Anais VIII Simpósio Brasileiro de Sensoriamento Remoto, Salvador, Brasil, 14-19 abril 1996, INPE, p. 567-572.

QUANTITATIVE ESTIMATION OF TROPICAL FOREST COVER BY SAR

SHAUN QUEGAN[1]
KEVIN D. GROVER[2]

[1]Sheffield Centre for Earth Observation Science, University of Sheffield, Sheffield S3 7RH, UK
[2]Department of Geography, University of Nottingham, Nottingham, UK

ABSTRACT- The capability of spaceborne SAR for monitoring tropical forest areas is assessed, using three ERS-1 images from the Tapajos region of Amazonia gathered in 1992 and a single JERS-1 image of the same area acquired in 1993. The multi-temporal ERS-1 data indicate that forest areas display a stable RCS, while in some case222 non-forest areas exhibit changes which appear to be associated with soil moisture variations, with greatest
discrimination between forest and non-forest after a dry period. Change detection techniques are compared and their ability to classify forest and non-forest are quantitatively assessed, assuming that a forest map inferred from a 1992 Landsat TM image is correct. Even in the best case, less than 50% of the non-forest region is detected using ERS-1. This figure may be improved by more frequent image acquisition, but there appear to be fundamental limitations in using C band data, since even comparatively low levels of standing biomass mask the soil moisture changes which drive the discrimination. At the longer wavelength of JERS-1, much better classification is possible, but the correction of topographic distortions is likely to present problems.

I INTRODUCTION The decreasing extent of the world's tropical forests due to anthropogenic activity is well documented (Skole and Tucker, 1993) and its possible impact on the global environment is an area of international concern. Spaceborne synthetic aperture radar (SAR) has the potential to provide timely information on the extent of these changes due to its near global coverage, regular revisit time and cloud cover penetration. This paper assesses the capabilities and limitations of SAR for tropical forest monitoring, using both multitemporal data from the ERS-1 C band SAR and also the longer wavelength L band data provided by JERS-1.

For ERS-1, a key observation is that areas of primary tropical forest exhibit very stable radar backscatter (Lecomte and Attema, 1993), while in some cases areas of pasture and bare soil show temporal variation associated with rainfall. A dataset from the Tapajos region of Brazilian Amazonia (described in Section II) is used to investigate the extent to which this behaviour can be exploited to provide a quantitative measurement of forest cover. This requires methods for identifying changes between images gathered at different times. In addition, pre-processing is necessary in order to combat image speckle. The first part of Section III describes and compares a variety of automatic methods to perform this task, and derives performance measures on the basis of a forest map extracted from Landsat TM data. The latter part of Section III provides a limited comparison with JERS-1, which makes clear the likely advantages to be gained by operating at longer wavelengths. Implications of these results and unresolved issues are discussed in Section IV.

II. TEST SITE AND DATA DESCRIPTION

The test site relevant to this paper lies to the west of the River Tapajos, at approximately 3° South and 55° West. It contains a range of ground cover types, with many cleared plots of land and regenerating forest of different ages adjacent to the primary forest in the protected Tapajos National Forest (FLONA Tapajos). Because of its use over a long period as a study area by INPE, there are large amounts of ground data available for this area, which were supplemented by measurements and an extensive land use survey carried out during fieldwork in August 1994.

The SAR analysis is based on three calibrated ERS-1 three-look PRI images gathered on 22/05/92, 31/07/92 and 18/12/92 and a single JERS-1 image from 26/06/93. The ERS-1 images are from the 35-day repeat cycle and hence have nearly identical geometry, which removes any severe registration problems. The data has resolution of the order of 9.7m × 18m (slant range × azimuth), with an equivalent ground range resolution lying between 22m at far range and 29m at near range, but were supplied at a pixel spacing of 12.5m × 12.5m (ground range × azimuth). JERS-1 provides three-look data with a resolution of 18m × 18m at the same pixel spacing as ERS-1. The incidence angle of the measurements is around 33° instead of the 23° of ERS-1. In addition, a relatively cloud-free Landsat TM image of Tapajos from 29/7/92 is available. All

Anais VIII Simpósio Brasileiro de Sensoriamento Remoto, Salvador, Brasil, 14-19 abril 1996, INPE, p. 567-572.

subsequent images of Tapajos have been too affected by cloud cover to be of use.

III. DATA ANALYSIS

A. ERS-1

The importance of multi-temporal data if ERS-1 is to be used in tropical forest monitoring is well illustrated by Figures 1(a) and 1(b). These show a section of the July and December images respectively, with radar cross-section (RCS) displayed in dB form. In the July image little structure is apparent, and the mean RCS exhibits only small variations across the region. Similar remarks apply to the May image. In December, large parts of the image have unchanged RCS, but other areas now show reduced backscatter (up to 3 dB), and structural components of the area are much clearer. Features to note are a large square area of farmland (bottom of the image, right of centre), the adjacent Santarem-Cuiaba highway running from top to bottom of the image and the two smaller highways perpendicular to it. A variety of cover types occurs along the highways, related to shifting agriculture and settlements. Most of the rest of the region shown is covered by forest.

Fieldwork established that the areas exhibiting reduced RCS in December were predominantly areas of pasture or low vegetation, and meteorological data indicate that the reduced RCS in December coincided with drier conditions (despite December being in the wet season and July being in the dry season for this region of Amazonia). Simple backscatter models are capable of explaining the observed behaviour (Grover et al, 1996), but in this paper we deal solely with the implications of these empirical observations for forest discrimination.

Figure 1 indicates that, in ERS-1 data, areas of non-forest may exhibit significant changes in RCS with time, which provides a means of discriminating forest. In order to detect change, differences of amplitude or intensity data are not useful because of the multiplicative nature of speckle, and image ratioing is to be preferred (Rignot and van Zyl, 1993). (Note, however, that this is equivalent to differencing of log images, which in practice is how the ratioing is carried out in the work described below.) This also has the advantage of removing the multiplicative modulation of RCS caused by variations in the incidence angle which can cause major problems in interpreting SAR data in hilly terrain. Since significant relief is a common feature in many tropical forest areas (though not a serious issue in Tapajos), this is important.

Direct ratioing of SAR images with only a small number of looks produces images which are unusable for classification, because of the effects of speckle (Rignot and van Zyl, 1993), and pre-processing is needed. We here compare four methods of reducing speckle before ratioing:

- block averaging;
- filtering using the local gamma maximum a posteriori (GMAP) filter (Lopes et al, 1993);
- global MAP estimation using simulated annealing (Geman and Geman, 1984; White, 1994);
- image segmentation (White, 1991; Caves and Quegan, 1995).

These four methods, described in more detail below, may be regarded as imposing progressively stronger models on the data.

The use of block averaging has been studied by Rignot and van Zyl (1993) who dealt with the simplest case, in which change occurs in two classes of equal area and the total error probability is minimised. In this case, a threshold of 1.5 dB is optimal for separating classes of which one remains constant and the other changes by 3 dB. Adopting the 1.5 dB threshold and using an 18×18 window (corresponding to approximately 240 independent samples after spatial correlation in the ERS-1 data is taken into account) would give a total error rate well below 1% for a 3 dB change in RCS. A more sophisticated calculation would allow for the fact that the areas of forest and non-forest are different, but this refinement is not pursued here. Figure 2(a) shows the effects of this operation, with white areas indicating where the July image exceeds the December image by at least 1.5 dB. Note the suppression of topographic features by this process, which is easily observed by comparing Figure 2(a) with 1(b).

A more refined approach to smoothing speckle is provided by the GMAP filter, which gives the local maximum a posteriori (MAP) estimate of the underlying RCS under the assumption that it is gamma distributed (Lopes et al, 1993). It is here applied in its locally adaptive form, which tests for structure (point targets, lines, edges) in the processing window before deciding which parts of the window are suitable for carrying out the RCS estimate. Where features are detected, less averaging is carried out, while lack of features and little underlying scene texture causes the filter to perform block averaging. Since its behaviour relies on theory developed for uncorrelated pixels, the images were pre-averaged by a factor of four to reduce spatial correlation between pixels. This gives rise to an approximately 5-look image containing nearly uncorrelated pixels. The pre-averaged images were then filtered using a 9×9 window, corresponding to 18×18 in the original data. The result of thresholding the difference of the log images at 1.5 dB is shown in Figure 2(b).

While the GMAP filter gives local estimates of the RCS, a constrained global MAP reconstruction of the RCS can be estimated using simulated annealing (Geman and Geman, 1984). Here we apply a version of this approach suitable for SAR images which was developed by White (1994). The July and December images were annealed separately after pre-averaging by a factor of four, as for the GMAP filter. The difference of the log of the July and December images after thresholding is shown in Figure 2(c).

A different approach to obtaining the underlying radar cross section is to first segment the image. The segmentation algorithm described in White (1991) has been modified to operate on multi-dimensional data (Caves and Quegan, 1995), and its modified form was used to provide a segmentation from all three available dates and also using just the July and December images. The operation is fully automatic and only requires the probability of an edge detection in pure speckle to be selected; this was set at 10^{-4}. The pixels within a segment were set to the mean backscatter within the segment and the effects of thresholding the difference of the ensuing log images is shown as Figure 2(d).

From Figure 2, it can be seen that all the methods pick out the large square area of farmland, as well as other areas of non-forest situated along the Santarem--Cuiaba highway and on the two other perpendicular roads, but there are marked differences in the results. Block averaging (Figure 2(a)) produces regions with little internal detail and smooth edges, as would be expected for such a large window size. Several regions of change are detected within the forested areas, possibly due to speckle fluctuations. The GMAP filter (Figure 2(b)) gives a very noisy and unsatisfactory result, with many detections of small isolated regions. This is qualitatively different from each of the other methods, where the change detections in the forest areas form larger regions. The performance of the filter did not improve markedly even when window sizes as large as 15×15 were used. Simulated annealing (Figure 2(c)) shows considerable detail in its detection of the main structural blocks, with retention of complex boundaries. Some of the smaller regions detected in the forested areas by averaging are also in the annealed image, but there are many differences in detail. Note the incomplete suppression of topographic effects; for example, the drainage feature to the right of the square of farmland (compare with Figure 1) gives rise to a number of small regions in Figure 2(c). Such effects could arise due to misregistration at the sub-pixel level. Segmentation (Figure 2(d)) gives results somewhere in between those of averaging and annealing as regards detailed structure, and appears to produce fewer isolated small features than any of the other methods. Note that it does not do as well as either averaging or annealing at detecting the area of change above the top right corner of the square pasture area.

B. Comparisons with TM

The performance of ERS-1 as a forest discriminator using these different approaches can be put on a quantitative basis by using the Landsat TM image from 29/7/1992. As a preliminary step, the TM and ERS-1 data were co-registered by resampling both images to a pixel size of 25m by 25m and then using ground control points to define a rotational and translational fit between the two images. The accuracy of the image registration has only a minor effect on the results discussed below, as long as the errors correspond to shifts by a small number of pixels. A forest/non-forest template was then prepared by thresholding band 5 of the TM data, to give Figure 3(a), in which 77.7% of the area is classified as forest. This allows a per-pixel comparison between the pre-processed ERS-1 data and Landsat TM to be performed, assuming that the TM map of forested areas is correct.

Results are summarised in Table 1. The first column shows, for each pre-processing method, the proportion of pixels detected as non-forest in the ERS-1 data classified as non-forest in the TM image. The second column shows the proportion of correctly detected forest pixels and the third column gives the area of false non-forest detections, as a proportion of the correct non-forest area. Certain non-forest areas detected by Landsat TM are not detected using ERS-1 for any of the pre-processing techniques (for example, the non-forest region protruding from the left side of the square area of farmland seen in the TM data of Figure 3(a) is not present in any of Figures 2(a) -- (d)). This is reflected in the low percentage of non-forest pixels found using ERS-1. Likely causes are significant soil roughness or secondary vegetation, both of which can lead to radar returns similar to primary forest at C band (Grover et al, 1996).

Averaging blocks of pixels causes loss of detail and poor definition of edges of the larger areas, as can be seen in Figure 2(a). Both these effects contribute to simple averaging giving the lowest percentage of non-forest detections (other than two-date segmentation). Annealing and three-date segmentation give similar performance; annealing does slightly better, but the percentage of non-forest pixels detected is only 45.7%.

Use of the GMAP filter appears from Table 1 to give better detection performance, but at the expense of significantly higher misclassification of forest pixels. Perhaps more important than the increased misclassification from GMAP is the nature of the errors. The fact that less than 50% of the non-forest pixels are detected by the other methods would seem to indicate a performance worse than randomly guessing at each pixel. However, the errors are clearly not spatially random, as is obvious from inspection of Figures 2(a), 2(c) and 2(d). By contrast, the errors in GMAP have much more of this spatially random quality, making this type of pre-filtering the least useful of those examined.

C. Processing Costs

An important issue in evaluating the pre-processing techniques is the computing power required. Table 2 shows the total CPU time for each of the techniques for a Sun Sparc10 machine. All methods started from images of 880 × 992 pixels; these were averaged by a factor of four for all methods except pixel averaging. Pixel averaging, GMAP and annealing all work on single frame images, but the times given correspond to processing two frames, which is the total time needed to form a difference of two log images. The segmentation algorithm works simultaneously on all channels of the multitemporal data, and inserts the mean value within each ensuing segment for each channel. The time given is for segmenting images from two dates. It is clear that the improved accuracy offered by annealing or segmentation is achieved at the cost of significantly increased machine time. Note that the time given for averaging is not for an optimised algorithm, and could be greatly improved. Recent work has indicated that the annealing algorithm could probably be made two or three times faster, but no ways of speeding up the segmentation code have yet been found.

In an operational context, the implications are that large-scale surveys to detect change can be carried out by using simple averaging, at the expense of lost detail and inability to detect small areas. Simulated annealing offers a more powerful approach which, after code optimisation, may be competitive with averaging, or which could be used to refine the classifications found by averaging.

D. Comparisons with JERS-1

A major limitation of ERS-1 is its short wavelength which prevents significant penetration into the vegetation canopy. As a result, even low vegetation or a comparatively young regenerating forest canopy may appear to the radar to be similar to primary forest. The age dependence of RCS has been studied for the Tapajos area by Yanasse et al (1995), who demonstrated that for large samples there are detectable differences, but that RCS is only weakly sensitive to age once biomass exceeds comparatively low levels (the saturation noted by several authors). At the longer wavelength of JERS-1, much greater penetration into the canopy occurs, and allows much more sensitive discrimination. The improved discrimination at L band is illustrated by Figure 3(b) which shows the JERS-1 data corresponding to Figure 1, after simulated annealing of a single image and thresholding at -9 dB. In the JERS-1 image, 25.8% of the area is detected as non-forest, as compared with 22.3% by TM. Comparison with Figure 3(a) indicates how closely related the detections of the non-forest areas are in the two sensors. The larger overall estimate of non-forest area by JERS-1 appears to be because many of the non-forest areas detected by JERS-1 are less fragmented than in the TM data (compare the highways at the top right and bottom left in the two images). The detection rates (corresponding to Table 1) are $p\{NF|NF\} = 0.841$ and $p\{F|F\} = 0.910$. False non-forest detections in JERS-1 are 12.4% of the true non-forest area determined from the TM data.

Figure 4 illustrates even more dramatically the different discriminating powers of the two SAR sensors. Figures 4(a) -- (c) show TM data, ERS-1 data and JERS-1 data respectively from a rougher area of pasture land than that dealt with in Figures (1) -- (3). The TM and JERS-1 data were treated as in Figure 3 and the ERS-1 image is the result of thresholding the difference of the July and December log images at 1 dB after simulated annealing. The most obvious feature to note is that the large area of pasture revealed in the TM and JERS-1 images is invisible in the corresponding ERS-1 images. This may reflect the presence of numerous small shrubs in this area of pasture, or different soil moisture conditions (see Grover et al, 1996). Also to be noted in this image is that JERS-1 finds areas of non-forest not in the TM image. The JERS-1 data was gathered a year later than the TM image, and this difference may correspond to real change, but we do not have information to verify this.

IV. DISCUSSION

The empirical analysis presented above illustrates a means of quantitatively comparing the performance of possible approaches to forest discrimination. Modelling suggests that for ERS-1 the physical parameter providing the discriminant is soil dielectric, and that the problems in its use as a forest discriminator reflect fundamental properties of the imaging process, namely that the radar signal is attenuated sufficiently, even by low canopies, to prevent detection of this soil signature. In addition, both observations and models indicate that RCS is only weakly sensitive to biomass, except at low biomass levels (Grover et al, 1996), so that even low vegetation canopies will give rise to returns very similar to those from a full forest canopy. Set against this, it must be noted that the ERS-1 images used in this study are the only ones available giving good coverage of the Tapajos test site, due to acquisition difficulties at the Cuiaba receiving station. It is not known whether a more complete set of data throughout the year might provide greater differentiation between the forest and non-forest areas. Acquisitions during more prolonged periods of dry weather may reveal greater effects. It is also possible that a regular time series of images may allow a satellite of the same type as ERS-1 to detect changes in forest boundaries, by locating regions where felling has taken place before significant regrowth occurs. Since images under wet and dry conditions would also be needed, the temporal sampling required would need to be based on RCS saturation time for regrowth, rainfall probability as a function of season and seasonal likelihood of forest clearance. This analysis is needed if an operational forest monitoring system based on C band is to be considered.

An important distinction between the use of JERS-1 and ERS-1 is that, at L band, modelling indicates that there is likely to be little variation in RCS as a result of rainfall. The use of ratio images to remove relief effects is therefore likely to be ineffective, and this could seriously affect the value of the data, unless other means of correcting for terrain effects on backscatter are used. The study area used in this paper has offered little chance to properly test how effective image ratioing is in removing relief effects. Assessment of both ERS-1 and JERS-1 in areas of higher relief would be very desirable.

V. ACKNOWLEDGEMENTS

We would like to thank colleagues at INPE for support and advice with this work and with data acquisition. We thank ESA for the supply of ERS-1 data and NASDA for the supply of JERS-1 data. We also thank Dr R G White of DRA Malvern, who is the originator of both the segmentation and simulated annealing algorithms used in this paper, Professor C J Oliver of DRA Malvern for permission to use the annealing software, and N A Software for providing this software. Kevin Grover was supported in this work by an NERC Studentship, reference no.ST/92/EOB/3.

References

Caves, R.G. and Quegan, S. The role of segmentation in multi-channel SAR image analysis. *Proceedings of RSS 95, 21st Annual Conference of the Remote Sensing Society, University of Southampton*, September, 1995, pp. 1171-1178.

Geman, G. and Geman, D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6): 721--741, 1984. Grover, K.D.; Quegan. S. and Yanasse, C. F. Use of C band SAR for quantitative estimation of tropical forest cover -methods and limitations. *IEEE Trans GRS (submitted).*

Lecomte, P. and Attema, E.P. Calibration and validation of the ERS-1 wind scatterometer. *Proc.First ERS-1 Symposium*, ESA SP-359, 19--29, 1993.

Lopes, A.; E. Nezry, E.; Touzi, R. and Laur, H. Structure detection and adaptive speckle filtering in SAR images. *International Journal of Remote Sensing, 14*: 1735--1758, 1993.

Rignot, E.J.M. and van Zyl, J. J. Change detection techniques for ERS-1 SAR data. *IEEE Transactions on Geoscience and Remote Sensing*, 31(4): 896--906, 1993.

Skole, D. and Tucker, C. Tropical deforestation and habitat fragmentation in the Amazon: Satellite data from 1978 to 1988. *Science*. *260*: 1905--1910, 1993.

White, R. G. Change detection in SAR imagery. *International Journal of Remote Sensing*. 12: 339-360,.

White, R.G. A simulated annealing algorithm for radar cross section estimation and segmentation. *Proceedings of SPIE International Conference on `Applications of artificial Neural Networks V*, April 1994.

Yanasse, C.C.F.; Sant'Anna, S.J.S.; Frery, A.C.; Kuplich, T.M.; Dutra, L.V. Mapeamiento de etapas de regeneracion y su relacion con datos de radar. In: *Simposio Latinoamericano de Perception Remota*, 7., Puerto Vallarta, MX, 05-10 Nov. 1995, Memorias, Selper, Instituto de Geografia, UNAM, MX, 1995, p.542-553.

	p{NF|NF|}	p{F|F}	p{F}p{NF|F}/p{NF}
Averaging (18 by 18)	0.420	0.942	0.156
GMAP (9 by 9)	0.467	0.931	0.180
Anneal	0.457	0.942	0.157
Segmentation (3 dates)	0.453	0.940	0.164
Segmentation (2 dates)	0.401	0.950	0.135

Table 1. Classification accuracies of the pre-processing methods

	CPU time (secs)
Averaging (18 by 18)	420
GMAP (9 by 9)	54
Anneal	2520
Segmentation (2 dates)	3467

Table 2. Computational cost

Figure Captions

Figure 1. (a) Section of ERS-1 PRI image acquired on July 31, 1992. (b) Corresponding image from December 18, 1992.

Figure 2. The effect of thresholding the difference of the July and December log images after four types of speckle reduction: (a) 18×18 block averaging; (b) GMAP filtering; (c) simulated annealing, and (d) segmentation.

Figure 3. (a) Forest/ non-forest map derived by thresholding band 5 of Landsat TM data acquired on 29 July, 1992. (b) Corresponding section of JERS-1 image acquired on June 26, 1993.

Figure 4. (a) Thresholded TM image containing an area of rough pasture land. (b) Corresponding ERS-1 image formed by thresholding the difference of the July and December log images at 1 dB after simulated annealing. (c) Corresponding JERS-1 image after simulated annealing and thresholding at -9 dB.