

Unsupervised polarimetric SAR image segmentation and classification using region growing with edge penalty

Peter Yu, A. K. Qin, *Member, IEEE*, David A. Clausi, *Senior Member, IEEE*

Abstract—A region-based unsupervised segmentation and classification algorithm for polarimetric SAR imagery that incorporates region growing and a Markov random field (MRF) edge strength model is designed and implemented. This algorithm is an extension of the successful Iterative Region Growing with Semantics (IRGS) segmentation and classification algorithm, which was designed for amplitude only SAR imagery, to polarimetric data. Polarimetric IRGS (PolarIRGS) extends IRGS by incorporating a polarimetric feature model based on the Wishart distribution and modifying key steps such as initialization, edge strength computation and the region growing criterion. Like IRGS, PolarIRGS oversegments an image into regions and employs iterative region growing to reduce the size of the solution search space. The incorporation of an edge penalty in the spatial context model improves segmentation performance by preserving segment boundaries that traditional spatial models will smooth over. Evaluation of PolarIRGS with Flevoland fully polarimetric data shows that it improves upon two other recently published techniques in terms of classification accuracy.

Index Terms—synthetic aperture radar (SAR), image segmentation, complex, polarimetry, Markov random field (MRF), Wishart, region adjacency graph (RAG), region-based

I. INTRODUCTION

POLARIMETRIC synthetic aperture radar (SAR) imagery provides useful information in a diverse number of applications from sea ice monitoring [1], [2] to land cover classification [3] and agricultural crop identification [4]. There is now an increasing volume of fully polarimetric data due to the launch of sensors capable of fully-polarimetric imaging such as RADARSAT-2. Therefore, automated image segmentation and classification methods are desired to replace manual interpretation, which is subjective and labour-intensive.

Automated segmentation and classification of fully polarimetric SAR imagery has been an ongoing field of research. Maximum likelihood classifiers based on the assumption of Wishart-distributed classes have been developed [5]–[7]. Another approach [8], which assumes complex Gaussian class distributions, also incorporates a Markov random field (MRF) spatial context model that overcomes some of the effects of noise by smoothing the segmentation according to local

interactions between pixel labels. A recent work [9] uses a non-parametric estimate of the class distributions with a modified spatial context model that attempts to prevent over-smoothing of the segmentation result across true class boundaries.

The previously mentioned papers perform segmentation on a per-pixel basis. Region-based segmentation methods also exist and are advantageous because they reduce the computation demand by working on regions instead of pixels, help the optimization procedure converge more effectively to the global solution and alleviate problems with noisy imagery by using region statistics instead of individual pixel values. Several recent papers advocate region-based image segmentation. Wu et al. [10] use a region-based Wishart MRF (Region WMRF) framework for segmentation. The image is oversegmented into square regions, which are then iteratively refined by the WMRF model and finally classified with training data. Wu et al.'s Region WMRF uses a basic MRF spatial context model which penalizes all class boundaries equally, including true class boundaries which should not be penalized.

Agglomerative hierarchical clustering [11] is another technique that oversegments the image into many regions by clustering over a polarimetric decomposition data space and merges regions to produce the final segmentation. The recent spectral graph partitioning [12] (SGP) technique first segments the image into regions with contour and spatial proximity information and then groups the regions by spectral clustering. Neither agglomerative hierarchical clustering [11] nor SGP [12] use a spatial context model to improve segmentation, although SGP makes use of image edge information.

The advantage of a region-based approach combined with an MRF spatial context model has been demonstrated for real-valued imagery by the recent Iterative Region Growing with Semantics (IRGS) algorithm [13], which has outperformed other algorithms when applied to amplitude-only SAR sea ice imagery [14], generic imagery [15] and optical imagery of savannah wetlands [16]. IRGS incorporates edge strength between regions in the spatial context model to ensure that true class boundaries are preserved as regions grow using a merging criterion. This is similar to the segmentation energy function defined in [17] but IRGS incorporates a different edge strength penalty function to aid in the segmentation process [13]. As IRGS is an unsupervised algorithm, it requires no training data prior to segmentation and classification.

This paper presents an extension to IRGS for polarimetric

The Natural Sciences and Engineering Research Council (NSERC) of Canada, NSERC International Polar Year (IPY) and GEOIDE (Geomatics for Informed Decisions, a Network of Centres of Excellence) financially supported this work.

The authors are members of the Vision and Image Processing (VIP) research group in the Department of Systems Design Engineering, University of Waterloo, Waterloo, ON, N2L 3G1 Canada.

SAR imagery. By combining the successful region-based approach and edge-strength spatial context model of IRGS with the Wishart distribution model, unsupervised polarimetric SAR segmentation and classification can be improved. The proposed technique is suitable for applications where the classes of interest have a reasonable spatial extent, such as agricultural classification. As with Region WMRF [10] and SGP [12], polarimetric IRGS is evaluated on an agricultural data set.

Section II outlines polarimetric SAR imagery, image segmentation, classification and the IRGS algorithm. Section III details the extension of IRGS to polarimetric IRGS (PolarIRGS). Section IV describes the fully polarimetric test data. Section V evaluates PolarIRGS with the test data. Conclusions and recommendations for future work follow in Section VI.

II. BACKGROUND

A. Polarimetric SAR images

A fully polarimetric radar measures the complex scattering matrix of the target medium, which in the monostatic case with a reciprocal medium reduces to a complex scattering vector [5]:

$$\mathbf{u} = [S_{hh} \quad \sqrt{2}S_{hv} \quad S_{vv}]^T \quad (1)$$

The terms S_{hh} , S_{hv} , S_{vv} correspond to the complex scattering coefficients of the HH, HV and VV polarizations, respectively. Multi-look processing is often performed for speckle reduction and data compression of SAR data by averaging several single-look outer products $\mathbf{u}\mathbf{u}^\dagger$ [18]:

$$\mathbf{Z} = \frac{1}{n} \sum_{k=1}^n \mathbf{u}_k \mathbf{u}_k^\dagger \quad (2)$$

where n is the number of looks, \mathbf{u}_k is the k^{th} single-look scattering vector and † is the conjugate transpose operator. The polarimetric covariance matrix \mathbf{Z} , or a measure that can be converted to it, is the measurement provided at each pixel from a multi-look, polarimetric image. The matrix $\mathbf{A} = n\mathbf{Z}$ has a complex Wishart [19] distribution [18]:

$$p(\mathbf{A}) = \frac{|\mathbf{A}|^{n-q} \exp(-\text{tr}(\mathbf{C}^{-1}\mathbf{A}))}{K(n, q) |\mathbf{C}|^n} \quad (3)$$

where $|\mathbf{A}|$ is the determinant of \mathbf{A} , $\text{tr}(\mathbf{C}^{-1}\mathbf{A})$ is the trace of $\mathbf{C}^{-1}\mathbf{A}$, $\mathbf{C} = \mathbf{E}[\mathbf{u}\mathbf{u}^\dagger]$ and:

$$K(n, q) = \pi^{\frac{1}{2}q(q-1)} \prod_{i=1}^q \Gamma(n-i+1) \quad (4)$$

The value of q is the number of elements in \mathbf{u} and Γ is the Gamma function. In the case of monostatic polarimetric radar, $q = 3$ under the reciprocity assumption [18]. For the Wishart distribution in Eq. (3) to be nonsingular, $n \geq q$. Eq. (3) will be used in Section III to develop PolarIRGS.

B. Region-based Image Segmentation Problem Definition

Let \mathcal{C} be the number of segmented classes in the image. Let $\mathcal{S} = \{(i, j) | 1 \leq i \leq M, 1 \leq j \leq N\}$ represent a discrete two dimensional rectangular lattice of size M rows \times N columns (i.e. an M row \times N column image) and let $s \in \mathcal{S}$ represent a site in the lattice (i.e. a pixel in the image). Let $\mathbf{X} = \{X_s | s \in \mathcal{S}\}$ be a set of discrete random variables forming a random field on \mathcal{S} , with each X_s taking a value from $\mathcal{L} = \{1, \dots, \mathcal{C}\}$ that indicates the class label of site s . Also, let $\mathbf{Y} = \{\mathbf{Y}_s | s \in \mathcal{S}\}$ be a random field on \mathcal{S} . Each \mathbf{Y}_s represents a measurement at site s and can take on a scalar, vector or matrix value. For the polarimetric data in this paper, each \mathbf{Y}_s takes on a positive definite Hermitian matrix which represents the complex covariance matrix \mathbf{Z}_s measured at site s (see Eq. (2)).

Label field configuration $\mathbf{x} = \{x_s | x_s \in \mathcal{L}, s \in \mathcal{S}\}$ and observed image $\mathbf{y} = \{\mathbf{y}_s | s \in \mathcal{S}\}$ are realizations of \mathbf{X} and \mathbf{Y} , respectively. The domain of \mathbf{y}_s depends on the type of data being considered: for a real-valued multichannel image with d bands, \mathbf{y}_s are d -dimensional vectors (i.e., $\mathbf{y} = \{\mathbf{y}_s | \mathbf{y}_s \in \mathbb{R}^d, s \in \mathcal{S}\}$), while for multi-look polarimetric data, $\mathbf{y} = \{\mathbf{Z}_s | s \in \mathcal{S}\}$, where the domain of \mathbf{Z}_s is the cone of positive definite Hermitian matrices [18].

Image segmentation involves finding an optimal \mathbf{x} according to some criteria on \mathbf{x} and \mathbf{y} . After segmentation, $\Omega_1, \dots, \Omega_C$ are disjoint subsets of \mathcal{S} that denote the \mathcal{C} classes in the label field configuration \mathbf{x} . Image segmentation, as defined here, allows classes that consist of unconnected regions. This definition of image segmentation is consistent with that given by Gonzalez and Woods [20, p. 690] and Li [21, p.188] and has been used in other published work [11] [9] [22]. By allowing classes that consist of unconnected regions and by not using any training data to find the optimal \mathbf{x} , the segmentation result is also an unsupervised classification result. Each region is a connected image segment with a classification label assigned. Under this definition of image segmentation, the proposed PolarIRGS algorithm is a *combined* unsupervised segmentation *and* classification algorithm. It is a region-based method based on the IRGS algorithm [13], which uses a region-adjacency graph (RAG) [23] to represent the image. Under a RAG representation, the image is first oversegmented into a number of disjoint and relatively homogeneous regions that consist of multiple image sites and the label field is defined over the RAG [24] instead of the lattice \mathcal{S} .

A RAG is represented as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} denotes the set of image regions as vertices of the graph and \mathcal{E} is the set of arcs that connect spatially adjacent regions. A region $v \in \mathcal{V}$ is an image region and \mathcal{S}_v is the set of image sites belonging to region v . An arc $e \in \mathcal{E}$ represents the shared boundary between two adjacent regions. The random field for the label configuration is now defined on \mathcal{G} [24] and is denoted by $\mathbf{X}^r = \{X_v^r | X_v^r \in \mathcal{L}, v \in \mathcal{V}\}$, where the superscript 'r' indicates a region-based definition of each term. The region-based label field configuration $\mathbf{x}^r = \{x_v^r | x_v^r \in \mathcal{L}, v \in \mathcal{V}\}$ is a realization of \mathbf{X}^r . Here, x_v^r is the label for all sites $s \in \mathcal{S}_v$.

All regions are mutually separated by region boundaries, which are comprised of image sites that are not part of any regions. The division of image sites into boundary sites and

sites that are part of regions follows several rules [25, Sec. 4.2.1]. All sites in the image are either boundary sites or part of a region and sites that are a part of one region cannot be a part of any other regions. Additionally, only sites that separate two regions are boundaries. An arc on the RAG connects regions v and w if they share a common boundary. These rules are formalized by the following definitions [25, Sec. 4.2.1]:

- 1) $\bigcup_{v \in \mathcal{V}} \mathcal{S}_v \bigcup_{v \in \mathcal{V}} \partial \mathcal{S}_v = \mathcal{S}$
- 2) $\forall v, w : \mathcal{S}_v \cap \partial \mathcal{S}_w = \emptyset$
- 3) $\forall v \neq w : \mathcal{S}_v \cap \mathcal{S}_w = \emptyset$
- 4) $\forall v : \partial \mathcal{S}_v \subseteq \bigcup_{w \in \mathcal{V}, w \neq v} \partial \mathcal{S}_w$

where $\partial \mathcal{S}_v$ is a one-pixel boundary outlining region \mathcal{S}_v . Regions v and w share a common boundary if $\partial \mathcal{S}_v \cap \partial \mathcal{S}_w \neq \emptyset$.

There are fewer possible configurations of the region-based label field \mathbf{x}^r than the pixel-based field as there are far fewer regions than pixels. In Section II-C, it will be seen that MRF-based segmentation problems are a combinatorial optimization problem. Therefore, the region-based approach reduces the solution search space, leading to better optimization results. Additionally, region-based segmentation reduces the negative effects of variation and noise within each region.

C. IRGS Algorithm

This section summarizes the details of the IRGS algorithm necessary for understanding the proposed polarimetric extension. Additional details are described in previous papers [13] [15]. Assuming that \mathbf{X}^r is an MRF with respect to a certain neighbourhood system on RAG \mathcal{G} , the unsupervised image segmentation and classification problem is formulated as finding the label field configuration \mathbf{x}^{r*} that satisfies [15] [21]:

$$\mathbf{x}^{r*} = \arg \max_{\mathbf{x}^r \in \mathbf{X}^r} p(\mathbf{y}|\mathbf{x}^r)P(\mathbf{x}^r) \quad (5)$$

where $p(\mathbf{y}|\mathbf{x}^r)$ is the conditional probability density function of the observed image \mathbf{y} given the specific label field configuration \mathbf{x}^r and $P(\mathbf{x}^r)$ is the probability of a specific label configuration and is defined by the MRF model chosen [21]. The operator $\arg \max$ selects the label field configuration \mathbf{x}^r that maximizes $p(\mathbf{y}|\mathbf{x}^r)P(\mathbf{x}^r)$. The selected configuration is denoted by \mathbf{x}^{r*} , where the superscript $*$ indicates that \mathbf{x}^{r*} is the optimal configuration with respect to Eq. 5. The term $p(\mathbf{y}|\mathbf{x}^r)$ is called the feature model because it models the distribution of the features (the observation values in \mathbf{y}) given \mathbf{x}^r and $P(\mathbf{x}^r)$ is called the spatial context model because it models the probability of the various possible configurations of \mathbf{x}^r , with some configurations being more likely under the chosen MRF model. Extending the IRGS algorithm to polarimetric data involves replacing the feature model and adapting the spatial context model to match the characteristics of polarimetric data.

Under a class-conditionally independent assumption, $p(\mathbf{y}|\mathbf{x}^r)$ can be written as:

$$p(\mathbf{y}|\mathbf{x}^r) = \prod_{i=1}^c \prod_{\mathcal{S}_v \in \Omega_i} \prod_{s \in \mathcal{S}_v} p(\mathbf{y}_s | x_v^r = i) \quad (6)$$

where $\mathcal{S}_v \in \Omega_i$ selects the set of sites \mathcal{S}_v for each region v that belongs to class i (i.e. $\{\mathcal{S}_v | \{v | x_v^r = i\}\}$). Although

the segmentation is region-based, (6) is written in terms of image sites because the calculations are based on individual site values. The term $p(\mathbf{y}_s | x_v^r)$ describes the probability of obtaining value \mathbf{y}_s given that the site belongs to the class specified by x_v^r . Note that under the region-based definition, $x_s = x_v^r$ for $s \in \mathcal{S}_v$.

In order to solve (5), it is converted into an energy function by taking the logarithm to change products into sums and changing the sign to give an equivalent minimization problem:

$$\mathbf{x}^{r*} = \arg \min_{\mathbf{x}^r \in \mathbf{X}^r} E_f(\mathbf{y}, \mathbf{x}^r) + E_s(\mathbf{x}^r) \quad (7)$$

where $E_f(\mathbf{y}, \mathbf{x}^r)$ is the feature model term corresponding to $p(\mathbf{y}|\mathbf{x}^r)$ and $E_s(\mathbf{x}^r)$ is the spatial context model term corresponding to $P(\mathbf{x}^r)$. $E_f(\mathbf{y}, \mathbf{x}^r)$ is derived for polarimetric data in Section III.

In IRGS, the neighbours for a given region v are all regions connected to it by one arc in the RAG. The spatial context energy $E_s(\mathbf{x}^r)$ for this MRF neighbourhood system is [15]:

$$E_s(\mathbf{x}^r) = \beta \sum_{i=1}^{c-1} \sum_{j=i+1}^c \sum_{s \in \partial \Omega_i \cap \partial \Omega_j} g(\nabla_s) \quad (8)$$

where $g(\nabla_s)$ is the edge penalty term and $\partial \Omega_i$ comprises all the boundary sites that separate regions assigned to class i from regions assigned to other classes. Hence, $\partial \Omega_i \cap \partial \Omega_j$ selects the shared boundary sites between classes i and j . The parameter β controls the degree to which the spatial context model is weighted, with larger β resulting in smoother segmentations.

At a conceptual level, (8) penalizes segmentations where adjacent regions are assigned to different classes. For every boundary site that separates a region of class i from a region of class j , $E_s(\mathbf{x}^r)$ is increased by $\beta \cdot g(\nabla_s)$. The edge penalty term $g(\nabla_s)$ is a monotonically decreasing function of ∇_s , where $\nabla_s \in [0, 1]$ is a measure of the edge strength at site s . The penalty is smaller when there is a strong edge between two regions assigned to different classes than when the edge is weak. This approach favours assigning adjacent regions to the same class only when the edge between them is weak, recognizing that strong edges indicate true class boundaries.

The spatial context energy used in IRGS is similar to the spatial context energy in the multi-level logistic (MLL) [26] model except for the addition of the edge penalty term $g(\nabla_s)$. The MLL model may over-smooth across true class boundaries since all boundaries are penalized equally.

The IRGS edge penalty function $g(\nabla_s)$ is [15, Eq. (15)]:

$$g(\nabla_s) = \exp \left[- \left(\frac{\nabla_s}{K} \right)^2 \right] \quad (9)$$

The only difference between $g(\nabla_s)$ for PolarIRGS and $g(\nabla_s)$ for IRGS [15] is the calculation of ∇_s for polarimetric data, which is described in Section III. The term K is a positive value that controls the strength of the edge penalty term (see [13, Sec. 3] for details).

IRGS iteratively searches for the configuration of labels that solves (7). At each iteration, IRGS performs region-merging on

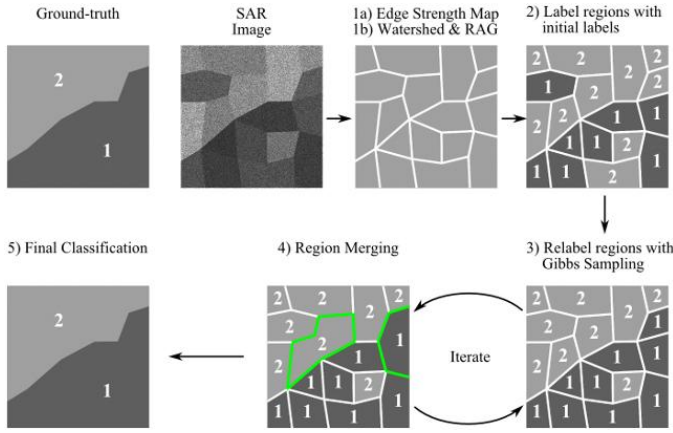


Fig. 1. Major steps of the IRGS algorithm. 1a) IRGS computes the edge strength map and 1b) generates the watershed regions and the region adjacency graph (RAG) from the image. Watershed regions shown here are simplified and have more arbitrary shapes in practice. 2) Watershed regions are given initial labels. 3) Regions are relabelled by Gibbs sampling [28]. 4) Regions with the same label are merged. Steps 3 and 4 repeat for a user-specified number of iterations. 5) The final segmentation is produced.

the intermediate segmentation to reduce the number of RAG nodes, which makes subsequent iterations more efficient and prevents the algorithm from becoming trapped in local minima in the solution space [13]. IRGS merges adjacent regions with the same class label in a greedy fashion. It computes a merging criterion ∂E [13] for each eligible pair of regions and merges the pair with the most negative ∂E . This continues until no more negative ∂E are found. Shared boundary sites between two regions become absorbed into the merged region.

Fig. 1 shows the major steps of the IRGS algorithm. The algorithm begins with an input image, which is oversegmented with a watershed algorithm [27] using an edge strength map $\nabla = \{\nabla_s | s \in \mathcal{S}\}$. The RAG is then created with each watershed region becoming a vertex in the graph. IRGS and PolarIRGS are not constrained to any particular oversegmentation algorithm for RAG construction; the Vincent and Soille watershed [27] was used because it produces a reasonable RAG for both IRGS and PolarIRGS. After RAG construction, each region is assigned an initial label to initialize IRGS. IRGS then enters its iterative phase. At each iteration, Gibbs sampling [28] assigns each region a label to move the label configuration toward the optimal solution. Regions are then merged and the next iteration of IRGS is executed.

The combined segmentation and classification nature of IRGS is apparent from the preceding description. Each of the regions is an image segment which is assigned a generic label from $\mathcal{L} = \{1, \dots, C\}$ in an unsupervised fashion as no *a priori* data is provided other than the number of classes. The final output is the label field configuration \mathbf{x}^{r*} , which is an unsupervised classification map of the scene.

III. PROPOSED METHOD

There are three main types of unsupervised classification algorithms for polarimetric SAR data [18]: statistical (e.g. clustering), physical and combined physical-statistical algorithms. The proposed PolarIRGS algorithm is an extension

of the statistical type since it combines a statistical model with a spatial context model. Physical model algorithms use domain knowledge such as the scattering characteristics of broad categories of surfaces (e.g. vegetation, buildings and water) to divide a scene into these categories. The combined physical-statistical algorithms use both physical and statistical models. For example, Lee et al. [29] used a Freeman-Durden decomposition to divide a scene into three categories of scattering characteristics. Each category is then separately classified with unsupervised statistical clustering. This paper is focused on the development of unsupervised statistical classification. The use of novel statistical techniques such as PolarIRGS in a combined physical-statistical classification framework could be investigated in future work.

To implement PolarIRGS, the following must be detailed:

- Derivation of the feature model energy $E_f(\mathbf{y}, \mathbf{x}^r)$ in (7) for polarimetric data.
- Definition of the region merging criterion ∂E .
- Calculation of the edge strength measure ∇_s , which is used by both the watershed algorithm and the edge penalty function $g(\nabla_s)$ in (8).
- Initialization method to assign initial labels to all regions.
- Choice of spatial context weighting parameter β .
- Model for labeling the boundary sites that remain after the region level label configuration \mathbf{x}^{r*} is generated to create the final classification image \mathbf{x} .

A. Feature Model Energy

In (3), $p(\mathbf{A})$ can be restated as $p(\mathbf{A}_s | x_v^r = i)$ [5]:

$$p(\mathbf{A}_s | x_v^r = i) = \frac{|\mathbf{A}_s|^{n-q} \exp(-\text{tr}(\mathbf{C}_i^{-1} \mathbf{A}_s))}{K(n, q) |\mathbf{C}_i|^n} \quad (10)$$

where \mathbf{C}_i is the mean polarimetric covariance matrix of class i . This corresponds to the $p(\mathbf{y}_s | x_v^r = i)$ terms in (6), assuming that $\mathbf{y}_s = \mathbf{A}_s = n\mathbf{Z}_s$. In the actual data, $\mathbf{y}_s = \mathbf{Z}_s$ but this can be ignored for the derivation. The feature model energy $E_f(\mathbf{y}, \mathbf{x}^r)$ can be derived by substituting (10) into (6), taking the natural logarithm and changing the sign:

$$E_{f_1}(\mathbf{y}, \mathbf{x}^r) = \sum_{i=1}^c \sum_{\mathcal{S}_v \in \Omega_i} \sum_{s \in \mathcal{S}_v} \{n \ln |\mathbf{C}_i| + \ln(K(n, q)) + \text{tr}(\mathbf{C}_i^{-1} \mathbf{A}_s) - (n - q) \ln |\mathbf{A}_s|\} \quad (11)$$

Eliminating all terms that are not a function of i and substituting $\mathbf{A}_s = n\mathbf{Z}_s$ into (11) gives [5]:

$$E_{f_2}(\mathbf{y}, \mathbf{x}^r) = \sum_{i=1}^c \sum_{\mathcal{S}_v \in \Omega_i} \sum_{s \in \mathcal{S}_v} \{n \ln |\mathbf{C}_i| + n \text{tr}(\mathbf{C}_i^{-1} \mathbf{Z}_s)\} \quad (12)$$

The constant multiplier n can be dropped without affecting the validity of the energy value, giving the final feature model energy in terms of \mathbf{C}_i and \mathbf{Z}_s :

$$E_f(\mathbf{y}, \mathbf{x}^r) = \sum_{i=1}^c \sum_{\mathcal{S}_v \in \Omega_i} \sum_{s \in \mathcal{S}_v} \{\ln |\mathbf{C}_i| + \text{tr}(\mathbf{C}_i^{-1} \mathbf{Z}_s)\} \quad (13)$$

The term in the sum is the Wishart distance measure derived by Lee et al. [5].

B. Region-merging criterion

The region-merging criterion ∂E determines whether two regions can be merged during region-merging. Region-merging considers a single pair of adjacent regions with the same class label at a time. Once a pair of eligible regions is selected, the algorithm determines the energy for a two region version of (7), where the energy is calculated with only the two regions in the pair. The energy for two cases is calculated [13]: when the two regions are separate classes (unmerged) and when the two regions are the same class (merged). If the merged case has lower energy, the regions are merged. Formally, IRGS merges regions v and w if $\partial E(v, w) < 0$, where:

$$\partial E(v, w) = E_{\text{merged}}(v, w) - E_{\text{unmerged}}(v, w) \quad (14)$$

The two energies in (14) have the following values for polarimetric data:

$$E_{\text{merged}}(v, w) = \sum_{s \in \mathcal{S}_v \cup \mathcal{S}_w} \{\ln |\mathbf{C}_{vw}| + \text{tr}(\mathbf{C}_{vw}^{-1} \mathbf{Z}_s)\} \quad (15)$$

$$\begin{aligned} E_{\text{unmerged}}(v, w) &= \sum_{s \in \mathcal{S}_v} \{\ln |\mathbf{C}_v| + \text{tr}(\mathbf{C}_v^{-1} \mathbf{Z}_s)\} + \\ &\quad \sum_{s \in \mathcal{S}_w} \{\ln |\mathbf{C}_w| + \text{tr}(\mathbf{C}_w^{-1} \mathbf{Z}_s)\} + \\ &\quad \beta \sum_{s \in \partial \mathcal{S}_v \cap \partial \mathcal{S}_w} g(\nabla_s) \end{aligned} \quad (16)$$

where \mathbf{C}_v is the average polarimetric covariance matrix of region v , and region vw is the union of regions v and w . Since $\mathbf{C}_v = \frac{1}{|\mathcal{S}_v|} \sum_{s \in \mathcal{S}_v} \mathbf{Z}_s$, where $|\mathcal{S}_v|$ is the cardinality of \mathcal{S}_v , $\partial E(v, w)$ can be simplified by making the following substitution in (15) and (16) for regions v , w and vw :

$$\begin{aligned} &\sum_{s \in \mathcal{S}_v} \{\ln |\mathbf{C}_v| + \text{tr}(\mathbf{C}_v^{-1} \mathbf{Z}_s)\} \\ &= |\mathcal{S}_v| \ln |\mathbf{C}_v| + \text{tr}(\mathbf{C}_v^{-1} \sum_{s \in \mathcal{S}_v} \mathbf{Z}_s) \\ &= |\mathcal{S}_v| \ln |\mathbf{C}_v| + \text{tr}(\mathbf{C}_v^{-1} |\mathcal{S}_v| \mathbf{C}_v) \\ &= |\mathcal{S}_v| \ln |\mathbf{C}_v| + |\mathcal{S}_v| \text{tr}(\mathbf{I}) \end{aligned} \quad (17)$$

where \mathbf{I} is the identity matrix of the same size as \mathbf{Z}_s . Since $|\mathcal{S}_{vw}| = |\mathcal{S}_v| + |\mathcal{S}_w|$, all terms involving $\text{tr}(\mathbf{I})$ sum to zero in the final merging criterion which yields:

$$\begin{aligned} \partial E(v, w) &= |\mathcal{S}_{vw}| \ln |\mathbf{C}_{vw}| - |\mathcal{S}_v| \ln |\mathbf{C}_v| - \\ &\quad |\mathcal{S}_w| \ln |\mathbf{C}_w| - \beta \sum_{s \in \partial \mathcal{S}_v \cap \partial \mathcal{S}_w} g(\nabla_s) \end{aligned} \quad (18)$$

C. Edge Strength Measure

Both the watershed [27] step and the edge penalty function in (8) require an edge strength measure ∇_s at each site s . Two approaches for calculating ∇_s were considered. The first is the vector field gradient (VFG) [30] that is already used in IRGS, which works on real-valued images with one or more channels and computes the gradient magnitude at each pixel. The second approach is a polarimetric edge strength calculation for complex polarimetric data, e.g. [31]–[33].

Since the VFG calculation works only on real-valued images, preliminary tests evaluated several real-valued polarimetric decompositions for use with VFG:

- (i) The amplitudes of the HH, HV and VV channels:

$$\mathbf{y}_s^{\parallel} = \begin{bmatrix} 20 \log |S_{hh}| \\ 20 \log |S_{hv}| \\ 20 \log |S_{vv}| \end{bmatrix} \quad (19)$$

- (ii) The Pauli decomposition [18]:

$$\mathbf{y}_s^{\text{Pauli}} = \begin{bmatrix} 20 \log \frac{|S_{hh} + S_{vv}|}{2} \\ 20 \log \frac{|S_{hh} - S_{vv}|}{2} \\ 20 \log (2|S_{hv}|) \end{bmatrix} \quad (20)$$

- (iii) The H/ α decomposition [18]:

$$\mathbf{y}_s^{\text{H}/\alpha} = [H \quad \alpha]^T \quad (21)$$

where H is the entropy and α is the polarimetric scattering parameter. The H/ α decomposition has been used to initialize other unsupervised polarimetric segmentation algorithms [7].

The amplitudes of the HH, HV and VV channels (i.e. \mathbf{y}_s^{\parallel}) produced the best overall classification accuracy in the tests and are therefore used in the VFG computation in this paper. The amplitude images have a large range of values, so each channel was clipped to the range of $[-40, -5]$ dB to prevent large values from dominating the computed gradient range. The VFG implementation in IRGS expects data to be normalized to $[0, 255]$ and the normalization from $[-40, -5]$ dB to $[0, 255]$ is done in floating point values to avoid quantization.

A polarimetric edge strength calculation method [31], which makes full use of the polarimetric information, was also evaluated but did not lead to consistently accurate classification results. This approach computes a likelihood ratio [34], known as the Bartlett distance [35], between pairs of pixels in different orientations to measure edge strength. Several parameters of the method [31] were adjusted, including the number of orientations used, the filter window size and whether number of looks estimation [36] in each filter window was used. As none of these adjustments produced consistent results, unlike the VFG gradient using \mathbf{y}_s^{\parallel} which did generate consistent results, the VFG gradient was chosen as the edge strength measure.

D. Initialization

IRGS requires that the statistics of each class (\mathbf{C}_i in (13)) be initialized before the algorithm begins. Several approaches are possible. The \mathbf{C}_i can be set to values from a classification of a related scene and used as the initial seed values for

a polarimetric K-means algorithm that uses a polarimetric distance measure [5]. Alternatively, the polarimetric matrices can be decomposed into H/ α parameters and classified by a pre-determined division of the H- α feature space to obtain a class map to calculate the initial \mathbf{C}_i [7]. Both of these approaches require initial set up with *a priori* knowledge: a classification of a related scene or a pre-determined division of the H- α plane. As IRGS aims to be as automated as possible, several fully automated approaches were considered:

- (i) Assign random labels to each region and compute the mean \mathbf{C}_i for each class.
- (ii) The initial \mathbf{C}_i from approach (i) are used as the initial seed for polarimetric K-means [5] to calculate a refined set of \mathbf{C}_i .
- (iii) Run the IRGS region-level K-means [15] on one of:
 - (a) The HH, HV and VV amplitude image from (19)
 - (b) The Pauli decomposition image from (20)
 - (c) The H/ α parameters from (21)
 and use the obtained classes to compute the initial \mathbf{C}_i .
- (iv) Use the initial \mathbf{C}_i from one of approaches (iii)(a) - (c) as the initial seed for polarimetric K-means [5] to find a second, refined set of initial \mathbf{C}_i .

Approach (iii)(a) with the amplitude of the HH, HV and VV (i.e. \mathbf{y}_s^{\parallel}) channels produced the best overall results across the image test cases described in Section IV. Thus, it is the approach used to initialize \mathbf{C}_i for all results in this paper.

E. Spatial context parameter β choice

At the beginning of each iteration (step 3 in Fig. 1), the value of β in (8) is automatically modified to change the weighting of the spatial context model in a data driven manner, which is described in previous literature [13] [15]. Standard IRGS requires the Fisher criterion J [37, Sec. 4.10] between two classes in the intermediate classification to compute a β that adapts to the class separability in the image. However, there is no Fisher criterion between classes with Wishart-distributed mean covariance matrices. Future work should investigate possible analogous measures for Wishart matrices. Currently, β in PolarIRGS must be calculated without the Fisher criterion:

$$\beta(\tau) = c_1 \beta_0(\tau) \quad (22)$$

where τ is the iteration number to emphasize that β changes at each iteration. The value of $\beta_0(\tau)$ is calculated at each iteration to maintain the expected class boundary length at the next iteration [38] to preserve the current detail level in the image. The multiplier c_1 controls the overall smoothness and level of merging of the algorithm [15]. In this paper, $c_1 = 5$ since it produced the best results. Previous experiments found that IRGS is not particularly sensitive to c_1 [39].

F. Boundary site labeling

After the region-based label field configuration is produced, the single pixel boundary between regions remains unlabelled. Each boundary site s is labelled by choosing the label that minimizes the following energy:

$$x_s = \arg \min_{i \in \mathcal{L}} \{ \ln |\mathbf{C}_i| + \text{tr}(\mathbf{C}_i^{-1} \mathbf{Z}_s) + \beta \sum_{t \in \mathcal{N}_s \cap \mathcal{S}_{\text{labelled}}} (1 - \delta(i, x_t)) \} \quad (23)$$

where \mathcal{N}_s is the eight pixel neighbourhood of site s , $\mathcal{S}_{\text{labelled}}$ is the set of all sites that have been labelled already, and δ is the Kronecker delta function. Eq. (23) labels pixels based on their similarity to the class statistics and the labels of its neighbours and is very similar to the MLL model [26].

G. Polarimetric IRGS Algorithm

Table I lists the full PolarIRGS algorithm. The maximum number of iterations τ in Step 2 is set to 100. This was found to be sufficient for convergence [14]. Convergence is achieved when the label field remains constant on further iterations. This is easily detected and for all tests in this paper, this occurred in fewer than 100 iterations.

TABLE I
ALGORITHMIC DESCRIPTION OF POLARIRGS

-
1. a) Compute edge strength map $\nabla = \{\nabla_s | s \in \mathcal{S}\}$ to measure edge strength at all sites using the image feature set defined in (19).
b) Oversegment image with the watershed algorithm [27] and construct RAG with one vertex for each watershed region.
 2. Let τ be the current iteration number and set $\tau = 1$. Assign a random label to each region and run a region-level K-means [15] on \mathbf{y}_s^{\parallel} from (19) to produce an initial classification. From this classification, compute initial \mathbf{C}_i (mean polarimetric covariance matrix) for each class i .
 3. a) Update β according to (22).
b) Each vertex $v \in \mathcal{V}$ is visited once in random order and assigned a label with Gibbs sampling [28] to minimize (7). This step moves the segmentation and classification result toward the optimal configuration. If $\tau > \tau_{max}$, go to Step 5, where τ_{max} is a user-specified maximum number of iterations.
 4. a) Compute merging criterion ∂E with (18) for every pair of regions with the same label that is connected by one arc in the RAG.
b) Merge the pair with the most negative ∂E . Repeat Step 4a until all $\partial E \geq 0$.
c) Go back to Step 3.
 5. Label boundary sites using (23) to produce \mathbf{x} , the final label field configuration.
-


IV. DATA

Four-look, fully polarimetric L-band data from NASA/JPL AIRSAR of Flevoland, The Netherlands [40] is used to evaluate PolarIRGS. A red-green-blue (RGB) composite of the Pauli decomposition is shown in Fig. 2(a). PolarIRGS' unsupervised classification accuracy was evaluated by four test images shown in Fig. 2(b):

- 1) Mask 1: This sub-image was used by Wu et al. [10] to evaluate their Region-based WMRF technique. The ground-truth used in this paper to evaluate Mask 1 classification results is shown in Fig. 4(d) and is an approximate recreation of Wu et al.'s ground-truth as the original authors could not be contacted. Evaluation of PolarIRGS on Mask 1 with a similar ground-truth as Wu et al. allows a comparison to be made between PolarIRGS and Region-based WMRF.



(a) Pauli RGB composite



(b) Sub-images for testing

Fig. 2. The Flevoland 4-look polarimetric L-band scene [40] (AIRSAR imagery courtesy SIR-C CDROM, A. Freeman - NASA/JPL). (a) The Pauli RGB composite is a partial visualization of the polarimetric data. (b) The sub-images correspond to images tested in other papers. Mask 1 was used to evaluate Region WMRF [10] while Mask 3 was used to evaluate SGP [12]. Mask 2 and the full image are additional test cases in this paper.

- 2) Mask 2: The ground-truth used to evaluate this sub-image is a portion of a ground-truth map (which will be referred to as GT2 in this paper) that combines information from [41] and [35]. The portion of GT2 used for Mask 2 is shown in Fig. 5(d).
- 3) Mask 3: This sub-image was used to evaluate SGP by Ersahin et al. [12]. The ground-truth used for Mask 3 is a portion of GT2 and is shown in Fig. 6(d). Ersahin et al. used the same ground-truth to evaluate SGP. Mask 3 allows for comparison with SGP.
- 4) Full image: The full image was evaluated by using the entire GT2 map as ground-truth. This ground-truth is shown in Fig. 7(b).

The ground-truth map used for Mask 1 and by Wu et al. [10] is a distinct version of the Flevoland ground-truth from the GT2 map (although the two maps do not contradict each other). Thus, Mask 1 is evaluated separately from the other test cases. None of the ground-truth maps of the Flevoland scene give a label for every single pixel in the image. Therefore, accuracy statistics are computed only for pixels of known class. The legend for the ground-truth maps and classification results is shown in Fig. 3

V. TESTING AND RESULTS

A. Evaluation Methodology

PolarIRGS was implemented in C++ as a module in the MAGIC image analysis system [42]. The algorithm was used to segment and classify the four test cases described in Section IV.

The classification result initially does not have meaningful labels (i.e., the generic labels assigned by the algorithm is not associated with the actual name of a ground-truth class). This is due to the fact that PolarIRGS is an unsupervised segmentation and classification algorithm (see Section II-B).

To compare with published ground-truth, each generic label in the unsupervised classification result (one of $\{1, \dots, C\}$) must be associated with a ground-truth class (Stembeans, Peas, etc.). This is accomplished by finding a mapping of generic labels to ground-truth class names (e.g. label 2 maps to Stembeans, etc.) that maximizes the overall accuracy, where overall accuracy is defined as the percentage of pixels with ground-truth information that are assigned to the correct class. This mapping can be found by exhaustively trying all possible mappings of generic labels to ground-truth labels and finding the one with highest overall accuracy. However, this is equivalent to finding the mapping that maximizes the trace of the classification confusion matrix, which is an assignment problem [43] that can be solved more efficiently by linear programming. The generic labels are mapped to ground-truth labels in this optimal fashion rather than labeling with training data so that the success of the unsupervised classification algorithm can be evaluated without introducing uncertainties due to an imperfect labeling process. As the image segments have already been given generic labels by PolarIRGS before a meaningful name is mapped to each generic label, the algorithm does not use any ground-truth information to generate the classification.

The overall accuracy and the individual class accuracies (the percentage of pixels of each class correctly classified) are used as measures of the quality of a classification algorithm. The ground-truth does not provide a label for each pixel of the entire image so the accuracy calculation is limited to only those pixels where the ground-truth provides a label. Since the Gibbs sampling in Step 3 is a stochastic process [28] and because the region-level K-means initialization [15] begins with a random seed, the algorithm can produce slightly different results on each execution on the same data. Therefore, PolarIRGS was run 10 times for each sub-image and the full scene to determine the stability and consistency of PolarIRGS.

To demonstrate the benefit of including edge strength information as part of the segmentation cost function, an MLL-based version of PolarIRGS (PolarMLL) was implemented and tested. This differs from PolarIRGS only in that $g(\nabla_s) = 1$, which penalizes all class boundaries equally.

B. Mask 1 Results

The mean accuracy over the 10 executions of PolarIRGS and the standard deviation of the accuracy are reported in Table II for Mask 1. PolarMLL results are also included. Results reported by Wu et al. [10] for their Region-based Wishart-MRF technique (Region WMRF) are also shown for comparison. The number of classes for Mask 1 was set to 8 to match the Region WMRF paper. Table II shows that PolarIRGS is highly successful for classifying Mask 1: the overall accuracy is 98.3%. This is higher than PolarMLL and Region WMRF, although Wu et al. [10] do not report sufficient statistics (i.e.

the confusion matrix) to determine whether this improvement is statistically significant. The consistency of PolarIRGS over 10 runs of the algorithm is also excellent, with very low standard deviations across the accuracy values for all classes. In terms of individual class accuracies, PolarIRGS compares favourably with Region WMRF, improving upon Region WMRF for 4 classes, although Region WMRF outperforms PolarIRGS for the other 4 classes. However, the overall accuracy shows that PolarIRGS has an advantage.

Fig. 4 shows a classification result from one of the ten runs of the algorithm alongside the Pauli decomposition RGB and the ground-truth. Fig. 4(b) and Fig. 4(e) show the same classification result but Fig. 4(e) is masked to remove pixels where the ground-truth does not provide a label, which is how the accuracy is calculated for Table II. The classification result in Fig. 4 matches very well with the ground-truth, with most pixels being correctly classified. Fig. 4 demonstrates the effectiveness of the IRGS MRF model: the segment boundaries correspond very well to the expected boundaries based on the ground-truth and each segment is contiguous and nearly noise-free. For this particular case, PolarMML and PolarIRGS produce results that are visually similar but PolarIRGS has higher classification accuracy.

C. Mask 2 Results

Table III reports the accuracy statistics for Mask 2 (as well as Mask 3 and the Full Image) using PolarIRGS and PolarMML. It is clear that including the edge penalty in PolarIRGS improves the accuracy compared to PolarMML.

Mask 2 was segmented and classified with 11 classes. Although this is larger than the number of ground-truth classes in the Mask 2 area, it was necessary because the ground-truth does not account for every class in the scene as not all pixels are labelled; the number of true classes is higher than the number of ground-truth classes. Since PolarIRGS is an unsupervised classification algorithm, setting the number of classes equal to the number of ground-truth classes produced unreasonable results because the classes found by the classification algorithm are not guaranteed to correspond only to the ground-truth classes. Specifying a larger number of classes allows the algorithm to find the true classes in the image, including any excess classes not shown in the ground-truth.

Table III indicates that the overall accuracy for Mask 2 is lower than that of Mask 1. The high standard deviation for the accuracy of Rapeseed and Grasses in Mask 2 are due to a single outlier execution of PolarIRGS. The consistency and accuracy of PolarIRGS are poorer for Mask 2 than for Mask 1. Increasing the number of classes used for classification (from 8 to 11) makes the solution search space for the optimal label configuration larger, which means that the algorithm is more likely to be trapped in various local minima of (7). This results in lower consistency (as different runs of the algorithm can be trapped in different local minima) and lower accuracy (as local minima are not the optimal solution). However, accuracy remains at over 90% for most classes and overall.

Fig. 5 shows the classification results in the same manner as Fig. 4. Black or dark gray indicates pixels assigned to

excess classes that cannot be mapped to a ground-truth class (due to setting the number of classes to 11 when there are only 8 ground-truth classes). These are classes found in the image but not represented in the ground-truth labeling. The classification result shown in Fig. 5(e) compares well with the ground-truth in Fig. 5(d). Segment boundaries match accurately and the segments are contiguous and free of noise. Grasses and Lucerne are confused, as well as Wheat and Rapeseed. However, the Pauli RGB image shows that such confusion may be unavoidable: Grasses and Lucerne appear similar, as do Wheat and Rapeseed. This suggests that there may be low separability between these classes in the polarimetric data space.

Fig. 5 shows examples of the role of the edge strength model. Fig. 5(b) is the result for PolarIRGS, which uses the edge strength model, and Fig. 5(c) is the result for PolarMML, which does not consider edge strength. The four corresponding circled regions in these two images demonstrate how visible class boundaries are ignored by PolarMML, leading to erroneous region merging:

- 1) In the centre of the PolarMML image, the Excess 1 class grows across the class boundary and erroneously merges across the boundary with the adjacent Rapeseed region.
- 2) In the top right, the Grasses segment is merged with Peas.
- 3) At the top left corner are two small regions identified as Stembeans which PolarMML merges across their boundaries with Potatoes.
- 4) In the bottom left, a large region of Excess 1 and a large region of Grasses are merged by PolarMML and their combined statistics relabel the class as Peas.

In each of these cases, using the edge strength model preserves the class boundaries.

D. Mask 3 Results

Mask 3 was segmented and classified with 9 classes, the same as the number of ground-truth classes and the same number that Ersahin et al. [12] used to evaluate their SGP technique using Mask 3. The overall accuracy of PolarIRGS for Mask 3 in Table III is higher than PolarMML. PolarIRGS achieved a higher overall accuracy of 84.4% compared to SGP's accuracy of 81.2% [12] for the same sub-image. Individual class accuracies were not reported for SGP. For PolarIRGS, several classes have low accuracy (Stembeans, Rapeseed and Lucerne). In Fig. 6, Rapeseed (orange in the ground-truth) was assigned the same class as Wheat (pink in the ground-truth). The Lucerne (cyan) was assigned the same class as Grass, which may be unavoidable due to the similar appearance of Grass and Lucerne in the Pauli RGB image (Fig. 6(a)). The Lucerne that does appear in the classification result is an unidentified extra class that is not in the ground-truth. It is similar to Potatoes in the Pauli RGB image but is darker. The algorithm is therefore finding Potatoes, dark Potatoes and Grass rather than Potatoes, Lucerne and Grass. Stembeans (red in the ground-truth) in Mask 3 is grouped together with Potatoes. The small patch of Stembeans in the upper right of Mask 3 does not have a strong enough edge with the adjacent Potato patch to remain separate in the IRGS model.



Fig. 3. Legend for all ground-truth maps and classification results. Excess classes are distinct classes in the image in regions without ground-truth information.




Fig. 4. Mask 1 classification results for PolarIRGS and PolarMLL. PolarMLL is identical to PolarIRGS except it ignores edge strength.

TABLE II

ACCURACY STATISTICS USING MASK 1 TO COMPARE POLARIRGS, POLARMLL AND REGION WMRF [10]. POLARIRGS GENERATES THE HIGHEST OVERALL CLASSIFICATION RATE AT 98.2% WHILE REGION WMRF AND POLARMLL GENERATE LOWER OVERALL ACCURACIES OF 95.5% AND 91.9%.

Algorithm	PolarIRGS		PolarMLL		Region WMRF [10]	
# of Runs	10		10		N/A	
# of Classes	8		8		8	
Class	Mean Acc.	Std. Dev.	Mean Acc.	Std. Dev.	Mean Acc.	Std. Dev.
Peas	96.7	0.1	96.3	0.75	98.4	N/A
Beet	87.6	0.2	93.6	4.3	85.3	N/A
Bare Soil	98.6	0	98.5	0.05	99.8	N/A
Rapeseed	100	0	100	0	99.2	N/A
Lucerne	99.8	~ 0	95.6	12.7	100	N/A
Potatoes	99.9	~ 0	100	0.02	100	N/A
Barley	99.1	0.05	93.2	18.1	90.2	N/A
Wheat	100	0	77.7	15.3	90.5	N/A
OVERALL	98.2	0.02	91.9	3.8	95.4	N/A

N/A = Not Applicable

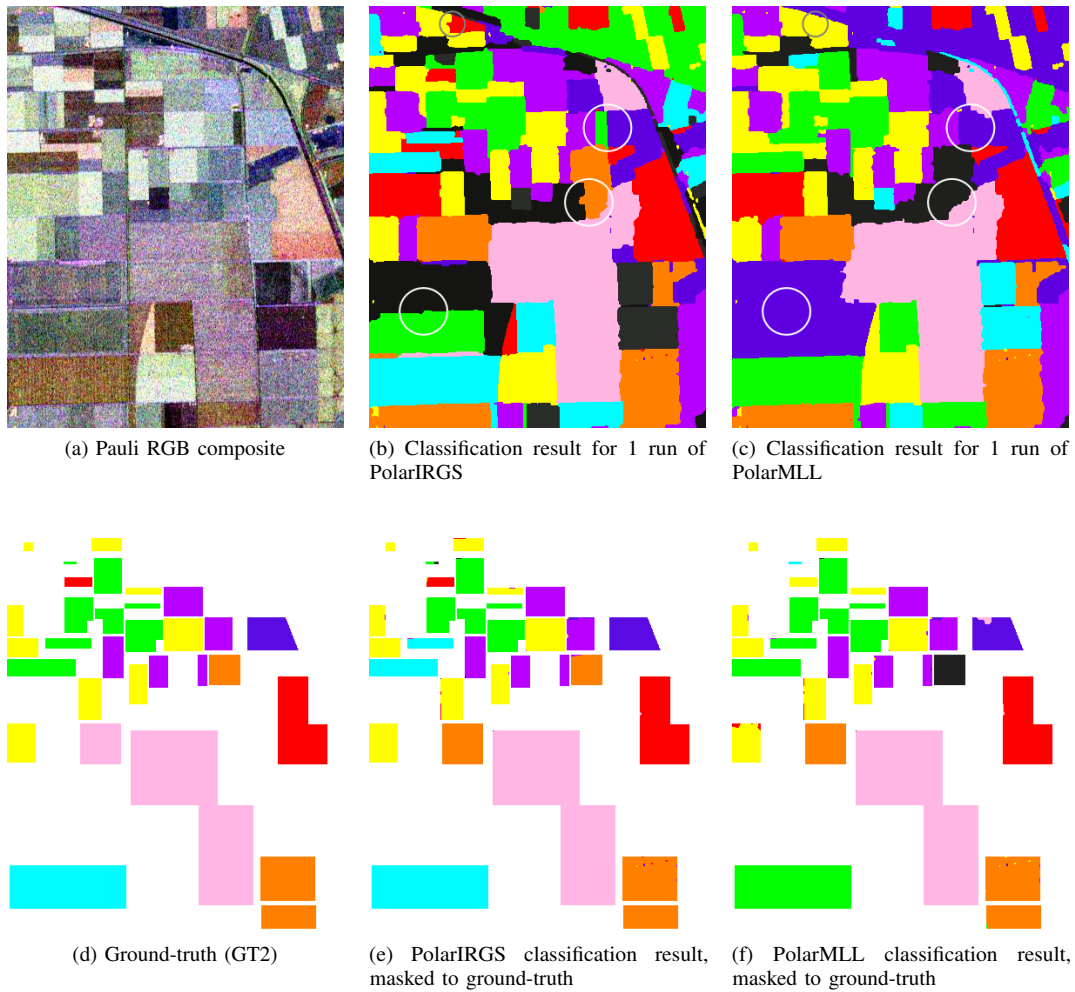


Fig. 5. Mask 2 classification results with 11 classes. Black or dark gray pixels are excess classes that are not associated with any ground-truth class. Circles show locations where the PolarMLL labeling smooths across visible class boundaries that the edge strength model of PolarIRGS preserves.




Fig. 6. Classification results for Mask 3 with 9 classes. Circles show locations where the PolarMLL labeling smooths across visible class boundaries that the edge strength model of PolarIRGS preserves.

TABLE III

ACCURACY STATISTICS FOR (a) POLARIRGS AND (b) POLARMLL FOR MASK 2, MASK 3 AND THE FULL IMAGE. A DASH (-) INDICATES THAT THE CLASS DOES NOT EXIST IN THE SUB-IMAGE. FOR MASK 3, POLARIRGS (84.4% OVERALL ACCURACY) IMPROVES OVER THE PUBLISHED SPECTRAL GRAPH PARTITIONING (SGP) METHOD [12] WHICH ACHIEVED 81.2% OVERALL ACCURACY (INDIVIDUAL CLASS ACCURACIES WERE NOT REPORTED IN [12]). THE FULL IMAGE CLASSIFICATION DOES NOT HAVE ANY KNOWN COMPARABLE PUBLISHED RESULTS FOR UNSUPERVISED CLASSIFICATION.

Sub-image	Mask 2		Mask 3		Full Image	
# of Runs	10		10		10	
# of Classes	11		9		17	
Class	Mean Acc.	Std. Dev	Mean Acc.	Std. Dev	Mean Acc.	Std. Dev
Peas	100	0	-	-	99.4	0.6
Stembeans	99.8	0.1	14.3	26.9	92	7.2
Beet	98.6	0.4	96.5	5.1	92.1	9.1
Forest	-	-	-	-	100	~ 0
Bare Soil	-	-	100	0	83.6	5.5
Rapeseed	91.4	25	0	0	33.6	7.1
Grasses	72.3	9.5	99.8	0	34.6	38.0
Lucerne	100	0	0.2	0	59.2	31.2
Potatoes	99.3	~ 0	83.3	~ 0	37.7	9.6
Barley	-	-	-	-	69.2	26.0
Wheat	89	3.9	100	0	97.6	5.2
Wheat 2	-	-	100	0	60	51.6
Wheat 3	-	-	-	-	79	23.7
Water	-	-	-	-	38.9	4.4
Buildings	-	-	-	-	93.9	1.4
Overall	91.9	0.2	84.4	0.9	69.8	4.4

(a) PolarIRGS

Sub-image	Mask 2		Mask 3		Full Image	
# of Runs	10		10		10	
# of Classes	11		9		17	
Class	Mean Acc	Std. Dev	Mean Acc	Std. Dev	Mean Acc	Std. Dev
Peas	97.7	1.8	-	-	70.2	36.8
Stembeans	92.6	0.1	11.1	27.9	71.3	16.2
Beet	91.8	6.5	81.5	16.9	65.8	15.8
Forest	-	-	-	-	99.9	0.1
Bare Soil	-	-	100	0	99.3	~ 0
Rapeseed	86.9	10.7	0	0	35.9	4.7
Grasses	85.1	26.5	99.8	0.1	38.4	28.2
Lucerne	30	48.3	0.2	0.2	66.6	23.4
Potatoes	98.7	0.4	75.7	4	14.2	15.2
Barley	-	-	-	-	80	1.6
Wheat	87.5	0.2	99.4	1.9	94.8	6.3
Wheat 2	-	-	97.4	8.2	20	42.0
Wheat 3	-	-	-	-	90	20.0
Water	-	-	-	-	43.6	6.7
Buildings	-	-	-	-	94	0.9
Overall	83.2	3.7	79.8	2.8	63.7	2.0

(b) PolarMLL

Although individual class accuracies are not reported in the SGP paper [12], examination of their visual results reveals that Rapeseed and Wheat were both mislabelled as Wheat 2, which would result in nearly 0% accuracy for those two classes. Lucerne is also confused with Grass in the SGP paper, although some of it is classified correctly. These classes (Rapeseed and Wheat, Lucerne and Grass) may be quite similar in the polarimetric data space since both recent algorithms (PolarIRGS and SGP) have trouble separating them.

Dramatically improving the accuracy of Mask 3 classification is possible by using a different method of edge strength calculation and initialization. For example, when the edge strength ∇_s was calculated on the H/α decomposition $\mathbf{y}_s^{H/\alpha}$ from (21) and PolarIRGS was initialized with approach (ii) in Section III-D, the overall accuracy improves from 84.3% to

93.4%, Stembean accuracy improves from 14.3% to 87.6%, Lucerne improves from 0.2% to 63.6% and Potato improves from 83.3% to 97.7% while all other classes remain similarly accurate. The standard deviations of the accuracies are also all reduced to less than 1%. These results are not emphasized to keep the edge strength calculation and initialization method consistent across all test cases in this paper. As mentioned in Section III-C and Section III-D, the edge strength measure and initialization method chosen for PolarIRGS are based on best overall performance across the three sub-images and the full-image. However, this test indicates that there is potential for much better performance provided the user is willing to tune the algorithm to a specific data set.

The circled regions in Fig. 6 illustrate additional cases where the edge strength model in PolarIRGS preserves class

boundaries that are ignored by PolarMLL because edge strength is not considered.

E. Full Image Results

The full image segmentation and classification was performed with 17 classes as the image appeared to contain more visually distinct classes than the previous sub-images. The full image results in Table III indicate lower classification accuracy and less consistency than previous cases. The authors know of no other papers that reported quantitative accuracies for an *unsupervised* segmentation and classification of the *full* Flevoland scene, so it is not possible to determine the relative performance of polarimetric IRGS against other algorithms for the full image.

There are several reasons for the lower performance in the full image results. As explained previously, an increased number of classes leads to a larger solution space with more possibilities for the algorithm to be trapped by local minima of (7). The full image is also much larger and contains more regions to be labelled, which again increases the solution search space. The larger spatial extent of the image encompasses more classes that can be confused with each other, such as the Potato and the Forest class, which appear very similar in the Pauli RGB image in Fig. 7(a). This is reflected in the classification result shown in Fig. 7(d), where many of the Potato pixels were grouped with the Forest class.

Table IV shows the mean confusion matrix [44] for the full image classification across 10 runs of PolarIRGS. This matrix was created by taking the element-wise mean of the individual confusion matrices for each run. This allows examination of which classes are frequently confused by PolarIRGS. Columns represent ground-truth classes. The numbers in each column indicate how many pixels from each ground-truth class are assigned to each class in the classification result. Table IV shows that Wheat 2 and Wheat 3 are frequently confused with each other. The high standard deviation of these two classes in Table III reflect the fact that under multiple executions of the algorithm, a pixel can be assigned to one or the other class, suggesting poor separability between them. Rapeseed is often erroneously assigned to the three Wheat classes, Grasses are assigned to Barley and Wheat 3, Potatoes are assigned to Forest and Water is assigned to Bare Soil and either one of the Excess classes.

Some of these segmentation and classification errors can be explained. There are several distinct signatures for Water in the scene. At the top right corner of the Pauli RGB image in Fig. 7(a), Water appears blue but dark elsewhere. The bright part appears to be a wind roughening effect that causes a higher backscatter return. Wheat 2 and Wheat 3 are poorly separable because they are actually one class (along with Wheat) that appear differently enough to be treated as different classes by the creator of the GT2 ground-truth image. Due to their similarity, there is no single classification that is the global minimum of the PolarIRGS energy function. As a result, any given region of Wheat 2 and Wheat 3 may be assigned to Wheat 2 or 3 in different runs of the stochastic algorithm and produces the confusion seen in the matrix. Finally, Potatoes

appear very similar to Forest. Du et al. [22] grouped them together as one class since they could not be distinguished in dual-polarization data. Other authors were able to distinguish between Potatoes and Forest in fully polarimetric data with a *supervised* algorithm [41] but supervised algorithms require training data that tell the system about the expected clusters, whereas PolarIRGS is fully unsupervised. Even in the more challenging case of full image unsupervised segmentation and classification, the segmentation boundaries in Fig. 7(c) visually match well with the image and there is little noise in the segments. This indicates the strength of the region-based spatial context model in IRGS.

When the classification results in the overlapping areas between Mask 1, Mask 2, Mask 3 and the full scene are compared, it can be seen that the classification results in these areas are not the same. This is because the labeling of each RAG region in a sub-image depends on the labeling of other regions in the sub-image due to the spatial context model. Therefore, when two images overlap, the classification result in the overlapping regions can be different due to the presence of other non-overlapping regions in the image, which affect the solution search space and the global cost function in (7). This is a fundamental property of spatial context models; the algorithm essentially considers the entire sub-image as a whole to find the classification result. This global approach helps PolarIRGS improve upon the results of other published techniques but does not guarantee that the overlapping areas are labelled the same way in each sub-image.

However, consistency is still maintained in a number of cases. Segments classified as Beet and some Grass segments are generally the same in overlapping regions in all three masks and the full image. Wheat is also consistent between the test cases if the three Wheat classes are considered one class. Segments classified as Potatoes are consistent between Masks 1, 2 and 3, although the Full Mask results identify the Potato segments as Forest. However, as noted earlier, Potato and Forest are difficult to distinguish. In areas where the full image classification is correct, Rapeseed, Bare Soil and Peas are consistent with the Mask 1 results.

F. Modified Full Scene Analysis

The observations made in Section V-E about the classification errors in the full scene classification motivate the need for a modified analysis. Poor separability and intra-class variation in polarimetric signatures make it difficult for an unsupervised algorithm like PolarIRGS to find the proper classification without *a priori* information such as training data that tell the algorithm what clusters to expect. The ground-truth image and the classification evaluation method were modified to account for these factors to obtain a more reasonable evaluation of PolarIRGS' performance.

The ground-truth image was modified so that only the Water in the top right corner of the image, where its signature is constant, was included in the evaluation. Wheat, Wheat 2 and Wheat 3 were grouped as a single Wheat class. Forest and Potatoes were also merged. Fig. 8(b) shows the modified ground-truth image. The classification accuracy was evaluated




Fig. 7. Classification results for the full image with 17 classes. Black or dark gray pixels are excess classes that are not associated with any ground-truth class.

on the same 17 class full image classification result from Fig. 7 but the three Wheat classes were manually merged into one class. Forest / Potatoes were also merged to form one class. The merging was done manually because there actually is more than one signature for the Wheat class and certain Potato and Forest patches are distinct. Thus, the algorithm cannot be controlled to automatically merge the classes. The classification accuracy was computed from the modified ground-truth and classification result.


Table V shows the result of this modified analysis. The accuracy has increased dramatically. Rapeseed, Grasses, Lucerne and Barley remain problematic but all other classes have high accuracies. The classification results are seen in Fig. 8(c)-(d). Accounting for the limitations inherent in the image data shows that PolarIRGS is capable of producing very reasonable classification results.

TABLE IV
 MEAN CONFUSION MATRIX [44] FOR THE FULL IMAGE CLASSIFICATION ACROSS 10 RUNS OF POLARIRGS, SHOWING THE GENERAL TRENDS REGARDING WHICH CLASSES ARE FREQUENTLY CONFUSED. COLUMNS REPRESENT GROUND-TRUTH CLASSES. THE NUMBERS IN EACH COLUMN INDICATE HOW MANY PIXELS OF THAT PARTICULAR GROUND-TRUTH CLASS WERE ASSIGNED TO EACH OF THE CLASSES IN THE CLASSIFICATION RESULT.

		Ground-truth Classes														
		Peas	Stembeans	Beet	Forest	Bare Soil	Rapeseed	Grasses	Lucerne	Potatoes	Barley	Wheat	Wheat 2	Wheat 3	Water	Buildings
Segmentation Classes	Peas	10044	0	51	0	0	1	0	0	0	0	0	0	0	0	0
	Stembeans	0	6298	1	1	0	2	0	0	17	0	0	0	0	0	36
	Beet	0	6	10703	0	0	21	8	0	29	0	0	0	0	0	0
	Forest	20	320	27	18543	0	1	1	2	11135	0	0	0	0	0	0
	Bare Soil	0	0	0	0	2765	0	0	0	0	0	0	0	0	5141	0
	Rapeseed	0	0	0	0	19	4706	12	11	55	0	0	0	0	0	0
	Grasses	0	0	0	1	115	0	2585	545	7	1610	0	0	1	383	0
	Lucerne	0	0	0	0	59	0	679	6224	3	173	0	0	419	0	0
	Potatoes	0	27	59	0	0	0	0	2	6859	0	0	0	0	0	0
	Barley	0	0	0	1	111	0	2505	2181	2	5574	0	0	20	0	0
	Wheat	30	6	11	0	0	2474	1	0	0	0	17612	0	0	0	0
	Wheat 2	5	0	0	0	177	4221	11	0	35	518	442	6949	4429	0	0
	Wheat 3	2	0	0	0	0	2359	1673	1555	15	0	0	4632	18279	0	0
	Water	0	0	0	0	0	0	0	0	0	0	0	0	0	5812	0
	Buildings	6	0	2	0	0	15	0	0	2	4	0	0	0	0	572
	Excess 1	0	192	667	2	0	0	1	0	2	0	0	0	0	2196	0
	Excess 2	0	0	105	0	61	193	0	1	9	173	0	0	0	1409	0




(a) Pauli RGB composite



(b) Ground-truth



(c) Classification result for 1 run of PolarIRGS



(d) PolarIRGS classification result, masked to ground-truth

Fig. 8. Classification results for the full image with 17 classes using the modified ground-truth and evaluation method of Section V-F. Black or dark gray pixels are excess classes that are not associated with any ground-truth class.

TABLE V

ACCURACY STATISTICS FOR POLARIRGS FOR THE FULL IMAGE USING THE MODIFIED GROUND-TRUTH AND EVALUATION METHOD OF SECTION V-F.

Algorithm	PolarIRGS	
Sub-image	Full Image	
# of Runs	10	
# of Classes	17	
Class	Mean Acc.	Std. Dev
Peas	99.4	0.6
Stembeans	92.0	7.2
Beet	92.1	9.1
Forest / Potatoes	99.5	0.1
Bare Soil	83.6	5.5
Rapeseed	33.6	7.1
Grasses	34.6	38.0
Lucerne	59.2	31.2
Barley	69.2	26.0
Wheat (combined)	99.2	2.5
Water	99.9	0.3
Buildings	93.9	1.4
OVERALL	85.8	3.0

VI. CONCLUSION

This paper has presented an extension of the IRGS algorithm to use a polarimetric feature model based on the Wishart distribution. PolarIRGS is unique among polarimetric segmentation and classification techniques in that it is a region-based algorithm that incorporates edge strength between regions as in its spatial context model and uses iterative region growing to aid in successful segmentation and classification. These properties allow it to compare favourably with and improve upon the results of other recent techniques. The region-based spatial context model is very successful at producing contiguous and noise free segments whose boundaries appear to match well with the ground-truth. PolarIRGS is well-suited for applications where classes have reasonable spatial coverage.

Future work could investigate improving the consistency of the algorithm between different executions for large scenes and scenes with a large number of classes. The bigger solution space associated with large scenes and more classes increases the likelihood for the current algorithm to be trapped in local minima. This can potentially be addressed by investigating whether the current Gibbs sampling optimization technique could be improved or replaced with a more robust technique if one exists. Another line of future work would determine whether improvements could be made by using measures analogous to the Fisher criterion for polarimetric data to set the value of the spatial context weighting parameter in the PolarIRGS energy function. Finally, the edge strength measure used in this paper only makes use of the amplitude information from the HH, HV and VV channels and future work could investigate whether polarimetric edge strength measures can offer improvements once a consistent set of parameters for the polarimetric edge strength filters can be found.

ACKNOWLEDGMENT

The GRSS-DFC provided the polarimetric image database. AIRSAR imagery courtesy SIR-C CDROM (A. Freeman - NASA/JPL). The authors would like to thank Prof. E. Pottier for some references and for suggesting the European

Space Agency's PolSARPro software, which was helpful for data conversion. The GT2 Flevoland ground-truth image was provided by Dr. S. N. Anfinson. The anonymous reviewers are acknowledged for their valuable feedback.

REFERENCES

- [1] B. Scheuchl, D. Flett, R. Caves, and I. Cumming, "Potential of RADARSAT-2 data for operational sea ice monitoring," *Canadian Journal of Remote Sensing*, vol. 30, no. 3, pp. 448 – 461, Jun. 2004.
- [2] K. Nakamura, H. Wakabayashi, K. Naoki, F. Nishio, T. Moriyama, and S. Uratsuka, "Observation of sea-ice thickness in the Sea of Okhotsk by using dual-frequency and fully polarimetric airborne SAR (Pi-SAR) data," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 11, pp. 2460 – 2469, Nov. 2005.
- [3] N. W. Park and K. H. Chi, "Integration of multitemporal/polarization C-band SAR data sets for land-cover classification," *International Journal of Remote Sensing*, vol. 29, no. 16, pp. 4667 – 4688, Jan. 2008.
- [4] B. Brisco and R. J. Brown, "Agricultural applications with radar," in *Principles & Applications of Imaging Radar*, 2nd ed., ser. Manual of Remote Sensing, F. M. Henderson and A. J. Lewis, Eds. New York: John Wiley & Sons, Inc., 1998, vol. 2, ch. 7, pp. 381 – 406.
- [5] J. S. Lee, M. R. Grunes, and R. Kwok, "Classification of multi-look polarimetric SAR based on complex Wishart distribution," *International Journal of Remote Sensing*, vol. 15, no. 11, pp. 2299 – 2311, Jul. 1994.
- [6] B. Scheuchl, R. Caves, D. Flett, R. De Abreu, M. Arnett, and I. Cumming, "The potential of cross-polarization information for operational sea ice monitoring," in *Proc. Envisat & ERS Symposium*, Salzburg, Austria, Sep. 2004.
- [7] J.-S. Lee, M. R. Grunes, T. L. Ainsworth, L.-J. Du, D. L. Schuler, and S. R. Cloude, "Unsupervised classification using polarimetric decomposition and the complex Wishart classifier," *IEEE Trans. Geosci. Remote Sens.*, vol. 37, no. 5, pp. 2249 – 2258, Sep. 1999.
- [8] E. Rignot and R. Chellappa, "Segmentation of polarimetric synthetic aperture radar data," *IEEE Trans. Image Process.*, vol. 1, no. 3, pp. 281 – 300, Jul. 1992.
- [9] Y. Wang, C. Han, and F. Tupin, "PolSAR data segmentation by combining tensor space cluster analysis and Markovian framework," *IEEE Geosci. Remote Sens. Lett.*, vol. 7, no. 1, pp. 210 – 214, Jan. 2010.
- [10] Y. Wu, K. Ji, W. Yu, and Y. Su, "Region-based classification of polarimetric SAR images using Wishart MRF," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 4, pp. 668 – 672, Oct. 2008.
- [11] F. Cao, W. Hong, Y. Wu, and E. Pottier, "An unsupervised segmentation with an adaptive number of clusters using the SPAN/H/α/A space and the complex Wishart clustering for fully polarimetric SAR data analysis," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 11, pp. 3454 – 3467, Nov. 2007.
- [12] K. Ersahin, I. Cumming, and R. Ward, "Segmentation and classification of polarimetric SAR data using spectral graph partitioning," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 1, pp. 164 – 174, Jan. 2010.
- [13] Q. Yu and D. A. Clausi, "IRGS: Image segmentation using edge penalties and region growing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 12, pp. 2126 – 2139, Dec. 2008.
- [14] —, "SAR sea-ice image analysis based on iterative region growing using semantics," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 12, pp. 3919 – 3931, Dec. 2007.
- [15] A. K. Qin and D. A. Clausi, "Multivariate image segmentation using semantic region growing with adaptive edge penalty," *IEEE Trans. Image Process.*, vol. 8, no. 19, Aug. 2010.
- [16] I. S. Barbosa and P. Maillard, "Mapping a wetland complex in the Brazilian savannah using an Ikonos image: assessing the potential of a new region-based classifier," *Canadian Journal of Remote Sensing*, vol. 36, no. S2, pp. S231–S242, 2010.
- [17] Y. Boykov, V. S. Lee, H. Rusinek, and R. Bansal, "Segmentation of dynamic N-D data sets via graph cuts using Markov models," in *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, ser. Lecture Notes in Computer Science, W. Niessen and M. Viergever, Eds. Springer Berlin / Heidelberg, 2001, vol. 2208, pp. 1058–1066.
- [18] J.-S. Lee and E. Pottier, *Polarimetric RADAR Imaging: From Basics to Applications*. Boca Raton: CRC Press, 2009.
- [19] M. S. Srivastava, "On the complex Wishart distribution," *Annals of Mathematical Statistics*, vol. 36, no. 1, pp. 313 – 315, 1965.
- [20] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. Upper Saddle River, NJ: Prentice-Hall, 2008.
- [21] S. Z. Li, *Markov Random Field Modeling in Image Analysis*, 2nd ed. Tokyo, Japan: Springer-Verlag, 2001.

- [22] L. Du, M. R. Grunes, and J. S. Lee, "Unsupervised segmentation of dual-polarization SAR images based on amplitude and texture characteristics," *International Journal of Remote Sensing*, vol. 23, no. 20, pp. 4383 – 4402, Oct. 2002.
- [23] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis and Machine Vision*, 3rd ed. Thompson Learning, 2008.
- [24] R. Kinderman and J. L. Snell, *Markov Random Fields and Their Applications*. Providence: American Mathematical Society, 1980.
- [25] Q. Yu, "Automated SAR sea ice interpretation," Ph.D. dissertation, Dept. of Systems Design Engineering, University of Waterloo, Waterloo, ON, Canada, 2006. [Online]. Available: <http://www.eng.uwaterloo.ca/~dclausi/students.html>
- [26] H. Derin and H. Elliott, "Modeling and segmentation of noisy and textured images using Gibbs random fields," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 9, no. 1, pp. 39 – 55, Jan. 1987.
- [27] L. Vincent and P. Soille, "Watersheds in digital spaces: An efficient algorithm based on immersion simulations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 6, pp. 583 – 598, Jun. 1991.
- [28] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 6, no. 6, pp. 721 – 741, Nov. 1984.
- [29] J.-S. Lee, M. R. Grunes, E. Pottier, and L. Ferro-Famil, "Unsupervised terrain classification preserving polarimetric scattering characteristics," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 4, pp. 722 – 731, Apr. 2004.
- [30] H.-C. Lee and D. R. Cok, "Detecting boundaries in a vector field," *IEEE Trans. Signal Process.*, vol. 39, no. 5, pp. 1181 – 1194, May 1991.
- [31] J. Schou, H. Skriver, A. A. Nielsen, and K. Conradsen, "CFAR edge detector for polarimetric SAR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 1, pp. 20 – 32, Jan. 2003.
- [32] W. Yang, C. He, Y. Cao, H. Sun, and X. Xu, "Improved classification of SAR sea ice imagery based on segmentation," in *Proc. IGARSS*, Jul. 2006, pp. 3727 – 3730.
- [33] C. J. Oliver, D. Blacknell, and R. G. White, "Optimum edge detection in SAR," *IEE Proc. Radar, Sonar Navig.*, vol. 143, no. 1, pp. 31 – 40, Feb. 1996.
- [34] K. Conradsen, A. A. Nielsen, J. Schou, and H. Skriver, "A test statistic in the complex Wishart distribution and its application to change detection in polarimetric SAR data," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 1, pp. 4 – 19, Jan. 2003.
- [35] S. Anfinsen, R. Jenssen, and T. Eltoft, "Spectral clustering of polarimetric SAR data with Wishart-derived distance measures," in *Proc. POLInSAR 2007*, Frascati, Italy, Jan. 2007.
- [36] F.-K. Li, C. Croft, and D. N. Held, "Comparison of several techniques to obtain multiple-look SAR imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 21, no. 3, pp. 370 – 375, Jul. 1983.
- [37] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. New York: Wiley, 2001.
- [38] X. Descombes, R. D. Morris, J. Zerubia, and M. Berthod, "Estimation of Markov random field prior parameters using Markov chain Monte Carlo maximum likelihood," *IEEE Trans. Image Process.*, vol. 8, no. 7, pp. 954 – 963, Jul. 1999.
- [39] P. Yu, "Segmentation of RADARSAT-2 dual-polarization sea ice imagery," Master's thesis, University of Waterloo, 2009. [Online]. Available: <http://hdl.handle.net/10012/4725>
- [40] "IEEE GRSS Data Fusion reference database, data set GRSS_DFC_0004," 2000. [Online]. Available: <http://www.dfc-grss.org/>
- [41] J. S. Lee, M. R. Grunes, and E. Pottier, "Quantitative comparison of classification capability: Fully polarimetric versus dual and single-polarization SAR," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 11, pp. 2343 – 2351, Nov. 2001.
- [42] D. A. Clausi, A. Qin, M. Chowdhury, P. Yu, and P. Malliard, "MAGIC: MAP-Guided Ice Classification System," *Canadian Journal of Remote Sensing*, vol. 36, no. Suppl. 1, pp. S11 – S13, 2010.
- [43] R. L. Rardin, *Optimization in Operations Research*. Upper Saddle River: Prentice Hall, 1998.
- [44] J. A. Richards and X. Jia, *Remote Sensing Digital Image Analysis*, 4th ed. Berlin: Springer, 2006.

PLACE
PHOTO
HERE

Peter Yu received the B.A.Sc. and M.A.Sc. degrees in Systems Design Engineering from the University of Waterloo (Waterloo, Ontario, Canada) in 2007 and 2009, respectively. From 2009 to 2010, he was a research associate with the Vision and Image Processing group, Department of Systems Design Engineering, University of Waterloo. Since 2010, he has been working in industry on image processing techniques for medical and remote sensing applications.

His research interests include image processing, pattern recognition, remote sensing and medical imaging. He was the recipient of the Canada Graduate Scholarship from the Natural Sciences and Engineering Research Council of Canada in 2007.

PLACE
PHOTO
HERE

A. K. Qin (M'05) received the B.Eng. degree from the Southeast University (Nanjing, Jiangsu, China) in 2001, and the Ph.D. degree from the Nanyang Technological University (Singapore) in 2007. During the Ph.D. study, he was with the Illinois Genetic Algorithms Laboratory at the University of Illinois at Urbana-Champaign (Urbana, IL, USA) as a visiting scholar for three months in 2005, and worked as a project officer at the Nanyang Technological University (Singapore) from 2006 to 2007.

From 2007 to 2009, he worked as a postdoctoral fellow in the Department of Systems Design Engineering at the University of Waterloo (Waterloo, Ontario, Canada). He is now with the MISTIS team at INRIA Grenoble Rhône-Alpes (Grenoble, France) as a postdoctoral researcher.

His research interests include pattern recognition, image processing, computer vision, evolutionary computation, neural networks, machine learning, remote sensing data analysis, biometrics and bioinformatics. He received the 2012 IEEE Computational Intelligence Society IEEE Transactions on Evolutionary Computation outstanding paper award. He is currently serving as the vice-chair of IEEE Emergent Technologies Task Force on Collaborative Learning and Optimization.

PLACE
PHOTO
HERE

David A. Clausi (S'93-M'96-SM'03) earned his B.A.Sc. (1990), M.A.Sc. (1992), and Ph.D. (1996) in Systems Design Engineering at the University of Waterloo (Waterloo, Ontario, Canada).

After completing his doctorate, Prof. Clausi worked in software medical imaging at Agfa (Waterloo). He started his academic career in 1997 as an Assistant Professor in Geomatics Engineering at the University of Calgary, Canada. In 1999, he returned to his alma mater and is now a Professor specializing in the fields of Intelligent and Environmental Systems as well as the departmental Graduate Chair. Prof. Clausi is an active interdisciplinary and multidisciplinary researcher. He has an extensive publication record, publishing refereed journal and conference papers in the diverse fields of remote sensing, computer vision, algorithm design, and biomechanics. His research efforts have led to successful commercial implementations including the building and selling of his own company. He was the Co-chair of IAPR Technical Committee 7 - Remote Sensing during 2004-2006.

He has received numerous scholarships, paper awards, and two Teaching Excellence Awards. In 2010, he received the award for "Research Excellence and Service to the Research Community" by the Canadian Image Processing and Pattern Recognition Society (CIPPRS).