faster than Gabor features, while H_pS_p are also significantly faster than both ab and HS features. The results demonstrate that the WHT features can achieve very similar performance to Gabor features at a fraction of the computational costs. In the next set of experiments, we increase the classification accuracy by using multiple classifiers.

We used a multiple classifier system consisting of two individual classifiers, one for the textural features and one for the chromatic features. Their outputs were combined through a *combine module* to make a final decision. Two *combination algorithms* were considered. In the first, *WAVG*, an unknown input sample was presented to both classifiers. Having obtained the respective responses, the combine module computed their weighted average, using two weighting factors called k_1 and $k_2 = 1 - k_1$, and then selected the class with the minimum MSE. In the second, *PRLSEC*, a primary/secondary scheme was applied in which the chromatic classifier was used as the main primary classifier due to its relative higher classification accuracy as already shown in Table 1. If this classifier was certain of its classification (i.e. the primary classifier's MSE was less than a certain threshold δ_1), then the secondary classifier was not used. Otherwise, the secondary textural classifier was deployed to aid in the final classification if it demonstrated high certainty in its decision (i.e. if the secondary MSE was less than another threshold, δ_2).

If the secondary classifier did not demonstrate high certainty, then the primary classifier’s decision is accepted anyway as the final classification. All other conditions were the same as with the previous experiments. To determine the optimum values for k_1 , δ_1 and δ_2 a linear search of these parameter spaces was carried out by measuring the classification performance on a validation set of 168 samples. In our experiments, the parameter space was $[0, 1]$ for k_1 (and $k_2 = 1 - k_1$), and $[0, 0.5]$ for δ_1 and δ_2 .

Test Features	W_AVG			PRLSEC		
	MSE	NI	CA%	MSE	NI	CA%
<i>ab</i> + Gabor	0.076	6	96.3%	0.068	6	96.3%
<i>HS</i> + Gabor	0.084	10	93.8%	0.085	9	94.4%
$H_p S_p$ + Gabor	0.081	8	95.0%	0.073	8	95.0%
<i>ab</i> + WHT	0.079	10	93.8%	0.090	9	94.4%
<i>HS</i> + WHT	0.101	12	92.5%	0.124	13	91.9%
$H_p S_p$ + WHT	0.089	9	94.4%	0.089	8	95.0%

Table 3: Multiple classification results for both W_AVG and PRLSEC schemes. (colour features are stated first since they carry larger weight)

All combinations of textural and chromatic feature sets were applied to measure their classification performance. Table 3 shows the results for both W_AVG and PRLSEC schemes. In comparison to the single classifier there is a marginal drop in accuracy with the WHT+*HS* features at 91.9% (down from 92.5% with a single classifier) but most significantly there is an increase in accuracy in every other case. The best performance is by the Gabor+*ab* features at 96.3% (up from 93.1%) since they capture the texture characteristics more powerfully. However, as before, the WHT+ $H_p S_p$ feature set follow extremely closely at 95.0% (up also from 93.1%) in the PRLSEC scheme and, as shown earlier in Table 1, they are much faster to compute. The overall performance of the PRLSEC multiple classifier is slightly better than the W_AVG scheme. The PRLSEC classifier allows a suitable combination rule as a way of giving primary importance to the chromatic features since they perform better than the texture features on their own. However, since the linear search of the parameter space for k_1 also gives a better weighting to the chromatic features in the W_AVG scheme the results between the two classification schemes are very close.

5 Conclusion

We described novel and faster methods for extraction of both directional texture features using the Walsh-Hadamard transform and simplified hue and saturation-like chromatic features. The methods were applied to high resolution outdoor scenes for colour texture classification. The WHT concept of sequency captures the lower and higher frequencies present in high resolution images very well. The performance of the proposed features were compared, for accuracy and speed, against Gabor and HLS/*Lab* features using both single and multiple classifiers. The multiple classification framework showed improvements in accuracy in all cases when texture features were combined with *ab* or the new proposed chromatic features $H_p S_p$. Furthermore, we demonstrated that the performance of the new features are highly comparable at a massively reduced computational cost.

Although high resolution images of outdoor scenes have been of interest in the framework of our experiments here, in future tests the proposed methods will be applied on a more standard texture test suite such as Vistex to implement a more reliable comparison with other texture analysis schemes.

Acknowledgement

A. Monadjemi is funded by the Iranian Ministry of Science, Research and Technology, (IMSRT).

References

- [1] T. Reed and J. Du Buf. A review of recent texture segmentation and feature extraction techniques. *CVGIP Image Understanding*, 57(3):359–372, 1993.
- [2] M. Unser. Texture classification and segmentation using Wavelet frames. *IEEE Trans. on Image Processing*, 4(11):1549–1560, 1995.
- [3] O. Pichler, A. Teuner, and B. Hosticka. Comparison of texture feature-extraction using adaptive Gabor filtering, pyramidal and tree-structured Wavelets. *PR*, 29(5):733–742, 1996.
- [4] G. Smith and I. Burns. Measuring texture classification algorithms. *Pattern Recognition Letters*, 18(14):1495–1501, December 1997.
- [5] A. Monadjemi, B. Thomas, and M. Mirmehdi. Experiments on high resolution images towards outdoor scene classification. In *7th Computer Vision Winter Workshop*, pages 325–334, 2002.
- [6] I. Fogel and D. Sagi. Gabor filters as texture discriminator. *Biological Cybernetics*, 61:102–113, 1989.
- [7] A. Jain and F. Farrokhnia. Unsupervised texture segmentation using Gabor filters. *Pattern Recognition*, 24(12):1167–1186, 1991.
- [8] S. Livens, P. Scheunders, G. Van de Wouwer, and D. Van Dyck. Wavelets for texture analysis, an overview. In *Proc IPA, IEE Pub. No. 443*, volume 1, pages 581–585, 1997.
- [9] A. Drimbarean and P.F. Whelan. Experiments in colour texture analysis. *Pattern Recognition Letters*, 22(10):1161–1167, 8 2001.
- [10] M. Mirmehdi and R. Perissamy. Perceptual image indexing and retrieval. Accepted for the Journal of Visual Communication and Image Representation, 2002.
- [11] O. Nestares, R. Navarro, J. Portilla, and A. Taberero. Efficient spatial-domain implementation of a multiscale image representation based on Gabor functions. *Journal of Electronic Imaging*, 7(1):166–173, 1998.
- [12] M. Unser. Local linear transform for texture measurements. *SP*, 11:61–79, 1986.
- [13] D. Kim and D. Cho. Texture segmentation using Walsh spectrum and modified ART2. <http://www.icspat.com/papers/92mfi.pdf>, 2000.
- [14] R. Cappelli, D. Maio, and D. Maltoni. Combining fingerprint classifiers. In *Multiple Classifier Systems*, pages 351–361. Springer-Verlag, 2000.
- [15] W. Wan and D. Fraser. A multiple self-organizing map scheme for remote sensing. In *Multiple Classifier Systems*, pages 300–309. Springer-Verlag, 2000.
- [16] W. Wang, P. Jones, and D. Partridge. Diversity between NN and decision trees for building multiple classifier systems. In *Multiple Classifier Systems*, pages 240–249. Springer, 2000.
- [17] M. Everingham, B. Thomas, T. Troscianko, and D. Easty. NN-VR mobility aid for the severely visually impaired. In *Proc of Disability, VR & Associated Technologies*, pages 183–192, 1998.
- [18] K. G. Beauchamp. *Applications of Walsh and Related Functions*. Academic Press, 1984.