INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

A Probabilistic Framework for Adaptive Texture Description

Karen Brady — Ian Jermyn — Josiane Zerubia

N° 4920

Septembre 2003

THÈME 3

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--4920--FR+ENG

A Probabilistic Framework for Adaptive Texture Description

Karen Brady , Ian Jermyn , Josiane Zerubia

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet Ariana

Rapport de recherche n° 4920 — Septembre 2003 — 44 pages

Abstract: This report details the development of a probabilistic framework for adaptive texture description. Starting with a probability distribution on the space of infinite images, we generate a distribution on finite regions by marginalisation. For a Gaussian distribution, the computational requirement of diagonalisation leads naturally to adaptive wavelet packet models which capture the principal periodicities present in the textures and allow long-range correlations while preserving the independence of the wavelet packet coefficients.

These models are then applied to the task of segmentation. Two data types are included in our test bed: synthetic Brodatz mosaics and high-resolution satellite images. For the case of the synthetic textures, undecimated versions of the wavelet packet transform are used to diagonalise the Gaussian distribution efficiently, albeit approximately. This enables us to perform a pixelwise classification of the mosaics. A regularisation step is then implemented in order to arrive at a smooth final segmentation. In order to obtain the best possible results for the real dataset, the mean of the distribution is included in the model. The approximation made for the classification of the synthetic texture mosaics is tested on the remote sensing images, but it produces unsatisfactory results. Therefore we introduce a heuristic classification technique for this dataset, based on a decimated wavelet packet transform. The resulting segmentation is then regularised using the same method as in the synthetic case. Results are presented for both types of data and a discussion follows.

Key-words: image segmentation, texture analysis, adaptive basis, wavelet packets

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex (France)
Téléphone : +33 4 92 38 77 77 — Télécopie : +33 4 92 38 77 65

Un Cadre Probabiliste pour la Description Adaptative de Texture

Résumé : Ce rapport présente le développement d'un nouveau cadre probabiliste cohérent pour la description adaptative de texture. En partant d'une distribution de probabilité sur un espace d'images infinies, nous générons une distribution sur des régions finies par marginalisation. Pour une distribution gaussienne, les contraintes de calcul imposées par la diagonalisation nous conduisent naturellement à des modèles utilisant des paquets d'ondelettes adaptatifs. Ces modèles reflètent les principales périodicités présentes dans les textures et permettent également d'avoir des corrélations à longue portée tout en préservant l'indépendance des coefficients des paquets d'ondelettes.

Nous avons appliqué notre méthode à la segmentation. Deux types de données figurent dans notre ensemble de test: des mosaïques synthétiques de Brodatz et des images satellitaires haute résolution. Dans le cas des textures synthétiques, nous utilisons la version non-décimée de la transformée en paquets d'ondelettes afin de diagonaliser la distribution gaussienne de manière efficace, bien qu'approximative. Cela nous permet d'effectuer une classification de la mosaique pixel par pixel. Une étape de régularisation est ensuite effectuée afin d'arriver à un résultat de segmentation final plus lisse. Afin d'obtenir les meilleurs résultats possibles dans le cas de données réelles, la moyenne de la distribution est ensuite introduite dans le modèle. L'approximation faite pour la classification des mosaiques de textures synthetiques a été testée sur des images réelles, mais les résultats obtenus n'étaient pas satisfaisants. C'est pourquoi nous avons introduit, pour ce type de données, une technique de classification heuristique basée sur la transformée en paquets d'ondelettes décimée. Les résultats de segmentation sont ensuite régularisés à l'aide de la même méthode que dans le cas synthétique. Nous présentons les résultats pour chaque type de données et concluons par une discussion.

Mots-clés : segmentation d'image, analyse de texture, base adaptative, paquets d'ondelettes

Contents

1 Introduction

The ability to describe and analyse texture is an integral part of many image processing applications. In the medical domain, for example, texture based features extracted from radiographic images provides specialists with an automatic aid for the diagnosis of malignant tissue [12]. Texture is also used in industry for the purposes of quality control [9] and can help in the automatic detection of defects, such as those found in lumber wood [7]. In many remote sensing images, texture is a salient cue for applications including cartography and damage assessment following natural disasters such as fires or floods.

Over the years numerous approaches have been proposed to tackle the problem of texture description. For a full overview, see [13, 21, 22, 23]. One research area which has been extremely active in recent years is the application of wavelets to texture analysis [17, 19, 24]. An example is the Hidden Markov Tree model developed in [4], which describes the interscale dependencies of a standard wavelet decomposition.

Many textures, however, possess intermediate and high frequency content. By repetitively decomposing only the low frequencies, the standard wavelet transform does not further analyse these other areas of activity. To address this problem, one can perform a wavelet packet analysis of the texture. In [16], packets were used in a classification experiment on natural textures. The problem with this is that performing a full wavelet packet decomposition results in a large number of subbands to be analysed and some of these will not contain much information. This was the motivation behind many of the efforts made to adapt the wavelet packet decomposition to the underlying structure of the texture. The pioneering work of [5] selected the best wavelet packet basis for an image by using an entropy criterion to decide if a subband should be further decomposed or not. Since then other selection criterion have been tested, some of which proved successful [1, 3] and others not so [20]. Interesting though these methods are, none of them were developed within a coherent probabilistic framework; they involve *ad hoc* choices that fail to reveal the assumptions underlying the models.

In this report we address the issue of texture description within a probabilistic framework. Starting from a probability measure on the infinite texture, which we assume here to be Gaussian, we derive the measure for the texture on a finite region. This leads naturally to a class of adaptive wavelet packet models that capture the structure of a given texture, for example its principle periodicities, in a manner analogous to the 2-D Wold decomposition [10]. These models are tested on both synthetic and real data sets. Two simple classification rules are described, both of which enable pixelwise classification of the image while retaining the advantages of more complicated prior models. We consider supervised segmentation because it is methodologically well-defined.

The report is organised as follows. In section 2 we introduce our texture model and provide an abstract development of it. This abstract development is then applied to Gaussian models in section 3. Section 4 explains how wavelet packets fit into our framework. Parameter estimation is discussed in section 5 and optimal representations are shown for several real and synthetic textures. The resulting texture models are then tested by segmenting textured scenes. Section 6 describes the segmentation procedures used. Results on mosaics of synthetic Brodatz [2] textures are obtained using an undecimated wavelet packet transform which results in a pixelwise classification of the mosaics. For the case of the remote sensing data, in order to obtain the best segmentation results possible we include the mean of the distribution in the model. The classification method used in the synthetic case is tested on the real data and is found to be inadequate. We therefore develop a heuristic classification technique for the segmentation of real images which is based on a decimated wavelet packet transform. Section 6.3 presents the segmentation results obtained with this heuristic method. The underlying characteristics of the texture models are discussed in section 7. A brief analysis of the adaptive wavelet packet subband statistics, along with a commentary on the choice of distribution, is presented in section 8. We conclude the report in section 9 with comments on possible future directions of this work.

2 A Coherent Framework for Texture Description

One of the characteristics of planar texture, perhaps its defining characteristic, is that it is infinitely extendable, so that images of such textures are functions on an infinite (or at least very large) domain D_∞. Thus in order to model such a texture accurately, one needs a probability distribution on the space of such images that describes the variations possible for images of that texture. We denote such a distribution by:

$$\Pr(\phi|\lambda \equiv l, B_l) \tag{1}$$

where ϕ is the infinite image; B_l is the set of parameters of the model of texture $l \in L$, the set of textures; and $\lambda : D_\infty \to L$ is the class map, which here takes every pixel to texture l.

In practice however, when we analyse and segment textures, we work with images of the texture defined on finite regions. This means that one needs, not the distribution on the space of infinite images, but that on the space of images defined on a finite region $R \subset D_\infty$. We thus need to marginalise equation (1) over the values of the pixels outside this region.

Let Φ be the space of infinite images and $R \subset D_\infty$ a region. $\bar{R}$ represents the complement of this region. We denote the space of images defined on R by Φ_R and on $\bar{R}$ by $\Phi_{\bar{R}}$. Figure 1 provides a visual representation of the maps that exist between Φ and these two subspaces. The surjective map π_R restricts an infinite image in Φ to

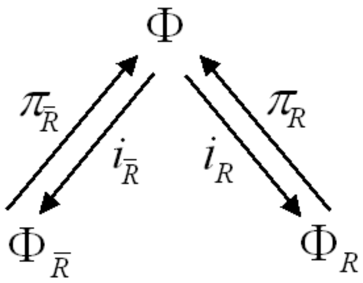

Figure 1: Maps Defined on Φ.

the region R; similarly $\pi_{\bar{R}}$ restricts an infinite image to $\bar{R}$. The injection i_R takes an image defined on the region and maps it back into Φ whereby all the values outside the region will be set to zero; similarly $i_{\bar{R}}$ takes an image defined on $\bar{R}$ and maps it back into Φ whereby the values of all the pixels inside the region will be set to zero. These four maps allow us to define two projections P_R and $P_{\bar{R}}$:

$$\begin{aligned} P_R &= i_R \pi_R \\ P_{\bar{R}} &= i_{\bar{R}} \pi_{\bar{R}} \end{aligned} \tag{2}$$

which generate an orthogonal decomposition of Φ such that an infinite image may be expressed as

$$\phi = P_R(\phi) + P_{\bar{R}}(\phi) \tag{3}$$

If we now marginalise the distribution in equation (1) over $\phi_{\bar{R}}$, we will end up with the probability measure for the image on the finite region R:

$$\Pr(\phi_R|\cdot) = \int_{\phi_{\bar{R}}} \Pr(\phi|\cdot) \tag{4}$$

which, in principle at least, solves the boundary problem for texture.

3 Model Choice: A Gaussian Distribution

We now implement the above procedure for a Gaussian distribution. In abstract notation, a Gaussian distribution can be expressed as:

$$\Pr(\phi|\cdot) = |F/\pi|^{\frac{1}{2}} e^{-\langle \phi-\mu|F|\phi-\mu\rangle} \tag{5}$$

where $\langle I|J\rangle$ is the inner product of the functions $|I\rangle$ and $|J\rangle$ in the space of images, $F|\phi\rangle$ denotes F acting on ϕ, μ is the mean of ϕ and $|F|$ is the determinant of the operator F. In the position basis, this distribution takes the form:

$$\Pr(\phi|\cdot) = |F/\pi|^{\frac{1}{2}} e^{-\sum_{(x,x')\in D_\infty}(\phi(x)-\mu)F(x,x')(\phi(x')-\mu)} \tag{6}$$

where $F(x,x^{'})$, the inverse covariance matrix, captures spatial correlations in the image ϕ. It is diagonal in the position basis only if the values of the pixels in the image are independent of each other.

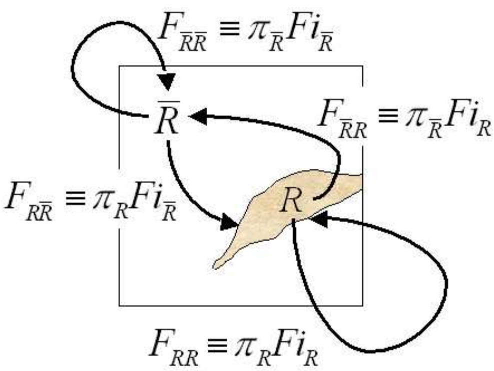

Figure 2: Splitting F according to the projections P_R and $P_{\bar{R}}$.

Once a finite region of interest R is specified for analysis, the operator F can be split up into four parts according to the orthogonal decomposition of Φ by the projections in equation (2). The meaning behind these new operators can be seen in figure 2: F_{RR} relates pixels inside the region to each other, $F_{R\bar{R}}$ relates pixels outside the region to pixels inside the region, $F_{\bar{R}R}$ relates pixels inside the region to pixels outside the region, and $F_{\bar{R}\bar{R}}$ relates pixels outside the region to each other.

$$F_{RR} = \pi_R F i_R$$

$$F_{R\bar{R}} = \pi_R F i_{\bar{R}}$$

$$F_{\bar{R}R} = \pi_{\bar{R}} F i_R$$

$$F_{\bar{R}\bar{R}} = \pi_{\bar{R}} F i_{\bar{R}} \tag{7}$$

Partitioning the operator F in this manner and marginalising over $\phi_{\bar{R}}$, as in equation (4), gives us the following probability measure for the image ϕ restricted to the finite region R:

$$\Pr(\phi_R|\cdot) \quad = \quad Z^{-1} e^{-\langle \phi_R - \mu_R | G_R | \phi_R - \mu_R \rangle} \tag{8}$$

where the operator $G_R \equiv F_{RR} - F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}$ and μ_R represents the mean of the infinite image ϕ restricted to the region R. Details of the calculations performed to arrive at this measure on the region can be found in appendix A.

Although in principle we can evaluate this operator and the exponent of equation (8), in practice computational complexity requires that we diagonalise G_R.

4 Computing the Probability of Texture on a Region

Calculating the probability in equation (8) requires us to diagonalise the operator G_R. One possible way to do this is to find a set of functions on the region

$$B = \{|a\rangle : a \in A\}, \; A \text{ an index set} \tag{9}$$

that satisfies the following two conditions:

1. The set $\{i_R|a\rangle : a \in A\}$ are eigenfunctions of the operator F (with eigenvalues f_a);

2. The set B forms an orthonormal basis for functions on R;

The first condition implies that the support of $F i_R|a\rangle = f_a i_R|a\rangle$ lies in the region R. Thus the second term in G_R is zero and the first term becomes $f_a|a\rangle$. The second condition then means that G_R is diagonalised by B, allowing us to write our measure as:

$$\Pr(\phi_R|\cdot) = \prod_{a\in A}\left(\frac{f_a}{\pi}\right)^{\frac{1}{2}} e^{-\sum_{a\in A} f_a \langle \phi_R - \mu_R|a\rangle\langle a|\phi_R-\mu_R\rangle} \tag{10}$$

Details of the calculations used in applying the conditions are left to appendix B.

So, in summary, if we can find such a set B that satisfies the above two conditions then we can diagonalise our operator G_R and hence have an expression for the probability of an image defined on a region R which is relatively quick to compute. The question that remains is how do we go about finding such a set B?

4.1 Using Wavelet Packets

As texture can appear in an image arbitrarily translated, we require that our probability measure on Φ is translation invariant. This condition has the following two consequences:

- The mean μ must be constant, so that its Fourier transform is a delta function:

$$\mu(k) \propto \delta(k)$$

- The operator F is diagonal in the Fourier basis:

$$F(k,k') = f(k)\delta(k,k')$$

This means that our measure is now characterized by a function f on the Fourier domain:

$$\Pr(\phi|\cdot) = |F/\pi|^{\frac{1}{2}} e^{-\int_k f(k)(\phi^*(k)-\mu^*(k))(\phi(k)-\mu(k))} \tag{11}$$

where $*$ denotes the complex conjugate. For an arbitrary function $f(k)$, it is very hard to find a set B that satisfies the conditions of section 4. We thus want to choose a set of functions f that is varied enough to capture the structure present in the texture, but limited enough that we can satisfy the conditions. To this end, we consider the set $\mathcal{T}$ of rooted trees in which each vertex has either four children, or two children and a label in the set $\{H, V\}$. Each such tree corresponds to a dyadic partition of the Fourier domain, where the leaves correspond to the elements of the partition. A vertex with four children corresponds to a standard quad tree split, while a vertex with two children corresponds to a two-way split in either the x (H) or the y (V) direction, but not both. The Fourier domain splits corresponding to each type of vertex are shown in figure 3.

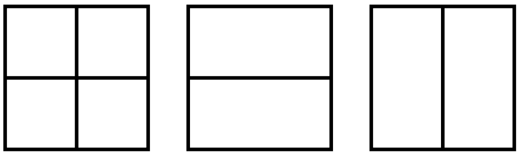

Figure 3: Fourier domain splits corresponding to quad and binary vertices.

Based on the set $\mathcal{T}$, we define a set of functions $\mathcal{F}$ by:

$$\mathcal{F} = \bigcup_{T \in \mathcal{T}} \mathcal{F}_T \tag{12}$$

where

$$\mathcal{F}_T = \{f : f \text{ is piecewise constant on } T\}\,. \tag{13}$$

Given a tree $T \in \mathcal{T}$, and a mother wavelet, we can define a wavelet packet basis B_T, each leaf of the tree corresponding to one subband. Each element of this basis has frequency support that lies approximately in the element of the partition T corresponding to its subband. Due to the piecewise constancy of the functions in $\mathcal{F}_T$, the basis elements are thus approximate eigenfunctions of the operators defined by those functions. Those basis elements whose support lies in R thereby satisfy condition 1 to a certain approximation. Our next task is to complete the set of wavelets inside the region R in order to make a basis for the region and in doing so satisfy condition 2. How we do this depends on the shape of R. We consider two possibilities: dyadic and arbitrarily shaped regions.

4.2 Dyadic Shaped Regions

For the first case, we can use a decimated wavelet packet decomposition to obtain a basis for R. Given a partition T and a function $f \in \mathcal{F}_T$, the distribution for a dyadic region R takes on the form:

$$\Pr(\phi_R|f,T) = \prod_{\alpha} \left[\left(\frac{f_\alpha}{\pi}\right)^{\frac{N_\alpha}{2}} e^{-f_\alpha \sum_{i\in\alpha}(\omega_{\alpha,i}^2 - \mu_\alpha)} \right] \tag{14}$$

where α is the index for the subbands (leaves) of T; f_α is the (constant) value of f on subband α; i is an index for the individual wavelets within each subband; $\omega_{\alpha,i}$ is the $\langle\alpha, i\rangle$ wavelet coefficient of the image ϕ_R; μ_α is the mean of subband α; and N_α is the number of coefficients in subband α.

4.3 Arbitrarily Shaped Regions

For arbitrarily shaped regions, dyadic wavelet packets no longer form a basis. There are two problems. First, the basis elements may not be aligned with the boundary and so include information from outside the region. Second, a shifting of the region with respect to the basis elements will produce a different representation of the same texture.

The problem of dealing with such regions comes into play in areas such as the segmentation of textured scenes, where we do not have control over the shape of the region being analysed. Later, in section 6, we detail two methods for treating this situation.

5 Describing Texture

Within our framework, training a texture model implies finding

- the wavelet packet basis, or equivalently the dyadic partition T, for the texture of interest and
- the values of the means, μ_α, and the inverse variances, f_α, associated with each subband $\alpha \in T$.

Note however that the constancy of the mean μ, and in consequence the region mean μ_R, means that the wavelet packet coefficients of μ_R are all zero except for the scaling coefficient subband, which we denote by $\alpha = 0$. Hence, μ_α, the mean of subband α, is zero unless $\alpha = 0$.

In order to estimate these parameters for a given texture we must examine the posterior probability of the parameters given the training data, d, available for that texture:

$$\Pr(f, \mu, T|d) \propto \Pr(d|f, \mu, T)\Pr(f, \mu, T) \tag{15}$$

The likelihood and prior terms of equation (15) will be discussed in turn.

- **The Likelihood Function:**
  As seen in section 4.2, it is easy to construct wavelet packet bases for dyadic shaped textured images within our proposed probabilistic framework. Patches used for model training are therefore chosen to be of this shape. Hence, the likelihood $\Pr(d|f, \mu, T)$ is given by equation (14), if a single training patch is used.

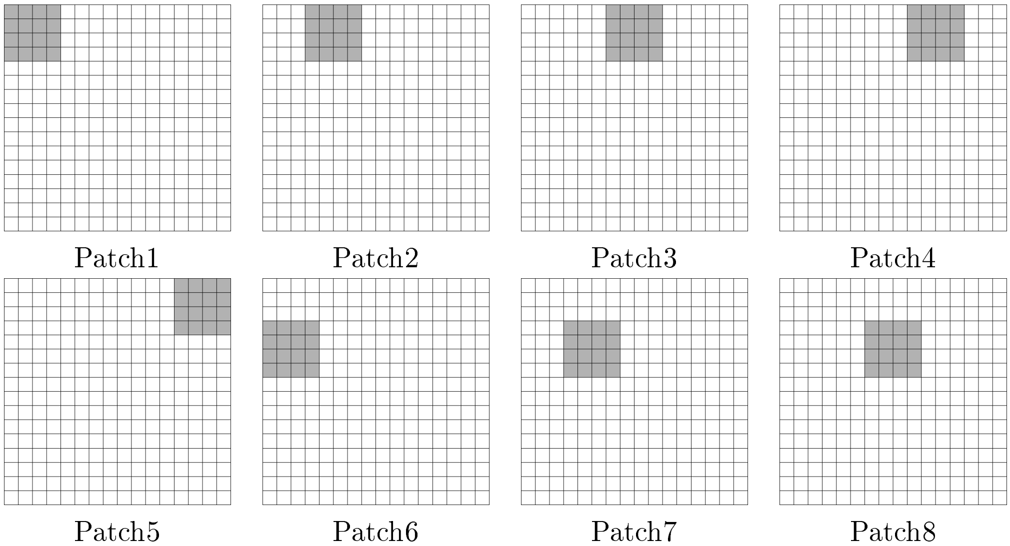

Figure 4: An example of the patch selection procedure. Here eight translated patches are chosen from the original textured image for use in training.

However as we would like our distribution to be translation invariant, we select M patches, $\phi^{(1)}, \ldots, \phi^{(M)}$, which represent shifted versions of the texture. The size of the patches are manually chosen so that the largest periodicity in the texture is present in the patch. Multiple patches are selected from the original texture using an evenly spaced process which starts from the top left corner and proceeds in a left to right, top to bottom manner. Figure 4 shows an example of this patch selection procedure.

The form of our likelihood in this case becomes:

$$\Pr(\phi^{(1)}, \ldots, \phi^{(M)} | f, \mu, T) = \prod_{\alpha} \left[\left(\frac{f_\alpha}{\pi} \right)^{\frac{MN_\alpha}{2}} e^{-f_\alpha \sum_{m=1}^{M} \sum_{i \in \alpha} [\omega_{\alpha,i}(m) - \mu_\alpha]^2} \right] \tag{16}$$

where $\omega_{\alpha,i}(m)$ is the $\langle \alpha, i \rangle$ wavelet coefficient of the m^{th} training patch, and μ_α is the mean of subband α.

- **The Prior Probability:** Our prior on the form of the distribution of f, μ and T can be further reduced to:

$$\Pr(f, \mu, T) = \Pr(f, \mu | T) \Pr(T) \tag{17}$$

The prior on T is given by

$$\Pr(T) = Z^{-1}(\beta)e^{-\beta|T|} \tag{18}$$

where $|T|$ is the number of elements in the decomposition. Such a prior will penalise large decompositions which is desirable given that we would like as sparse a representation of the signal as is possible.

For a fixed decomposition, we assume a Jeffrey's prior for $\Pr(f, \mu|T)$ [14].

5.1 Parameter Estimation

With these distribution choices in mind, we can make invariant maximum *a posteriori* (MAP) estimates [15] for f and μ which for a fixed T yield:

$$\hat{f}_\alpha = \frac{1}{2\bar{e}_\alpha} = \frac{MN_\alpha}{2\sum_{m=1}^{M}\sum_{i\in\alpha}\left[\omega_{\alpha,i}(m) - \mu_\alpha\right]^2} \tag{19}$$

$$\hat{\mu}_\alpha = \frac{2}{MN_\alpha}\sum_{m=1}^{M}\sum_{i\in\alpha}\omega_{\alpha,i}(m) \tag{20}$$

where $\bar{e}_\alpha$ is the average energy per wavelet coefficient in subband α. In order to find the optimal decomposition and its corresponding f_α and μ_α values we use an algorithm which performs a depth-first search through the space $\mathcal{T}$ to find the exact MAP estimates for T, f, and μ.

5.2 Synthetic Data

When training models for textures from the Brodatz album we omitted the mean from the calculations in order to test the texture properties of the models. This was an attempt to make our model luminosity invariant. To account for this in the training phase we subtracted the spatial mean from each of the patches before using them to estimate the optimal values of T and f. Figure 5 contains some optimal decompositions. The reason for not displaying the raw e_α values is because most of these values are zero when scaled and quantized to eight bits. Instead the images show the logarithm of e_α, shifted and rescaled to the range $[0, 255]$. Also shown for comparison are the logs of the magnitudes of the Fourier coefficients.

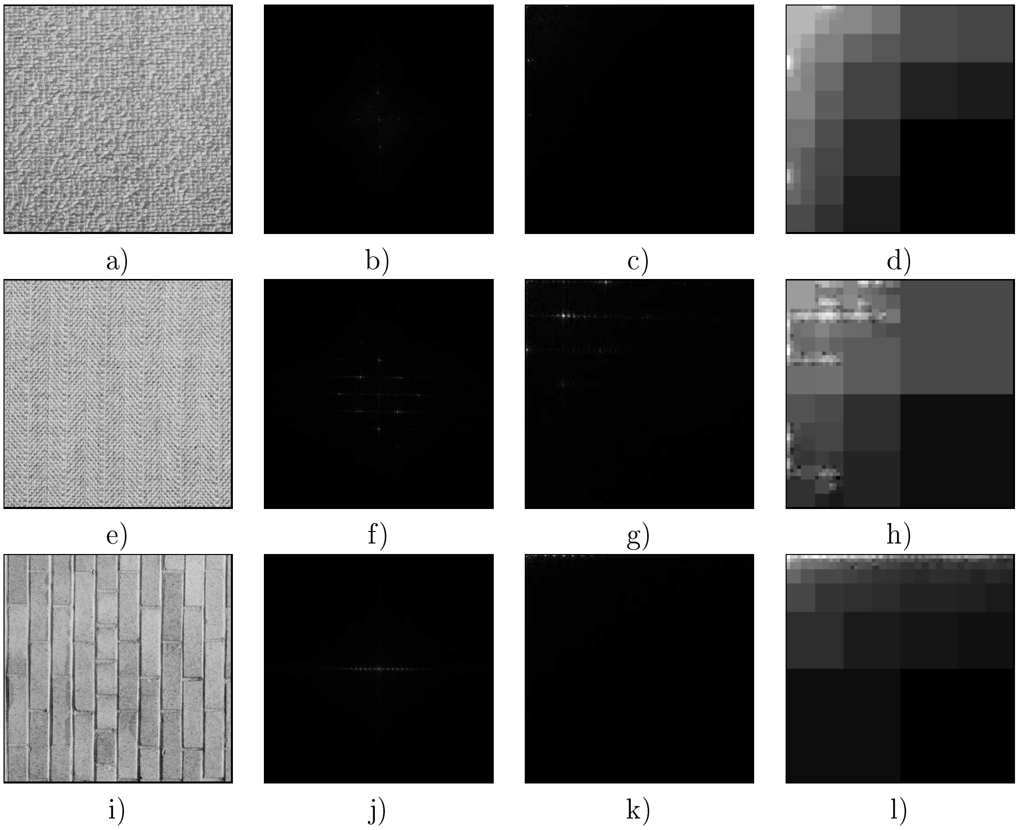

Figure 5: Optimal decompositions for several synthetic textures. Column 1: Original textures. Column 2: The log display of their Fourier magnitude images. Column 3: The log magnitudes of their Fourier coefficients, combined so as to correspond to a wavelet decomposition. Column 4: The log display of their optimal decompositions.

5.3 Real Data

For textures found in high resolution aerial and satellite images we train up the models with the mean included. This is due to the fact that pixel intensity is an extremely important feature in the analysis of remote sensing data. Several textures of this type were trained and their optimal decompositions are shown in figure 6.

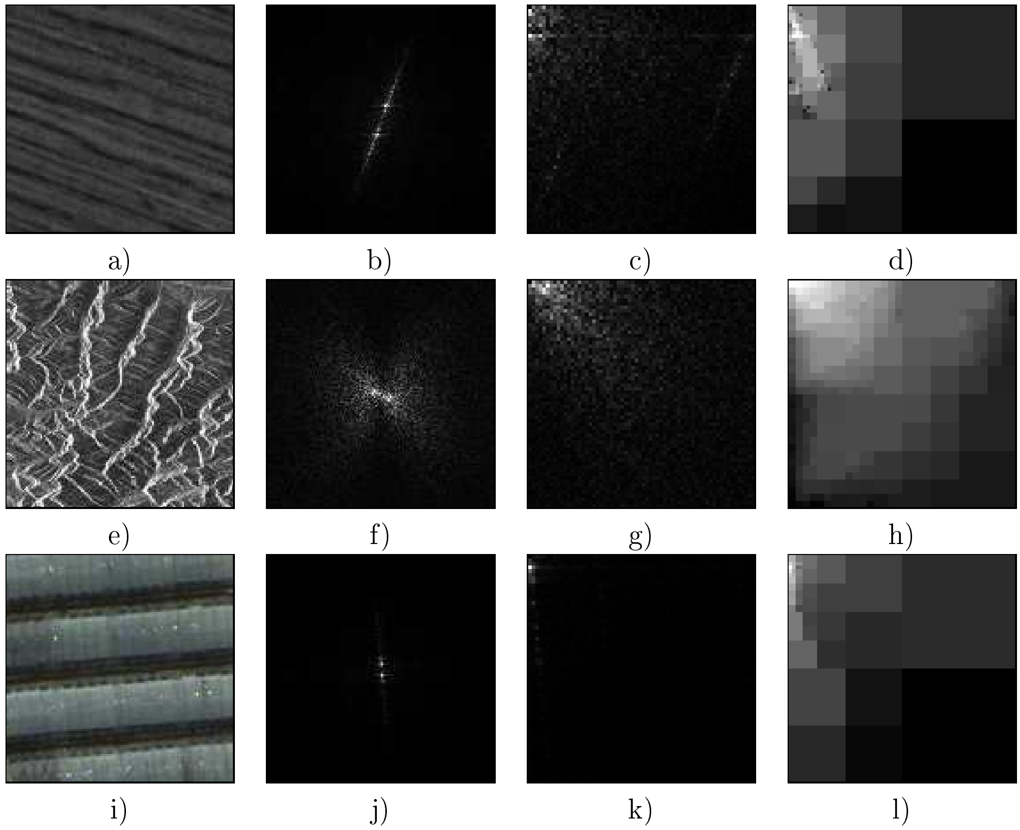

Figure 6: Optimal decompositions for several real textures. Column 1: Original textures. Column 2: The log display of their Fourier magnitude images. Column 3: The log display of their cropped Fourier magnitudes. Column 4: The log display of their optimal decompositions.

6 Application: Segmentation of Textured Scenes

Now that we have developed a framework for adaptively describing texture we would like to examine how our resulting texture models work for the purpose of image segmentation. The task at hand is to segment a given scene into regions of constant texture and to assign to each of these regions a label corresponding to the real world entity it represents. Let us consider a toy example of this general segmentation task. Figure 7 shows the original finite composite image ϕ which we would like to segment [1]. We define a *label set*, denoted by L, to be the finite collection of labels

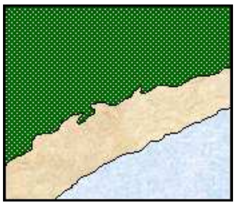

Figure 7: Original image, ϕ, to be segmented into regions of constant texture.

representing the textures present in an image. The three real world entities that gave rise to this particular scene are forest, sand and water. Hence our label set is:

$$L = \{F, S, W\} \tag{21}$$

In classifying an image we define a *class map*, $\lambda \colon D \to L$, which assigns a label to each pixel in the finite image domain D. Three possible class maps of the image ϕ are presented in figure 8. Clearly some class maps represent more accurate results than others. The goal of our analysis is to find the best class map, λ^*, and hence obtain the most accurate classification of our textured scene.

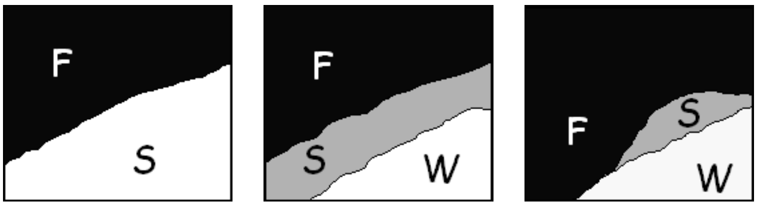

Figure 8: Examples of class maps for the image ϕ.

6.1 A Probabilistic Statement of the Problem

In terms of probabilities, our segmentation problem can be stated as the task of finding the class map λ with the highest probability $\Pr(\lambda \mid \phi, B)$, where $\phi \colon D \to C$

[1]For notational simplicity, we redefine unsubscripted ϕ and λ to refer to the image and class map defined on the finite domain D, instead of as before when they were defined on D_∞.

is the observed image (with the image co-domain $C = \mathbb{R}$ or $\mathbb{Z}$) and B is the set containing any prior information we may have on textures in the scene, which includes the trained values of T and f for each texture class. Hence we need to find:

$$\lambda^* = \arg\max_{\lambda} \Pr(\lambda \mid \phi, B) \tag{22}$$

Using Bayes' Theorem, we can expand this posterior probability to get:

$$\Pr(\lambda \mid \phi, B) = \frac{\Pr(\phi, B \mid \lambda)\Pr(\lambda)}{\Pr(\phi, B)} \tag{23}$$

which, after some manipulation, becomes

$$\Pr(\lambda \mid \phi, B) \propto \Pr(\phi \mid \lambda, B)\Pr(\lambda \mid B) \tag{24}$$

where $\Pr(\phi \mid \lambda, B)$ is the likelihood function and $\Pr(\lambda \mid B)$ is the prior probability of the class map λ.

We make a couple of assumptions before proceeding with our analysis. Firstly, we assume that given a classification, pixel values inside a region with a fixed label do not depend on pixel values outside that region. This first assumption allows us to simplify our likelihood function so that the posterior probability can be expressed as:

$$\Pr(\lambda \mid \phi, B) \propto \left[\prod_{l \in L} \Pr(\phi_{R_l} \mid \lambda, B)\right] \Pr(\lambda \mid B) \tag{25}$$

where $R_l \subset D$ is the set of points whose label is l. This assumption is almost always made in segmentation methods because, while the assumption of independence is certainly incorrect, the dependencies which are present in practice are extremely complex to model.

Our next assumption, which is also usually made in segmentation schemes, states that the probability of the pixel values inside the region labelled l does not depend on the class map outside this region. This simplifies the conditioning in our likelihood function to give:

$$\Pr(\lambda \mid \phi, B) \propto \left[\prod_{l \in L} \Pr(\phi_{R_l} \mid \lambda_{R_l}, B_l)\right] \Pr(\lambda \mid B) \tag{26}$$

where λ_{R_l} is the label map restricted to R_l (which by definition is constant and equal to l) and B_l is our prior information on texture l.

So in order to perform our segmentation task we need:

- for each texture represented in the label set, a probability distribution on an arbitrarily shaped finite region, and
- a prior probability on the class map λ.

In this report we consider the segmentation of two different data types: synthetic Brodatz texture mosaics and remote sensing images. Due to the difference in component textures in these images, we treat the two cases separately. The segmentation methods proposed will be described in turn in sections 6.2 and 6.3, respectively.

6.2 Synthetic Data

Suppose that the region in question, labelled l, is dyadic in shape. If this is the case, we would proceed in a manner analogous to that of section 4.2. That is, we would apply the optimal wavelet packet decomposition of the model for texture l to the region ϕ_{R_l} and, using the resulting coefficients from the decimated subbands, calculate the probability that ϕ_{R_l} belongs to texture class l, *i.e.* $\Pr(\phi_{R_l} \mid \lambda_{R_l}, B_l)$, which can be computed using equation (14). But what happens if the region is arbitrarily shaped? The decimated wavelet packet decomposition described above no longer forms a basis for the region and using it will produce the problems described in section 4.3. One solution would be to use an iterative optimisation method in an attempt to satisfy condition 2 of section 4 exactly. However this would be extremely computationally intensive. Instead, we try to ameliorate the situation by completing the basis on the region using the following approximate scheme. For each subband in the *ideal* decimated decomposition described above, we calculate all its possible shifts within the region. This should overcome the first problem of not including all the information inside the region due to a misalignment of the wavelet basis function with the region boundary. We then approximate the energy in the decimated subband by taking a geometric mean of the energies computed from the shifts. This should give us a more translation invariant measure of the energy which can then be used to compute an approximation of the probability for the region. This procedure is equivalent to creating an undecimated wavelet packet decomposition of the region and incorporating the resulting wavelet coefficients in the following expression:

$$\begin{aligned}
\Pr(\phi_{R_l} \mid \lambda_{R_l}, B_l) &= \prod_{\alpha} \prod_{x \in R_l} \left[\left(\frac{f_\alpha}{\pi} \right)^{\frac{1}{2M_\alpha}} e^{-\frac{f_\alpha}{M_\alpha}(\omega^2_{\alpha,x} - \mu_\alpha)} \right] && (27) \\
&= \prod_{x \in R_l} \prod_{\alpha} \left[\left(\frac{f_\alpha}{\pi} \right)^{\frac{1}{2M_\alpha}} e^{-\frac{f_\alpha}{M_\alpha}(\omega^2_{\alpha,x} - \mu_\alpha)} \right] && (28) \\
&= \prod_{x \in R_l} H_l \left[\{\omega_{\alpha,x}\}_{\alpha \in T_l} \right] && (29)
\end{aligned}$$

where M_α, the *redundancy factor* for subband α, captures the redundancy present in the undecimated representation of the region. Note that this distribution is not the same as that found by pretending that the coefficients in the undecimated wavelet decomposition are independently distributed.

For the prior probability on the class map, we choose a trivial prior which gives each texture class equal probability:

$$\begin{aligned}
\Pr(\lambda \mid B) &= \Pr(\lambda = 1)^{|R_1|} . \Pr(\lambda = 2)^{|R_2|} \ldots \Pr(\lambda = L)^{|R_L|} && (30) \\
&= \left(\frac{1}{|L|} \right)^{|R_1|} . \left(\frac{1}{|L|} \right)^{|R_2|} \cdots \left(\frac{1}{|L|} \right)^{|R_L|} && (31) \\
&= \left(\frac{1}{|L|} \right)^{|R_1| + |R_2| + \ldots + |R_L|} && (32) \\
&= \left(\frac{1}{|L|} \right)^{|R|} && (33) \\
&= \prod_{x \in R} \frac{1}{|L|} && (34)
\end{aligned}$$

From the form of equation (27), it is clear that by choosing such a trivial prior we can perform a pixelwise Maximum Likelihood (ML) classification of the image:

$$\Pr(\lambda \mid \phi, B) \propto \prod_{x \in R} \frac{1}{|L|} H_l \left[\{\omega_{\alpha,x}\}_{\alpha \in T_l} \right] \qquad (35)$$

In practice, we know that λ is likely to be somewhat regular. One option, of course, is to define a Potts prior and use simulated annealing to make a MAP estimate of λ. However, as is well known, simulated annealing is very slow, and it turns out that another approach produces results that are as good, if not better, while being considerably more computationally efficient. We use the following classification rule:

$$\lambda^*(x) = \arg \max_{l \in L} \prod_{x' \in V(x)} H_l \left[\{\omega_{\alpha,x'}\}_{\alpha \in T} \right] \tag{36}$$

where $V(x)$ is the set of neighbours of pixel x, which includes x itself. Examples of neighbourhoods can be seen in figure 9. This rule has a similar effect to the Potts prior, but it still allows a pixelwise classification because it uses the data at the neighbours of a pixel but not their unknown classes. In consequence, one can use larger neighbourhoods with little extra penalty.

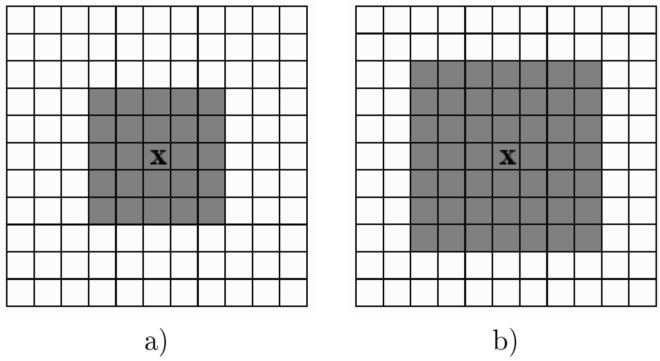

Figure 9: a) Neighbourhood-5 and b) Neighbourhood-7 schemes for pixel x

This segmentation method was tested on several 512×512 Brodatz texture mosaics. The results can be seen in figure 10. The misclassification percentages were 2.6%, 1.7%, 2.5%, and 2.4%, respectively.

We have tested the performance of our method for synthetic data against another supervised texture segmentation scheme developed by Gimel'farb [11]. This approach, based on a conditional Gibbs model, describes the spatial structure of a grayscale texture mosaic via the pixelwise and pairwise pixel interactions in a fixed neighbourhood. These interactions appear in the form of Gibbs potential functions in the exponent of the conditional probability distribution of the region map, given its corresponding grayscale mosaic. The model captures not only the texture properties of a single homogeneous textured region, but also the inter-region structure present in the mosaic.

When presented with a new test mosaic to segment, a likelihood maximization algorithm (Controllable Simulated Annealing) is used to minimise the probabilistic

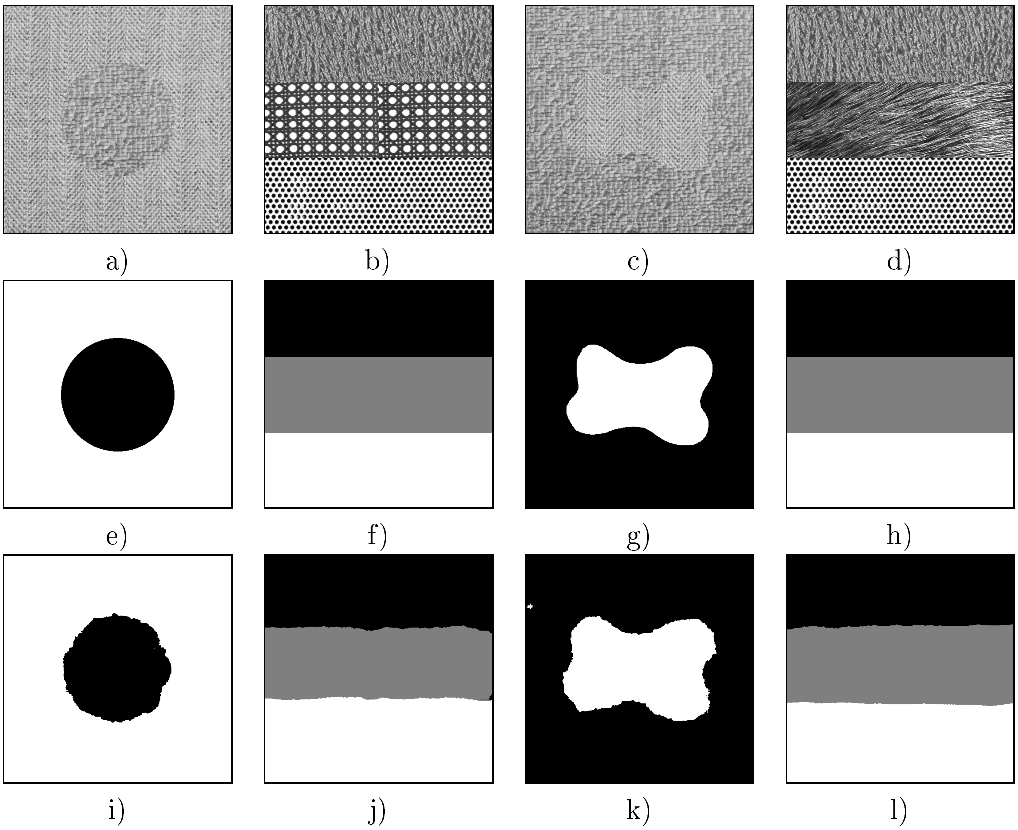

Figure 10: Row 1: Original mosaics: a) Circular mosaic of Herring and Raffia; b) Rectangular mosaic of Calf, D102, and Hexholes152; c) Freehand mosaic of Herring and Raffia; d) Rectangular mosaic of Calf, Fabric0004, and Hexholes152. Row 2: Their ground truth images. Row 3: Segmentation results.

distance between the first-order and second-order statistics for the training data (mosaic and region map) and the mosaic to be segmented. Some comparison results can be seen in figure 11.

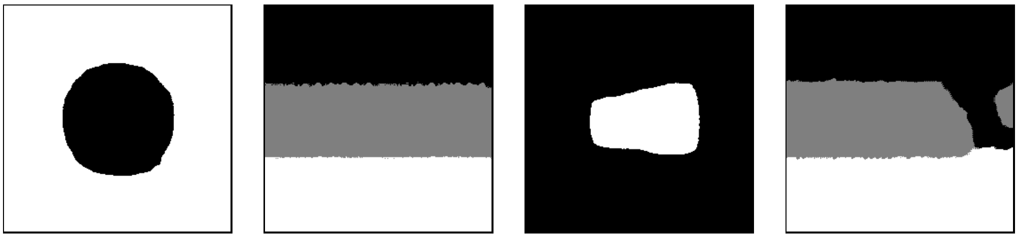

a) b) c) d)

Figure 11: Comparison results from the method presented in [11].

It should be noted that the method of [11] was tested on scaled down (by a factor of two) versions of the textures used in our experiments. The training mosaics used can be seen in figure 12. Results for this method could probably be improved if full scale images are used.

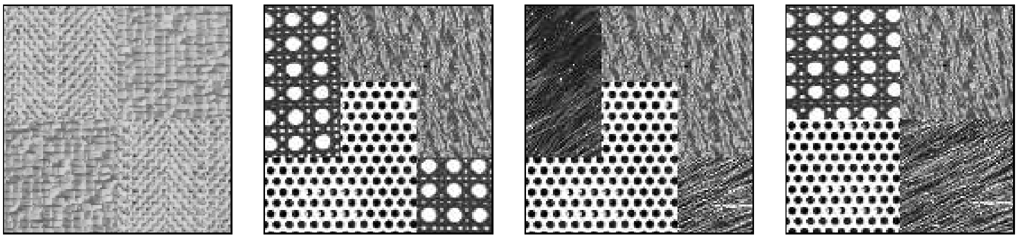

a) b) c) d)

Figure 12: Training mosaics used in the method presented in [11].

However, even with the results on the scaled down test images, one can see where the two methods differ. Firstly, the presence of a texture which varies in brightness and/or directionality poses a problem for the method of [11]. The fourth mosaic of figure 10 has a centre strip containing the texture Fabric0004. We can see where this texture changes toward the right hand edge of the mosaic. It is apparent from the segmentation results of figure 11 that the algorithm had difficulty with this particular area of the mosaic. The training procedure of our method, which uses a collection of sample patches from each texture, allows us to avoid such problems at the segmentation phase.

Secondly, the method of [11] relies heavily on the geometric structure of the training mosaics which consist of rectangular regions. This structure created particular problems for the third test mosaic. However, we can also see how the inclusion of structure in the training process has perhaps enhanced the segmentation of the rectangular shaped test mosaics. As our method trains each texture separately, there is no inter-region structure present in our models. This might explain the lack of precision at the boundaries in our results on the rectangular mosaics, but it also allows us the freedom to segment mosaics with all kinds of inter-region structure present, such as the freehand and circular test mosaics presented above.

6.3 Real Data

We tested the classification method described in section 6.2 on scenes cropped from remote sensing images. It did not produce very satisfactory results and so we investigated the use of a simple and heuristic classification method which is described below. This method has one major advantage: it is faster than the iterative optimization algorithms that would be necessary if condition 2 were to be satisfied.

First note that each texture model has a largest effective filter size, corresponding to the smallest elements in the Fourier domain partition. This filter size is dyadic. Given a set of texture models then, there is a largest effective filter size among them all. To classify pixel x in the image domain, we examine an image patch centred on x of size equal to this largest filter size among all models. We can calculate the probability of the image restricted to this dyadic patch for each trained texture model using equation (14). Assuming that the prior on the class map λ is uniform on L, these data probabilities can be normalized to give the posterior probabilities for the class of the patch. The MAP estimate of the class of the patch is then given, tautologically, by the class with maximum posterior probability. Rather than assign this class to every pixel in the patch however, we heuristically assign this class only to the pixel x at the centre of the patch. We then repeat this procedure for every pixel in the image.

The danger of such a pointwise procedure is that the class map λ can be irregular. However, as in section 6.2, we use the heuristic procedure that has a similar effect to a regularizing prior: it smooths the class map and shortens boundaries. This is done by combining the posterior probability distributions on L at each pixel, as in equation (36).

We carried out a number of experiments using the method described above. The test images were cropped from real aerial and satellite optical and SAR images. Examples of the results obtained are shown in figure 13. Kappa values [6] for these segmentation results are given in table 1. These values are based on the manually created ground truth images presented in the second row of figure 13. The classes used to generate each ground truth image are shown in table 2.

	Figure 13(a)	**Figure 13(b)**	**Figure 13(c)**	**Figure 13(d)**
Nbd-1	0.81	0.70	0.74	0.78
Nbd-7	0.81	0.68	0.75	0.78

Table 1: Kappa values for the segmentation results presented in figure 13.

A more insightful look into the nature of the segmentations can be gained from examining the entropy $S(\pi_x)$ of the posterior probability distribution:

$$\pi_x = \frac{1}{|L|} H_l \left[\{\omega_{\alpha,x}\}_{\alpha \in T_l} \right] \tag{37}$$

In figure 14, the first row shows the entropy which lies in the interval $[\,0,\ |L|\,]$, rescaled to $[\,0,\ 255\,]$. The second row of this figure shows an entropy-weighted misclassification map. In these images, the value of a pixel is given by:

$$e^{-S(\pi_x)}[2\delta(\lambda(x), \lambda^*(x)) - 1] \tag{38}$$

where λ is the algorithm output, and λ^* is the manually created ground truth (note that e^S is analogous to the number of classes that were confused by the algorithm). Thus values greater than 128 represent 'correct', with certainty increasing with increasing intensity, while values less than 128 represent 'incorrect', with certainty increasing with decreasing intensity. These images show that the method is uncertain

	Classes (in order of increasing grey level)
Figure 13(a)	unploughed field, scrubby field, ploughed field, field border
Figure 13(b)	flood, trees
Figure 13(c)	mountains, plain
Figure 13(d)	scrub, ploughed field, woodland, unploughed field, field border

Table 2: Classes defined for the ground truth images in figure 13.

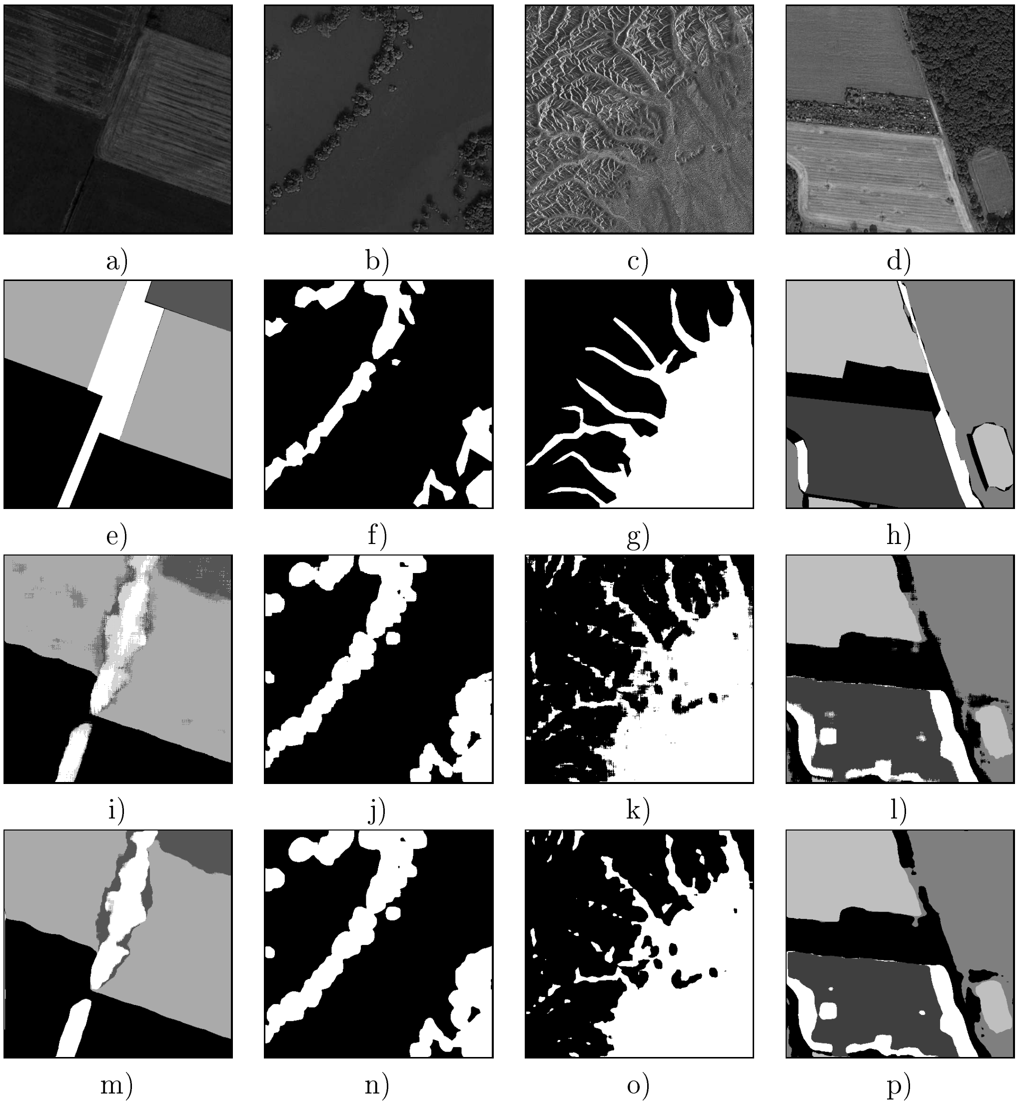

Figure 13: Original images: a) Aerial, 1m resolution ©IGN, b) IKONOS satellite, 1m resolution ©Space Imaging, c) SAR, d) IKONOS satellite, 1m resolution ©Space Imaging; Manually created ground truth images: e) – h); Neighbourhood-1 classifications: i) – l); Neighbourhood-7 classifications: m) – p).

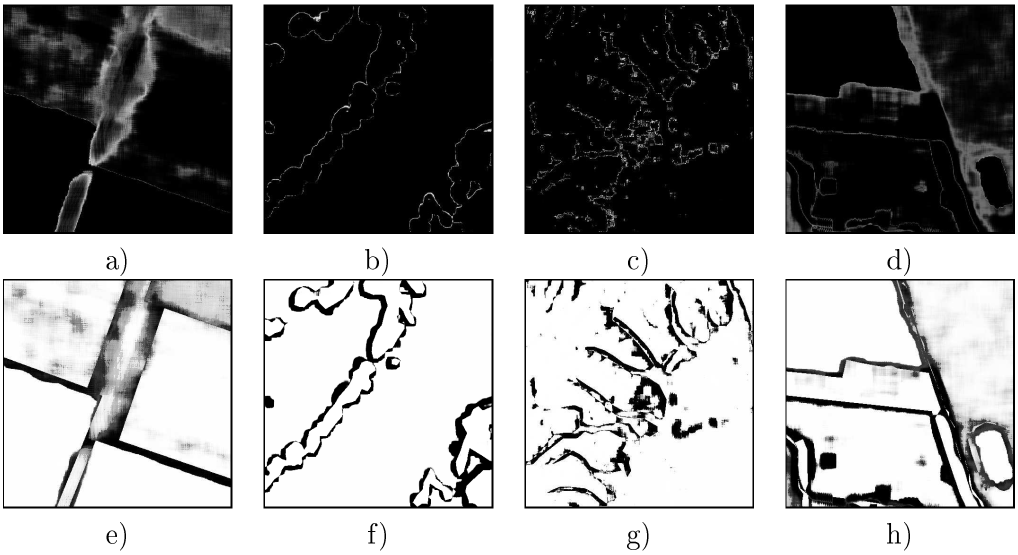

Figure 14: Entropy images for the neighbourhood-1 classifications of figure 13: q) – t); Corresponding entropy-weighted misclassification maps: u) – x).

about its segmentation primarily at the boundaries between entity regions, as would be expected. In the interior of entity regions, the entropy is generally extremely low.

7 Interpretation

The above models could equally well be expressed in terms of a standard wavelet basis, and considering what this means throws light on what the models achieve. Although the inverse covariance operator of a model is (approximately) diagonal in the adapted wavelet packet basis chosen by the above procedure, it will not in general be diagonal in the standard wavelet basis. Mathematically, this is clear. Intuitively, the reason is the following. Many textures possess periodic correlations between pixels at intermediate frequencies that extend over many wavelengths [8]. The standard wavelet basis, however, links the notions of scale and frequency, so that the range of correlations is of the order of the inverse of the frequency of the wavelet (i.e. one wavelength). Correlations between pixels over a range corresponding to many wavelengths can thus only be captured in a standard wavelet basis by introducing interactions between coefficients. This is typically achieved through the use of tree models, in which coefficients interact via their parents, grand-parents, and so on.

In contrast, the adaptive bases presented here capture these long-range correlations using wavelet packet bases that can represent correlations between pixels over many wavelengths while remaining independent, with no need for interactions between the coefficients themselves. For Gaussian models, this type of description is always possible, since the inverse covariance can always be diagonalized, at least in principle. The problem is that the basis elements may not have compact spatial support, which renders them unsuitable for segmentation. The above procedure can thus be thought of as finding the basis that diagonalizes an inverse covariance that must simultaneously be learnt, subject to probabilistic restrictions on its form that ensure the more or less compact spatial support of the basis elements.

8 Discussion on the Choice of Model

In this report we have proposed an adaptive probabilistic framework for dealing with the problem of texture description. The development of this framework was our main focus, and so we did not address the issue of model choice and simply used a Gaussian distribution to model texture.

However, an analysis of the wavelet coefficients of the resulting adaptive packet decompositions points towards an inaccuracy in our initial model assumption. Examples of subband histograms are shown in figure 15 where we can see three distinct types of statistics emerging.

Previous work in the area of texture analysis has found that marginal wavelet statistics can be modelled accurately by a Generalized Gaussian distribution [18]. The first row in our figure shows coefficients which follow this rule. These coefficients belong to high frequency subbands and the form of the distribution can easily be attributed to the fact the coefficients are a response to the edge content in the texture. Those coefficients which lie in or around zero represent the homogeneous area in the texture and some noise, while the deviations are coefficients representing the edges in the textures. The second row shows histograms from subbands in broad intermediate to low frequency ranges which follow a Gaussian distribution. Finally, in the third row, we can see a new bimodal form. Histograms of this shape arise from the subbands which adapted to the periodic content of the texture, and represent much narrower intermediate to low frequency ranges.

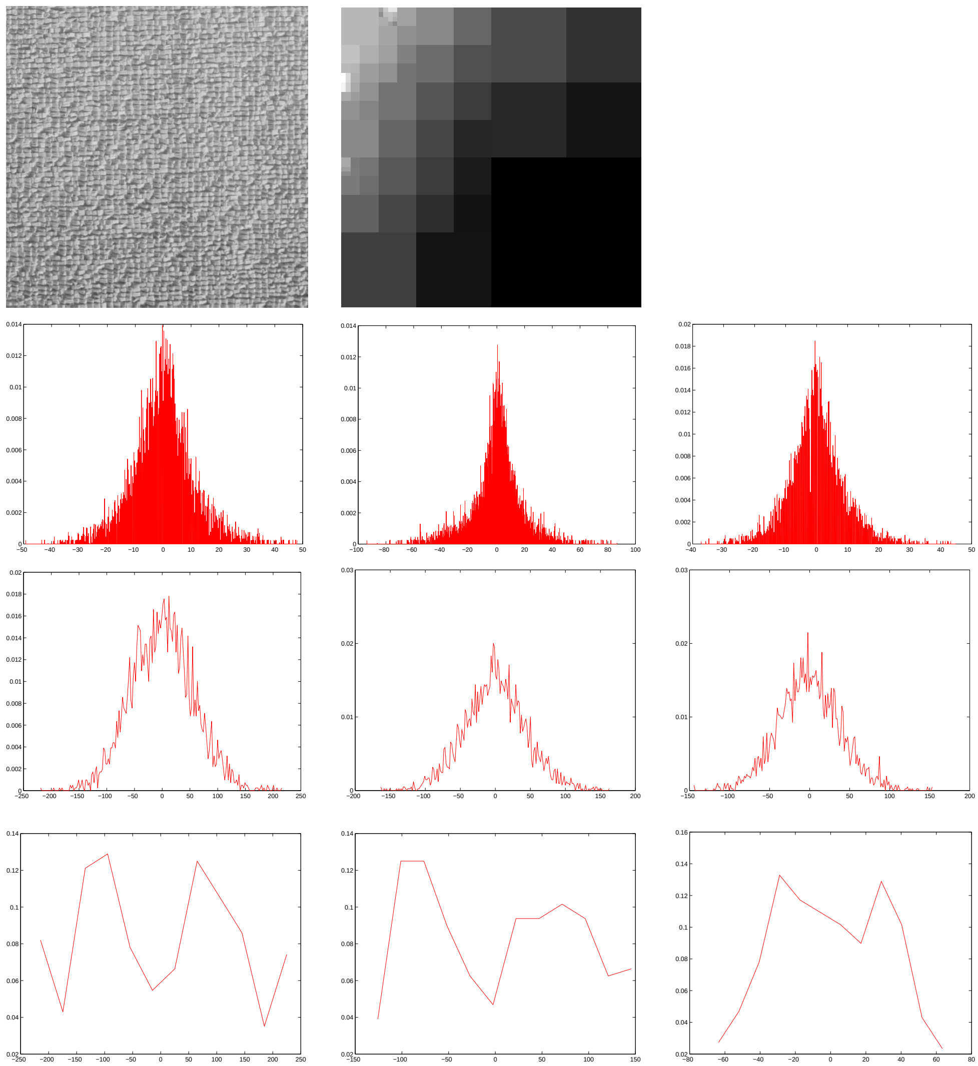

Figure 15: Subband histograms from an adaptive wavelet packet decomposition. Row1: The original Raffia texture patch and its corresponding optimal adaptive decomposition. Row2: Subband histograms which follow a Generalised Gaussian distribution. Row3: Subband histograms which follow a Gaussian distribution. Row4: Subband histograms which have a bimodal form.

In an attempt to model this new form, we have chosen to use the following quartic model:

$$\Pr(\phi|\cdot) = \frac{1}{Z(\mu_\alpha, \sigma_\alpha)} \exp \sum_{\alpha} \sum_{i \in \alpha} \left[\frac{-(\omega_{\alpha,i}^2 - \mu_\alpha)^2}{2{\sigma_\alpha}^2} \right] \tag{39}$$

Under the assumption that μ_α is large relative to σ_α, an approximation was made for the partition function $Z(\mu_\alpha, \sigma_\alpha)$. The resulting Maximum Likelihood estimates for the parameters are given by:

$$\hat{\mu}_\alpha = \sqrt{\frac{\sum_{i \in \alpha} \omega_{\alpha,i}^4}{N_\alpha}} \tag{40}$$

and

$$\hat{\sigma}_\alpha = \sqrt{\frac{2}{N_\alpha} \left(\sum_{i \in \alpha} \omega_{\alpha,i}^2 - \sum_{i \in \alpha} \omega_{\alpha,i} \; \hat{\mu}_\alpha \right)} \tag{41}$$

where N_α is, as before, the number of pixels in the subband α.

Figure 16 shows some type three histograms plotted against a quartic distribution. Several of these fit the distribution quite well. We have also fitted a Gaussian distribution to this data to highlight where our model may be going wrong.

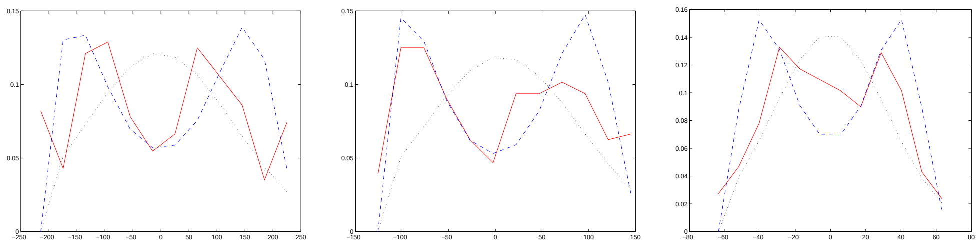

Figure 16: Subband histograms from the adaptive wavelet packet decompositions. Gaussian - black (dotted line); Quartic - blue (dashed line); Data histogram - red (solid line);

From this analysis of the coefficients of the adaptive wavelet packet decompositions we can see that choosing a distribution is a very important issue. Taking a Generalized Gaussian as a starting point allows us to cover both type one and type two histograms, but the type three histograms need to be described separately. All this indicates that there should be some form of model selection process at each of the nodes in the tree, perhaps with the prior probability of a model depending on the frequency range concerned.

9 Conclusions and Future Work

Texture is an important attribute of the appearance of many semantic entities of interest in a wide variety of images. In this paper, we have described a new adaptive probabilistic model for texture description and applied it to the supervised segmentation of both synthetic and real texture data. Wavelet packet bases, which arise naturally within our probabilistic framework, allow the model to adapt to an individual texture class, and in doing so to capture its structure, and in particular the long-range correlations that create its principal periodicities. The method works well on a variety of textures, despite its simplicity. Two outstanding issues to be resolved are the estimation of the parameter β in the prior on the set of trees, and the choice of a neighbourhood size. In addition, there are two sets of issues that point to future work. First, in some cases the models are too simple, and second, the segmentation method is too rough. We discuss these points in turn.

As noted in Section 8, the presence of a narrow band of frequencies that runs throughout a texture means that the statistics of the wavelet packet coefficients with support in that band are no longer clustered around zero (although they have mean zero). This is in contrast to the statistics reported for standard wavelet coefficients [19]. Gaussian models for such coefficients are thus not accurate for modelling the statistics of adaptive bases.

The models described in this paper are not rotation invariant, although in principle this is just as much of a requirement as translation invariance. One can impose rotation invariance on a Gaussian distribution, but the results are not that interesting since the inverse covariance must now be a function only of the magnitude of the frequency. Non-trivial rotation invariance can be achieved by mixtures over rotations or by otherwise increasing the complexity of the model.

Finally, translation-invariant Gaussian models have no way of taking into account the relative phase of Fourier components, and yet phase is critical for texture structure.

All three of these modelling issues will be treated in future work by extending the Gaussian model through the addition of quartic terms to the exponent, as discussed in section 8.

The segmentation method used above also possesses a drawback. For small neighbourhoods the results are noisy, while for larger neighbourhoods the smoothing effect distorts the shape and size of entity regions. Future work will focus on the satisfaction of condition 2, and the use of iterative optimization algorithms with regularizing priors to improve the accuracy of the segmentations.

Acknowledgments

This work was partially supported by European Union project MOUMIR, HP-99-108 (`www.moumir.org`). The authors would like to thank Dr. Georgy Gimel'farb for his interesting discussions and for the images he provided in figures 11 and 12. The authors would also like to thank the French Mapping Agency, IGN (`www.ign.fr`) for the donation of some of the remote sensing images used in this paper, and Space Imaging (`www.spaceimaging.com`) for permission to use others.

A Deriving a Probability Measure on a Finite Region

This appendix details the calculations needed to arrive at a probability measure for a texture on an arbitrarily-shaped finite region.

In order to model texture in an accurate manner we must capture the notion of its infinite extendability by describing it via a distribution on Φ, the space of infinite images. Without loss of generality, we assume this distribution to be Gaussian with zero mean and so the probability measure takes on the following form:

$$\Pr(\phi|\cdot) = |F|^{\frac{1}{2}} e^{-\langle\phi|F|\phi\rangle} D\phi \tag{42}$$

where $\phi \in \Phi$; F is the inverse variance-covariance operator whose determinant is given by $|F|$; and $D\phi$ is the underlying measure on the space Φ and represents the infinite product $\prod d\phi(x)$ over all points $x \in D_\infty$, the infinite image domain.

In practice, when we analyse and segment textures, we work with images which are defined on a finite region, $R \subset D_\infty$. This means that, starting from the distribution on the space of infinite images, we must derive a probability measure on R. This requires us to marginalize equation (42) over all points in $\phi_{\bar{R}}$, the image defined on $\bar{R}$, the complement of the region:

$$\Pr(\phi_R|\cdot) = \int_{\phi_{\bar{R}}} \Pr(\phi|\cdot) \tag{43}$$

To ease this marginalisation process we express our distribution on Φ in terms of Φ_R, the space of images defined on the region, and $\Phi_{\bar{R}}$, the space of images defined on the complement of the region.

A.1 Expressing the Energy in Terms of ϕ_R and $\phi_{\bar{R}}$

The energy of the distribution on the space Φ is defined to be:

$$E(\phi) = \langle\phi|F|\phi\rangle \tag{44}$$

In section 2 of this report, we introduced two projections, P_R and $P_{\bar{R}}$, which form an orthogonal decomposition of Φ. With this in mind we can rewrite the energy as:

$$E(\phi) = \langle P_R\phi + P_{\bar{R}}\phi|F|P_R\phi + P_{\bar{R}}\phi\rangle \tag{45}$$

which when expanded yields:

$$E(\phi) = \langle P_R\phi|F|P_R\phi\rangle + \langle P_R\phi|F|P_{\bar{R}}\phi\rangle + \langle P_{\bar{R}}\phi|F|P_R\phi\rangle + \langle P_{\bar{R}}\phi|F|P_{\bar{R}}\phi\rangle \tag{46}$$

The projections, P_R and $P_{\bar{R}}$, are defined according to surjective and injective maps on Φ. In terms of these maps the energy is given by:

$$\begin{aligned}
E(\phi) &= \langle i_R \pi_R \phi | F | i_R \pi_R \phi \rangle + \langle i_R \pi_R \phi | F | i_{\bar{R}} \pi_{\bar{R}} \phi \rangle + \\
&\quad \langle i_{\bar{R}} \pi_{\bar{R}} \phi | F | i_R \pi_R \phi \rangle + \langle i_{\bar{R}} \pi_{\bar{R}} \phi | F | i_{\bar{R}} \pi_{\bar{R}} \phi \rangle \\
&= \langle i_R \pi_R \phi | F i_R \pi_R \phi \rangle + \langle i_R \pi_R \phi | F i_{\bar{R}} \pi_{\bar{R}} \phi \rangle + \\
&\quad \langle i_{\bar{R}} \pi_{\bar{R}} \phi | F i_R \pi_R \phi \rangle + \langle i_{\bar{R}} \pi_{\bar{R}} \phi | F i_{\bar{R}} \pi_{\bar{R}} \phi \rangle \\
&= \langle \pi_R \phi | i_R^\dagger F i_R \pi_R \phi \rangle + \langle \pi_R \phi | i_R^\dagger F i_{\bar{R}} \pi_{\bar{R}} \phi \rangle + \\
&\quad \langle \pi_{\bar{R}} \phi | i_{\bar{R}}^\dagger F i_R \pi_R \phi \rangle + \langle \pi_{\bar{R}} \phi | i_{\bar{R}}^\dagger F i_{\bar{R}} \pi_{\bar{R}} \phi \rangle
\end{aligned} \tag{47}$$

where $i_R^\dagger$ and $i_{\bar{R}}^\dagger$ are the adjoints of i_R and $i_{\bar{R}}$. As $i_R^\dagger = \pi_R$ and $i_{\bar{R}}^\dagger = \pi_{\bar{R}}$ the energy becomes:

$$\begin{aligned}
E(\phi) &= \langle \pi_R \phi | \pi_R F i_R \pi_R \phi \rangle + \langle \pi_R \phi | \pi_R F i_{\bar{R}} \pi_{\bar{R}} \phi \rangle + \\
&\quad \langle \pi_{\bar{R}} \phi | \pi_{\bar{R}} F i_R \pi_R \phi \rangle + \langle \pi_{\bar{R}} \phi | \pi_{\bar{R}} F i_{\bar{R}} \pi_{\bar{R}} \phi \rangle \\
&= \langle \pi_R \phi | \pi_R F i_R | \pi_R \phi \rangle + \langle \pi_R \phi | \pi_R F i_{\bar{R}} | \pi_{\bar{R}} \phi \rangle + \\
&\quad \langle \pi_{\bar{R}} \phi | \pi_{\bar{R}} F i_R | \pi_R \phi \rangle + \langle \pi_{\bar{R}} \phi | \pi_{\bar{R}} F i_{\bar{R}} | \pi_{\bar{R}} \phi \rangle
\end{aligned} \tag{48}$$

Note here that $\pi_R F i_R$, $\pi_R F i_{\bar{R}}$, $\pi_{\bar{R}} F i_R$ and $\pi_{\bar{R}} F i_{\bar{R}}$ in equation (48), are in fact the four operators of section 2 obtained by splitting F up according to the orthogonal decomposition formed by the projections P_R and $P_{\bar{R}}$. These operators therefore enter into energy as follows:

$$\begin{aligned}
E(\phi) &= \langle \pi_R \phi | F_{RR} | \pi_R \phi \rangle + \langle \pi_R \phi | F_{R\bar{R}} | \pi_{\bar{R}} \phi \rangle + \\
&\quad \langle \pi_{\bar{R}} \phi | F_{\bar{R}R} | \pi_R \phi \rangle + \langle \pi_{\bar{R}} \phi | F_{\bar{R}\bar{R}} | \pi_{\bar{R}} \phi \rangle
\end{aligned} \tag{49}$$

Using the definition of the surjective maps π_R and $\pi_{\bar{R}}$ we can further reduce this to:

$$E(\phi) = \langle \phi_R | F_{RR} | \phi_R \rangle + \langle \phi_R | F_{R\bar{R}} | \phi_{\bar{R}} \rangle + \langle \phi_{\bar{R}} | F_{\bar{R}R} | \phi_R \rangle + \langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} \rangle \tag{50}$$

As our goal is to obtain an expression for the energy in terms of ϕ_R and $\phi_{\bar{R}}$, we must manipulate the expression in equation (50) to separate out completely the action on R from the action on $\bar{R}$. To do this we must complete the square in $\phi_{\bar{R}}$.

A.1.1 Completing the Square in $\phi_{\bar{R}}$

As it stands our expression for the energy can be viewed as containing the following three terms in $\phi_{\bar{R}}$:

$$\begin{aligned}
\textit{Quadratic Term} &: \langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} \rangle \\
\textit{Linear Term} &: \langle \phi_R | F_{R\bar{R}} | \phi_{\bar{R}} \rangle + \langle \phi_{\bar{R}} | F_{\bar{R}R} | \phi_R \rangle \\
\textit{Constant Term} &: \langle \phi_R | F_{RR} | \phi_R \rangle
\end{aligned}$$

To complete the square in $\phi_{\bar{R}}$ we take the following steps:

- Preserve the quadratic term $\langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} \rangle$
- Introduce a linear term $\langle \phi_{\bar{R}} + F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} + F_{\bar{R}R}\phi_R \rangle$
- Cancel the extra $F_{\bar{R}\bar{R}}$ term $\langle \phi_{\bar{R}} + (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} + (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R \rangle$

Expanding this last expression yields:

$$\begin{aligned}
E(\phi) &= \langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} \rangle + \\
&\quad \langle (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} \rangle + \\
&\quad \langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} | (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R \rangle + \\
&\quad \langle (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} | (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R \rangle \\
&= \langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} \rangle + \\
&\quad \langle (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} \phi_{\bar{R}} \rangle + \\
&\quad \langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R \rangle + \\
&\quad \langle (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R \rangle \\
&= \langle \phi_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} \rangle + \\
&\quad \langle (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} \phi_{\bar{R}} \rangle + \\
&\quad \langle \phi_{\bar{R}} | F_{\bar{R}R} | \phi_R \rangle + \\
&\quad \langle (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}R}\phi_R \rangle \qquad (51)
\end{aligned}$$

Examining the second term of equation (51) in detail gives:

$$
\begin{aligned}
\langle {F_{\bar{R}\bar{R}}}^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}}\phi_{\bar{R}}\rangle &= \langle \pi_{\bar{R}} F {i_{\bar{R}}}^{-1} \pi_{\bar{R}} F i_R \phi_R | \pi_{\bar{R}} F i_{\bar{R}} \phi_{\bar{R}}\rangle \\
&= \langle \pi_{\bar{R}} F i_R \phi_R | i^{\dagger}_{\bar{R}} F {\pi^{\dagger}_{\bar{R}}}^{-1} \pi_{\bar{R}} F i_{\bar{R}} \phi_{\bar{R}}\rangle \\
&= \langle \pi_{\bar{R}} F i_R \phi_R | \pi_{\bar{R}} F {i_{\bar{R}}}^{-1} \pi_{\bar{R}} F i_{\bar{R}} \phi_{\bar{R}}\rangle \\
&= \langle \pi_{\bar{R}} F i_R \phi_R | \phi_{\bar{R}}\rangle \\
&= \langle F i_R \phi_R | \pi^{\dagger}_{\bar{R}} | \phi_{\bar{R}}\rangle \\
&= \langle i_R \phi_R | F \pi^{\dagger}_{\bar{R}} | \phi_{\bar{R}}\rangle \\
&= \langle \phi_R | i^{\dagger}_R F \pi^{\dagger}_{\bar{R}} | \phi_{\bar{R}}\rangle \\
&= \langle \phi_R | \pi_R F i_{\bar{R}} | \phi_{\bar{R}}\rangle \\
&= \langle \phi_R | F_{R\bar{R}} | \phi_{\bar{R}}\rangle \qquad (52)
\end{aligned}
$$

which is exactly the second term in equation (50).

Hence, the extra piece of the energy puzzle is the final term of equation (51) which when expanded becomes:

$$
\begin{aligned}
\langle F_{\bar{R}\bar{R}}{}^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}R}\phi_R \rangle &= \langle \pi_{\bar{R}} F i_{\bar{R}}{}^{-1} \pi_{\bar{R}} F i_R \phi_R | \pi_{\bar{R}} F i_R \phi_R \rangle \\
&= \langle \pi_{\bar{R}} F i_R \phi_R | i^{\dagger}_{\bar{R}} F \pi^{\dagger}_{\bar{R}}{}^{-1} \pi_{\bar{R}} F i_R \phi_R \rangle \\
&= \langle \pi_{\bar{R}} F i_R \phi_R | \pi_{\bar{R}} F i_{\bar{R}})^{-1} \pi_{\bar{R}} F i_R \phi_R \rangle \\
&= \langle \phi_R | i^{\dagger}_R F \pi^{\dagger}_{\bar{R}} \pi_{\bar{R}} F i_{\bar{R}}{}^{-1} \pi_{\bar{R}} F i_R | \phi_R \rangle \\
&= \langle \phi_R | \pi_R F i_{\bar{R}} \pi_{\bar{R}} F i_{\bar{R}}{}^{-1} \pi_{\bar{R}} F i_R | \phi_R \rangle \\
&= \langle \phi_R | F_{R\bar{R}} F_{\bar{R}\bar{R}}{}^{-1} F_{\bar{R}R} | \phi_R \rangle \qquad (53)
\end{aligned}
$$

We can therefore complete the square in the following manner:

$$
\begin{aligned}
E(\phi) = \; & \langle \phi_{\bar{R}} + F_{\bar{R}\bar{R}}{}^{-1} F_{\bar{R}R}\phi_R | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} + (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R \rangle - \\
& \langle \phi_R | F_{R\bar{R}} (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R} | \phi_R \rangle + \langle \phi_R | F_{RR} | \phi_R \rangle \qquad (54)
\end{aligned}
$$

Letting

$$
(F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R}\phi_R = \mu_{\bar{R}} \qquad (55)
$$

allows us to finally write our energy as desired, in terms of ϕ_R and $\phi_{\bar{R}}$, separately:

$$
E(\phi) = \langle \phi_{\bar{R}} + \mu_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} + \mu_{\bar{R}} \rangle + \langle \phi_R | F_{RR} - F_{R\bar{R}} (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R} | \phi_R \rangle \qquad (56)
$$

A.2 Marginalising over Elements in $\phi_{\bar{R}}$

Now that we have an expression for the energy of our distribution in terms of ϕ_R and $\phi_{\bar{R}}$ separately, we can write our distribution on the space of infinite images as:

$$
\Pr(\phi|\cdot) = |F|^{\frac{1}{2}} e^{-\langle \phi_{\bar{R}} + \mu_{\bar{R}} | F_{\bar{R}\bar{R}} | \phi_{\bar{R}} + \mu_{\bar{R}} \rangle - \langle \phi_R | F_{RR} - F_{R\bar{R}} (F_{\bar{R}\bar{R}})^{-1} F_{\bar{R}R} | \phi_R \rangle} D\phi_R D\phi_{\bar{R}} \qquad (57)
$$

To obtain a measure on the region R one must marginalize over all points in $\phi_{\bar{R}}$ as follows:

$$\begin{aligned}\Pr(\phi_R|\cdot) &= \left[\int_{-\infty}^{+\infty} e^{-\langle \phi_{\bar{R}}+\mu_{\bar{R}}|F_{\bar{R}\bar{R}}|\phi_{\bar{R}}+\mu_{\bar{R}}\rangle} D\phi_{\bar{R}}\right] * \\ &\quad \left[|F|^{\frac{1}{2}} e^{-\langle \phi_R|F_{RR}-F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1}F_{\bar{R}R}|\phi_R\rangle} D\phi_R\right] \end{aligned} \tag{58}$$

Make the following change of variable:

$$\phi'_{\bar{R}} = \phi_{\bar{R}} + \mu_{\bar{R}} \tag{59}$$

to get:

$$\begin{aligned}\Pr(\phi_R|\cdot) &= \left[\int_{-\infty}^{+\infty} e^{-\langle \phi'_{\bar{R}}|F_{\bar{R}\bar{R}}|\phi'_{\bar{R}}\rangle} D\phi'_{\bar{R}}\right] * \\ &\quad \left[|F|^{\frac{1}{2}} e^{-\langle \phi_R|F_{RR}-F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1}F_{\bar{R}R}|\phi_R\rangle} D\phi_R\right] \end{aligned} \tag{60}$$

Here we note that:

$$\left[\int_{-\infty}^{+\infty} e^{-\langle \phi'_{\bar{R}}|F_{\bar{R}\bar{R}}|\phi'_{\bar{R}}\rangle} D\phi'_{\bar{R}}\right] \tag{61}$$

is a non-normalized Gaussian distribution. Hence the correct normalization constant is given by $|F_{\bar{R}\bar{R}}|^{-\frac{1}{2}}$ and so our probability measure on the region becomes:

$$\Pr(\phi_R|\cdot) = \frac{|F|^{\frac{1}{2}}}{|F_{\bar{R}\bar{R}}|^{\frac{1}{2}}} e^{-\langle \phi_R|F_{RR}-F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1}F_{\bar{R}R}|\phi_R\rangle} D\phi_R \tag{62}$$

For notational simplicity we shall denote the operator on the region R by G_R:

$$G_R = F_{RR} - F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1}F_{\bar{R}R} \tag{63}$$

So the probability distribution for images defined on the finite region R is given by:

$$\boxed{\Pr(\phi_R|\cdot) = Z^{-1} e^{-\langle \phi_R|G_R|\phi_R\rangle} D\phi_R}$$

B Diagonalising the Operator on the Region

Diagonalising the operator on the region requires us to find a set of functions on the region

$$B = \{|a\rangle : a \in A\} \tag{64}$$

where A is an index set, such that this set of functions satisfies the two conditions outlined in section 4. The details of this Diagonalisation procedure are given below. Expressing the distribution on Φ_R in terms of the elements of the set B we have that:

$$\begin{aligned}
\Pr(\phi_R|\cdot) &= Z^{-1}e^{-\sum_{a,a' \in A} \langle\phi_R|a'\rangle\langle a'|G_R|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a,a' \in A} \langle\phi_R|a'\rangle\langle a'|F_{RR}-F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1}F_{\bar{R}R}|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a,a' \in A} \langle\phi_R|a'\rangle\langle a'|F_{RR}-F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1}\pi_{\bar{R}}Fi_R|a\rangle\langle a|\phi_R\rangle}
\end{aligned} \tag{65}$$

Applying the injective map i_R to each element in B yields a set of functions defined on Φ:

$$\{i_R|a\rangle : a \in A\} \tag{66}$$

By definition of the injective map, each function in this set is equal to $|a\rangle$ in R and to zero in $\bar{R}$. Condition 1 states that each such function should be an eigenfunction of the operator F (with eigenvalue f_a). This implies that the support of $Fi_R|a\rangle = f_a i_R|a\rangle$ lies in the region R. Applying the surjective map $\pi_{\bar{R}}$ to this therefore sets the second term in the operator G_R to zero:

$$\begin{aligned}
\Pr(\phi_R|\cdot) &= Z^{-1}e^{-\sum_{a,a' \in A} \langle\phi_R|a'\rangle\langle a'|F_{RR}-F_{R\bar{R}}(F_{\bar{R}\bar{R}})^{-1}\pi_{\bar{R}}f_a i_R|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a,a' \in} \langle\phi_R|a'\rangle\langle a'|F_{RR}-0|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a,a' \in A} \langle\phi_R|a'\rangle\langle a'|\pi_R Fi_R|a\rangle\langle a|\phi_R\rangle}
\end{aligned} \tag{67}$$

As $\pi_R i_R = id$, applying the surjective map π_R to $f_a i_R$ results in the first term of the operator G_R becoming $f_a|a\rangle$.

$$\begin{aligned}
\Pr(\phi_R|\cdot) &= Z^{-1}e^{-\sum_{a,a'\in A}\langle\phi_R|a'\rangle\langle a'|\pi_R F i_R|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a,a'\in A}\langle\phi_R|a'\rangle\langle a'|\pi_R f_a i_R|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a,a'\in A}\langle\phi_R|a'\rangle\langle a'|f_a \pi_R i_R|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a,a'\in A}\langle\phi_R|a'\rangle\langle a'|f_a|a\rangle\langle a|\phi_R\rangle} \qquad (68)
\end{aligned}$$

This last calculation shows that the elements of the set B are also eigenfunctions of the operator F_{RR}.
The second condition, which states that the set B should form an orthonormal basis for function on the region R, means that G_R is diagonalised by B, allowing us to write our measure as:

$$\begin{aligned}
\Pr(\phi_R|\cdot) &= Z^{-1}e^{-\sum_{a,a'\in A}\langle\phi_R|a'\rangle f_a\langle a'|a\rangle\langle a|\phi_R\rangle} \\
&= Z^{-1}e^{-\sum_{a\in A} f_a\langle\phi_R|a\rangle\langle a|\phi_R\rangle} \qquad (69)
\end{aligned}$$

C Translation Invariance for a Gaussian Distribution

As part of our model development we require that our distribution on Φ for a given texture be translation invariant. This means that images such as those in figure 17 will have equal probability of belonging to a given texture class.

a) b)

Figure 17: a) ploughed field texture patch, b) a shifted version of the patch.

However this requirement of translation invariance and our choice of a Gaussian distribution on Φ makes the operator F of our distribution diagonal in the Fourier basis. This can be seen in what follows.

C.1 Translation Operator

We define our translation operator T_δ on the spatial domain as:

$$[T_\delta \phi](x) = \phi(x+\delta) \text{ , for all } x \in D_\infty \tag{70}$$

The Fourier transform of $[T_\delta\phi](x)$ is given by:

$$\begin{aligned}
[\hat{T_\delta\phi}](k) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{-ikx}[T_\delta\phi](x)dx \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{-ikx}\phi(x+\delta)dx \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{-ik(x'-\delta)}\phi(x')dx' \\
&= e^{ik\delta}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{-ikx'}\phi(x')dx' \\
&= e^{ik\delta}\hat{\phi}(k) \qquad (71)
\end{aligned}$$

C.2 Requiring Translation Invariance

Saying that we require our distribution to be translation invariant means that we want:

$$\Pr(\phi|\cdot) = \Pr(T_\delta\phi|\cdot) \qquad (72)$$

Expressing these two probabilities in the Fourier basis yields the requirement that:

$$Z^{-1}e^{-\int_k\int_{k'}\hat{\phi}^*(k)F(k,k')\hat{\phi}(k')dkdk'} = Z^{-1}e^{-\int_k\int_{k'}[\hat{T_\delta\phi}]^*(k)F(k,k')[\hat{T_\delta\phi}](k')dkdk'} \qquad (73)$$

Using equation (71) we can further develop the right hand side of this expression to give the requirement that:

$$\begin{aligned}
Z^{-1}e^{-\int_k\int_{k'}\hat{\phi}^*(k)F(k,k')\hat{\phi}(k')dkdk'} &= Z^{-1}e^{-\int_k\int_{k'}e^{-ik\delta}\hat{\phi}^*(k)F(k,k')e^{ik'\delta}\hat{\phi}(k')dkdk'} \\
&= Z^{-1}e^{-\int_k\int_{k'}e^{i(k'-k)\delta}\hat{\phi}^*(k)F(k,k')\hat{\phi}(k')dkdk'} \qquad (74)
\end{aligned}$$

The only way this can be true is if the exponents are equal, namely if:

$$\begin{aligned}
\int_k \int_{k'} \hat{\phi}^*(k) F(k,k') \hat{\phi}(k') dk dk' &= \int_k \int_{k'} e^{i(k'-k)\delta} \hat{\phi}^*(k) F(k,k') \hat{\phi}(k') dk dk' \\
\Rightarrow e^{i(k'-k)\delta} &= 1 \\
\Rightarrow i(k'-k)\delta &= 0, \ \forall \delta \\
\Rightarrow k' &= k
\end{aligned} \tag{75}$$

So the condition of translation invariance on our Gaussian distribution causes it to have the following form in the Fourier basis:

$$\begin{aligned}
\Pr(\phi \mid \cdot) &= Z^{-1} e^{-\int_k \int_{k'} \hat{\phi}^*(k) F(k,k') \hat{\phi}(k') dk dk'} \\
&= Z^{-1} e^{-\int_k \int_{k'} \hat{\phi}^*(k) f(k) \delta(k,k') \hat{\phi}(k') dk dk'} \\
&= Z^{-1} e^{-\int_k \hat{\phi}^*(k) f(k) \hat{\phi}(k) dk} \\
&= Z^{-1} e^{-\int_k f(k) [\hat{\phi}(k)]^2 dk} \\
&= Z^{-1} e^{-\int_k f(k) E(k) dk}
\end{aligned} \tag{76}$$

Hence our distribution is characterised by a function, f, on the Fourier domain.

References

[1] M. Acharyya and M. K. Kundu. Adaptive basis selection for multi texture segmentation by M-band wavelet packet frames. In *Proc. IEEE International Conference on Image Processing*, pages 622–625, Thessaloniki, Greece, 2001.

[2] P. Brodatz. *Textures - A Photographic Album for Artists and Designers.* Dover, 1966.

[3] T. Chang and C.C.J. Kuo. Texture analysis and classification with tree-structured wavelet transform. *IEEE Transactions on Image Processing*, 2(4):429–441, 1993.

[4] H. Choi and R. Baraniuk. Multiscale image segmentation using wavelet-domain hidden markov models. *IEEE Transactions on Image Processing*, 10(9):1309–1321, 2001.

[5] R. R. Coifman and M. V. Wickerhauser. Entropy-based methods for best basis selection. *IEEE Transactions on Information Theory*, 38(2):713–718, 1992.

[6] R.G. Congalton. A review of assessing the accuracy of classifications of remotely sensed data. *Remote Sensing of Environment*, 37:35–46, 1991.

[7] R.W. Conners, C.W. McMillin, K. Lin, and R.E. Vasquez-Espinosa. Identifying and locating surface defects in wood: Part of an automated lumber processing system. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-5:573–583, 1983.

[8] G. Van der Wouwer. *Wavelets for Texture Analysis.* PhD thesis, University of Antwerp, May 1998.

[9] P. Dewaele, P. Van Gool, and A. Oosterlinck. Texture inspection with self-adaptive convolution filters. In *Proc. International Conference on Pattern Recognition*, pages 56–60, Rome, Italy, November 14-17 1988.

[10] J. M. Francos, A. Z. Meiri, and B. Porat. A unified texture model based on a 2-d Wold like decomposition. *IEEE Transactions on Signal Processing*, 41:2665–2678, 1993.

[11] G. Gimel'farb. Supervised texture segmentation by maximising conditional likelihood. In *Proc. Energy Minimization Methods in Computer Vision and Pattern Recogntion*, Sophia Antipolis, France, September 2001. Springer. LNCS 2134.

[12] R. Gupta and P. Undrill. The use of texture analysis to delineate suspicious masses in mammography. *Physics in Medicine and Biology*, 40:835–855, 1995.

[13] R.M. Haralick. Statistical and structural approaches to texture. *Proceedings of the IEEE*, 67:786–804, 1979.

[14] H. Jeffreys. *Theory of Probability*. Clarendon Press, Oxford, 1939.

[15] I. Jermyn. On Bayesian estimation in manifolds. Inria Research Report 4607, 2002. (www-sop.inria.fr/ariana/personnel/Ian.Jermyn/).

[16] A. Laine and J. Fan. Texture classification by wavelet packet signatures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1186–1190, 1993.

[17] S. Livens, P. Scheunders, G. Van de Wouwer, and D. Van Dyck. Wavelets for texture analysis, an overview. In *Proc. IEE International Conference on Image Processing and its Applications*, pages 581–585, Dublin, Ireland, 1997.

[18] S. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-11(7):674–693, July 1989.

[19] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.

[20] O. Pichler, A. Teuner, and B.J. Hosticka. A comparison of texture feature extraction using adaptive Gabor filtering, pyramidal and tree structured wavelet transforms. *Pattern Recognition*, 5(29):7333–742, 1996.

[21] T. Randen and J. H. Husoy. Filtering for texture classification: A comparative study. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):291–310, 1999.

[22] T.R. Reed and J.M. Hans du Buf. A review of recent texture segmentation and feature extraction techniques. *Computer Vision, Graphics and Image Processing: Image Understanding*, 57(3):359–372, 1993.

[23] M. Tuceryan and A. K. Jain. *Texture Analysis*, chapter 2.1. World Scientific, 1998.

[24] M. Vetterli and J. Kovacevic. *Wavelets and Subband Coding*. Prentice Hall, 1995.

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)
Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399