Texture Analysis of SAR Images: A Comparative Study

Anne H. Schistad Solberg[1] *and Anil K. Jain*[2]

[1] Norwegian Computing Center, P.O.Box 114 Blindern, 0314 Oslo, Norway
Phone +47 22 85 25 00 Fax +47 22 69 76 60
e-mail: Anne.Solberg@nr.no

[2] Department of Computer Science, Michigan State University, East Lansing, Michigan 48824
e-mail: jain@cps.msu.edu

Abstract

Texture provides useful information in the interpretation of SAR images. The SAR signature of a spatial region is subject to changes depending on the weather conditions. Several published studies have indicated that a classification of SAR images based on texture is more robust than a classification based on gray values alone. In this paper, we compare the performance of a number of different texture features. These features are derived from the gray-level co-occurrence matrix, local statistics, lognormal random field models, and fractal models. The motivations behind this comparative study are twofold: (i) to investigate the robustness of texture features for SAR image classification; and (ii) to combine texture features derived from different models to determine if texture fusion improves the classification accuracy. The invariance of the texture features with respect to changes in the mean backscatter values is investigated on a data set of thirteen ERS-1 SAR images of the same scene, captured under different conditions. The effect of speckle filtering on the discrimination ability of the texture features is examined. The relative classification accuracy using different texture models is compared. Finally, feature selection is applied to the pooled set of texture features to select the optimal feature combinations for several images. Three main conclusions can be drawn from this study: (i) the texture features are not invariant with respect to natural changes in the mean backscatter values; (ii) speckle filtering improves the discrimination ability of texture features derived from the co-occurrence matrix and from a lognormal random field model; and (iii) texture fusion and selection by combining texture features obtained by different models significantly improve the classification accuracy.

1 Introduction

Synthetic aperture radar (SAR) imaging provides an all-weather coverage of a large number of different phenomena (e.g., area usage, snow coverage, ice monitoring in polar regions, and oil spill detection in the ocean). In the interpretation of SAR images, texture provides important information, in addition to image gray levels or the backscatter values alone. Several studies have shown that classification based on texture features can improve the accuracy of the interpretation [2, 3, 38, 49]. The SAR signature of an area (spatial region) is dependent on the weather conditions. Factors like temperature, soil moisture content, and wind conditions for water areas highly influence the SAR signatures. The same area may have different backscatter values when imaged repeatedly. For operational use (automatic scene interpretation) of SAR imagery for monitoring a phenomenon, this dependency is a complicating factor. Some studies [3, 42] have indicated that classification based on texture might be more robust than classification based on gray values alone. Based on these initial results, it is important to study multiple SAR images of the same scene carefully to investigate the reliability of the texture features.

SAR images are degraded by speckle noise due to interference between individual scatterers within each resolution cell. The presence of speckle noise presents a challenge to most texture models. It is common to first filter the SAR images to reduce speckle. The filtering process, however, will also affect the degree of textural information present in the image. Thus, the effect of speckle filtering on the discrimination ability of different texture features needs to be analyzed.

A large number of approaches for computation of image texture has been proposed in the literature. Some studies have compared the performance of different methods [3, 14, 52] for SAR image analysis, but we are not aware of any study which has combined texture features derived from different models and then applied feature selection to find the optimal feature combination for a given classification task.

In this paper, we investigate the performance of texture features derived from the gray-level co-occurrence matrix [21], features based on local statistics, fractal features [15, 29], and features derived from a lognormal random field model [17] for land-use classification of ERS-1 SAR images and agricultural classification of MAESTRO SAR images. The invariance of the texture features with respect to natural changes in backscatter values within each pattern class is studied using a set of 13 ERS-1 SAR images of the same scene. The effect of speckle filtering and image transformations prior to the computation of the texture features is investigated. The discrimination ability of the different methods for texture computation is examined and compared. Feature selection methodology and discriminant analysis are applied to find the optimal combination of texture features. All experiments are conducted on several SAR images to be able to provide generalizations of the results.

The remainder of the paper is organized as follows. Section 2 reviews previous approaches to computation of texture features in general and, in particular, texture features suited for SAR images. In section 3 the methods used for texture computation and evaluation are described. The experimental data and results are given in section 4. Finally, a discussion of the results and conclusions are provided in section 5.

2 Related work

Although no precise definition of texture exists, certain intuitive concepts of texture can be defined. Texture is a property of a local region in an image; the texture of a point is undefined. Texture involves the spatial distribution of gray levels. In the description of texture, the following qualitative properties

are important [30]: uniformity, density, coarseness, roughness, regularity, linearity, directionality, frequency, and phase. Because there are so many qualitative properties which may describe the texture of a region, a generic texture model applicable to all types of images has evaded us.

Methods for texture analysis and modelling are often grouped into the following four categories: statistical methods, geometrical methods, model-based methods, and signal processing methods [48]. Statistical methods are based on various models of the spatial distribution of gray levels in an image region. The most widely used method is based on gray-level co-occurrence matrices which describe how the gray levels in a spatial neighborhood are related [21]. Other statistical measures of texture include measures of first-order and higher-order statistics and autocorrelation features [32, 43, 46].

Geometrical methods are characterized by the fact that texture is assumed to be composed of "texture elements" or primitives. Given these primitives, the texture features can be derived from their statistical properties [1, 47], or by extracting features based on syntactic or structural methods [5, 44, 56]. Structural methods involve two major steps: extraction of the texture elements and inference of their placement rules.

The basic concept in model-based methods is to formulate a generative model that can be used not only to describe a texture, but also to synthesize it. This category of methods include Markov random field (MRF) models [8, 11] and fractal models [34, 40]. MRF models can be used for texture synthesis, classification, segmentation, and image restoration.

Signal processing methods decompose an input image into component images which capture different frequency and orientation information. The filtering of the input image can be done either in the frequency domain or in the spatial domain. Simple texture properties like "edgeness" can be computed by applying edge operators to the input image [30, 33]. Recent approaches have also been based on a bank of Gabor filters with different frequencies and orientations [27].

For texture analysis of SAR images, the most common approach is to use features derived from the gray-level co-occurrence matrix [36, 37, 38, 42, 49]. Among model-based approaches to analyze texture in SAR images is the well-known model proposed by Ulaby et al., where the observed backscatter value is assumed to be a product of three factors: the mean intensity of a field, a variable which represents the within-field natural variations due to texture, and a speckle variable [49]. In this scheme, the degree of texture of a region can be found from the deviation of backscatter values from pure speckle statistics. Markov random fields and auto-regressive models have also been used to model SAR texture [13, 17, 28]. Fractal models are also used to analyze SAR image texture [4, 6, 35]. Bourissou [6] found that the fractal dimension was constant for a scene imaged using three different polarizations. Examples of the use of texture primitives can be found in [3, 36].

Several studies have compared the performance of different texture extraction methods. Barber et. al [3] compared structural, statistical, and frequency-based methods for discrimination of sea-ice types. They achieved the best performance with statistical features derived from the gray-level co-occurrence matrix (GLCM) and the gray-level dependence matrix (GLDM). They found discrimination of sea-ice types based on GLCM and GLDM features to be more robust than discrimination based on pixel gray levels alone. This observation was previously made by Shokr [42]. Murni et. al [36] compared the segmentation performance of texture features based on the concept of texture units [51], features based on the co-occurrence matrix, and features from local statistics.

Du [14] compared the discrimination ability of GLCM features and Gabor filtering [7] and found that while the co-occurrence features preserved the finer texture detail, the Gabor method preserved the overall locational distributions of ice types and was computationally faster. Weszka et al. [52], in their comparison of first-order statistics, second-order co-occurrence probabilities and frequency

domain analysis, found that the spatial-domain methods provided more precise classification than the frequency-domain methods.

A few studies have investigated the effect of adaptive filtering on texture classification. Barber and LeDrew [2] found no significant difference in classifier performance when processing a filtered or unfiltered SAR image. In a previous small-scale study, we found a slight increase in the classification accuracy when texture was computed after adaptive filtering with a 3x3 window [41].

The studies reviewed above were all performed on a limited data set, consisting of one or two images. None of these studies has employed feature selection to find the optimal subset of texture features from a combined or pooled set of features derived from different models, although Murni et. al [36] achieved good performance by manually selecting a combination of two features from each of three categories (co-occurrence features, features based on a concept of texture units, and features based on local statistics). In this paper, we will compare the performance of various approaches to texture computation on a larger data set. Feature selection and discriminant analysis will be conducted to see if the accuracy of the interpretation of SAR images can be increased. This follows the methodology used by You and Jain [55] who compared the performance of various features for shape matching.

In this comparative study, we include four different texture feature extraction methods: features computed from local statistics and GLCM matrices, fractal features, and texture modelling based on a lognormal random field model. Features based on local statistics define texture in SAR images as the deviation from fully developed speckle. GLCM features have been selected based on their popularity and performance in a number of prior texture studies. Fractal features are included because fractal models often are believed to be adequate models of natural scenes. Lognormal random fields have performed well for texture synthesis and classification in SAR images [17]. We have chosen not to include frequency-based methods in this study, because SAR images of the land-use classes considered here do not contain periodic or orientation-specific information. In our preliminary experiments, we found that the SAR signatures of some of the land-use classes studied here gave very similar power spectra. The estimated spectra were also sensitive to within-class variations when they were computed in small windows. In addition, several comparative studies reported inferior performance of frequency-based methods compared to statistical methods [14, 52].

For land-use classification in Norway, the average size of typical regions is rather small. Thus, we need to limit the window size used in the texture computations. In this study, a window of size 9×9 pixels will be used. Shork [42] compared the effect of varying the window sizes from 5×5 pixels to 9×9 pixels and found no significant differences in GLCM texture features. However, if the window size is increased to, e.g., 30×30 pixels or larger, fractal features and/or frequency-based features might have better performance.

3 Texture features and classification methodology

A brief description of the selected texture features is given below. The classification methodology and the principles of feature transformation in terms of supervised discriminant analysis are then described. The input to the texture models will be an $N \times N$ image $[f(x, y)]$ where $f(x, y)$ denotes the gray level of pixel (x, y).

3.1 Features derived from gray-level co-occurrence matrices

The co-occurrence texture features are based on gray-level spatial dependencies [21]. A co-occurrence matrix, computed in a local window, contains the relative frequencies of all pairwise combinations of backscatter values at a certain distance d and direction α within the window. From this matrix, a number of features can be computed. With N_g gray levels in the image, the dimension of the matrix will be $N_g \times N_g$. The (i,j)th element, P_{ij}, of this matrix is defined by

$$P_{ij} = \frac{p_{ij}^{d,\alpha}}{\sum_{i,j}^{L} p_{ij}^{d,\alpha}}, \tag{1}$$

where $p_{ij}^{d,\alpha}$ is the frequency of occurrence of gray levels i and j, separated by a distance d and direction $\alpha = 0^o, 45^o, 90^o$ and 135^o. The summation is over the total number of pixel pairs L, given d, in the window.

We compute the following texture features from the co-occurrence matrix: *angular second moment (ASM), contrast (CONT), entropy (ENT), cluster shade (CLSH), inertia (INER)* , and *inverse difference moment (IDM)*. These features are among the most commonly used co-occurrence features.

$$ASM = \sum_{i}^{N_g} \sum_{j}^{N_g} P_{ij}^2 \tag{2}$$

$$CLSH = \sum_{i}^{N_g} \sum_{j}^{N_g} (i + j - \mu_i - \mu_j)^3 P_{ij} \tag{3}$$

$$CONT = \sum_{n=0}^{N_g - 1} n^2 \{ \sum_{i}^{N_g} \sum_{j}^{N_g} P_{ij} \} \tag{4}$$

$$ENT = \sum_{i}^{N_g} \sum_{j}^{N_g} P_{ij} \log P_{ij} \tag{5}$$

$$INER = \sum_{i}^{N_g} \sum_{j}^{N_g} (i - j)^2 P_{ij} \tag{6}$$

$$IDM = \sum_{i}^{N_g} \sum_{j}^{N_g} \frac{1}{1 + (i - j)^2} P_{ij} \tag{7}$$

In the above equations, $n = |i - j|$, $\mu_i = \sum_i i \sum_j P_{ij}$ and $\mu_j = \sum_j j \sum_i P_{ij}$. Prior to the computation of the co-occurrence matrix, the number of gray levels in the image needs to be reduced to a small number to get reliable estimates of the relative frequencies when texture is computed in a small window (9×9 pixels). In this study, histogram equalization [19] is first applied to the images to produce an image with gray levels which span the entire 8-bit interval of pixel values, followed by a quantization into $N_g = 8$ gray levels.

To check the dependence of distance d on the performance of texture features, the following experiment was conducted. The GLCM texture features listed above were computed for distances ranging from $d = 1, ..., 30$ for a five-class classification problem involving the following five classes: *water, urban areas, forests,* and two types of agricultural areas. The resulting entropy feature values are plotted in Figure 1. In this case, the computed entropy was reasonably constant for different values of d. This was also the case for the other GLCM features. This confirmed the results obtained by Barber and LeDrew [2]. Based on this experiment, we set the value of $d = 1$. Furthermore, the co-occurrence matrices are computed as averages over the four directions $\alpha = 0^o, 45^o, 90^o, 135^o$.

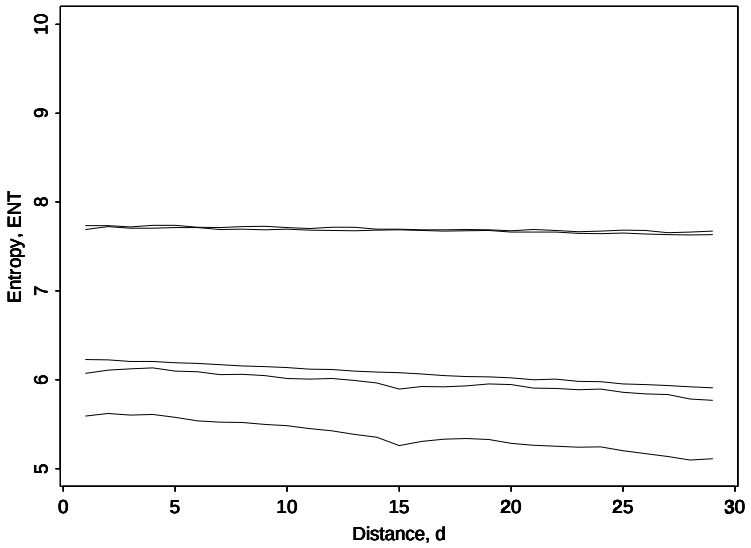

Figure 1: *Entropy values for five different classes as a function of distance d.*

3.2 Features computed from local statistics

Based on local statistics, the following features are computed:

- Power-to-mean ratio, $PMR = \sigma/\mu$, where σ is the local standard deviation and μ is the local mean of the backscatter values.
- Skewness, $SKW = \frac{E[(f(x,y)-\mu)^3]}{\sigma^3}$, where $E[]$ denotes the expectation operator.
- Kurtosis, $KUR = \frac{E[(f(x,y)-\mu)^4]}{\sigma^4}$.
- Contrast [43], $CNT = \frac{1}{W-1}\sum_{i,j\in\mathcal{N}}[(f(x,y) - f(i,j))/\mu]^2$, where W is the number of pixels in a local neighborhood $\mathcal{N}$ around the pixel of interest (x, y) .
- Homogeneity [43], $HOM = \frac{1}{W-1}\sum_{i,j\in\mathcal{N}}\frac{1}{1+[(f(x,y)-f(i,j))/\mu]^2}$.

Subsets of these features have been used for texture analysis in SAR images in several previous studies [20, 43].

3.3 Fractal features

Many natural textured surfaces can be described as a fractal surface. The key parameter describing a fractal surface is its fractal dimension. Several methods for computation of the fractal dimension of a surface exist. The most common methods are the variation method [15], the ϵ-blanket method [39], and the box-counting method [29]. In previous comparative studies [16, 24] of the performance of these three methods, the variation method was found to have the smallest variance, indicating the best stability. Based on this result, we will use the variation method for computation of the fractal dimension.

The variation method is based on the observation that the variation of a fractal surface within small boxes as a function of the box size produces good estimates of the fractal dimension [15]. The method computes the maximum variation, $V(\epsilon)$, of the image surface $f(x, y)$ in a square of side ϵ. The fractal dimension D is estimated as 3 minus the slope of the least square fit of the data $\{\ln(\epsilon), \ln(V(\epsilon))\}$ as ϵ varies.

Several studies have shown that the fractal dimension alone does not capture all the textural properties (see, e.g., [29]). Another measure, called *lacunarity*, is often used in addition to the fractal dimension. The lacunarity value is small when texture is fine and it is large when the texture is coarse. An estimate of lacunarity can be derived from the box counting algorithm [29]. Consider the image to be a set S of points in a three-dimensional space, as defined by the two spatial coordinates of a pixel (x, y) and the associated gray level $f(x, y)$. Following Voss [50], let $P(m, L)$ be the probability that m points will fall within a box of side length L centered at a point.

The lacunarity for a box of size L is defined as:

$$\Lambda(L) = \frac{M^2(L) - (M(L))^2}{(M(L))^2}, \tag{8}$$

where

$$M(L) = \sum_{m=1}^{N} mP(m, L), \tag{9}$$

$$M^2(L) = \sum_{m=1}^{N} m^2 P(m, L), \tag{10}$$

and N is the number of possible points within the box. In our experiments, a 5×5 box ($L = 5$) is used. The amount of time needed to compute the fractal features was an order of magnitude higher than the computation times for the other texture features. Because of this, fractal features were only included in selected experiments, and not in the full-scale study on the large set of test images.

3.4 Parameters of lognormal field models

Following Frankot and Chellappa [17], we will model the SAR image using a multiplicative autoregressive random field (MAR). The parameters of the model will be used as texture descriptors.

Let the observed image $f(x, y)$ be represented by a white-noise-driven multiplicative system, where $g(x, y) = \ln f(x, y)$ follows a Gaussian autoregressive (AR) model

$$g(x, y) = \sum_{r,s \in \mathcal{N}} \theta_{rs}(g(x + r, y + s) - \mu_g) + u(x, y). \tag{11}$$

$\mathcal{N}$ is the neighborhood system, and μ_g is the mean value of the logarithmic image g. The noise process $u(x,y)$ is an uncorrelated white noise with variance σ^2.

We will use the least squares estimates for the parameters θ_{rs}, σ^2, and μ_g [17]. The parameter estimates are computed in windows of size 9×9 and with $\mathcal{N} = \{(0,-1),(-1,-1),(-1,0)\}$.

3.5 Classification procedure

To test the performance of the texture features with respect to their discrimination ability, a classifier must be chosen. We can use either a non-parametric classifier, like the K-Nearest Neighbor classifier, which does not assume any particular statistical distribution of the texture features, or a parametric classifier based on a statistical model. To be able to perform a large number of classification experiments in a reasonable amount of time, we chose a simple quadratic classifier based on the multivariate normal distributions. Finding the best classifier for SAR texture feature vectors is a topic of future research and will not be considered in this paper. Assuming that class k has a probability distribution $p_k(\mathcal{F}(x,y))$, where $\mathcal{F}(x,y)$ is the vector of texture features for pixel (x,y), the classification rule is defined by

$$\text{Allocate pixel } (x,y) \text{ to class } k \text{ if } p_k(\mathcal{F}(x,y)) = \max_{l=1,..,K} p_l(\mathcal{F}(x,y)),$$

where

$$p_l(\mathcal{F}(x,y)) = \frac{1}{|2\pi\Sigma_l|^{\frac{1}{2}}} e^{-\frac{1}{2}(\mathcal{F}(x,y)-\mu_l)^T \Sigma_l^{-1}(\mathcal{F}(x,y)-\mu_l)}. \tag{12}$$

We use the maximum likelihood estimates of the class-conditional mean vectors μ_l and covariance matrices Σ_l, $l = 1,..K$. The error rates reported for the classifiers are computed by the leave-one-out method using all the available ground-truth data.

3.6 Feature transformations by discriminant analysis

Discriminant analysis is a well-known approach for dimensionality reduction and multivariate data projection [18]. The input features are projected into a space having fewer dimensions. The pattern class labels for each pattern are utilized in the projection to maximize the between-class scatter. For a data set involving K classes and d features, the multiple discriminant analysis is based on the following linear transformation

$$\mathcal{G}(x,y) = \mathcal{H}_0 \mathcal{F}(x,y), \tag{13}$$

where the transform matrix $\mathcal{H}_0$, with dimension $(K-1) \times d$, is defined in terms of the within-class ($\mathcal{S}_W$) and the between-class ($\mathcal{S}_B$) scatter matrices. For the definition of the scatter matrices, see, e.g., [26]. This transformation has the important property that the patterns are projected from d dimensions to $K-1$ dimensions while maintaining the scatter ratio to be constant. The rows of $\mathcal{H}_0$ are the eigenvectors corresponding to the $(K-1)$ non-zero eigenvalues of $\mathcal{S}_W^{-1}\mathcal{S}_B$, and the new feature vector $\mathcal{G}(x,y)$ is (K-1) dimensional. To get a projection into a b-dimensional space ($b \leq K-1$), the b eigenvectors corresponding to the b largest eigenvalues are chosen.

3.7 Feature selection

To find the best subset of texture features from the set of d available features, Whitney's algorithm for feature selection will be used [53]. The algorithm works in the following manner. First, the feature with the smallest probability of misclassification is selected. Then the feature which, in combination with the previously selected features, gives the smallest probability of misclassification, is selected, and so on. This feature selection process is suboptimal, because it does not perform an exhaustive evaluation of all the possible feature subsets. Whitney used a K-NN classifier with the leave-one-out method for error estimation. We have moditied this algorithm to use a quadratic classifier and the leave-one-out method.

4 Experimental results

Two test data sets were used to evaluate the performance of the texture features. Evaluation was based on the following two objectives: (i) land-use classification based on ERS-1 SAR images, and (ii) crop classification based on airborne SAR images from the MAESTRO survey [9].

ERS-1 SAR images A data set consisting of 13 ERS-1 SAR images from the fall of 1991 of Kjeller, Norway was used. The images were captured under different weather conditions, so the SAR signature varies greatly from one image to another. The images were taken at different dates, but from exactly the same path of the satellite track. Thus, they covered exactly the same area, and no co-registration or geometric correction was necessary. Figure 2 shows a small part of the scene for all the 13 images. Large variations in SAR signature for the same ground-cover class were observed with different wind conditions (for water areas), temperature, and soil moisture content (for forest and agricultural areas).

A five-class classification problem was considered, consisting of the following common ground cover classes: water, urban areas, forests, and two classes of agricultural fields (plowed and unplowed). The data set was acquired in a project for monitoring soil erosion due to runoff from agricultural fields tilled in the fall. Between the acquisition of the SAR images, a large number of the agricultural fields has been tilled. Ground control samples were available from each acquisition date. The size of the training areas was typically 4,000 pixels per class for each image. For agricultural areas, the training set varied from 300 pixels to 12,000 pixels as the plowing season progressed.

Based on selected areas from the five ground-cover classes, we computed the mean values for each class for each image. Figure 3 shows a plot of the mean backscatter values for each image and each category. For the water class, we can identify high mean values for acquisition dates with windy conditions (image nos. 3, 7, and 12). Of the five classes, the plot for the urban areas shows the smallest variation, and the plowed agricultural class had the largest variation. Plowed areas can be easily separated from the unplowed areas when the soil moisture content is high and the soil is not frozen. Such conditions occurred for image nos. 3, 7, 10, 11, and 12. On the acquisition dates for image nos. 8 and 9, the soil was frozen.

MAESTRO SAR images A 4-look C-band multipolarization SAR image from the MAESTRO survey [9] from August 16, 1989, of Feltwell, United Kingdom, was used. Only the 3-band amplitude image consisting of the three polarizations HH, VV, and HV was available for our study. Figure 4 shows the SAR image corresponding to HH polarization. The image was calibrated, but it still contained intensity variations in the range direction due to varying incidence angles. For classification

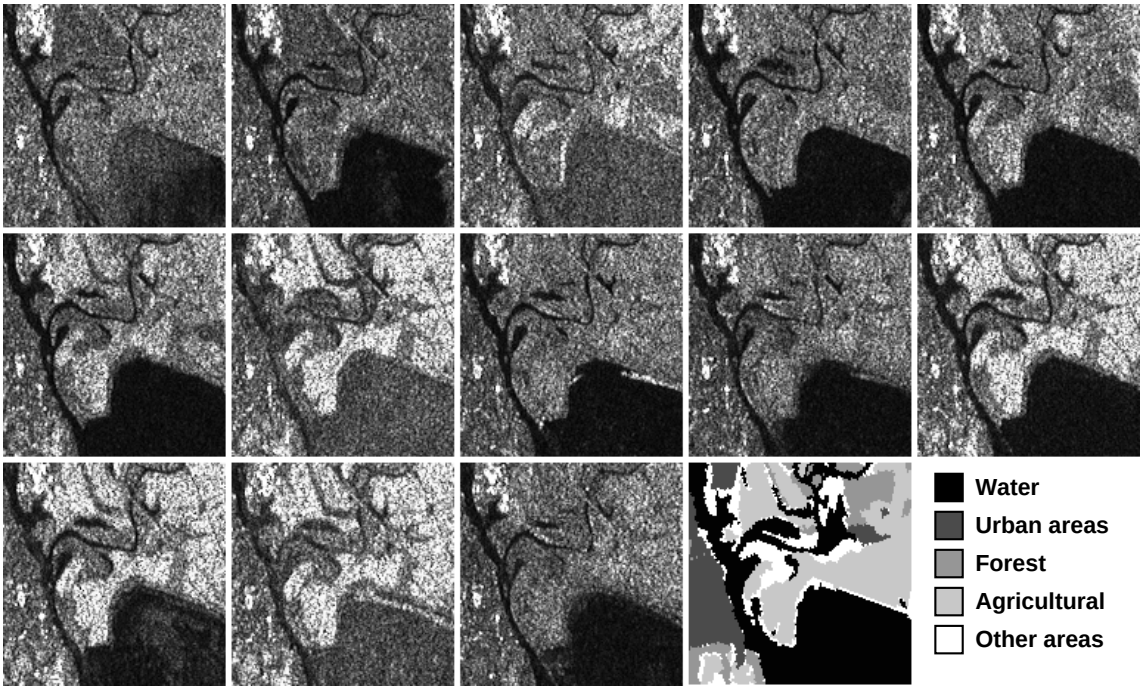

Figure 2: *ERS-1 SAR images.*
Images from August 24 and 27, October 17, 20, 23, and 26, November 4, 16, 19, 22, 25, and 28, and December 7, 1991 of Kjeller, Norway are shown. The SAR signature changes are due to wind conditions (water areas), and soil moisture content and temperature (forest and agricultural areas). The large bright areas (see, e.g., image nos. 1-2 in the third row) are plowed agricultural fields. These same fields appear darker when the soil is dry or frozen (see, e.g., image no. 3 in the third row). A general ground cover map is included for reference purposes.

purposes, we partitioned the image into five subimages corresponding to different ranges, and then processed each subimage separately to reduce misclassifications due to the range variations.

The classification problem considered here involved the following six agricultural classes: wheat, sugar beet, potatoes, carrots, grass, and stubble. A crop map (Figure 5) was used to train and test the classifier [54]. The size of the training areas varied between 4,000 pixels for the grass class and 25,000 pixels for the sugar beet class.

4.1 Invariance of texture features

The first question we attempted to answer, based on the set of 13 ERS-1 SAR images, was whether any of the texture features included in our study was invariant to the natural changes in SAR signature observed within the same land-use category.

Invariance studies were performed for GLCM features, and features computed from local statistics and from the lognormal random field model. Due to the very high computation times needed for fractal features, they were not included in the invariance study.

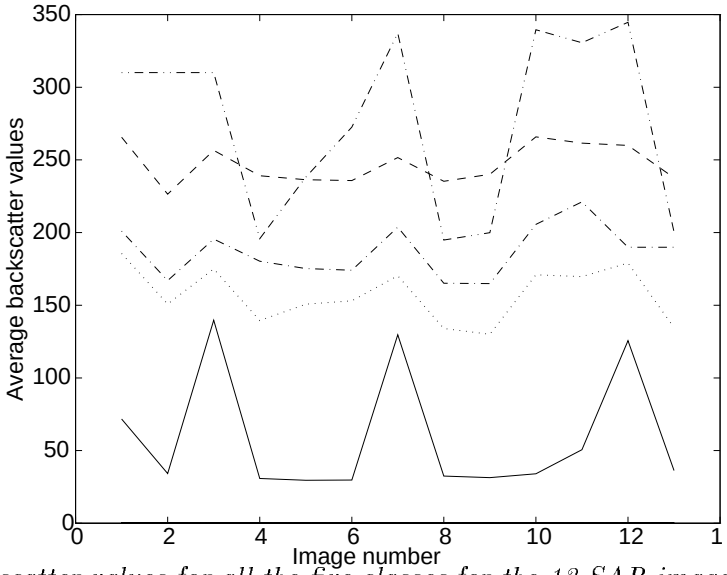

Figure 3: *Average backscatter values for all the five classes for the 13 SAR images. Water: —–; Urban areas: - - -; Forest .....; Unplowed agricultural areas: -.-.-; Plowed agricultural areas: ..- - .*

4.1.1 Invariance of GLCM features

The invariance of GLCM features relies on the invariance of the gray-level co-occurrence matrix itself. The co-occurrence matrix for an image which has been quantized into N_g gray levels can be visualized as an $N_g \times N_g$ image. In our study, $N_g = 8$. To get an indication of the invariance of the co-occurrence matrix, we have computed average co-occurrence matrices corresponding to the five classes for each of the 13 SAR images. The averaged matrices have been generated by first computing the co-occurrence matrices in windows of size 9×9, and then averaging all these matrices for each class. A mosaic of the resulting images is given in Figure 6. The first row shows the 13 8×8 co-occurrence matrices for water, while the second row corresponds to urban areas, the third row to forest, and the fourth and fifth rows to unplowed and plowed agricultural areas, respectively. As expected, the first row (water) shows nearly constant values, except for windy days. For urban aras, non-zero co-occurrence values are found for nearly all the gray levels. The matrices for forest regions are reasonably similar for the 13 images. Agricultural areas show large variations in co-occurrence values. From this preliminary analysis, finding invariant texture features from the co-occurrence matrices seems difficult, particularly for agricultural fields.

The speckle noise present in SAR images is multiplicative in the sense that the backscatter standard deviation is proportional to the mean backscatter value. When the weather conditions cause an increase in the mean backscatter value, the corresponding standard deviation will also increase. An effect of this is that the co-occurrence matrix will show an increase in the off-diagonal values, and a smaller number of elements in the matrix will have close-to-zero values. By inspecting Eqs. (2)-(7), we can predict the changes in feature values when the mean backscatter value increases as follows:

- ASM and IDM features are measures of homogeneity. An increase in the mean backscatter value, which results in an increase in the number of significant co-occurrence values, should lead to a decrease of the ASM and IDM values.
- *Entropy, Inertia, Contrast,* and *Cluster Shade* features are measures of contrast and should increase with an increase in the mean backscatter value.

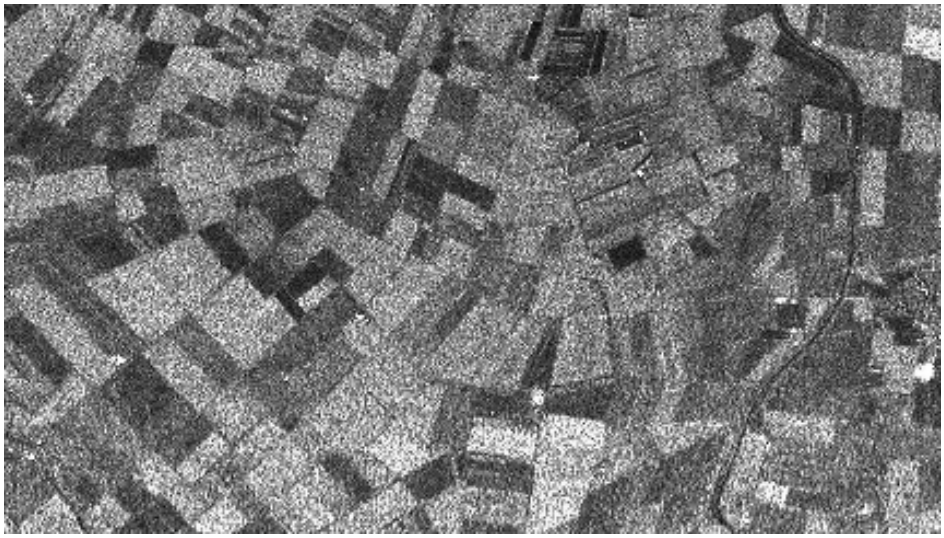

Figure 4: *MAESTRO SAR image of Feltwell, United Kingdom, August 16, 1989. (Only the HH polarization is shown).*

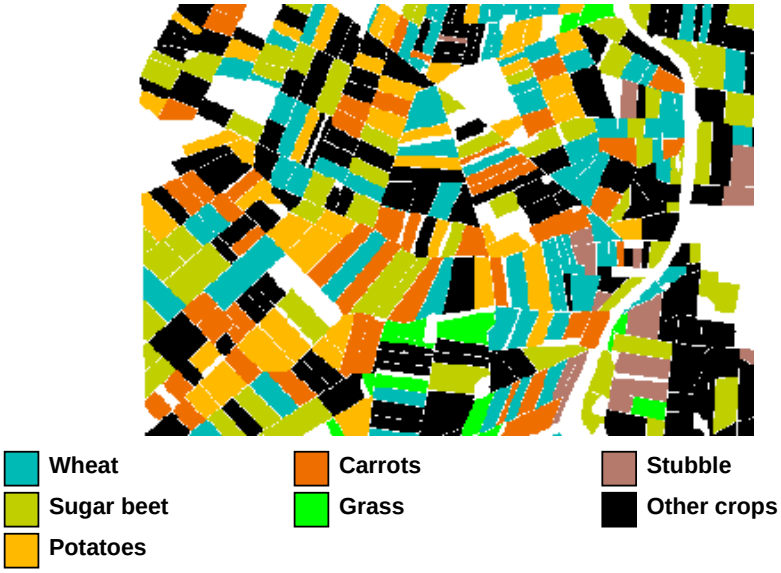

Figure 5: *Crop map of Feltwell, United Kingdom.*

Figure 7 shows plots of the mean values for each of the five classes for the 13 images for the four GLCM features: ASM, IDM, Entropy, and Cluster Shade (the corresponding mean backscatter values are plotted in Figure 3). The predicted effects of an increase in the mean backscatter value can easily be observed for the water class by comparing, e.g., the feature values for image nos. 2 and 3. The average values for each of these features shows large variations with varying mean backscatter values. The texture features do not show better invariance properties than the average backscatter value. Inertia and Contrast features produced curves which were less stable and had larger variations than the curves shown in Figure 7.

The multiplicative speckle noise can, in principle, be transformed into an additive noise by applying a logarithmic transform to the images. To check whether this could lead to a better invariance of the texture features, we computed the GLCM features from logarithmically transformed images. The resulting average feature values for all the GLCM features were found to be highly correlated with the original GLCM features, and these features did not show any better invariance.

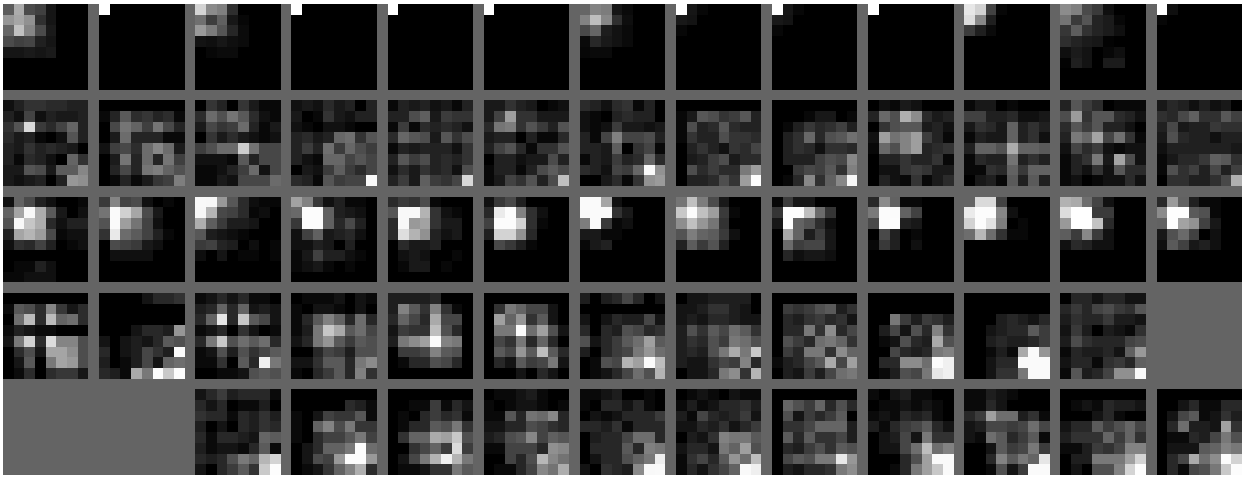

Figure 6: *Co-occurrence matrices.*
A visualization of the average 8×8 *co-occurrence matrices for five different classes for the set of 13 SAR images. The first row shows the 13 co-occurrence matrices for water, while row two corresponds to urban areas, row three to forest, and rows four and five to unplowed and plowed agricultural areas, respectively. Dark values indicate low co-occurrence probabilities, while bright areas indicate high probabilities.*

Since the GLCM features were highly influenced by changes in the mean backscatter value, a normalization which subtracts the local mean backscatter value and divides the result by the local standard deviation was applied prior to computation of the GLCM features to see if the invariance properties improved. No significant improvements in invariance were found, while the discrimination ability of the resulting features was significantly decreased.

4.1.2 Invariance of features computed from local statistics

From speckle statistics, we expect the power-to-mean ratio (PMR) of a homogeneous region to be independent of the mean backscatter value [49]. Figure 8 shows the average values for the five classes for the 13 images for the four features *power-to-mean ratio, skewness, contrast,* and *homogeneity.* We find some variations in the water category for the PMR plot, which we would expect to be constant because water is often considered as pure speckle or without texture. These variations are due to local, light wind variations within the test areas. All the classes show some variations in PMR values, and the variations for unplowed agricultural fields are quite large, but these variations are not highly correlated with the average backscatter values. The higher order moments, kurtosis and skewness, were quite unstable for our images, compared to the power-to-mean ratio, homogeneity, and contrast. Again, we also computed these features from a logarithmically transformed image to see if the noise transformation improved the invariance properties. This transformation resulted in slightly worse performance with respect to invariance.

4.1.3 Lognormal field models

Plots of the average values for each of the five classes for the parameters of the lognormal random field model are shown in Figure 9. Among the five model parameters (μ, θ_1, θ_2, θ_3, and ρ), θ_2 and θ_3 show the largest variations.

In summary, none of the texture features were found to be invariant to natural backscatter changes. A

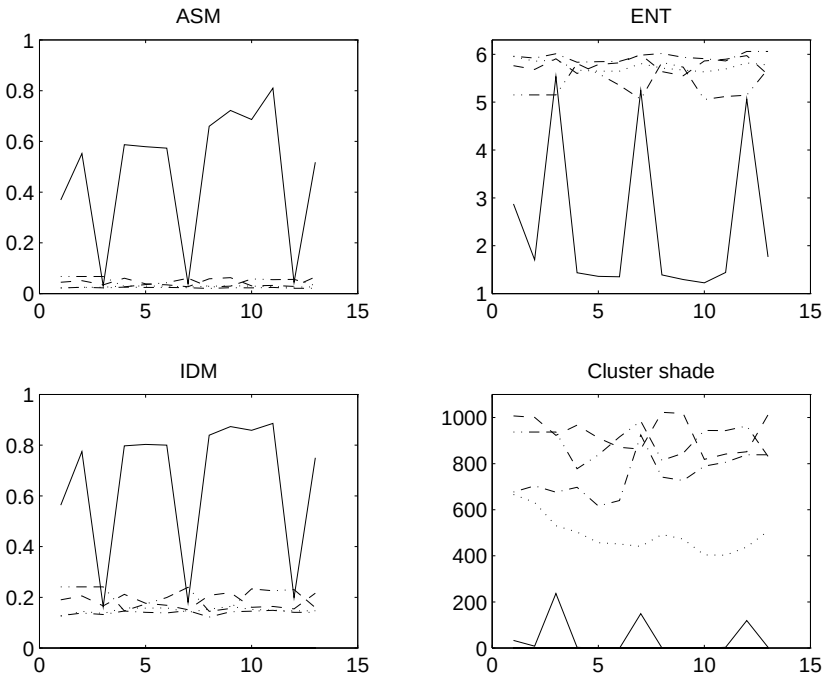

Figure 7: *Average values for the GLCM features for the 13 SAR images for the five ground-cover types. Water, —– ; Urban areas, - - -; Forest, ...; Unplowed agricultural areas, -.-.-; Plowed agricultural areas, ..- - .*

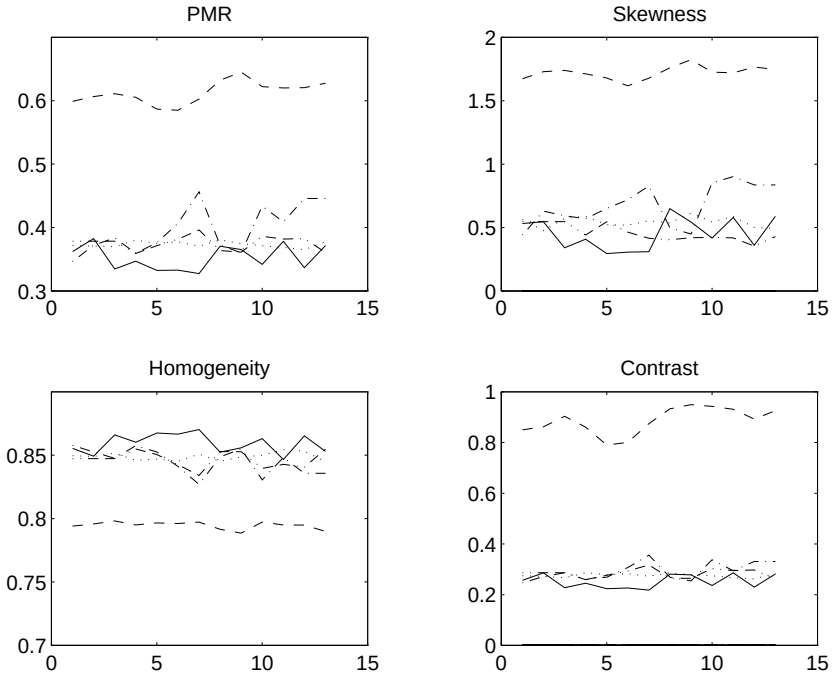

Figure 8: *Average local statistics features for five classes for the 13 SAR images. Water: —– ; Urban areas: - - -; Forest ....; Unplowed Agricultural areas: -.-.-; Plowed agricultural areas: ..- - .*

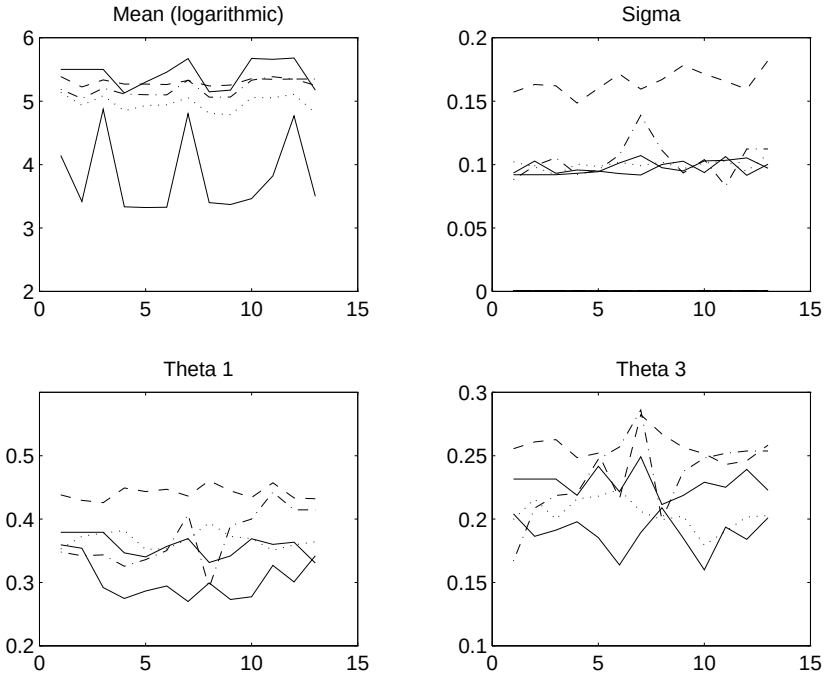

Figure 9: *Average values of the parameters in the random field model for five classes for the 13 SAR images. Water: —— ; Urban areas: - - -; Forest: .....; Unplowed agricultural areas: -.-.-; Plowed areas: .. - -.*

brief discussion on how to circumvent the problem of changing SAR signatures is given in section 5.

4.2 The effect of speckle filtering and noise transformation

In order to investigate the effect of speckle filtering and noise transformation, we addressed the following two questions: (i) Should the texture features be computed from the original images, or do their discrimination ability improve when an image transformation is applied prior to their computation? (ii) Should speckle filtering, which not only reduces the speckle noise, but also smooths the image texture, be applied to the images prior to the computation of the texture features? For this test, five SAR images representing different conditions were selected.

GLCM features The GLCM features were computed from the raw ERS-1 SAR images, and from logarithmically transformed images. The images were then classified, and the average classification accuracy for each image was computed (Table 1). Averaged over five different images, the error rate for classification based on texture features computed from raw images was 32.0%, compared to 31.8 % for texture features computed from logarithmically transformed images. Thus, no significant improvement in classification accuracy was achieved by applying a logarithmic transformation prior to the texture computations.

To study the effect of speckle reduction filters on the discrimination ability of the GLCM features, mean filtering and Lee filtering [31] with varying window sizes were applied to the raw images. By computing GLCM features from smoothed images, the water areas were, in some cases, quantized into the same gray value, resulting in singular covariance matrices. Because of this, the water class was excluded from the filter study, and the reported error rates for the raw image in Table 2 are different from those given in Table 1. The resulting average error rates (average performance over five images)

Acquisition date	Classification error (in %)	
	No preprocessing	Logarithmically transformed images
August 24	32.1	32.5
October 17	39.3	39.6
October 23	35.3	35.6
November 22	24.8	23.9
November 28	28.6	27.4
Average error rates	32.0	31.8

Table 1: *The effect of a logarithmic transform on the classification error for GLCM features. The reported numbers are overall error rates for classification into five classes based on six GLCM features computed from the raw SAR image and from a logarithmically transformed image.*

Acq. date	Classification error (in %)						
	No filtering	3x3 mean	5x5 mean	7x7 mean	3x3 Lee	5x5 Lee	7x7 Lee
August 24	54.4	46.2	48.4	37.0	47.0	42.7	30.9
October 17	48.2	50.2	47.9	49.3	52.0	49.0	44.8
October 23	53.0	50.2	46.5	54.2	53.4	54.0	52.8
November 22	32.9	32.3	29.9	30.6	38.3	41.5	32.1
November 28	33.0	33.4	32.2	32.2	38.9	42.1	35.8
Average	44.4	42.5	41.0	40.7	45.9	45.9	39.3

Table 2: *The effect of speckle filtering on the classification error for GLCM features. The reported numbers are overall error rates for classification into four classes based on six GLCM features computed from the raw SAR image and from speckle-reduced images. The four-class classification error rates are higher than the five-class error rates from Table 1 because the water class, which normally is classified highly accurately, has been removed because speckle filtering of smooth water resulted in co-occurrence features with a zero variation.*

for the four-class problem was 44.4% without speckle filtering, 41.0% after applying a 5×5 mean filter, 40.7% with a 7×7 mean filter, and 39.3% with a 7×7 Lee filter (Table 2). However, with a mean filter or Lee filter of size 7×7, the performance decreased for some images. For further analysis of ERS-1 images, we decided to compute the GLCM features from 5×5 mean filtered images.

Features computed from local statistics The local-statistics features were also computed from the raw ERS-1 SAR images and from filtered images produced by mean filtering and Lee filtering with filters of different window sizes (3×3, 5×5, and 7×7). Without filtering, the average classification error rate for the five-class classification problem was 38.6% (averaged over four images), while after filtering the average error rate varied between 52.2% and 55.8% (Table 3). As both the speckle and the texture are smoothed by applying a filter, a reduction in the discrimination ability of the power-to-mean ratio, as observed, is not unreasonable from a theoretical viewpoint. For further analysis, we computed the statistical features from raw images only.

Features from the lognormal random field model The parameters of the lognormal random field model were computed from the raw SAR images, and from mean-filtered and Lee-filtered images.

Acq. date	Classification error (in %)						
	No filtering	3x3 mean	5x5 mean	7x7 mean	3x3 Lee	5x5 Lee	7x7 Lee
October 17	46.5	52.1	53.1	52.4	51.8	52.1	51.3
October 23	44.9	53.2	52.9	52.3	53.4	52.4	55.8
November 22	31.6	52.2	52.1	51.9	51.9	53.4	58.5
November 28	31.3	52.0	53.1	55.3	52.1	54.2	57.5
Average	38.6	52.4	52.8	53.0	52.3	53.0	55.8

Table 3: *The effect of speckle filtering on the classification accuracy of features computed from local statistics. The reported numbers are overall error rates for classification into five classes based on the features PMR, $SKEW$, $KURT$, and CNT computed from the raw SAR image and from speckle-reduced images.*

Acq. date	Classification error (in %)						
	No filtering	3x3 mean	5x5 mean	7x7 mean	3x3 Lee	5x5 Lee	7x7 Lee
October 17	33.9	32.1	32.3	34.9	32.2	31.9	31.3
October 23	28.1	30.9	31.9	35.6	26.7	27.9	28.4
November 22	20.7	16.5	17.9	22.7	17.7	18.9	21.2
November 28	22.6	18.7	19.9	31.2	19.8	18.9	19.2
Average	26.3	24.5	25.5	31.1	24.1	24.4	25.0

Table 4: *The effect of speckle filtering on the classification accuracy of parameters of the lognormal random field model. The reported numbers are overall error rates for classification into five classes based on the five features θ_1, θ_2, θ_3, σ^2, and μ_x computed from the raw SAR image and from speckle-reduced images.*

The average error rate without filtering was 26.4% (Table 4). The best average overall performance obtained by filtering was achieved by a 3×3 Lee filter, which gave 24.1% error. Based on this, we decided to apply the Lee filter to the ERS-1 SAR images prior to computation of the parameters of the lognormal random field model in the remaining experiments.

Classification based on MAESTRO SAR images To see if the effects of filtering on the texture computation found for land-use classification of ERS-1 SAR images were valid for a different classification problem and a different SAR sensor, the filtering experiments were also performed on the MAESTRO SAR images. The MAESTRO SAR images were classified into six agricultural classes. The resulting error rates for GLCM features and parameters of the lognormal random field model are given in Table 5. For this classification task involving only the agricultural classes, the classification accuracy was not improved by performing speckle filtering prior to texture computation. Agricultural areas are found to be less textured than urban areas and forest regions.

4.3 Selection of texture features with optimal discrimination abilitily

In order to determine the discrimination ability of different subsets of texture features, we addressed the following questions: (i) Which texture feature extraction model performs the best when applied to different images? (ii) When fusing the features computed from different texture models, which feature combination is optimal? (iii) Can texture fusion improve the classification accuracy?

A comparison between the classification performance of the GLCM features, the local-statistics features,

	Classification error (in %)			
Texture method	Filter	HH-polarized image	HV-polarized image	VV-polarized image
GLCM	None	60.2	66.4	60.1
GLCM	5x5 mean	63.2	71.4	64.6
GLCM	7x7 Lee	70.2	75.2	70.5
Lognormal	None	56.2	58.8	53.4
Lognormal	3x3 mean	56.2	58.8	54.8
Lognormal	3x3 Lee	59.6	65.2	59.2

Table 5: *The effect of speckle filtering on the classification error for six agricultural classes of texture features computed from MAESTRO SAR images.*

	Classification error (in %)			
Acquisition date	Mean only	GLCM features	Local statistics	Lognormal fields
October 17	43.8	36.5	46.5	32.2
October 23	33.3	31.0	44.9	26.7
November 22	20.8	22.5	31.6	17.7
November 28	23.9	26.6	31.3	19.8
MAESTRO SAR Image	43.1	40.9	60.7	32.4
Average performance	33.0	31.5	43.0	25.8

Table 6: *A comparison of the performance of different methods for textural feature extraction for SAR images*

and features from the lognormal random field model is provided in Table 6. On an average, the error rate was 25.8% for classification based on lognormal-field features, compared to 31.5% for GLCM features, and 43.0% for features from local statistics. Classification accuracy based on a speckle-reduced backscatter image, with no texture features, was 33.0% on an average. In all the experiments, the features derived from the lognormal field model performed significantly better than the other texture features. The classification accuracy was always significantly improved by using the lognormal field parameters compared to using the speckle-filtered mean backscatter value alone. On the other hand, classification based on features from local statistics only always performed worse than classification based on only the mean backscatter value.

To see if the classification performance could be further improved by combining the features derived from the co-occurrence matrix, local statistics, fractal models, and lognormal field models, a feature selection experiment was performed. The 6 GLCM features (*angular second moment, cluster shade, contrast, entropy, inertia,* and *inverse difference moment*), the five features from local statistics (*power-to-mean ratio, skewness, kurtosis, homogeneity,* and *contrast*), the five parameters of the lognormal field model (θ_1, θ_2, θ_3, σ^2, and μ_x), the two fractal features (*lacunarity* and *fractal dimension*), and the local mean value were merged into one large feature vector, containing a total of 19 features. A feature selection was performed on this set of 19 features, using Whitney's method for feature selection [53]. As mentioned earlier, we replaced the K-NN classifier in the original Whitney's algorithm by a quadratic classifier.

In our experiments, the "best" subset with b features out of the total of 19 features was found for $b = 1, ..., 19$. By best subset of features we mean the feature subset which gives the lowest classification error as determined by Whitney's modified method. Figure 10 shows the classification error rate as a function of the number of texture features included in the classification for the original texture features and for features transformed based on discriminant analysis as described in section 3.6. The results after applying a feature transform as defined by discriminant analysis were not significantly different from classification based on the features selected by Whitney's method using the same number

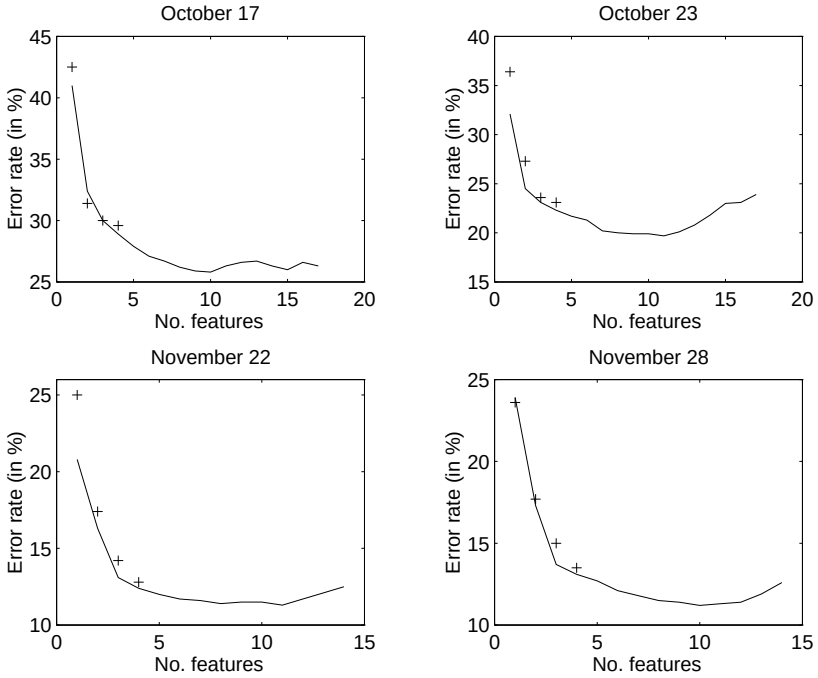

Figure 10: *Classification error rate vs. number of features used.*
Original features: —; Discriminant analysis for a five-class problem gives four new features whose performance is marked as '+'. The total number of features considered for each image varies because features with a high correlation were excluded to avoid singular covariance matrices in the classification.

Texture feature name	Explanation
Mean	Average backscatter value
μ_x, σ^2, θ_1, θ_2, θ_3	Lognormal field parameters
ASM, CONT, CLSH, ENT, IDM, INER	GLCM features
PMR, SKEW, KURT, CNT, HOM	Features from local statistics

Table 7: *Texture feature abbreviations.*

of features. We also note that the lowest classification error was achieved with approximately ten features. Increasing the number of features beyond a certain point resulted in a higher error rate (this phenomenon is called the "curse of dimensionality", see [25]).

Table 7 explains the feature names, while Table 8 lists the 10 best features using Whitney's selection scheme for different ERS-1 SAR images. The features are listed in the order in which they were selected. Among the ten best features, we find reasonably consistent results for different ERS-1 SAR images, though with different individual rankings. Among the frequently selected features are the mean value, the lognormal parameters θ_1, θ_2, θ_3, and σ^2, the GLCM features *Inertia* and *Cluster Shade*, the statistical features *power-to-mean ratio* and *kurtosis*, and the fractal *lacunarity*. Fusion of the texture features also reduced the classification error rate from 24.1% using only the lognormal features to 17.1% on an average.

Image	Best features	Error rate (in %)
October 17	CLSH, SKEW, μ, θ_1, INER, LAC, θ_2, PMR, SKEW, CONT	25.8
October 23	CLSH, SKEW, KURT, θ_3, θ_2, Mean, σ^2, LAC, θ_1, σ^2, IDM	19.9
November 22	Mean, PMR, INER, θ_1, θ_2, CLSH, LAC, θ_3, KURT, σ^2	11.4
November 28	Mean, θ_1, θ_2, θ_3, σ^2, INER, CLSH, PMR, KURT, LAC	11.2

Table 8: *The 10 best features for ERS-1 SAR images and the corresponding error rates for a subset of texture features using Whitney's feature selection scheme. The features are listed in the order in which they were selected.*

5 Discussion and conclusions

In this paper, we have investigated the performance of four different and commonly used methods for texture computation with respect to three important questions pertaining to the interpretation of SAR images. The following methods for texture computations were included in this comparative study: (i) gray-level co-occurrence matrices, (ii) local image statistics, (iii) lognormal field models, and (iv) fractal features. Due to the computational complexity of fractal features, they were only included in some selected experiments.

We first studied the invariance of the texture features with respect to natural changes in SAR signature to check if the texture features could be used in an operational monitoring system. No invariant texture features was found. Secondly, the effect of speckle filtering on the discrimination ability of the texture features was examined. Finally, the relative classification accuracy using different texture models was compared, and feature selection was applied to the pooled set of texture features to select the optimal feature combinations for several images.

Invariance of texture features None of the texture features tested were found to be invariant. Some features were found to be very unstable (in the sense that their values showed substantial variations, and these variations were not significantly correlated with the weather conditions). Other features behaved in a predictable manner when the average backscatter value changed. For operational use of SAR images in a monitoring system, a possible approach might be to avoid using the most unstable features, to establish a library of texture signatures taken under different conditions, and then choosing the appropriate database entry for classification. A further improvement may be to use methods for automatic parameter updating of the unclassified pixels prior to classification to adjust the database entries [12, 23].

The effect of speckle filtering and noise transformations Applying a logarithmic transformation to convert the multiplicative noise to additive noise did not improve the discrimination ability of the texture features. For land-use classification of ERS-1 images, an average reduction in error rate of 3% was found by computing GLCM features from a 5×5 mean-filtered image instead of the input image. Speckle filtering decreased the performance of texture features from local statistics. Lee filtering with a 3×3 filter improved the classification accuracy of features derived from the lognormal field model by only 2.3%, on an average. For classification of MAESTRO SAR images involving only agricultural classes, the classification accuracy was not improved by performing speckle filtering. A possible explanation for the difference in the effect of filtering observed with general land-use classification based on ERS-1 SAR images and agricultural classification based on MAESTRO SAR images might be that the filtering smoothes the moderate texture of agricultural areas too much, while the texture in more

heterogeneous land-use areas is preserved. Since the effect of speckle filtering is a function of both the texture model and the classification problem, we recommend that speckle filtering be used with care in an operational SAR image analysis system.

Texture fusion and feature selection The performance of classifiers based on speckle-filtered backscatter values alone, GLCM features, features from local statistics, and features estimated from the lognormal random field model were compared. The lognormal random field model consistently produced the best results, with an average classification error of 25.8% , compared to 31.5% for GLCM features and 33.0% for speckle-filtered backscatter values.

Feature selection was applied to the pooled set of texture features to find the optimal subset of features. Combining features derived from different models significantly improved the classification accuracy by 7% compared to using features from the lognormal field model alone. A feature transformation defined by discriminant analysis produced similar classification accuracies as feature selection using Whitney's modified method.

By applying feature selection to various SAR images, reasonably consistent classification results were found, although with different individual rankings of the selected features. The images used for this evaluation showed quite a large variation in backscatter values. Although the texture features were not invariant to natural SAR signature variations, the same features were selected in the best subset of features for each image. This robustness in feature selection is a very important property of an operational SAR image analysis system. If a library of texture signatures corresponding to different weather conditions is established, the robustness should be further checked on the large set of images in the data base.

We have compared the texture features in terms of their resulting overall classification accuracy. This approach can be called goal-directed evaluation [45], in which the performance of low-level vision tasks is evaluated in terms of how well the subsequent image analysis steps will perform. Using the overall classification accuracy as the basis for the comparisons is just one of several possible choices (see, e.g., [10]). If discrimination between certain classes is more important than others, then a different measure of classification performance could be used, and the ranking of the selected features might be different.

Combining texture features computed using different models is an important topic for further research. Additional approaches to linear or non-linear feature transformation based on, e.g., flexible discriminant analysis [22] or neural nets should be investigated.

Conclusions Three main conclusions can be drawn from this study: (i) The texture features are not invariant with respect to natural changes in the mean backscatter values. (ii) Speckle filtering improves the discrimination ability of texture features derived from the co-occurrence matrix and from a lognormal random field model. (iii) The most significant result of this study is that texture fusion and selection by combining texture features obtained by different models significantly improve the classification accuracy.

Acknowledgement
This work was supported by a grant from The Research Council of Norway. The comments of Dr. Torfinn Taxt are highly appreciated. We are grateful to the Norwegian Defense Research Establishment, the State Pollution Control Authority (SFT) and the Norwegian Space Center for providing the ERS-1 SAR images. The MAESTRO SAR images were kindly made available to us by NORUT-IT and JRC/ESA.

References

[1] N. Ahuja. Dot pattern processing using Voronoi polygons. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12:336–343, 1982.

[2] D. G. Barber and E. F. LeDrew. SAR sea ice discrimination using texture statistics: A multivariate approach. *Photogrammetric Engineering & Remote Sensing*, 57:385–395, 1991.

[3] D. G. Barber, M. E. Shokr, R. A. Fernandes, E. D. Soulis, D. G. Flett, and E. F. LeDrew. A comparison of second-order classifiers for SAR sea ice discrimination. *Photogrammetric Engineering & Remote Sensing*, 59:1397–1408, 1993.

[4] J. M. Blackledge and E. Fowler. Fractal dimensions segmentation of synthetic aperture radar images. In *Proceedings of the International Conference on Image Processing and its Applications*, pages 445–449, Maastricht, Netherlands, 1992.

[5] D. Blostein and N. Ahuja. Shape from texture: Integrating texture-element extraction and surface estimation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11:1233–1251, 1989.

[6] A. Bourissou, K. Pham, W. C. Au, J. A. Kong, and T. Le Toan. Classification of a polarimetric SAR image using fractal concepts. In *Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 18–21, Tokyo, Japan, 1993.

[7] A. C. Bovik. Analysis of multichannel narrow-band filters for image texture segmentation. *IEEE Trans. Signal Proc.*, 39:2025–2043, 1991.

[8] R. Chellappa and S. Chatterjee. Classification of texture using Gaussian Markov random fields. *IEEE Trans. Acoust. Speech Signal Process.*, 33:959–963, 1985.

[9] P. N. Churchill. MAESTRO 1 proceedings of the data quality workshop 5-6/7-1990. Technical report, number IRSA/MWT/4.90, Version 1.0, July 1990, CEC-JRC/ESA, 1990.

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

[11] G. C. Cross and A. K. Jain. Markov random field texture models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 5:25–39, 1983.

[12] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society, Series B*, 39:1–38, 1977.

[13] H. Derin and H. Elliot. Modeling and segmentation of noisy and textured images using Gibbs random fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 9:39–55, 1987.

[14] L. J. Du. Texture segmentation of SAR images using localized spatial filtering. In *Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 1983–1986, Washington, DC, 1990.

[15] B. Dubuc, S. W. Zucker, C. Tricott, J. F. Quiniou, and D. Wehbi. Evaluating the fractal dimension of surfaces. In *Proceedings of the Royal Society of London A, 235*, pages 113–127, 1989.

[16] M.-P. Dubuisson and R. C. Dubes. Efficacy of fractal features in segmenting images of natural textures. *Pattern Recognition Letters*, 15:419–431, 1994.

[17] R. T. Frankot and R. Chellappa. Lognormal random-field models and their applications to radar image synthesis. *IEEE Trans. Geosc. Remote Sensing*, 25:195–206, 1987.

[18] H. P. Friedman and J. Rubin. On some invariant criteria for grouping data. *J. American Statistical Association*, 62:1159–1178, 1967.

[19] R. C. Gonzales and P. Wintz. *Digital Image Processing.* Addison-Wesley, 1977.

[20] B. Guindon, R. K. Hawkins, and D. G. Goodenough. Spectral-spatial analysis of microwave sea ice data. In *Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 755–758, Vancouver, Canada, 1989.

[21] R. M. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *IEEE Trans. Syst. Man Cyber.*, 3:610–621, 1973.

[22] T. Hastie, R. Tibshirani, and A. Buja. Flexible discriminant analysis by optimal scoring. *J. American Statistical Association*, 89:1255–1270, 1994.

[23] M. Holden, L. Eikvil, and G. Storvik. Methods for updating of model parameters applied within the area of symbol recognition. In *Proceedings of the 8th Scandinavian Conference on Image Analysis*, pages 1291–1298, Tromsø Norway, May, 1993.

[24] Q. Huang, J. R. Lorch, and R. C. Dubes. Can the fractal dimension of images be measured? *Pattern Recognition*, 27:339–349, 1994.

[25] A. K. Jain and B. Chandrasekaran. Dimensionality and sample size considerations in pattern recognition practice. In P. R. Krishnaiah and L. N. Kanal, editors, *Handbook of Statistics, Vol. 2*, pages 835–855. North-Holland, 1982.

[26] A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data.* Prentice Hall, Englewood Cliffs, N. J., 1988.

[27] A. K. Jain and F. Farrokhnia. Unsupervised texture segmentation using Gabor filters. *Pattern Recognition*, 24:1167–1186, 1991.

[28] R. Jha and M. E. Jernigan. Classification and segmentation of SAR sea ice imagery using stochastic image models. In *Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 1881–1884, Washington D.C., 1990.

[29] J. M. Keller, S. Chen, and R. M. Crownover. Texture description and segmentation through fractal geometry. *Computer Vision, Graphics, and Image Processing*, 45:150–166, 1989.

[30] K. I. Laws. Textured image segmentation. Ph.D. thesis, University of Southern California, 1980.

[31] J-S Lee. Speckle analysis and smoothing of SAR images. *Computer Graphics and Image Processing*, 7:24–32, 1981.

[32] J. D. Lyden, B. A. Burns, and A. L. Maffett. Characterization of sea ice types using synthetic aperture radar. *IEEE Trans. Geosc. Remote Sensing*, 22:431–439, 1984.

[33] J. Malik and P. Perona. Preattentive texture discrimination with early vision mechanisms. *J. Opt. Soc. Am. Series A*, 7:923–932, 1990.

[34] B. B. Mandelbrot. *The Fractal Geometry of Nature.* Freeman, San Francisco, 1983.

[35] B. Moghaddam, C. V. Stewart, and K. J. Hintz. Fractal dimension segmentation of synthetic aperture radar. In *Proceedings of the Third International Symposium on Signal Processing and Its Applications, ISSPA'92*, pages 455–458, Gold Coast, Australia, 1992.

[36] A. Murni, N. Darwis, M. Mastur, and D. Hardianto. A texture classification experiment for SAR radar images. In *Proceedings of Pattern Recognition in Practice IV*, pages 213–224, Vlieland, Netherlands, June 1994.

[37] J. A. Nystuen and F. W. Garcia. Sea ice classification using SAR backscatter statistics. *IEEE Trans. Geosc. Remote Sens.*, 30:502–509, 1992.

[38] D. R. Peddle and S. E. Franklin. Image texture processing and data integration for surface pattern discrimination. *Photogrammetric Engineering & Remote Sensing*, 57:413–420, 1991.

[39] S. Peleg, J. Naor, R. Hartley, and D. Avnir. Multiple resolution texture analysis and classification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 6:518–523, 1984.

[40] A. Pentland. Fractal-based description of natural scenes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 9:661–674, 1984.

[41] A. H. Schistad and A. K. Jain. Texture analysis in the presence of speckle noise. In *Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 884–886, Houston, Texas, May 1992.

[42] M. E. Shokr. Texture measures for sea-ice classification from radar images. In *Proceedings of the International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 763–768, Vancouver, Canada, 1989.

[43] H. Skriver and P. Gudmansen. Study on the use and characteristics of SAR for sea ice and polar ice applications. Tech. Report, Electromagnetic Institute, Technical University of Denmark, R-307, Final Report, ESA Contract No. 5442/83/D/IM(SC), June 1986.

[44] F. Tomita and S. Tsuji. *Computer Analysis of Visual Textures.* Kluwer Academic Publishers, Boston, 1990.

[45] Øivind Due Trier and Anil K. Jain. Goal-directed evaluation of binarization methods. In *NSF/ARPA Workshop on Performance Versus Methodology in Computer Vision*, pages 206–217, Seattle, Washington, June 1994.

[46] M. Tuceryan. Moment-based texture segmentation. *Pattern Recognition Letters*, 15:659–668, 1994.

[47] M. Tuceryan and A. K. Jain. Texture segmentation using Voronoi polygons. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12:211–216, 1990.

[48] M. Tuceryan and A. K. Jain. Texture analysis. In C. H. Chen, L. F. Pau, and P. S. P. Wang, editors, *Handbook of Pattern Recognition and Computer Vision*, pages 235–276. World Scientific Publishing Company, 1993.

[49] F. T. Ulaby, F. Kouyate, B. Brisco, and T. H. Lee Williams. Textural information in SAR images. *IEEE Trans. Geosc. Remote Sens.*, 24:235–245, 1986.

[50] R. Voss. Random fractals: Characterization and measurement. In R. Pynn and A. Skjeltorp, editors, *Scaling Phenomena and Disordered Systems*, pages 1–11. Plenum Press, New York, 1986.

[51] L. Wang and D. C. He. A statistical approach for texture. *Photogrammetric Engineering & Remote Sensing*, 56:61–66, 1990.

[52] J.S. Weszka, C. R. Dyer, and A. Rosenfeld. A comparative study of texture measures for terrain classification. *IEEE Trans. Syst., Man, Cybern.*, 6:269–285, 1976.

[53] A. Whitney. A direct method of non-parametric measurement selection. *IEEE Trans. Computers*, 20:1100–1103, 1971.

[54] M. G. Wooding. Crop survey of the Feltwell area - August 1989, short report and crop map prepared within the MAESTRO-1 experimental programme. Technical report, British National Space Centre, January 1990.

[55] Z. You and A. K. Jain. Performance evaluation of shape matching via chord length distribution. *Computer Vision, Graphics, and Image Processing*, 28:185–198, 1984.

[56] S. W. Zucher. Toward a model of texture. *Comput. Graph. Image Process.*, 5:190–202, 1976.