Using airborne laser altimetry to improve river flood extents delineated from SAR data

D.C. Mason, J.T. Dall'Amico, T.R. Scott
Environmental Systems Science Centre,
University of Reading,
Reading, UK
dcm@mail.nerc-essc.ac.uk

M.S. Horritt, P.D. Bates
School of Geographical Sciences,
University of Bristol,
Bristol, UK

Abstract— Flood extent maps derived from SAR images are a useful source of data for validating hydraulic models of river flood flow. The accuracy of such maps is reduced by a number of factors, including changes in returns from the water surface caused by different meteorological conditions and the presence of emergent vegetation. The paper describes how improved accuracy can be achieved by modifying an existing flood extent delineation algorithm to use airborne laser altimetry (LiDAR) as well as SAR data. The LiDAR data provide an additional constraint that waterline (land-water boundary) heights should vary smoothly along the flooded reach. The method was tested on a SAR image of a flood for which contemporaneous aerial photography existed, together with LiDAR data of the un-flooded reach. Waterline heights of the SAR flood extent conditioned on both SAR and LiDAR data matched the corresponding heights from the aerial photo waterline significantly more closely than those from the SAR flood extent conditioned only on SAR data.

Keywords-component; Synthetic aperture radar, Laser applications, Water resources

I. INTRODUCTION

Flood extent maps derived from remotely sensed data are of considerable use in hydrology, providing spatially distributed data for validating hydraulic models of river flood flow, for emergency flood relief management and for developing spatially accurate hazard maps. The all-weather day-night capability of SAR sensors gives these a considerable advantage for flood mapping over sensors operating at visible or infrared wavelengths, as the latter are unable to penetrate the cloud that often accompanies flood events. Flood water usually appears dark compared to the surrounding land because the smooth water surface acts as a specular reflector. However, wind or rain may cause roughening of the water surface, such that the backscatter from the water may rise to similar or greater levels than the surrounding land. A further complicating factor may be the presence of emergent vegetation at the flood edge, leading to substantial increase in backscatter due to multiple reflections. These factors tend to reduce the accuracy of SAR-derived flood extent maps.

Methods for the automatic delineation of flood extent in SAR images of both fluvial and tidal environments have been developed by several authors (see e.g. [1, 2]). A technique for delineating a fluvial flood using a statistical active contour model (or snake) applied to a single-frequency single-polarisation SAR image of the flood to identify regions of homogeneous speckle statistics is described in [3, 4]. This was aimed at producing an observed flood extent against which to validate the flood extent predicted by a hydraulic model. The technique was applied to a flood imaged using the ERS-1 SAR sensor, and proved capable of identifying 75% of the flooded area correctly, with 70% of the waterline coinciding with ground data to within 20m. The main error in waterline position was found to be due to un-flooded short vegetation adjacent to the flood giving similar radar returns to open water, causing an over-estimate of flood extent. The loss of flood extent due to emergent vegetation was found to be a secondary source of error.

However, further investigation revealed that, as a result of these error sources and the relatively large size of the ERS pixel, the heights of the SAR waterline along a flooded reach could exhibit considerable noise. One reason for this was that the snake algorithm did not constrain the SAR waterline height to vary smoothly along the reach, and thus did not exploit the fact that changes in height along a reach are usually very gradual. With the level of height errors exhibited, the SAR image often became much less useful for flood model validation than it could otherwise be.

This paper describes how improved accuracy can be achieved by modifying the existing flood extent delineation algorithm to use airborne LiDAR as well as SAR data. The algorithm is modified to look not only at SAR image space but also at LiDAR DTM and vegetation height maps, so that the snake can be conditioned to be smoothly-varying in ground height as well as in SAR intensities and textures.

II. TEST DATA SET

The event used to test the method was a 1-in-5 year flood on the Thames west of Oxford that occurred in December 1992. An ERS-1 overpass occurred just after the flood peak, at a time when there was no wind or rain in the area. Aerial photographs were acquired at the time of the overpass [5], so that the SAR snake waterlines conditioned without and with the LiDAR data could be compared with the waterline from the aerial photos. The floodplain over this reach is semi-rural, with the majority of fields being used at the time for pasture or having been ploughed. The flood waterline was delineated by

eye from the aerial photos and georeferenced [4]. The error in the waterline position was assessed from waterline segments where the waterline was observed to lie alongside a hedgerow or field boundary that could be located on a 1:25000 scale map, and was found to be less than 20m.

LiDAR data of the un-flooded reach were subsequently acquired by the Environment Agency of England and Wales (EA). By comparing them with GPS readings, the LiDAR heights were found to have a random error of 10.2 cm and a systematic error of 2.6cm. The post-processed LiDAR DTM and vegetation height maps were obtained from the EA. These were degraded to 2m pixel size to avoid too large a mismatch with the SAR pixel size of 12.5m. Fig. 1 shows a section of the DTM (1.5km square) with the aerial photo waterline overlain.

III. FLOOD EXTENT EXTRACTION FROM SAR DATA

The algorithm to delineate a flood in a SAR image using an active contour model is described in [3], and only an overview is presented here. Active contour models or snakes are useful for converting incomplete or noisy edge maps into smooth continuous vector boundaries [6]. The edge image space is searched using a dynamic curvilinear contour that is driven to be attracted to edge pixels using an energy minimisation function, so that the contour can link together unconnected edge segments. The contour (snake) is represented in a piecewise linear fashion as a set of nodes (i.e. the coordinates of the snake points) linked by straight line segments.

Ivins and Porrill [7] developed a statistical snake that operates on the image itself rather than an edge image, dispensing with the need for a prior edge detection stage. The statistical snake is formulated as an energy minimisation problem with the total snake energy $E(\boldsymbol{u}(s))$ given by –

$$E(\boldsymbol{u}(s)) = E_{tension} + E_{curvature} - \iint G(I(x, y)) \, dx \, dy \qquad (1)$$

where $\boldsymbol{u}(s) = (x(s), y(s))$ describes the contour position (x, y) in the 2-D image space as a vector function of an arc length parameter s. $E_{tension}$ and $E_{curvature}$ are energies generated by the model's internal tension and stiffness constraints, which favour a smooth uncrenellated contour made up of evenly spaced nodes. G is a goodness functional that measures how homogeneous in tone and texture is a set of image pixels $I(x,y)$.

G provides an estimate of how similar the mean intensity μ' and variance of the set of n pixels on the contour immediately on either side of a particular snake node are to the mean intensity and variance of those pixels contained within the interior of the contour. The dependence of G on μ' is given by –

$$G(\mu') = 1 - n(\mu' - \mu)^2/vk^2 \qquad (2)$$

where μ and v are respectively the mean and variance of the seed population enclosed within the contour, and k is a parameter that can be adjusted to tune algorithm performance. G is equal to 1 for a set of pixels with $\mu' = \mu$, but falls to zero if μ' differs from μ by $k\sqrt{(v/n)}$ (i.e. k standard deviations, with k usually set at around 2 or 3).

The roles of the tension and curvature constraints are to produce a contour of appropriate smoothness with evenly spaced nodes, by a consideration of the balance between image and curvature forces. Consider adjacent snake nodes at positions $\boldsymbol{u}_{i-1}$, $\boldsymbol{u}_i$ and $\boldsymbol{u}_{i+1}$, linked by vectors $\boldsymbol{v}_i$ and $\boldsymbol{v}_{i+1}$. The contribution to the total curvature energy is [3] –

$$\Delta E_{curvature} = \gamma \, |\boldsymbol{v}_{i+1} - \boldsymbol{v}_i|^2/ a_i \qquad (3)$$

where a_i is the distance between the midpoints of $\boldsymbol{v}_i$ and $\boldsymbol{v}_{i+1}$, and γ is a curvature energy weighting parameter. The contribution to the tension energy is given by –

$$\Delta E_{tension} = \lambda(|\boldsymbol{u}_{i+1} - \boldsymbol{u}_i|^2 + |\boldsymbol{u}_i - \boldsymbol{u}_{i-1}|^2) \qquad (4)$$

where λ is the tension energy weighting parameter. The magnitudes of these energies can be adjusted using the weighting parameters. Too large a value for the curvature parameter will make the curvature term dominate the model energy and produce an unrealistically smooth contour. Too large a value of the tension parameter will favour a short contour and stifle the growth of the snake. The scheme used to adjust the snake to minimise its energy is the algorithm of Williams and Shah [6].

Fig. 2 shows a section of snake waterline generated using SAR data only, overlain on the SAR image. The snake was seeded manually as a narrow strip lying along the course of the un-flooded river channel, ensuring that it contained only flooded pixels. The snake shows a tendency to leak to higher ground (see the DTM in Fig. 1). This is likely to be due to the presence of vegetated fields, which correspond to areas of low SAR backscatter and are likely to be misclassified as flooded.
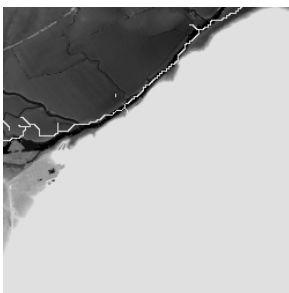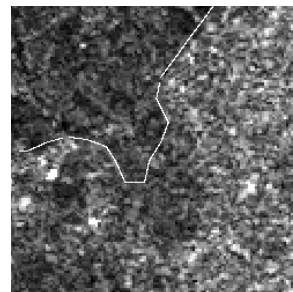

Figure 1. A section of the aerial photo waterline (white) overlain on the LiDAR DTM (lighter = higher). The low-lying north-west part of the image was flooded

Figure 2. Waterline conditioned only on SAR data overlain on SAR data

IV. Flood Extent Extraction from SAR and LiDAR Data

The snake algorithm was modified so that the snake is conditioned not only on the SAR image, but also on the LiDAR DTM, so that it becomes smoothly-varying in ground height as well as in SAR intensities and textures. The LiDAR DTM is able to provide a ground height at each pixel, so that each position $u(x, y)$ becomes $u(x, y, z)$. The modification involves using the LiDAR heights to measure curvatures and tensions at snake nodes in 3-D rather than 2-D space. Consider an instance where an un-flooded field with low SAR backscatter is adjacent to a flood edge, such that the field is included in the SAR waterline determined by the snake (Fig. 3). As there will likely be a rise in height (dh) across the field perpendicular to the true flood edge, the error in the waterline will give rise to a significant component of curvature in the vertical plane, which will not be present in the waterline segments adjacent to the field. To reduce the vertical curvature component at node i in Fig. 3, the snake will try to contract to drag node i back to be collinear with nodes $i-1$ and $i+1$, which will also reduce the horizontal curvature. The 3-D tension energy will also be reduced by this move. In order to take account of the fact that a change in height at a node should in general cause different changes in 3-D curvature and tension compared to the same magnitude change of node position in the xy plane, the LiDAR heights were scaled by a weighting factor w_l with respect to the (x, y) coordinates.

For the modified algorithm, the LiDAR image was degraded to the same pixel size as the SAR image (12.5m) by averaging the LiDAR heights within each SAR pixel. Fig. 4 shows a section of the snake conditioned on both SAR and LiDAR data, overlain on the SAR data. It is clear that the tendency of the snake to leak onto higher ground has been much reduced compared to Fig. 2 (see also Fig. 1 DTM). This behaviour is repeated at a number of other points over the whole waterline. It turns out that the main errors in waterline position corrected using the LiDAR are due to un-flooded short vegetation adjacent to the flood giving similar returns to open water.

V. Quantitative Comparison of Methods

A quantitative estimate of the improvement in performance using LiDAR data was made by comparing height differences between the snake and aerial photo waterlines. These were heighted by superimposing them on the LiDAR DTM. Only snake nodes on low slopes and in areas of short vegetation in the LiDAR vegetation height map were selected for heighting, as these were the ones likely to be heighted most accurately. For each snake node selected, the aerial photo height to associate with the snake height was found by finding the height of the closest aerial photo waterline point. To ensure that adjacent pairs of heights were largely uncorrelated, the pairs selected were thinned further so that no pair was closer than 200m to another. Parameters were optimised by minimising the sum of the squared height differences between the snake nodes and their corresponding aerial photo waterline points. For the snake conditioned on only SAR data, the parameter optimised was k. For the snake conditioned on SAR and LiDAR data, k and w_l were optimised.

=   waterline determined from SAR only
    (O = snake nodes)

- - - - -   =   waterline determined from SAR and LIDAR

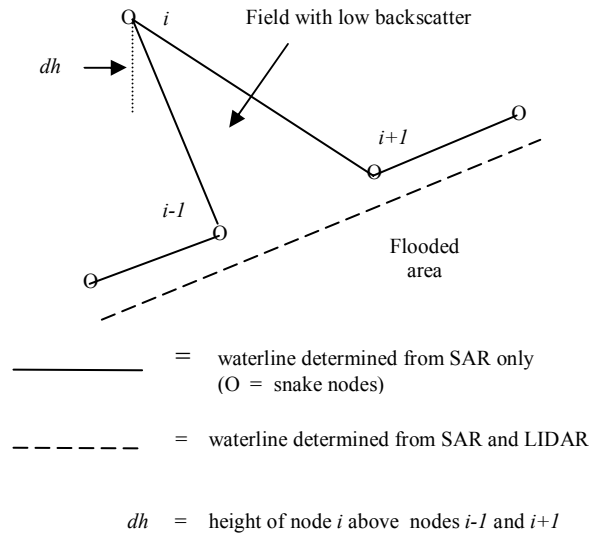dh =   height of node i above nodes $i-1$ and $i+1$

Figure 3. An example error that might be corrected using LiDAR.

For the snake conditioned only on SAR data, the minimum r.m.s. error was obtained for $k = 2.0$, for which the snake waterline heights turned out to be significantly higher than those of the aerial photo. For the snake conditioned on SAR and LiDAR data, the minimum r.m.s. error was obtained for $k = 3.0$ and $w_l = 0.15$. In this case, there was no significant difference between the snake and aerial photo waterline heights. There was also a factor of four reduction in r.m.s. error compared to that for the snake conditioned only on SAR data. The snake sections shown in Figs. 2 and 4 are for the relevant optimum parameters.

VI. Discussion

The emphasis in the foregoing has been on ERS SAR data because of the availability of simultaneous ERS SAR and aerial photography of the 1992 Oxford flood. While the ERS SAR sensor has single polarisation and a fixed viewing angle,
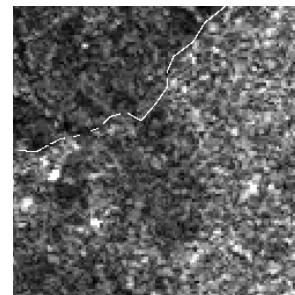

Figure 4. Waterline conditioned on SAR and LiDAR data overlain on SAR data.

the advent of later sensors with higher resolutions, multiple polarisations and variable viewing angles (e.g. RADARSAT, ENVISAT ASAR) has allowed improved flood delineation (e.g. [8]). ERS SAR data also have poorer resolution than airborne LiDAR data, whereas the high resolution satellite SAR sensors due for launch shortly (e.g. RADARSAT-2, TerraSAR and the Cosmo-Skymed constellation) will have resolutions more similar to that of LiDAR. The technique should also be applicable in this case.

Production of a more smoothly-varying waterline may allow the development of improved performance measures for flood extent validation based on patterns of height differences between observed and modelled waterlines rather than on patterns of wet or dry pixels as at present [9]. This aspect is currently being studied.

VII.   CONCLUSION

An algorithm has been developed for the automatic extraction of flood extent using a snake generated from combined SAR and LiDAR data, and the resulting waterline compared to that generated using SAR data alone. The conclusion is that the use of LiDAR data has resulted in an observed waterline that varies more smoothly along the reach and is a better match to our best estimate of the true waterline heights.

ACKNOWLEDGMENT

REFERENCES

[1] G. Schumann, R. Hostache, C. Puech, L. Hoffmann, P. Matgen, F. Pappenberger and L. Pfister, "High-resolution 3D flood information from radar imagery for flood hazard management," IEEE Trans. Geoscience and Remote Sensing (in press).

[2] A. Niedermeier, E. Romaneessen and S. Lehner, "Detection of coastlines in SAR images using wavelet methods," IEEE Trans. Geoscience and Remote Sensing, vol. 38, no. 5, pp.2270-2281, 2000.

[3] M.S. Horritt, "A statistical active contour model for SAR image segmentation," Image and Vision Computing, vol. 17, pp. 213-224, 1999.

[4] M.S. Horritt, D.C. Mason and A.J. Luckman, "Flood boundary delineation from synthetic aperture radar imagery using a statistical active contour model," Int. J. Remote Sensing, vol. 22(13), pp. 2489-2507, 2001.

[5] D.S. Biggin and K. Blyth, "A comparison of ERS-1 satellite radar and aerial photography for river flood mapping," Journal of the Chartered Institute of Water Engineers and Managers, vol. 10, pp. 1777-1795, 1996.

[6] D.J. Williams and M. Shah, "A fast algorithm for active contours and curvature estimation," CVGIP: Image Understanding, vol. 55, pp. 14-26, 1992.

[7] J. Ivins and J. Porrill, "Statistical snakes: active region models," Proc. 5th British Machine Vision Conference, York, September, 1994 (London, BMVA), vols 1 and 2, ch. 79, pp. 377-386, 1994.

[8] Y. Yu and S.T. Acton, "Automated delineation of coastline from polarimetric SAR imagery," Int. J. Remote Sensing, vol. 25, no. 17, pp.3423-3438, 2004.

[9] G. Aronica, P.D. Bates and M.S. Horritt, "Assessing the uncertainty in distributed model predictions using observed binary pattern information within GLUE," Hydrological Processes, vol.16, pp. 2001-2016, 2002.