A SIMPLE AND EXTENDABLE SEGMENTATION METHOD FOR MULTI-POLARISATION SAR IMAGES

Anthony P. Doulgeris

University of Tromsø, The Auroral Observatory, 9037 Tromsø, Norway

anthony.p.doulgeris@uit.no

ABSTRACT

This work presents a simple feature-based multi-channel SAR segmentation method that produces good, smooth, fast and robust results for image segmentation and interpretation. A basic set of six features produce good class distinction for general applications, and the method is easily extendable with new features and for multi-source data fusion. Optional steps of sub-sampling and Markov Random Field based contextual smoothing are presented. The different stages of the basic approach are demonstrated with real data examples from several sensors and for several applications and is fast enough for operational use.

Key words: Polarimetric; Synthetic Aperture Radar; Product Model; Segmentation; Contextual Smoothing.

1. INTRODUCTION

The objective of image segmentation is to segment, or partition, the image into several groups, of possibly disjoint regions, that have the same properties and presumably represent distinct land-cover classes. Polarimetric SAR images are difficult to interpret because the main information is obtained from a complex matrix with statistical properties that may be difficult to evaluate. This work tries to simplify the information extraction from this complex matrix into a general approach that may be used for effective image segmentation purposes. In addition, and possibly after the segmentation phase, the features may explored as a means of characterising, and potentially labelling, the different regions.

The general approach is to extract features with a sliding window technique that represent the polarimetric and statistical information in the multi-channel SAR image. The choice of window size affects the degree of smoothing and class distinguishing capabilities. Simple transformations, such as the logarithm, are applied during extraction to remove any obvious non-linear spreading of the data to produce features that appear well spread with simple visible clusters. This transformation stage is important because it allows the application of a simple segmentation or clustering method to identify separate classes based on data similarity. For example, a mixture of Gaussian clustering algorithm has achieved good results and is particularly fast when applying sub-sampling before full image classification. The addition of Markov random field based contextual smoothing greatly simplifies the segmented image result, although at the expense of some processing speed.

The principle behind the choice of features, the feature extraction and the feature transformation will be explained in detail. We demonstrate a basic set of five polarimetric features to extract from quad-pol data, which have produced good general results for many applications [1]. These features are an absolute multivariate brightness measure, the cross-pol fraction, the co-pol ratio, and the correlation magnitude and phase. The approach works consistently for dual-pol or single-pol data, although fewer features are available in those cases. We demonstrate how to include additional features such as non-Gaussianity (radar texture) measures, polarimetric decomposition parameter features, as well as advanced correlation and entropy measures. If the initial transformation of the parameters is carefully chosen, then the resulting simple segmentation achieves good results. In addition, the optional stages of sub-sampling and contextual smoothing will also be discussed and demonstrated.

This approach has been performing well for many real data applications such as urban, agriculture, forestry and sea ice analysis. Examples are shown for San Francisco and Vancouver Radarsat-2 scenes, because they may be familiar to many people, as well as for an ALOS PALSAR sea ice image from Barrow, Alaska. These demonstrate good, visually qualified, segmentation, but without validation. Validation would require ground truth data and be specific for each application or problem. However, the proposed segmentation and features may prove suitable as a starting point for such investigations.

The simple nature of the approach means that a new scene can be imported, features extracted and segmented with contextual smoothing in less than half an hour and therefore falls well within the near-real-time processing desired for operational use.

2. BASIC APPROACH

The basic workflow, summarised in Fig. 1, is to extract features with a sliding window method over the multi-channel single-look complex (SLC) PolSAR Image Data. It is assumed that all calibrations and terrain corrections are already applied to the data-sets. The polarimetric information is extracted from the local covariance matrix and can, thus, be extracted from MLC data-sets too. Statistical information regarding the radar textural properties are measured on the local neighbourhood collection of SLC vectors, with respect to the local mean covariance matrix. For MLC data, such a statistical property would have to be estimated from a local neighbourhood of MLC pixels.

SLC vector data: $\mathbf{s} = [S_{hh}, S_{hv}, S_{vh}, S_{vv}]^T$

MLC matrix data: $\mathbf{C} = \frac{1}{L}\sum_{i=1}^{L} \mathbf{s}_i \mathbf{s}_i^H$

2.1. Texture and the Product Model

Radar texture describes the concept of random variation around the class mean that is in excess of that expected for the simple Gaussian model for the complex vector, or equivalently, the Wishart model for the covariance matrix. This additional variation is evident in the non-Gaussian statistics often observed in real radar images and, thus, texture, or non-Gaussianity, may help with class distinction in certain terrain types . Radar texture is usually accounted for with the so-called "product model" that describes the vector or matrix data as the product of two independent random variables, one for *texture* and one for *speckle*, and may be found in many standard text books on radar, e.g. [1]:

$$texture \times speckle: \quad \mathbf{s} = \sqrt{\tau}\mathbf{g} \quad , \quad \mathbf{C} = \tau\mathbf{W}$$

where $\mathbf{g}$ multivariate complex Gaussian distributed, $\mathbf{W}$ matrix-variate complex Wishart distributed, and the τ texture random variable has its own positive-only scalar distribution that determines the final non-Gaussian model for the SAR data.

2.2. Extended Polarimetric Feature Space (EPFS)

We first introduced six basic features in [2] that have a reasonably general application for PolSAR analysis. They are referred to as the extended polarimetric features, because we have extended five basic polarimetric features from the covariance matrix with a feature for non-Gaussianity. All features are texture model independent, because the non-Gaussianity feature is an empirical measure of the sample distribution and the covariance matrix will be the same for all models under the product model scheme.
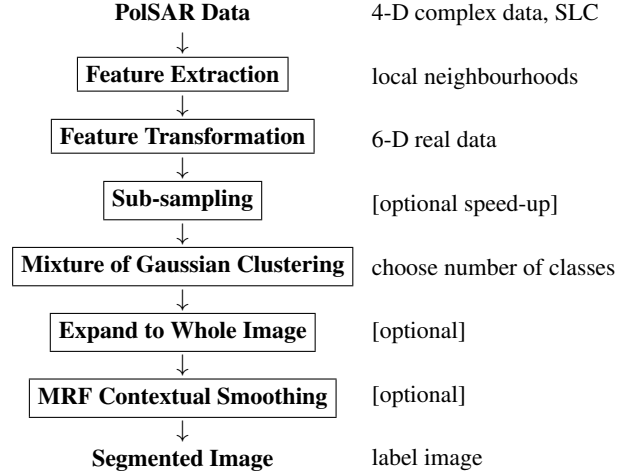
Basic set of six real features:

PolSAR Data	4-D complex data, SLC
Feature Extraction	local neighbourhoods
Feature Transformation	6-D real data
Sub-sampling	[optional speed-up]
Mixture of Gaussian Clustering	choose number of classes
Expand to Whole Image	[optional]
MRF Contextual Smoothing	[optional]
Segmented Image	label image

Figure 1. The basic workflow, with comments.

1. A non-Gaussianity measure: relative kurtosis RK

$$RK = \frac{1}{Nd(d+1)}\sum_{i=1}^{N}[\mathbf{s}_i^H \mathbf{C}^{-1}\mathbf{s}_i]^2$$

2. An absolute backscatter value: multi-variate radar cross section

$$MRCS = \sqrt[d]{\det(\mathbf{C})}$$

3. A cross-polarisation fraction or ratio:

$$R_{cr} = \mathbf{C}_{hvhv}/MRCS$$

4. A co-polarisation ratio: $R_{co} = \mathbf{C}_{vvvv}/\mathbf{C}_{hhhh}$

5. The co-polarisation correlation magnitude: $|\rho|$

$$\rho = \mathbf{C}_{hhvv}/\sqrt{(|\mathbf{C}_{hhhh}|\,|\mathbf{C}_{vvvv}|)}$$

6. The co-polarisation correlation angle:

$$\angle\rho = <\phi_{hh} - \phi_{vv}>$$

These features were chosen for several reasons: from a purely mathematical point of view, given some basic symmetry properties of backscatter from natural targets, and features that have shown good potential from the general literature. The absolute brightness measure for the covariance matrix was chosen to be the determinant rather than the span, because we felt that it better suited the geometric scaling of the product model. Presumably, the span would work in a similar way, but we have not tried to compare the two since the determinant has shown good results. This backscatter brightness feature is usually one of the most important features. The natural symmetries and scattering mechanism principles indicate that the cross-pol terms, HV and VH, are usually reasonably independent from the co-pol terms, HH and VV. Hence
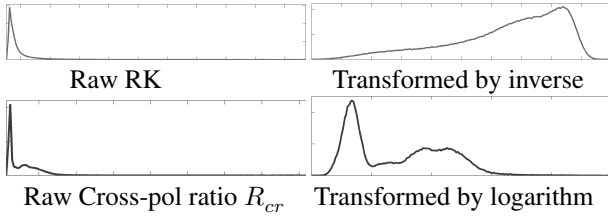
Raw RK | Transformed by inverse

Raw Cross-pol ratio R_{cr} | Transformed by logarithm

Figure 2. Example feature transformations to improve visualisation of the data spread and natural groupings.

we separate out this fraction. The co-pol ratio has relevance to single or multiple surface scattering and hence is a useful feature. Similarly, the co-polarisation correlation holds important dielectric information and has been found useful in many different studies. It is complex and both the magnitude and phase, or real and imaginary parts, may hold information.

2.3. Feature transforms and histograms

Exploring histograms of these features soon indicated that many of them were heavily skewed and it was difficult to see whether they held good class distinguishing information. Applying some simple monotonic transformations, such as the logarithm, inverse or square-root, soon revealed a wealth of hidden detail and made visualisation of natural groupings within the data space much clearer. The choice of transform was guided by any physical interpretation of the feature, e.g., power is best displayed on a log-scale, or from mathematical considerations, as in inverting the non-Gaussianity to compact the asymptotic trend from infinity to zero. Another consideration is that complex magnitude and phase do not separate well nor suit distance based discrimination measures because of the phase wrapping. By reverting the complex values back to the real and imaginary parts allows a 2-D distance to be meaningful near the phase wrapping boundary, such that a phase of 2π and 0 are similar, and near the origin, where phase has little meaning on distance.

The San Francisco Radarsat-2 sample scene from 2008 will be used as an example, because many people are familiar with it. Marginal histograms (normalised 0-1 scale) of all features are shown at the top of Fig. 3 and a 2-D scatter density plot is shown underneath for the features MRCS vs. co-pol ratio. The histograms show several clear groupings within many of the features and the scatter plot shows obvious globular clusters that can indicate at least four major class divisions, with details of several more sub-divisions. Images of the transformed features are shown in Fig. 4 and indicate good contrast and distinct areas that represent natural classes in the images.
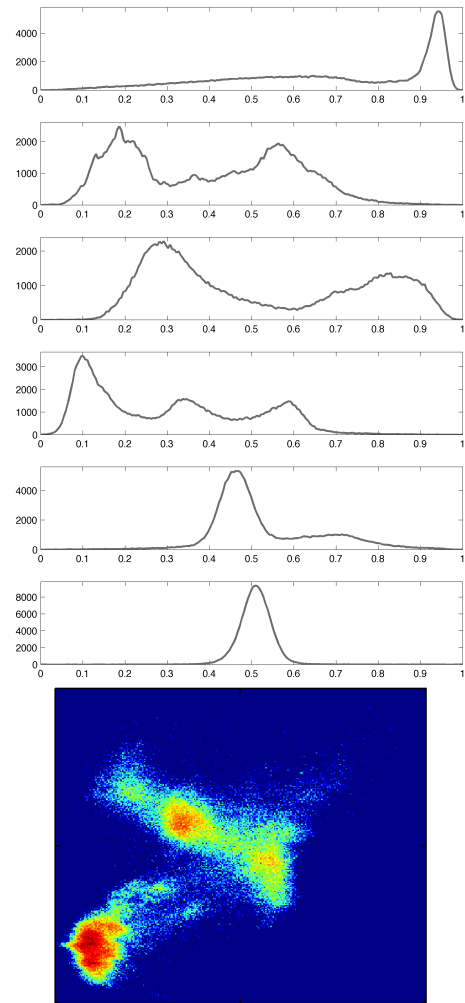

Figure 3. Simplified Feature Space: marginal histograms (top) and example 2-D scatter density plot, for MRCS vs. co-pol ratio, showing good natural class groupings in simple globular clusters.

2.4. Simple Mixture of Gaussian Clustering

The appearance of nice globular clusters in the feature space influenced the choice of a simple segmentation algorithm and suggests that a mixture of Gaussian clustering should achieve good results. Mixture of Gaussian clustering is also very fast, simple to understand and built-in to many analysis software packages. It has an advantage over the simpler k-means in that it can account for elliptical clusters, i.e., dissimilar and possibly correlated variation in different feature dimensions, and will certainly capture the observed density centres in the feature domain. The number of classes must be given in advance and may be guided by the visual peaks in the feature domain, or repeatedly applied (since it is so fast) for different numbers of clusters and a suitable resulting image chosen by visual inspection. The algorithm result is influenced by the initial conditions, so it is advisable to use multiple starts if implemented. Fig. 5 shows the
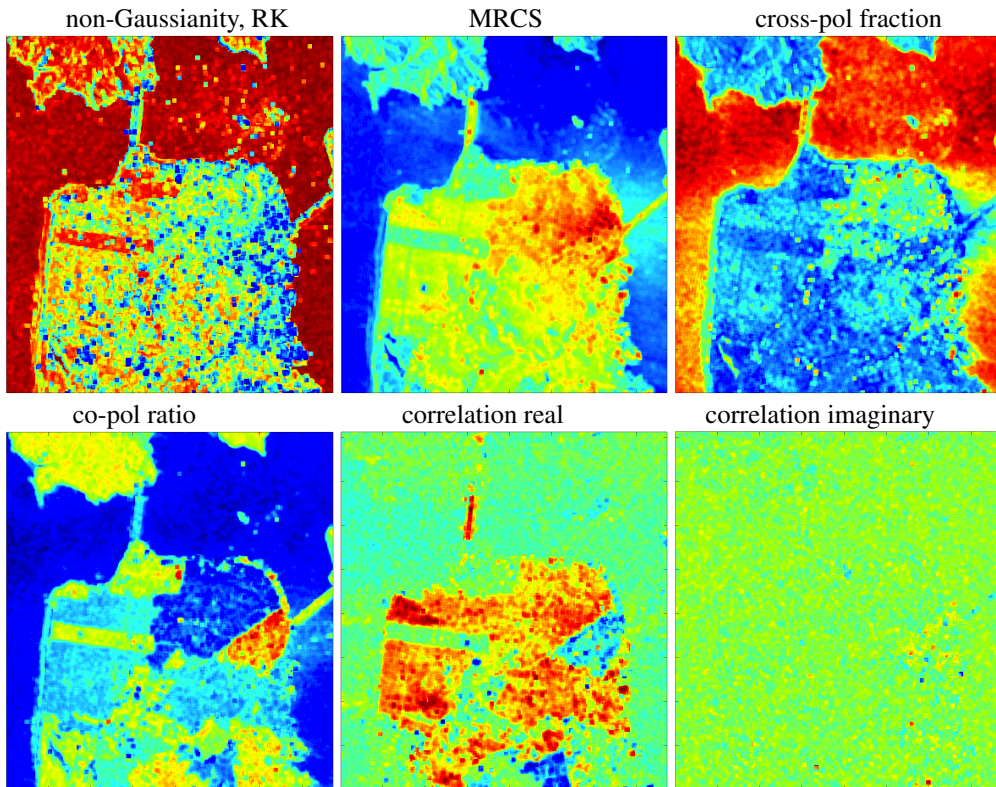
Figure 4. Example Feature Images - San Francisco. Many features have good distinguishing ability as noted by differently coloured areas. The last parameter shows little distinction for this image.

clustering results in both the feature space and the image space. It may be clearly seen that the coloured clusters in the feature domain match the natural density centres and that they also correspond to clear regions, and likely ice type classes, in the image domain. The window size is 24×3 and we chose 10 classes.

It is worth noting that the clusters are not perfectly Gaussian, nor symmetric, in profile, but at a coarse level, and particularly with sub-sampling of the data, the Gaussian model nevertheless achieves reasonable results. With increasing the number of clusters, it certainly finds the major, dominant classes before splitting clusters due to shape mis-match. Future work may explore non-Gaussian clustering or kernel methods, but the speed of the simple mixture of Gaussian clustering is highly desirable and worth some compromise. The mis-match in cluster profiles also means that some advanced methods to determine the number of classes through goodness-of-fit testing [3] cannot be applied.

2.5. Contextual Smoothing - MRFs

Contextual smoothing with Markov random fields (MRFs) is a rigorous image smoothing approach that accounts for the pixel-wise class probability of each pixel in the context of its neighbours. We use the basic isotropic, 8 neighbour MRF with a globally optimised smoothing parameter and the mean-field approximation [4, 5]. Essentially, the global class prior probabilities are replaced by local prior probabilities that depend on the neighbouring probabilities and can boost the likelihood of the local majority class throughout the image. This can easily be applied after the main mixture of Gaussian clustering phase, once the cluster parameters have been obtained, and is an iterative smoothing phase that allows the Markov property to spread throughout the image towards the globally optimum result. This stage may take many iterations, but can still process large scenes in only seconds or minutes.

The effect of the contextual smoothing can be dramatic for classes with highly overlapping probabilities. An example is shown in Fig. 6 for the San Francisco Radarsat-2 sample scene, with an 8×4 window size. Note the effect on the red and green classes, and between the brown, purple and cyan classes. The smooth, more solid regions are also more pleasing to the eye and simplify the interpretation of the segmented image.

2.6. Generic and Extendable to Data Fusion

This approach is generic in the sense that it is independent of the specific texture distribution or model, uses five generic features from polarisation matrix given basic symmetries that work well for many applications, can
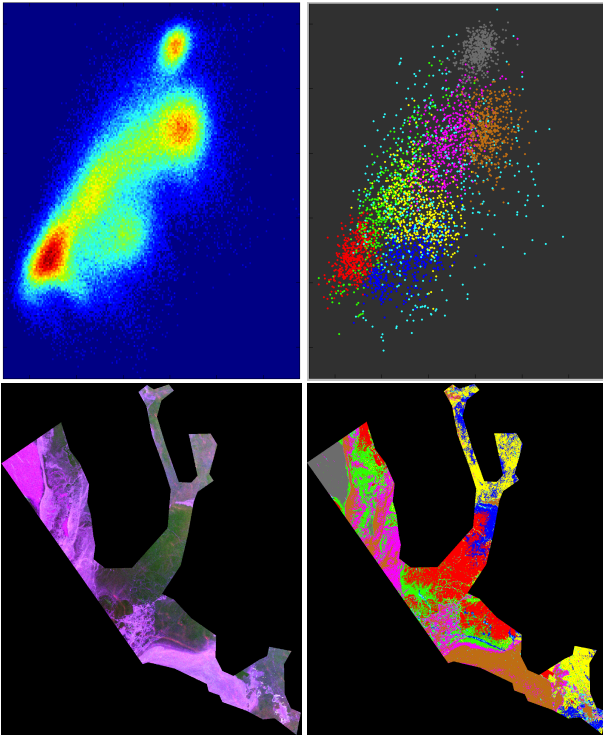
Figure 5. ALOS PALSAR, 2010, example of sea ice around Svalbard, with the land masked out. Feature space density (top-left) and the coloured clusters (top-right), image space Pauli RGB (bottom-left) and corresponding coloured classes (bottom-right). The feature domain clusters do segment visual regions in the image domain.

be applied to quad, dual or mono-pol data (although with reduced features, see Tab. 1), and can be extended in a consistent manner with new features.

Can use any suitably transformed real valued features, for example:

- Log-cumulants κ_2, κ_3 for texture

- Optical data (e.g., Intensity, NDVI)

- Directional / image texture

- Multi-scale / wavelet measures

- Polarimetric decomposition parameters

- Model based decompositions (e.g., RVOG)

3. RESULTS: FURTHER EXAMPLES

We have applied this method to many scenes and include three examples here to demonstrate that it works well for a variety of cases. Please note that we have not consistently applied geocoding or performed any radiometric
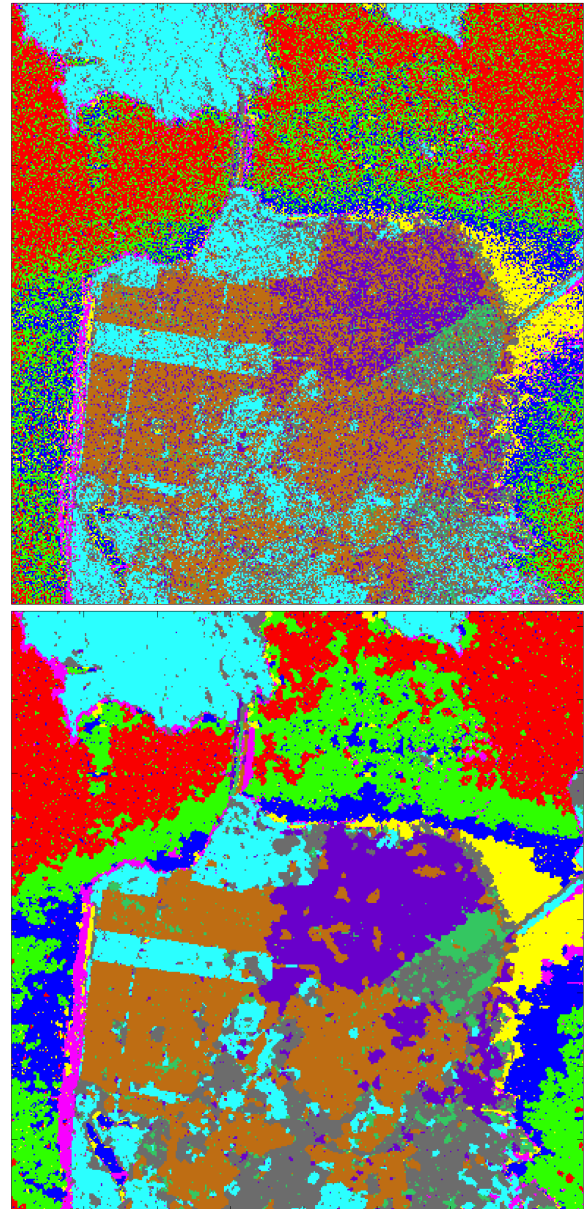

Figure 6. Contextual Smoothing Example with the San Francisco, Radarsat-2 scene, 8×4 window, 10 classes. After clustering (top) and with MRF smoothing (bottom).

terrain correction to these examples, they are simply presented as quick demonstrations. Different window sizes and number of classes would, of course, influence all of these results.

The first is again the San Francisco image, Fig. 7, but here with a slightly larger window size, 24×12, to achieve smoother, regional, urban results. The obvious water, vegetation and urban are well separated, and different urban classes are distinguished, which appear to correlate with either density or block orientation.

The second example, Fig. 8, is the Vancouver sample im-

	Quad	dual-co/cross	dual-co/co	mono
1. RK	+	+	+	+
2. MRCS	+	+	+	+
3. R_{cr}	+	+	-	-
4. R_{co}	+	-	+	-
5. $|\rho|$	+	-	+	-
6. $\angle\rho$	+	-	+	-

+ indicates available, - indicates not computable.

age, Radarsat-2 scene from 2008, with 8 classes and a 16×8 window. The central city region and the large airport structures are clearly distinguished in brown and dark-blue, the sub-urban areas are a mixture of cyan and grey, the agricultural fields and forestry are segmented as pink and yellow, and the water is primarily red and green with some dark-blue where the side-lobe ghosting effects of the city high-rise buildings are spread over the water.

A third example is of sea ice near Barrow, Alaska and is from an ALOS PALSAR scene from 2009. The lower quarter of the scene is land and includes frozen ponds and the township of Barrow. There is a region of land-fast ice above this and then some highly complicated dynamic ice floes. Even though it is quite complex, and is predominantly one material type, "ice", this simple approach has clearly distinguished many regions that we can visually distinguish by eye on the Pauli RGB image.

The fourth example, Fig. 10, is for a Sigma-nought, wide-swath, HH-HV, Radarsat-2 image of sea ice around Svalbard, with the land masked out. This demonstrates that although the dual-channel data has less features, three in this case, the general approach still segments the main regions and certainly the ice and water classes in the lower corner.

The final example, Fig. 11, depicts only the three entropy-alpha-anisotropy decomposition parameters and their segmentation. Obviously, three parameters does not describe the full polarimetric information (which has five free parameters) and these normalised terms do not include the important information of absolute backscatter intensity, therefore, they should not be expected to achieve optimal results by themselves. They are segmented here, by themselves, simply to demonstrate that they can easily be included in this overall approach and that any class distinction that they contain can be utilised for segmentation. They should, of course, be included in addition to other useful parameters and may improve some aspect of the overall image segmentation.

4. CONCLUSIONS

We have demonstrated a simple feature-based multi-channel SAR segmentation method that produces good,
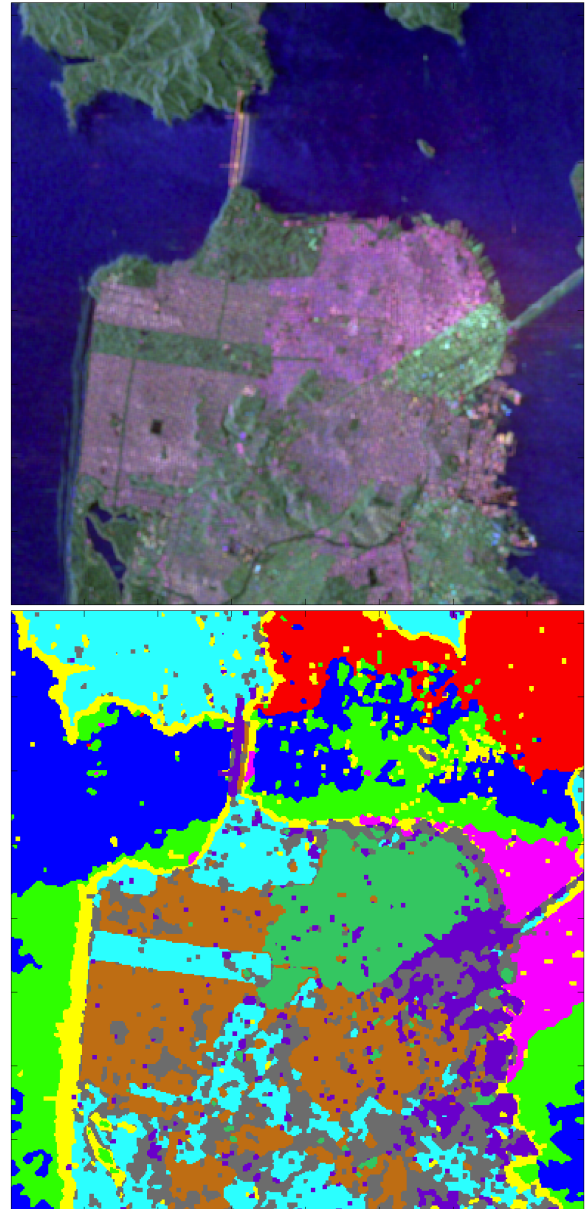

Figure 7. San Francisco, Radarsat-2 example: Pauli RGB (top) and 10 class segmentation with MRF, 24×12 window.

smooth, fast and robust results for image segmentation and interpretation. The method is quite generic and easily extendable with new features and for multi-source data fusion. The basic approach was demonstrated to work well for several different sensors and land-cover types. Of course, the best choice of features for different applications may be improved by further application specific studies, and be extended to use the class properties to identify and label the segmented image.

The main message of this presentation is that a good choice of features, hopefully capturing all of the polarimetric information, and some smart transformations, al-
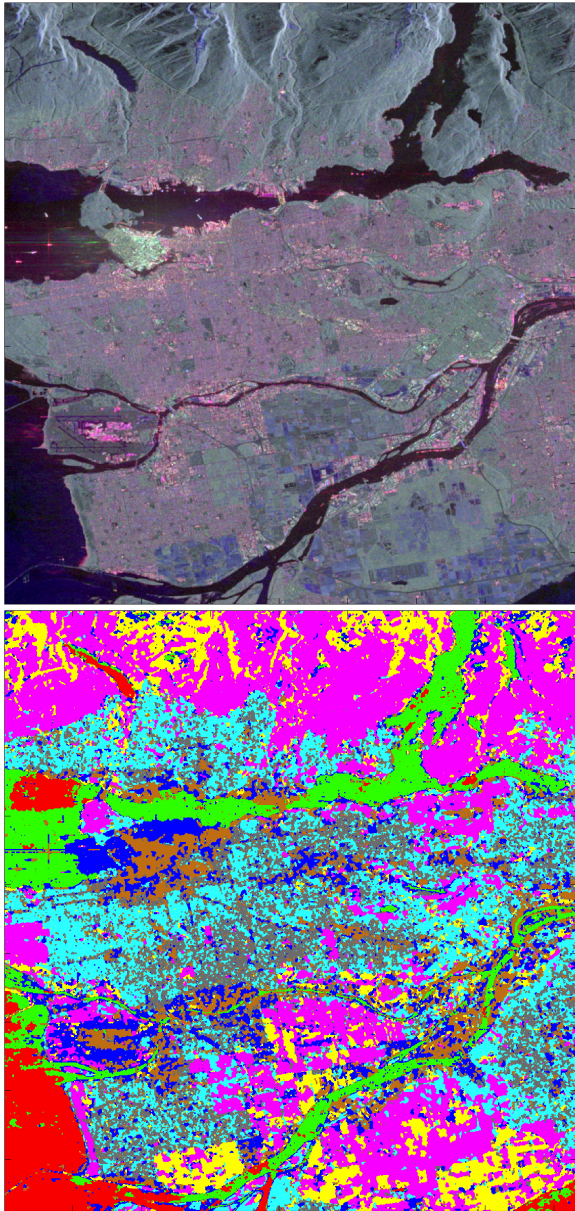
Figure 8. Vancouver, Radarsat-2 example: Pauli RGB (top) and 8 class segmentation with MRF, 16×8 *window.*

Figure 9. Barrow, sea ice example, ALOS PALSAR, 2009: Pauli RGB (left) and 9 class segmentation with MRF (right), 16×8 *window.*

lows a very simple segmentation approach to achieve satisfactory results. The advantage of this approach is its simplicity and speed, being able to segment full scenes into meaningful class segments in a matter of minutes, and is generic and consistent enough for operational use.

REFERENCES

[1] C. Oliver and S. Quegan. *Understanding Synthetic Aperture Radar Images*. SciTech Publishing, Raleigh, USA, 2nd edition, 2004.

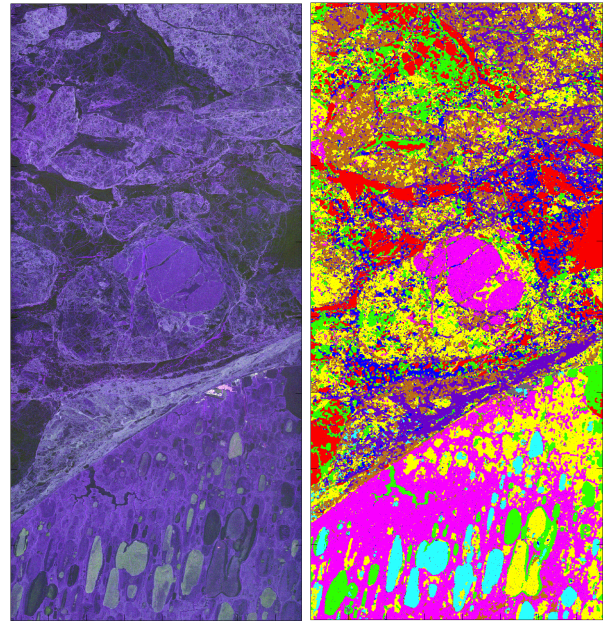[2] A. P. Doulgeris and T. Eltoft. Scale Mixture of Gaussian Modelling of Polarimetric SAR Data. *EURASIP Journal on Advances in Signal Processing*, 2010(874592):1–12, 2010.

[3] A. P. Doulgeris, S. N. Anfinsen, and T. Eltoft. Automated non-Gaussian clustering of polarimetric synthetic aperture radar images. *IEEE Trans. Geoscience and Remote Sensing*, 49(10):3665–3676, 2011.

[4] S. Li. *Markov random field modeling in image analysis*. Springer, 2009.

[5] G. Celeux, F. Forbes, and N. Peyrand. Em procedures using mean field-like approximations for markov model-based image segmentation. *Pattern Recognition*, 36(1):131–144, 2003.
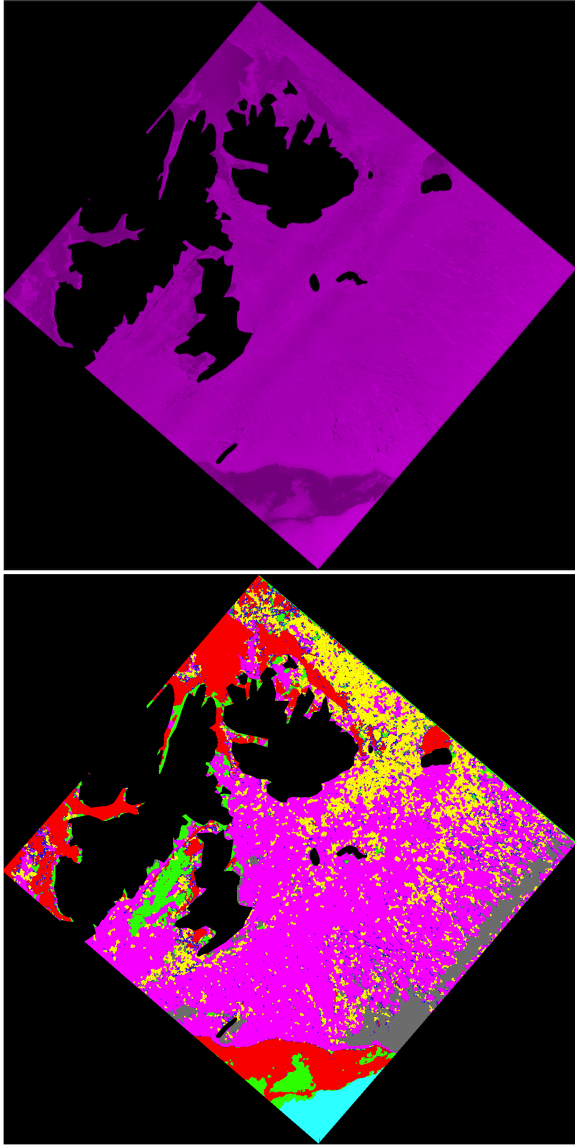
Figure 10. Wide-swath, sigma-nought example, Radarsat-2, 2012: Pseudo-Pauli (2 channel) RGB (top) and 7 class segmentation with MRF, 5×5 window.
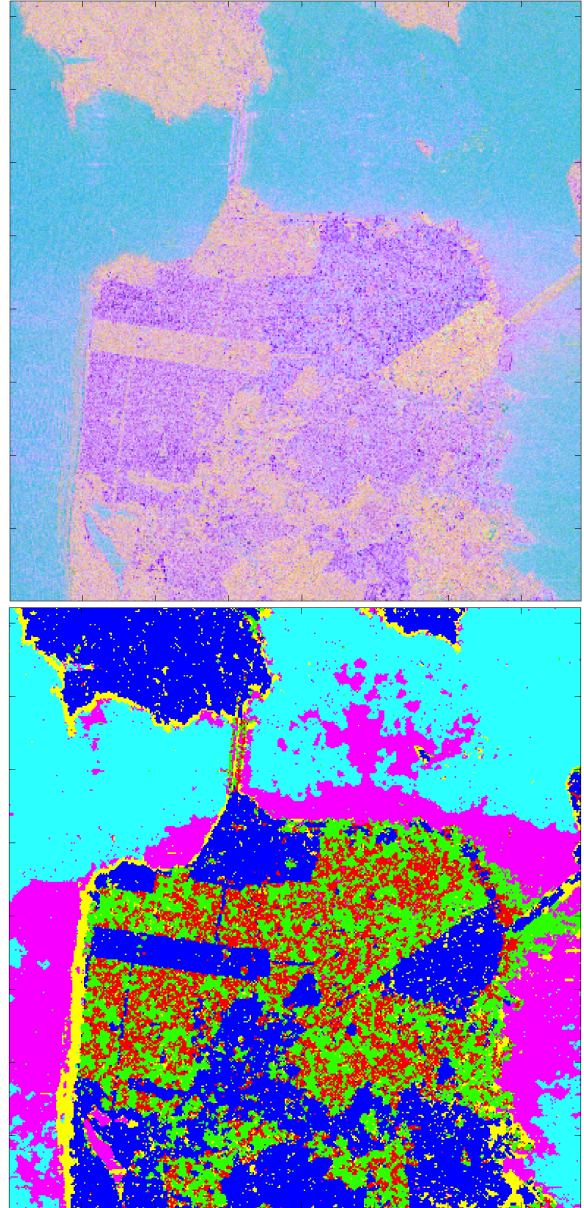

Figure 11. H-A-alpha example, Radarsat-2, San Francisco: H-A-alpha RGB (top) and 6 class segmentation with MRF, 5×5 window. Segmentated uses only the three decomposition parameters to highlight that they are suitable for this scheme.