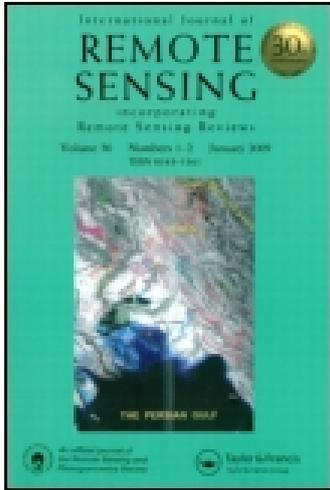


This article was downloaded by: [University of Western Ontario]
On: 07 March 2015, At: 08:55
Publisher: Taylor & Francis
Informa Ltd Registered in England and Wales Registered Number:
1072954 Registered office: Mortimer House, 37-41 Mortimer Street,
London W1T 3JH, UK



International Journal of Remote Sensing

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/tres20>

Flood boundary delineation from Synthetic Aperture Radar imagery using a statistical active contour model

M. S. Horritt , D. C. Mason & A. J. Luckman
Published online: 25 Nov 2010.

To cite this article: M. S. Horritt , D. C. Mason & A. J. Luckman (2001) Flood boundary delineation from Synthetic Aperture Radar imagery using a statistical active contour model, International Journal of Remote Sensing, 22:13, 2489-2507

To link to this article: <http://dx.doi.org/10.1080/01431160116902>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Flood boundary delineation from Synthetic Aperture Radar imagery using a statistical active contour model

M. S. HORRITT†, D. C. MASON

Environmental Systems Science Centre, University of Reading, 3 Earley Gate, Whiteknights, P.O. Box 238, Reading RG6 6AL, England, UK

and A. J. LUCKMAN

Department of Geography, University of Wales, Swansea, Wales, UK

(Received 28 September 1998; in final form 22 November 1999)

Abstract. Flood extent maps derived from remotely sensed data can provide distributed validation data for hydraulic models of fluvial flow, and can be used for flood relief management and to develop spatially accurate hazard maps. A statistical active contour model is used to delineate a flood from the first European Remote Sensing satellite Synthetic Aperture Radar (ERS-1 SAR) imagery as a region of homogeneous speckle statistics. The segmentation uses both local tone and texture measures and is capable of accurate feature boundary representation. The results are assessed by comparison with simultaneous aerial photography, the SAR segmentation scheme classifying 75% by area of the shoreline region correctly. Seventy per cent of the shoreline coincides with the ground data to within 20 m. The main error is due to unflooded vegetation giving similar radar returns to open water.

1. Introduction

Recent research has identified the potential of remotely sensed data to improve our understanding of hydrological systems, especially surface phenomena such as flooding. Reviews of the use of remote sensing in flood hydrology are given in Bates *et al.* (1997) and Smith (1997), both identifying the all weather capability of Synthetic Aperture Radar (SAR) imagery as a considerable advantage over other sensors operating at visible or infrared wavelengths of the electromagnetic spectrum which cannot penetrate the cloud cover often associated with flood events. Flood extent maps derived from remote sensing data can provide valuable distributed calibration and validation data for hydraulic models of river flow processes (Bates and Anderson 1995, Horritt 1998) and can be used for emergency flood relief management and to develop spatially accurate hazard maps.

The simplest model of SAR returns from flooded areas assumes that the water surface is generally much smoother than the surrounding dry land, and acts as a specular reflector, giving low backscatter. This simplistic model is reflected in many previous studies of flood mapping through SAR imagery, which generally adopt a qualitative

†e-mail: matt.horritt@bristol.ac.uk

approach (Badji and Dautrebande 1995, Calabresi 1995, Kannen 1995, Noyelle *et al.* 1995). Pre-flood, in-flood and post-flood imagery are combined in the red, green and blue channels of a false colour composite image, in which the low returns from the flood waters generate a magenta hue, which is identified and delineated manually. Imhoff *et al.* (1987) uses a simple thresholding technique for automatic classification of flooded regions in SAR imagery, and Oberstadler *et al.* (1997) use an evidential reasoning technique for automatic classification of the flood area. This technique has the drawback that the flood boundary is poorly defined, and so small scale flood features which the high resolution of SAR systems is capable of detecting are ignored.

The imaging of the water surface is complicated by wind roughening and the effects of protruding vegetation, both of which may produce significant radar returns. Multiple reflections between the water surface and upright vegetation enhance back scattering giving flooded vegetation a bright radar return in the image, the magnitude of this effect being a function of radar look angle, wavelength and polarization. Inundated forests have been identified with high returns in the L-band (Richards *et al.* 1987), but this effect is reduced at shorter radar wavelengths due to increased volume scattering in the canopy. Ormsby *et al.* (1985) found that the X-band gave bright returns for flooded marshland but no backscatter enhancement in forests. Henderson (1995) detected high radar returns from flooded forest and swamp and Ramsey (1995) found an inverse relationship between SAR returns and the depth of marsh flooding. Solomon (1993) describes backscatter from rivers in tropical forest areas and its variation with river orientation, with the water acting mainly as a specular reflector when looking along the river but giving higher returns when looking across due to multiple water/vegetation reflection. Wang *et al.* (1995) have modelled these interactions mathematically, their results predicting that for C-band SAR at an incidence angle of approximately 20° (the characteristics of ERS-1 and 2 SARs) flooding under Amazonian forest canopies should increase backscatter by approximately 2.6 dB. The effects of wind on ocean surface roughness and radar backscatter are also well known (Ulaby *et al.* 1986) but the effect on terrestrial water bodies is complicated by the variable fetch and effects of local topography. Due to the multiplicative nature of SAR noise, any increase in backscatter from the water surface (whether due to wind roughening or vegetation) will also produce a higher level of noise in the flooded region.

Despite these complications, the use of SAR imagery compares favourably with other remote sensing systems. Imhoff *et al.* (1987) compare Landsat and SIR-B (Shuttle imaging radar) images with a control set of colour and infrared aerial photographs of monsoon flooding in Bangladesh. SAR data processed with a simple density slicing and threshold technique provided a better identification of the flooded area (85% correct) compared to Landsat (64% correct), when compared to the air photograph control. Biggin and Blyth (1996) mapped a flood from SAR imagery of the River Thames, UK, validating the result against simultaneous aerial photographs. The images were analysed by eye with the SAR image correctly identifying approximately 80% of the flooded area.

The phase information in SAR imagery also has the potential to be used for flood mapping. Wegmuller *et al.* (1995) found that water could be identified as regions of low backscatter and low interferometric phase correlation between 2 SAR scenes. Corr *et al.* (1995) used coherence between ERS-1 images taken 3 days apart to distinguish between fields in winter and deciduous and coniferous forest, but achieved poor results for 35 day repeat cycles. Crop growth and farming activities tend to reduce phase

correlation, making classification more difficult. Coherence mapping also tends to generate results at a lower spatial resolution than the SAR image itself.

This research aims to rectify some of the current problems associated with flood mapping from SAR imagery. The development of an automatic segmentation algorithm allows rapid development of objective flood extent maps with accurately defined (to within ≈ 1 pixel) boundaries in order to make best use of the spatial resolution of SAR imagery. Comparison with simultaneous aerial photography data generates a measure of confidence in the maps derived from SAR imagery. The use of coherence mapping for flood detection is also assessed as an alternative to working with intensity data alone.

This paper is organized as follows. In §2.1, a statistical active contour model for SAR image segmentation is described, and the possibility of using coherence measurements for flood mapping is discussed in §2.2. Section 3.1 gives the results of applying these techniques to imagery of a flood on the river Thames, and compares them with flood maps derived from simultaneous aerial photographs. A Bayesian method of estimating classification accuracy is then developed in §3.2. Finally, the results are summarized and future research needs outlined in §4.

2. Methodology

2.1. Segmentation using a statistical active contour model

Active contour models or snakes (Cohen 1991, Williams and Shah 1992) have recently gained popularity as a means of turning incomplete and noisy edge maps into smooth continuous vector segment boundaries. The snake method uses a dynamic curvilinear contour to search the edge image space until it settles upon image region boundaries, driven by an energy function that is attracted to edge points, so that a snake will still settle on a line feature even if it is incomplete. The contour is usually represented as a series of nodes linked by straight line segments. The technique has been applied to SAR images of the North Sea shoreline on the east coast of the UK (Mason and Davenport 1996). Edges are detected using a constant false alarm rate contrast ratio technique to counteract the multiplicative nature of SAR speckle, the snake being used as the edge linking stage in a much larger image processing chain. Ivins and Porrill (1994) developed a statistical snake that operates on the image itself, rather than an edge image, dispensing with the need for a previous edge detection stage. It is suggested in their paper that mean image intensity can be computed for a number of pixels lying along the contour, rather than at a number of discrete points, which would give a level of noise reduction but still result in accurate delineation of segment boundaries. This would seem an ideal strategy for this problem, combating the effects of SAR speckle but fulfilling the aim of mapping the shoreline to within 1 pixel. This idea can be taken further, as variance can be calculated along the contour, allowing local texture to be taken into account in the segmentation process. It is stated in Luckman *et al.* (1997) that, 'At C-band, SAR imagery often exhibits little variation in mean amplitude between different types of natural landcover', and the failure of SAR amplitude to discriminate strongly between different surface features and cover types has led to texture being used as the main classifier in SAR imagery in many applications (Collins *et al.* 1997, Henebry and Kux 1997, Soares *et al.* 1997). The capability of the statistical snake to calculate image statistics over a number of pixels while maintaining well defined segment boundaries makes it a potentially powerful tool for SAR image processing for this application (and others), and worthy of further development.

A detailed description of the algorithm used here can be found in Horritt (1999),

and only an overview of the main features follows. The statistical snake is formulated as an energy minimization problem (Ivins and Porrill 1994), with the total snake energy $E(\mathbf{u}(s))$ being given by:

$$E(\mathbf{u}(s)) = E_{\text{tension}} + E_{\text{curvature}} - \iint_{\text{Inside Contour}} G(I(x, y)) dx dy \quad (1)$$

where $\mathbf{u}(s) = (x(s), y(s))$ describes the contour position (x, y) in the two-dimensional image space as a vector function of an arc length parameter, s . The contour is represented in a piecewise linear fashion as straight lines linking a series of nodes. E_{tension} and $E_{\text{curvature}}$ are energies generated by the model's internal tension and stiffness constraints, which favour a smooth contour made up of evenly spaced nodes (this enhances the numerical stability of the model). G is a goodness functional that assesses how well a set of image pixels $(I(x, y))$ meet certain criteria. The total energy is minimized if the contour encloses a large area of good pixels, and in this respect the model behaves as a region growing algorithm.

The driving term of the algorithm is the goodness functional which in this case is based on the statistical properties of small samples taken from a random population of image pixel intensity values. Image statistics are assumed to obey the simple speckle model (Touzi *et al.* 1988, Oliver *et al.* 1994), with the order parameter equal to the number of looks of the SAR system (three in the case of ERS-1 SAR PRI (PRrecision Image) data). The mean of small samples are assumed to be normally distributed (according to the central limit theorem), and numerical experiment validates this assumption for sample size $n \geq 5$ (see figure 1). The measured variance is assumed to be gamma distributed, with mean and order parameter found from empirical models again determined by numerical experiment (see figure 2).

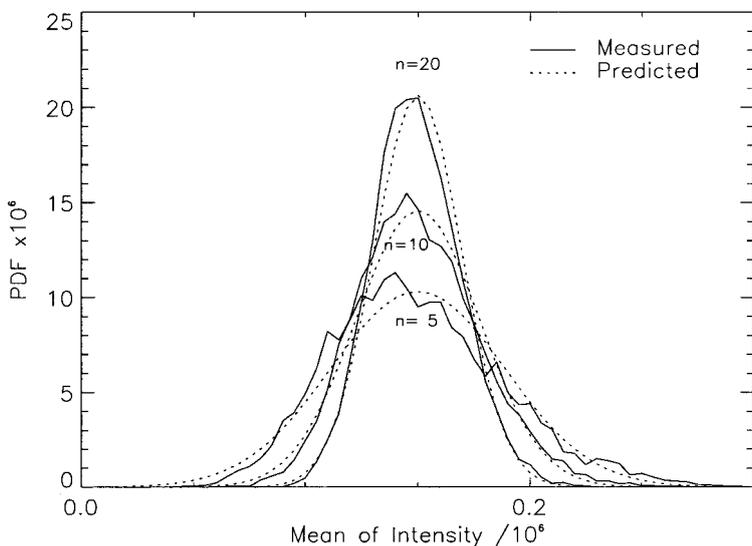


Figure 1. Results of numerical experiment to find the distribution of the means of samples of $n = 5, 10$ and 20 pixels taken from a gamma-distributed population, along with the normal distribution predicted by the central limit theorem. The mean and order parameter of the original population were chosen to mimic typical ERS-1 three-look amplitude imagery.

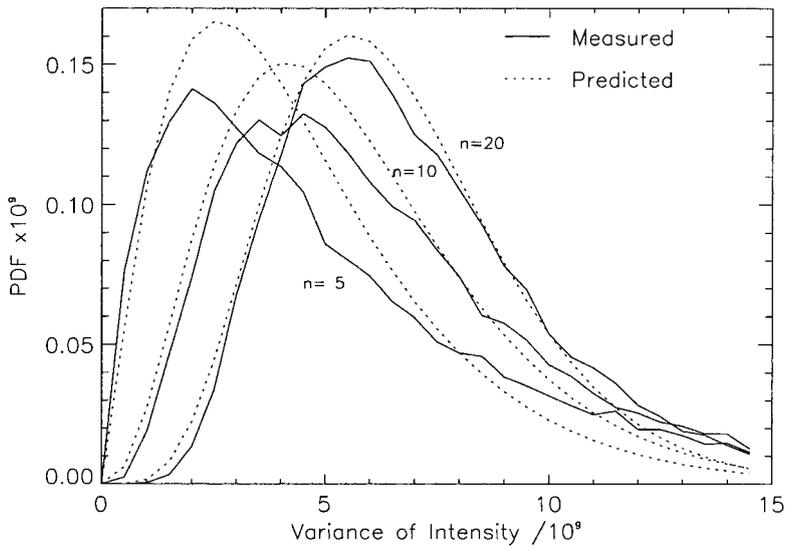


Figure 2. Results of numerical experiment to find the distribution of the variances of samples of $n=5$, 10 and 20 pixels taken from a gamma-distributed population, along with the distribution predicted by the empirical model.

If the mean and variance of pixel intensities along a segment of the contour are measured, we can use this knowledge of how the mean and variance behave to assess the probability that these pixels belong to the same region as that already inside the contour. The goodness functional G is then based on the log of this probability, the dependence on the measured mean μ' , for example, having the following form:

$$G(\mu') = 1 - \frac{n(\mu' - \mu)^2}{vk^2} \quad (2)$$

where μ is the mean of the seed population already enclosed within the contour, v is the variance of the seed population, n is the sample size and k is a parameter that can be adjusted to tune algorithm performance. This gives a measure which is equal to 1 for a group of pixels with the expected mean, falling to 0 for pixels with a mean differing by $k\sqrt{(v/n)}$ (i.e. k standard deviations) from the expected value. Taking the energy to equal the log of the probability has parallels with the simulated annealing formulation, where probability is exponentially related to the energy. The parameter k is set at around 2 or 3 but may be increased further to allow for a level of statistical inhomogeneity in the region being segmented. The dependence of the energy function on the sample size, n , means that the false alarm rate (a false alarm occurring when $G < 0$ even for a population with mean μ and variance v) is independent of sample size. A large value of n , however, does narrow the region for which $G > 0$, and so enhances the model's ability to discriminate between regions of different speckle statistics. Given this enhanced power of discrimination, the false alarm rate can always be reduced by increasing k . The overall goodness functional (with components based on both the measured mean and variance) is limited to a minimum value of -1 , as this makes interpretation of the model energy balance easier.

The roles of the tension and curvature constraints differ between the usual edge driven active contour model and the statistical version used here. If an edge detecting

snake 'snags' on an isolated (noise) edge point, the resulting high curvature, tension and inflation forces conspire to allow the snake to pass over the point. In this algorithm, the curvature and tension constraints are viewed as a means of producing a contour of appropriate smoothness with evenly spaced nodes, by a consideration of the balance between image and curvature forces. The magnitudes of these energies, which can be adjusted with a pair of weighting parameters, are usually determined by a process of trial and error (Cohen 1991, Williams and Shah 1992, Ivins and Porrill 1994), but here an analysis of the relationship between image and curvature/tension forces in the model yields appropriate values for the weighting parameters. The analysis places an upper bound on the curvature parameter for features of a certain length scale r (curvature energy scales as $1/r$), if the parameter is any larger, the curvature will dominate the model energy and smooth out features of this length. Similar analysis places an upper bound on the tension parameter (too large a value will favour a short contour and stifle the growth of the snake), although a much smaller value can be used to good effect.

Since the flooded region may not be simply connected (islands and isolated water bodies form holes and outliers), the algorithm incorporates a method for dealing with complex topology and snake self-intersection. An example of the algorithm behaviour is shown in figure 3, where the snake spawns a smaller sub snake to represent an island.

The snake is represented as a series of nodes, each having a discrete location on the image grid. The energy minimization scheme used is the greedy algorithm (Williams and Shah 1992). At each iteration, change in energy, dE , is computed for moves to all eight neighbours of each node in turn:

$$dE = -GdA + dE_{\text{tension}} + dE_{\text{curvature}} \quad (3)$$

Obviously dE is equal to zero for no node movement. The lowest (most negative) dE is chosen. The goodness functional G is calculated for pixels along the lines linking the node with its two neighbours, and dA is the local change in area. If G is positive, the snake is in a region of good pixels, a positive dA is favoured and the snake expands. If G is negative, a negative dA is favoured, and the snake retreats from the bad pixels. The statistics of the seed population, μ and σ , are calculated from all pixels lying inside the contour, every 10 iterations. Where the contour is artificially constrained by the image boundary, the curvature constraint is switched off, allowing the snake to expand freely along the edges of the image.

2.2. Coherence mapping

An algorithm based on backscatter intensity may have difficulty in situations where wind- or turbulence-induced waves increase backscatter from water to the

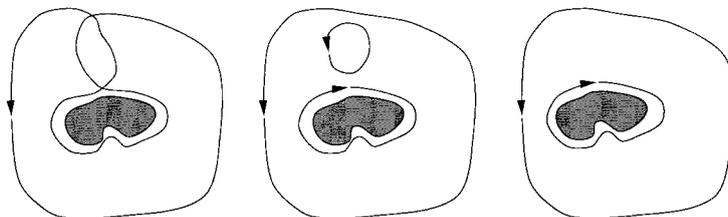


Figure 3. The contour adapts its topology to describe an island. The snake meets itself, spawning two new sub-snakes, one of which is recognized as an artefact and destroyed.

level of that from surrounding non-flooded areas. However, because the water surface is in constant motion, it should exhibit strong temporal decorrelation of interferometric phase. A phase coherence map constructed from a pair of SAR images taken at different times should exhibit very low coherence over water bodies. Provided that coherence from surrounding non-flooded areas remains sufficiently high, this could provide an additional means of distinguishing flooded from non-flooded areas (Wegmuller *et al.* 1995, Smith and Alsdorf 1997). Surfaces which change little with time, such as bare soil, rock and buildings, tend to exhibit high coherences independent of the time between overpasses. The coherence measured for vegetated surfaces, however, decays strongly as the time between overpasses increases as growth destroys the correlation between the two scenes. Corr *et al.* (1995) measured coherence between ERS-1 images taken 3 days apart (in winter, when vegetation growth is minimal) and found high coherence values for fields, which could be distinguished from forest because volume scattering gave much lower coherences for the latter. However, images with a 35-day repeat cycle also gave much lower coherences for vegetation.

3. Results and analysis

3.1. The snake algorithm and coherence mapping: comparison with airphoto data

Two reaches of the river Thames, both approximately 15 km in length, were selected as test sites for the snake algorithm (see figure 4). The dataset comprises three ERS-1 SAR images, covering both test sites and acquired on three dates (see table 1), with the image of 4 December 1992 capturing a flood event. The pre-flood and flood SAR pair have a baseline appropriate for interferometric processing. A set of oblique colour aerial photographs (Biggin and Blyth 1996), simultaneous with the flood SAR image (to within 2 hours), has also been acquired.

The flood shoreline was delineated by eye from the aerial photographs and vectorized. The shoreline vectors were then georeferenced using an orthographic

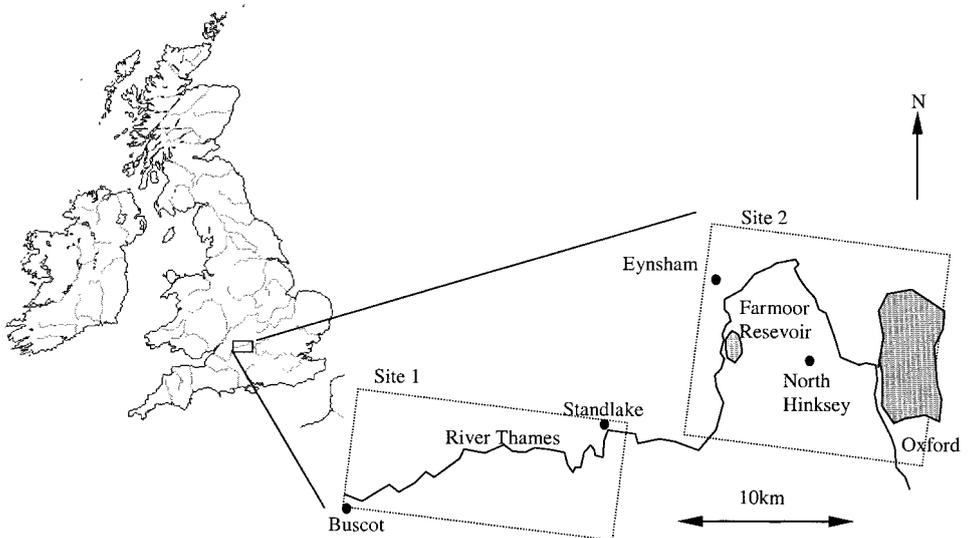


Figure 4. Schematic map of the Thames upstream of Oxford, showing the two SAR test sites.

Table 1. Details of ERS-1 SAR imagery used in the investigation.

Date	Time (UTC)	Frame	Orbit	Track	Type
30 Oct 1992	11:01	2565	06752	00094	descending
4 Dec 1992	11:00	2565	07253	00094	descending
12 Feb 1993	11:00	2565	08255	00094	descending

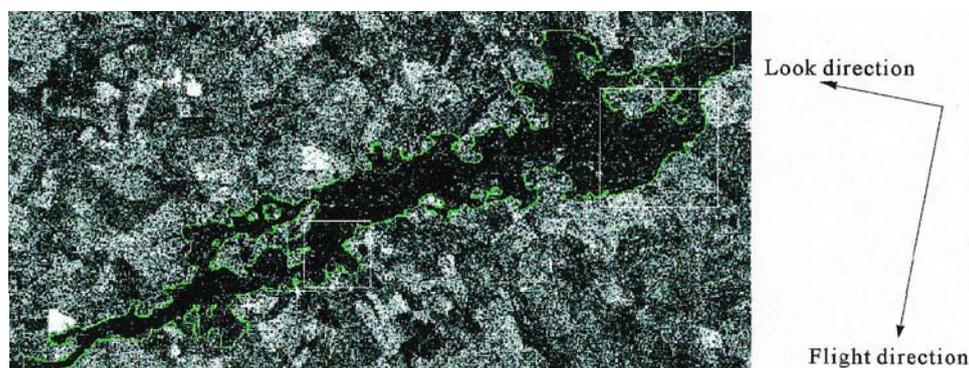


Figure 5. Results of applying the snake algorithm to the first test site. The sub-regions shown in figures 7 and 8 are also shown (upper right and lower middle, respectively). © European Space Agency.

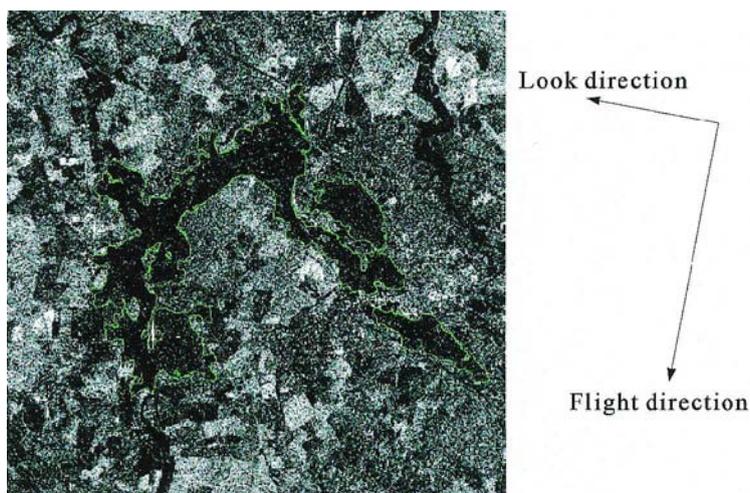


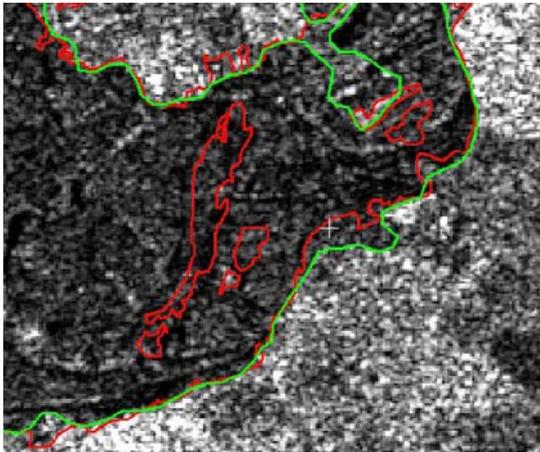
Figure 6. Results of applying the snake algorithm to the second test site. © European Space Agency.

transform parametrized by a least-squares method (Slama 1980) from 15–20 ground control points for each photograph. Since some segments of the shoreline were visible from more than one photograph, these could be used to assess the error (from both georeferencing and delineation) in the shoreline location. Where this occurred, the shoreline taken from the nearest field view was taken as the most accurate. The error could also be verified from segments where the shoreline was observed to lie

alongside a hedgerow or field boundary which could be located on 1:25 000 scale maps, and was found to be less than 20 m.

The results of applying the snake algorithm to the in-flood SAR imagery are shown in figures 5 and 6 for the two test reaches. In applying the algorithm, a two-stage strategy was adopted, the first stage using a nodal spacing of 12 pixels, no tension or curvature constraints and with the texture component of the goodness functional switched off, lasting for 200 iterations. The snake was started as a narrow strip lying along the course of the river channel, ensuring that it contained only flooded pixels. The first stage results in an approximate depiction of the flood boundary. The second stage used a nodal spacing of six pixels, and tension and curvature constraints that allow the snake to depict features with a radius of half the nodal spacing (this is generally the lowest useful value), and stabilized after 170 iterations for the first reach and 260 for the second. The constant k was set at 3.0 in both stages to allow for some statistical inhomogeneity in floodplain SAR statistics.

Figure 7 shows the 3 km \times 3 km detail highlighted in figure 5, with the shorelines



(a)



(b)

Figure 7. Shorelines derived from aerial (red) and SAR (green) imagery for a 3 km \times 3 km area, overlaid on the SAR imagery itself (a). Aerial photograph of the region, looking north (b). © European Space Agency.

found from SAR imagery and the aerial data overlaid on the SAR imagery and an aerial view of the area. The scene shows well the sort of backscatter increase from flooded vegetation reported in the literature, for example in the system of hedgerows present in the meander region (top right). Another problem is that some areas of dry land give the same returns as the flood, the island in the middle of the flood, for example. Figure 8 (the second region highlighted in figure 5) illustrates this effect for another dry field, along with an aerial photograph of the area. Bare or sparsely vegetated fields (seen as brown in the photograph) exhibit the high backscatter associated with wet soils (Giacomelli *et al.* 1995, Ulaby *et al.* 1996). A vegetated canopy will reduce these returns, simple water cloud models (Prevot *et al.* 1993, Taconet *et al.* 1996) predicting that the effect will increase with canopy water content. This is borne out by the green fields in the scene which correspond with areas of low SAR backscatter, and are likely to be misclassified as flooded.

Figure 9 shows a coherence map of the first test site (i.e. corresponding with the image shown in figure 5) with the vector shoreline found from the aerial photography. The coherence map (≈ 80 m pixel size) was formed from the flood and pre-flood

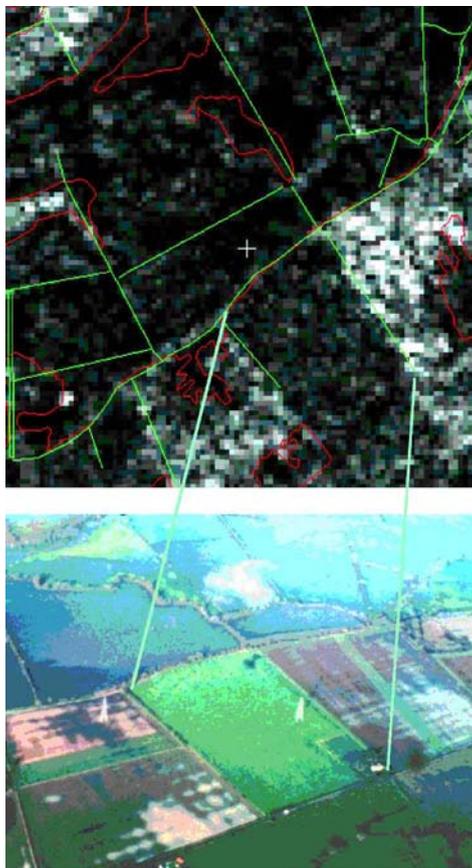


Figure 8. SAR and aerial view of a shoreline region, showing a vegetated field giving similar radar returns to those from the flood. The shoreline derived from aerial photography (red) is shown on the SAR imagery, along with field boundaries (green). © European Space Agency.

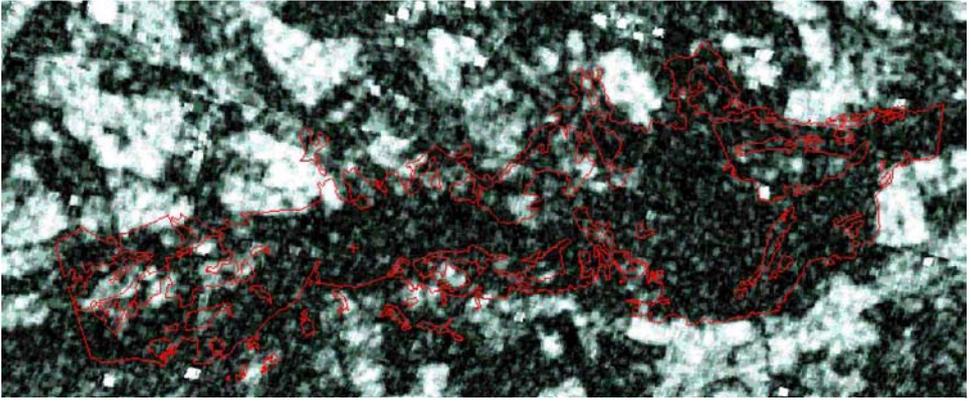


Figure 9. Coherence map for the first test site with the shoreline derived from the aerial photography in red. © European Space Agency.

SLC (Single Look Complex) images, which were found to have a baseline suitable for coherence mapping. Interferometric processing was carried out using the Atlantis Earthview InSAR software. The flooded region does indeed show consistently low levels of coherence, but so do many other regions of the floodplain, in agreement with the findings of Corr *et al.* (1995) for imagery acquired under a 35-day repeat cycle regime. Thus coherence mapping is of little use for this event.

3.2. Development of quantitative measures of algorithm performance

We first adopt a region based Bayesian approach in comparing the aerial and SAR-derived vector shorelines, evaluating the probability of correct classification from the SAR imagery using the aerial photography as ground data. A Bayesian method allows us to determine classification accuracy given additional information, such as landcover type and pre- or post-flood imagery, and to give a spatially variable measure of confidence in the SAR classification.

Care is needed, however, in fixing the region over which this probability of correct classification is defined. Considering the whole test site, for example, may bias the probability of dry regions having been identified correctly towards 1. For a small flood in a large domain, classifying the whole region as dry is a reasonable guess and would be considered a good classification in terms of the probabilities. This problem can be avoided if only the shoreline region is considered by limiting the analysis to a strip containing the flood boundary, but this raises the question of how wide to make this strip. A narrow strip would exclude many misclassified areas, and a strip that was too wide would suffer the same problems as using the whole domain in the analysis. A sensible compromise would be to make the strip width equal to typical floodplain field scales, so that it can include misclassified fields near the shoreline but without introducing too much bias. Results are presented in table 2 for strips 300 m and 600 m wide along the shoreline and the whole domain for the first test site and a strip 300 m wide for the second test site. These show only small differences between the two sizes of strips, but a larger difference between statistics for the shoreline region and the whole domain, which are biased towards 100% as expected. The results imply that the statistics are reasonably stationary with respect to strip width in the region 300–600 m, and can be used to predict the classification accuracy of the SAR segmentation algorithm in the shoreline region. They also show

Table 2. Statistics of SAR algorithm performance

Class	SAR status	Aerial status	Site no. 1300 m strip (%)	Site no. 1625 m strip (%)	Site no. 1, full domain (%)	Site no. 2, 300 m strip (%)
1	Wet	Wet	37.5	38.5	22.6	33.3
2	Dry	Dry	37.5	42.8	67.9	40.8
3	Wet	Dry	15.5	11.2	5.4	19.7
4	Dry	Wet	9.6	7.6	4.0	6.2
Correct			75.0	81.2	90.5	74.1

that the algorithm performs similarly in the two test reaches, identifying $\approx 75\%$ of the area correctly.

It is possible that certain landcover types are more prone to misclassification than others (because forest canopies may be impenetrable at C-band radar frequencies, for example). Landcover classification is available for most domestic floodplains in the form of the landcover map of Great Britain, developed from multitemporal Landsat Thematic Mapper data (Fuller *et al.* 1994). The virtually guaranteed availability of this data would allow the application of this technique to many other test sites in the UK.

To test this hypothesis, conditional probabilities of the form $p(\text{class}|\text{landcover type})$ were calculated using:

$$p(\text{class}|\text{landcover type}) = \frac{p(\text{class, landcover type})}{p(\text{landcover type})} \quad (4)$$

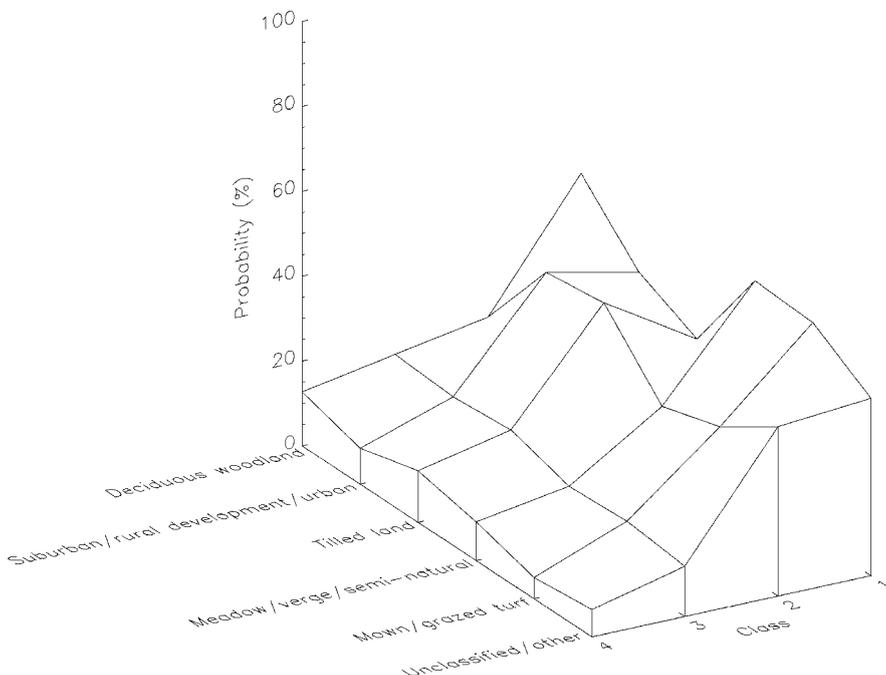


Figure 10. Class probability against landcover type for the first test site.

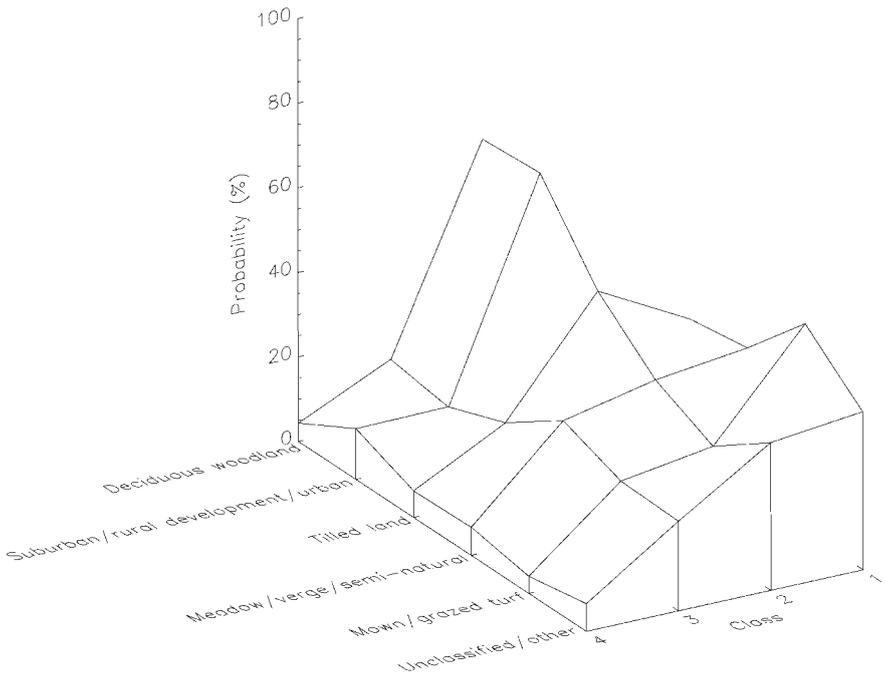


Figure 11. Class probability against landcover type for the second test site.

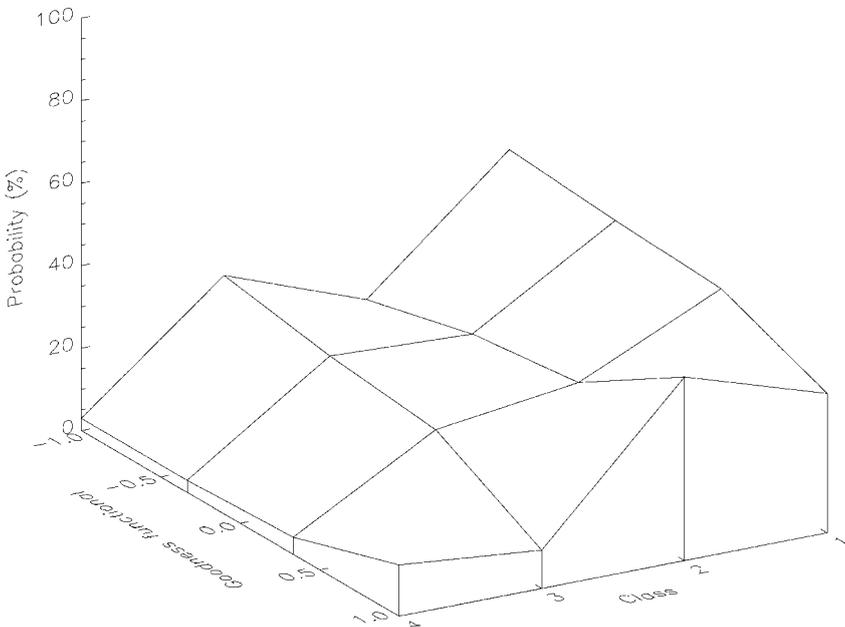


Figure 12. Class probability against goodness functional for the pre-flood SAR imagery for the first test site.

This expression measures the probability of getting a particular class (in the range 1–4 as shown in table 2) and landcover type, but normalized by the prevalence of that landcover type. Statistics were measured from the 300m shoreline strip and compiled into 2-D histograms and the results for the two test reaches are shown in figures 10 and 11. The results are inconsistent between the two sites, implying that the results cannot be extended to other reaches. This is probably due to the lack of correlation between landcover types and vegetation depth and density for the three most common landcover types on the floodplain: mown/grazed turf, meadow/verge/semi-natural and tilled land. Tilled land, for example, will support a range of vegetation depths depending on agricultural practices and the time of SAR image acquisition.

Pre-flood SAR imagery may help in identifying areas prone to misclassification, such as fields that give low returns even when dry, which could be erroneously identified as flooded. A methodology identical to that used to analyse correlation between misclassified areas and landcover type was adopted. The goodness functional, as used in the snake algorithm, was calculated for pre-flood imagery for both test sites, thresholded into four classes, and statistics compiled from the 300 m strip. The results are shown in figures 12 and 13. The seed mean and standard deviation used in calculating the goodness functional were taken from the flood as found by the snake algorithm for each reach, the goodness functional being calculated over 12 pixels as in the snake algorithm used to delineate the flood (two line segments of six pixels each). The results are similar for the two test sites, which implies that the results should be applicable to other reaches, and the dependence on the goodness functional is what would be expected. Areas with a high value of the goodness functional in pre-flood imagery are more likely to be identified as flooded and fall into classes 1 and 3.

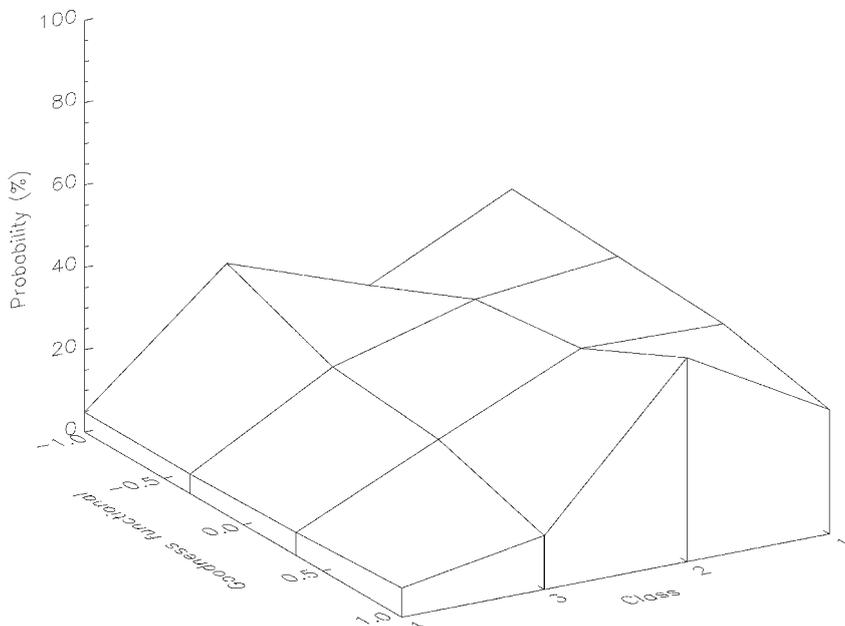


Figure 13. Class probability against goodness functional for the pre-flood SAR imagery for the second test site.

Probabilities like the ones displayed in figures 12 and 13 can be inverted to give a measure of confidence in the segmentation produced by the snake algorithm based on the goodness functional of the pre-flood imagery. For example, if an area has been identified as wet by the algorithm in a region with goodness functional G , the probability that this identification is correct will be given by:

$$p(\text{wet} | \text{SAR wet}, G) = \frac{p(\text{wet}, \text{SAR wet}, G)}{p(\text{SAR wet}, G)} \quad (5)$$

The probabilities on the right-hand side give the ratio of area identified correctly as wet from the SAR imagery to the total area identified as wet from SAR, with the given goodness functional. A similar expression is derived for areas identified as dry in the SAR scene. These probabilities derived from SAR imagery with ground data can now be applied to imagery without referral to the airphoto data. The result is a map of flood probability, which is shown in figure 14 for the first reach. The probabilities of equation (5) were taken from both test reaches. The map shows lower probabilities for regions identified as flooded but corresponding with darker areas in the pre-flood imagery.

The shorelines derived from SAR and aerial photographic imagery can also be compared directly in their vector form. The distance between each node of the contour describing the SAR shoreline and the nearest point on the aerial shoreline was measured and compiled into the histogram shown in figure 15. The mean distance between the two shorelines is 48 m, but this value is strongly influenced by the large tail of the distribution, which corresponds to regions where the SAR segmentation has misplaced the shoreline by the order of a field length. The mean of the distribution for distances less than 50 m is 20 m, which is a more representative estimate of the shoreline error for 69% of its length.

Pre-flood imagery could be combined directly into the classification process (as in the multitemporal techniques referred to in the introduction), for example by inhibiting the growth of the snake over areas that give low radar returns when dry. The extension of the snake model to deal with multidimensional imagery is outlined in Ivins and Porrill (1994). This would, however, ignore the uncertainty inherent in classifying areas with similar SAR returns to the flooded region, where there is little contrast between flooded and dry land. The region based approach adopted here

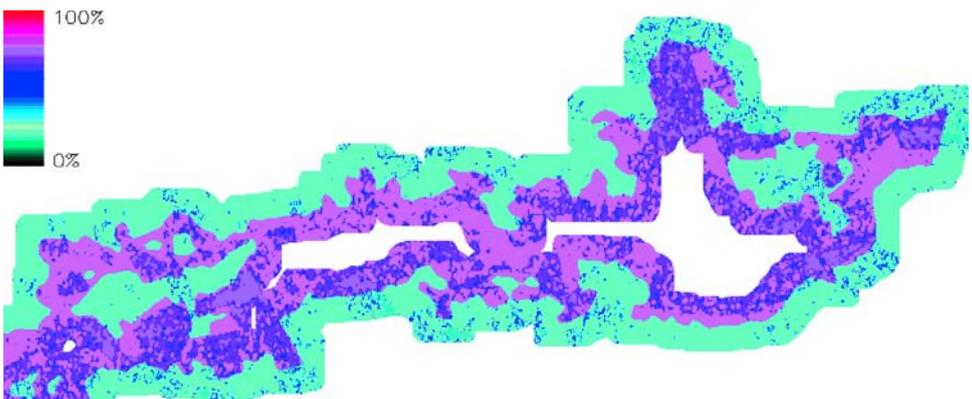


Figure 14. Map of flood probability derived from in- and pre-flood SAR imagery.

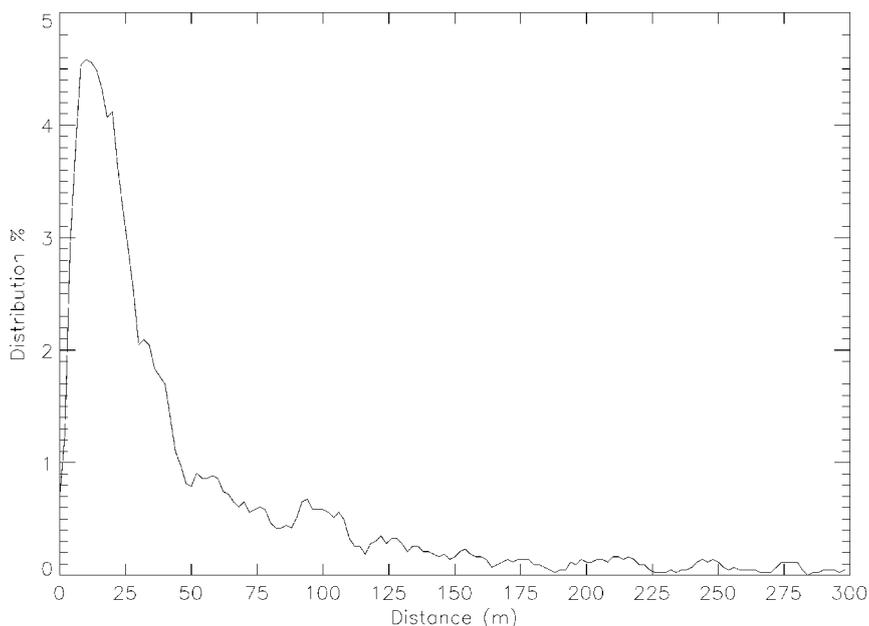


Figure 15. Distribution of distance error between SAR and aerial shoreline.

can deal with such areas, in a worse case scenario assigning a flood probability of 50% to regions where the state of flooding is unknown. The technique could also be adapted to incorporate other remotely sensed or ground surveyed information into the classification process. There may well be a correlation between SAR classification accuracy and vegetation depth which could be obtained from airborne LiDAR (Light Detection And Ranging) survey (Lin 1997).

4. Conclusions and future research needs

A statistical active contour model has been applied to the problem of flood boundary delineation from SAR imagery. Two 15 km reaches of the Thames are used to assess the performance of this algorithm by comparing the results with quasi-simultaneous aerial photography. The effects of flooded vegetation explored in the literature were confirmed over the test reaches, flooded hedgerows giving high returns, probably due to a double reflection mechanism. A more serious practical difficulty not previously identified is that some regions of unflooded vegetation give very low returns similar to open water, and are likely to be misclassified as flooded. These regions can be identified in imagery from before or after the flood, and this multitemporal data used to generate a measure of confidence in the derived shoreline. These low returns were identified with thickly vegetated areas, areas of bare wet soil giving relatively high backscatter. Despite these difficulties, the SAR segmentation algorithm classifies 75% of the shoreline region correctly when compared with simultaneous aerial photographic data. Furthermore, $\approx 70\%$ of the SAR shoreline is coincident with the aerial shoreline to within 20 m.

Phase information in SAR imagery was also used in an attempt to map the flood as a region of low coherence. The 35 day interval between image acquisition does, however, produce low coherence levels in many other regions of the image, in agreement with the findings of Corr *et al.* (1995). This would seem to limit the use

of coherence mapping for flood detection, unless imagery with a shorter repeat cycle can be acquired, from the ERS-1/2 tandem mission, for example (Smith and Alsdorff 1997).

The accuracy of the flood mapping procedure may be improved by the adoption of a more model based approach, rather than the somewhat heuristic method used here, whereby the flood is identified as a region of (reasonably) homogeneous speckle statistics. The real world situation is complicated by the vegetation cover on the floodplain. A qualitative comparison between the SAR imagery and aerial photographs has shown how bare soils produce high levels of radar backscatter, with much lower returns coming from thickly vegetated areas. Enhanced backscatter from flooded vegetation is also observed, principally for linear features such as hedgerows and the tree lined river banks. Vegetation mapping through airborne LiDAR and multispectral scanners would enable these processes to be quantified and catered for explicitly in the segmentation process, perhaps via some form of physically based model. This could permit the inversion of SAR returns into not only a binary classification of flooded/unflooded, but also give information on water depths. Flooding below vegetation may not be visible until the water depth approaches the canopy depth. A region of enhanced backscatter may then occur, followed by open water backscatter as the canopy is overtopped. Given the relative rarity of SAR imagery of fluvial flood events, it may be efficacious to expand future studies to coastal environments such as salt marshes where inundation occurs on a more regular basis.

Acknowledgments

This work was done as part of a PhD studentship funded by the National Remote Sensing Centre Ltd, Farnborough, UK. Thanks go to Ken Blyth of the Institute of Hydrology, Wallingford, UK, for providing the SAR and photographic data, without which this work would not have been possible.

References

- BADJI, M., and DAUTREBANDE, S., 1995, Characterization of flood inundated areas and delineation of poor drainage soil using ERS-1 SAR imagery. *Hydrological Processes*, **11**, 1441–1450.
- BATES, P. D., and ANDERSON, M. G., 1995, Issues of hydraulic model validation and design using remotely sensed data. *Proceedings of the 1st ERS thematic working group meeting on flood monitoring, 26–27 June 1995* (Frascati, Italy: ESA-ESRIN).
- BATES, P. D., HORRITT, M. S., SMITH, C. N., and MASON, D. C., 1997, Integrating remote sensing observations of flood hydrology and hydraulic modelling. *Hydrological Processes*, **11**, 1777–1795.
- BIGGIN, D. S., and BLYTH, K., 1996, A comparison of ERS-1 satellite radar and aerial photography for river flood mapping. *Journal of the Chartered Institute of Water Engineers and Managers*, **10**, 59–64.
- CALABRESI, G., 1995, The use of ERS data for flood monitoring: an overall assessment. *Proceedings of the 2nd ERS applications workshop, 6–8 December 1995* (London: ESA) pp. 237–241.
- COHEN, L. D., 1991, On active contour models and balloons. *CVGIP: Image Understanding*, **53**, 211–218.
- COLLINS, M. J., LIVINGSTONE, C. E., and RANEY, R. K., 1997, Discrimination of sea ice in the Labrador marginal ice zone from synthetic aperture radar image texture. *International Journal of Remote Sensing*, **18**, 535–571.
- CORR, D. G., KEYTE, G. E., and WHITEHOUSE, S., 1995, Studies of decorrelation in multi-

- temporal SAR imagery. *International Geoscience and Remote Sensing Symposium, 1995* (Piscataway, NJ: IEEE), pp. 1026–1028.
- FULLER, R. M., GROOM, G. B., and WALLIS, S. M., 1994, The landcover map of Great Britain: an automated classification of Landsat Thematic Mapper data. *Photogrammetric Engineering and Remote Sensing*, **15**, 1357–1362.
- GIACOMELLI, A., BACCHIEGA, U., TROCH, P. A., and MANCINI, M., 1995, Evaluation of surface soil moisture distribution by means of SAR remote sensing techniques and conceptual hydrological modelling. *Journal of Hydrology*, **166**, 445–459.
- HENDERSON, F. M., 1995, Environmental factors and the detection of open surface water using x-band radar imagery. *International Journal of Remote Sensing*, **16**, 2423–2437.
- HENEERY, G. M., and KUX, H. J. H., 1997, Lacunarity as a texture measure for SAR imagery. *International Journal of Remote Sensing*, **16**, 565–571.
- HORRITT, M. S., 1998, Enhanced flood flow modelling using remote sensing techniques. PhD thesis, University of Reading.
- HORRITT, M. S., 1999, A statistical active contour model for SAR image segmentation. *Image and Vision Computing*, **17**, 213–224.
- IMHOFF, M. L., VERMILLION, C., STORY, M. H., CHOUDHURY, A. M., GAFOOR, A., and POLCYN, F., 1987, Monsoon flood boundary delineation and damage assessment using space borne imaging radar and Landsat data. *Photogrammetric Engineering and Remote Sensing*, **53**, 405–413.
- IVINS, J., and PORRILL, J., 1994, Statistical snakes: active region models. *Proceedings of the 5th British Machine Vision Conference, September 1994, York* (London: BMVA), vols 1 and 2, ch. 79, pp. 377–386.
- KANNEN, A., 1995, Use of SAR data for flood mapping and monitoring in Thuringen, Germany. *Proceedings of the 2nd ERS applications workshop, 6–8 December 1995* (London: ESA) pp. 237–241.
- LIN, C. S., 1997, Waveform sampling LiDAR applications in complex terrain. *International Journal of Remote Sensing*, **18**, 2087–2104.
- LUCKMAN, A. J., FREY, A. C., YANASSE, C. C. F., and GROOM, G. B., 1997, Texture in airborne SAR imagery of tropical forest and its relationship to forest regeneration stage. *International Journal of Remote Sensing*, **18**, 1333–1349.
- MASON, D. C., and DAVENPORT, I. J., 1996, Accurate and efficient determination of the shoreline in ERS-1 images. *IEEE Transactions on Geoscience and Remote Sensing*, **34**, 1243–1253.
- NOYELLE, J., DELIMIERE, S., and MARINELLI, L., 1995, Identification of flooded areas in the Rhone. *Proceedings of the 2nd ERS applications workshop, 6–8 December 1995* (London: ESA), pp. 237–241.
- OBERSTADLER, R., HONSCH, H., and HUTH, D., 1997, Assessment of the mapping capabilities of ERS-1 SAR data for flood mapping: a case study in Germany. *Hydrological Processes*, **11**, 1415–1425.
- OLIVER, C. J., BLAKE, A., and WHITE, R. G., 1994, Optimum texture analysis of synthetic aperture radar images. *Algorithms for Synthetic Aperture Radar Imagery*, **2230**, ch. 34, 389–398.
- ORMSBY, J. P., BLANCHARD, B. J., and BLANCHARD, A. J., 1985, Detection of lowland flooding using active microwave systems. *International Journal of Remote Sensing*, **5**, 317–328.
- PREVOT, L., DECHAMBRE, M., TACONET, O., VIDALMADJAR, D., NORMAND, M., and GALLE, L., 1993, Estimating the characteristics of vegetation canopies with airborne radar measurements. *International Journal of Remote Sensing*, **14**, 2803–2818.
- RAMSEY, E. W., 1995, Monitoring flooding in coastal wetlands by using radar imagery and ground based measurements. *International Journal of Remote Sensing*, **16**, 2495–2502.
- RICHARDS, J. A., WOODGATE, P. W., and SKIDMORE, A. K., 1987, An explanation of enhanced radar backscattering from flooded forests. *International Journal of Remote Sensing*, **8**, 1093–1100.
- SLAMA, C. C., 1980, *Manual of Photogrammetry*, 4th edn (Falls Church, VA: American Society of Photogrammetry).
- SMITH, L. C., 1997, Satellite remote sensing of river inundation area, stage and discharge: a review. *Hydrological Processes*, **11**, 1427–1439.
- SMITH, L. C., and ALSDORFF, D. E., 1997, Flood mapping from phase decorrelation of tandem ERS data: Ob river, Siberia. *Proceedings of the 3rd ERS symposium on space at the service of our environment, Florence, Italy, 17–21 March 1997* (Paris: ESA), pp. 537–539.

- SOARES, J. V., RENNO, C. D., FORMAGGIO, A. R., YANASSE, C. C. F., and FREY, A. C., 1997, An investigation of the selection of texture features for crop discrimination using SAR imagery. *Remote Sensing of Environment*, **59**, 234–247.
- SOLOMON, S. I., 1993, Methodological considerations for the use of ERS-1 SAR imagery for the delineation of river networks in tropical forest areas. *Proceedings of the 1st ERS-1 symposium, Frascati, Italy, 26–27 June 1995* (London: ESA), pp. 595–600.
- TACONET, O., VIDALMADJAR, D., EMBLANCH, C., and NORMAND, M., 1996, Taking into account vegetation effects to estimate soil moisture from C-band radar measurements. *Remote Sensing of Environment*, **56**, 52–56.
- TOUZI, R., LOPES, A., and BOUSQUET, P., 1988, A statistical and geometrical edge detector for SAR images. *IEEE Transactions on Geoscience and Remote Sensing*, **26**, 764–773.
- ULABY, F. T., MOORE, R. K., and FUNG, A. K., 1986, *Microwave Remote Sensing*, vol. 3 (Norwood, MA: Artech House).
- ULABY, F. T., DUBOIS, P. C., and VAN ZYL, J., 1996, Radar mapping of surface soil moisture. *Journal of Hydrology*, **184**, 57–84.
- WANG, Y., HESS, L. L., FILOSO, S., and MELACK, J. M., 1995, Understanding the radar backscattering from flooded and non-flooded amazonian forests: results from canopy backscatter modelling. *Remote Sensing of Environment*, **54**, 324–332.
- WEGMULLER, U., WERNER, C. L., NUESCH, D., and BORGEAUD, M., 1995, Forest mapping using ERS repeat-pass SAR interferometry. *Earth Observation Quarterly*, **49**, 4–7.
- WILLIAMS, D. J., and SHAH, M., 1992, A fast algorithm for active contours and curvature estimation. *CVGIP: Image Understanding*, **55**, 14–26.