Super-resolution Target Mapping from Soft-Classified Remotely Sensed Imagery

Peter M. Atkinson

Department of Geography, University of Southampton, Highfield, Southampton SO17 1BJ, UK

Email: pma@soton.ac.uk

Abstract. A simple, efficient algorithm is presented for super-resolution target mapping from remotely sensed images. Following an initial random allocation of pixel proportions to binary 'hard' sub-pixel classes, the algorithm works in a series of iterations, each of which contains two stages. For each iteration, a distance weighted function of neighbouring pixels is computed for all sub-pixels. Then, on a pixel-by-pixel basis, the '1' with the minimum value of the function is swapped with the '0' with the maximum value of the function, if the swap results in an increase in some objective function. The algorithm is demonstrated to work reasonably well with simple images, opening the way for further research to explore the algorithm, to extend the algorithm to multiple classes and to develop more efficient, but equally simple algorithms.

1. Introduction

Land cover is a fundamental variable that underpins much scientific research. For example, data on land cover are required to provide boundary conditions for climate models (e.g., global circulation models) and hydrological and hydraulic models. Although important, informative and accurate data on land cover are both difficult and expensive to provide. Therefore, much of the land cover data currently being used in scientific research are of inadequate quality. For example, much land cover data may be (i) incomplete spatially, (ii) out-of-date or (iii) inaccurate. Given the problems associated with land cover data it is not surprising that remote sensing has been of great value for land cover mapping. For example, remote sensing is capable of providing synoptic and complete coverage with a single image.

Despite the obvious utility of remote sensing for land cover mapping, many of the above data problems remain. For example, it is often difficult to ensure appropriate spatial and, in particular, temporal coverage. Further, increasing the accuracy of land cover classification has been the subject of intensive research for many years (e.g., Justice and Townshend, 1981). The issue that forms the focus of the present research is that land cover data provided by remote sensing are limited by the *spatial resolution* of the sensor. While complete *cover* may be provided for an area of interest, the *sample* can never be complete: increasing the spatial resolution will always reveal greater detail.

Spatial resolution has been the subject of research in remote sensing for many years because it forms a fundamental scale of measurement (Woodcock and Strahler, 1987; Atkinson and Tate, 2000). The spatial variation observed in remotely sensed imagery is a function of both the property of interest (i.e., the real world) and the sampling framework (i.e., the ensemble of sensor characteristics including the spatial resolution). Researchers have sought to evaluate the effect of spatial resolution on detectable spatial variation as characterised by functions such as the scale variance (Woodcock and Strahler, 1987) and variogram (Curran and Atkinson, 1998). Further, researchers have attempted to find a suitable means of selecting a spatial resolution given knowledge of functions such as the variogram (e.g., Curran and Atkinson, 1998).

Just as spatial resolution affects the spatial variation observable in remotely sensed imagery, it also affects land cover classification based on that imagery. For example, researchers realized early on that an increase in spatial resolution may lead to a decrease in classification accuracy (Justice and Townshend, 1981). This paradoxical outcome may be explained by an increase in within-class variance (between ground resolution elements) with an increase in spatial resolution. However, the result relates mainly to hard classification and assumes no post-classification analysis. The more important (and often overlooked) property is that the spatial information increases dramatically with spatial resolution, making spatial resolution the most fundamental limit to quantitative land cover information from remotely sensed imagery.

The earliest research on techniques for land cover classification from remotely sensed imagery focused on hard classification (both supervised and unsupervised) in which each pixel is allocated to one class. Around fifteen years ago (Adams *et al*., 1985) researchers began to realize that for most remotely sensed scenes hard classification is inappropriate. Many pixels in remotely sensed images represent more than one land cover class on the ground. Such 'mixed pixels' occur where the frequency of spatial variation in land cover is greater

than or equal to the frequency of sampling afforded by the sensor's spatial resolution (Woodcock and Strahler's (1987) L-resolution case). However, a proportion of pixels will be mixed even where the spatial resolution is fine relative to the land cover variation (H-resolution case) because some pixels inevitably straddle boundaries between land cover 'objects'.

The existence of mixed pixels led to the development of several approaches for soft (often termed fuzzy in the remote sensing literature) classification in which each pixel is allocated to all classes in varying proportions. Examples of techniques for soft classification applied to remotely sensed imagery include

(i) the linear mixture model, which is based on least squares approximation to solve simultaneous linear equations (e.g., Adams *et al*., 1985; Garcia-Haro *et al*., 1996; Atkinson *et al*., 1997),
(ii) fuzzy *c*-means classification, which is based on distance metrics expressed in feature (waveband) space (e.g, Bezdek, 1981; Bezdek *et al*., 1984; Atkinson *et al*., 1997), and
(iii) feed-forward, back-propagation (FFBP) neural networks trained on class proportions (e.g., Atkinson and Tatnall, 1997; Atkinson *et al*., 1997)

Recently, support vector machines have also become popular (e.g., Brown *et al*., 1999).

All of the above techniques may be used to provide a soft classification of land cover that is both more informative and potentially more accurate than the equivalent hard classification. However, while the proportions of each land cover within each pixel may be predicted, the spatial location of each land cover within each pixel is not. For example, the soft classifier may predict 60% woodland within a pixel. This is undoubtedly more informative than the '100% woodland' predicted by a hard classifier. However, it would also be useful to know *where*, within the pixel, the woodland is located spatially. That goal, referred to as super-resolution mapping, is the subject of this paper. It amounts to transforming multispectral (multiple attribute) data into spatial (single attribute) data. While no new information is created, it does result in an increase in spatial resolution above that achieved with hard (and soft) classification of the original remotely sensed imagery. In the next section, previous research on super-resolution mapping is reviewed.

2. Super-resolution mapping

There exist many different potential techniques for super-resolution mapping from remotely sensed imagery. A simple approach involves converting a hard-classified image into the vector data model by replacing class object boundaries with vectors. Generalizing these vectors will produce sub-pixel spatial information on land cover. However, not withstanding the problems associated with hard classification, such an approach is under-constrained. Foody (1998) evaluated an interpolation-based technique for predicting the boundary of a lake with sub-pixel geometric precision. However, this approach was similarly under-constrained. In both of the above cases, the algorithm may be subject to effects such as conditional bias and smoothing which may affect the final vector boundary. More recently, Aplin et al., (2001) developed a technique for converting the output from a per-pixel soft-classification of land cover into a per-parcel hard classification of land cover objects. Land-line vector data from the Ordnance Survey were used to constrain the placement of the soft proportions within each pixel. This requirement for vector data makes the technique redundant for (i) less developed areas of the world and (ii) updating the vector database.

Several authors have attempted super-resolution mapping directly from multispectral remotely sensed imagery. For example, in a series of papers, Schneider (1993, 1999) and Steinwendner *et al*. (1998) document a technique for sub-pixel mapping of linear features based on a 3 pixel by 3 pixel moving window. The approach was extended to include neural network prediction of vector boundaries, but is restricted to remotely sensed images and the detection of linear features. Flack et al. (1994) developed a technique based on the Hough transform for, first, detecting linear features in remotely sensed images of agricultural scenes and, second, unmixing the signal on either side of the boundary. Again, the technique is suitable for application to linear features in unprocessed remotely sensed images.

Atkinson (1997) was the first to suggest super-resolution mapping based solely on the output from a soft classification. The idea proposed was to convert soft land cover proportions to hard (per-sub-pixel) land cover classes (that is, at a finer spatial resolution). The solution that is most intuitive (most visually appealing) is attained by maximizing the spatial correlation or spatial dependence between neighbouring sub-pixels. Spatial dependence is the likelihood that observations close together are more alike than those that are further apart (Matheron, 1965; Goovaerts, 1997). The basic idea was, therefore, to maximize the spatial correlation between neighbouring sub-pixels under the constraint that the original pixel proportions were maintained (Atkinson, 1997). This objective is reasonable where the land cover target of interest is larger than the pixels in the imagery.

While successful for simple shapes such as a circle and a torus, the algorithm failed to handle satisfactorily the

complex arrangement of multiple-class land cover objects found in a System Pour L'Observation de la Terre (SPOT) High Resolution Visible (HRV) image of the New Forest, Hampshire. Two problems relating to the data were identified in the paper: (i) generalisation caused by the point-spread function of the sensor and (ii) classification error. However, a more fundamental problem that was not identified was that sub-pixels to be assigned a class were compared to neighbouring *pixel* proportions, thus, mixing scales of measurement (sub-pixel v. pixel). As a consequence, land cover was allocated up against pixel boundaries creating linear artefacts in the final super-resolution map.

Vehoeye *et al*. (2000), building on the earlier work of Atkinson (1997), attempted to allocate sub-pixel hard classes using a technique similar to the spectral mixture model. The pixel proportion constraints were built into the mixture model and a solution was achieved by least squares approximation. Unlike the computational solution (Atkinson, 1997) Vehoeye *et al*.'s solution was not iterative. Unfortunately, the non-iterative solution was achieved only by comparing sub-pixels to pixel proportions, thereby mixing scales. The result was that linear artefacts were once again produced in the final map.

A solution to the super-resolution problem may be achieved by comparing sub-pixels to sub-pixels meaning that the problem is non-linear and the technique adopted to solve it is likely to be iterative. Recently, Tatem *et al*. (2001a) developed a Hopfield neural network (HNN) technique (Hopfield and Tank, 1985) for super-resolution target mapping. The HNN is similar to the more common FFBP neural network except that all neuron outputs are connected to all neuron inputs. Whereas the FFBP network is used to model the relation between two sets of variables (in the same sense as regression or classification), the HNN is used as an optimization tool. To solve the super-resolution mapping problem, with the pixel proportions as initial conditions, the HNN architecture must be arranged as in Figure 1. The sub-pixel class allocations are initialised randomly with $p.n$ sub-pixels having the value 0.55 and $n\text{-}p.n$ value 0.45, where p is the proportion and n the number of sub-pixels. The HNN is the set up to minimise an energy function which comprises a goal and constraints:

$$E = G + C$$

where the goal G is to increase the spatial correlation between neighbouring sub-pixels (sub-pixels are compared directly to other sub-pixels) and the constraint C is that the original class proportions per-pixel are maintained in the super-resolution land cover map.

The HNN was applied initially to detect targets (two-class problem) (Tatem *et al*., 2001a), but eventually extended to super-resolution land cover mapping (multiple class problem) (Tatem *et al*., 2001b). Further, the simple spatial clustering goal G was extended to a variogram-matching goal SV for which the objective was not to maximise spatial correlation, but to match the spatial correlation in the output map to that known *a priori* (i.e., provided by a sample variogram) (Tatem *et al*., 2001c). This latter objective is similar in concept to the techniques of conditional simulation and simulated annealing in geostatistics (see Journel, 1996; Goovaerts, 1997). It means that super-resolution mapping is possible, even for objects that are smaller than a pixel.

The HNN has been demonstrated to be a successful tool for super-resolution mapping. However, alternative classes of algorithm remain to be investigated. One, in particular, stands out as an obvious candidate. It is the extension of the work in Atkinson (1997) to the comparison of sub-pixels with sub-pixels (rather than sub-pixels to pixels).

In the present paper, a simple algorithm, similar in character to simulated annealing, is described which is capable of producing super-resolution maps from binary input images. This simple 'pixel-swapping' optimization algorithm allocates randomly hard binary classes to sub-pixels initially. Thereafter, the spatial *location* of the hard classes is altered, rather than the attribute value at each location (c.f. the HNN). . The algorithm is similar in concept to the set of techniques known as simulated annealing in a geostatistical framework (this is especially obvious given the use of the variogram in Tatem *et al*. (2001c)). Thus, in the future, readily available geostatistical algorithms (Gooverts, 1997) may provide alternative solutions to the HNN.

The new pixel-swapping optimization algorithm for super-resolution mapping is described in the next section.

3. Method

The simple optimization algorithm presented here is both simple and efficient. It is designed to take, as input, an image of land cover proportions in c=2 classes (probably obtained by application of a soft classifier to a remotely sensed image). First, the pixel proportions are transformed into sub-pixel hard land cover classes allocated randomly within each pixel. Once allocated, only the spatial arrangement of the pixels can vary, not the actual attribute values. Further, the number of sub-pixels allocated within each pixel remains fixed (corresponding to the pixel proportion constraint in the HNN approach). However, this is not a constraint: the proportions at the pixel level cannot vary.

Given the above random initialization, the objective is to vary the spatial arrangement of the sub-pixels in such a way that the spatial correlation between neighbouring sub-pixels (both *within* and, perhaps more importantly, *between* pixels) is maximised. There are many possible

approaches. The one adopted here is described below. It comprises two basic steps. First, for every sub-pixel the attractiveness O_i of the location is predicted as a distance weighted function of its neighbours:

$$O_i = \sum_{j=1}^{n} l_{ij} z(\mathbf{x}_j)$$

where n is the number of neighbours, $z(\mathbf{x}_j)$ is the value of the (binary) class z at the jth pixel location $\mathbf{x}_j$, and l_{ij} is a weight predicted as:

$$l_{ij} = \exp\left(\frac{-h_{ij}}{a}\right)$$

where h_{ij} is the distance between the pixel location for which the attractiveness is desired $\mathbf{x}_i$ and the location of the neighbour $\mathbf{x}_j$, and a is the non-linear parameter of the exponential model. The exponential weighting function chosen here is essentially arbitrary and several alternatives such as a simple inverse distance weighting function or the Gausian model could be used instead. The choice of non-linear parameter and the number of nearest neighbours are both important considerations and these will be revisited in the discussion.

Once the attractiveness of each sub-pixel location has been predicted based on the current arrangement of sub-pixel classes the optimization algorithm ranks the scores on a *pixel-by-pixel* basis. For each pixel, the least attractive location currently allocated to a '1' (i.e., a '1' surrounded mainly by '0's) is stored. Similarly, the most attractive location currently allocated to a '0' (i.e., a '0' surrounded mainly by '1's) is also stored. If the attractiveness of the least attractive location is less than that of the most attractive location then the classes are swapped for the pixel in question. If it is more attractive, no change is made.

The above two-stage process is repeated such that a solution is approached *iteratively*. The process can be stopped either at a fixed number of iterations or when the optimization algorithm fails to make a change.

4. Examples

In all of the examples discussed in this section the number of nearest neighbours used was 2 (i.e., the clique was second-order) and the non-linear parameter of the exponential model a was set to 5.

4.1 Simple shapes

To test the performance of the algorithm two simple geometric shapes (circle and linear feature) were simulated using the Splus™ software. These shapes are shown in Figures 2a and 3a. These images were aggregated into pixels of size 7 sub-pixels by 7 sub-pixels to form pixel-level images of 5 pixels by 5 pixels (Figures 2b and 3b). These images of proportions were used to simulate (i.e., used in place of) real soft-classified remotely sensed images of land cover proportions. These images represent the sole input to the optimization algorithm.

Given the images of pixel-level proportions the optimization algorithm allocated randomly each sub-pixel to a binary hard class initially such as to maintain the original pixel proportions (Figures 2c and 3c). Thereafter, the optimization algorithm made a maximum of one swap per-pixel for each iteration. This maximum helps to avoid local minima. For each of the simple images shown in Figures 2 and 3 less than 10 iterations was sufficient to achieve convergence. The circle was reproduced well (Figure 2g). This result is comparable to that of Tatem *et al*. (2001a). The circle is the most compact 2-dimensional shape and thus it is the easiest for the algorithm to reproduce. For the linear feature (Figure 3) the result was also acceptable visually, although there was a problem near the boundaries of the image (edge effect). Again, this problem was also identified by Tatem *et al*. (2001a). Thus, it appears that the simple optimization algorithm is capable of producing acceptable results, at least for simple geometric shapes.

The root mean square error (RMSE) is plotted in Figures 2h (circle) and 3h (linear feature). Clearly, the RMSE decreased with each iteration of the optimization algorithm until convergence. The RMSEs of the predicted images (Figures 2g and 3g) are much less than for the initial random allocations (Figure 2c and 3c).

4.2 Remotely sensed target identification

The most obvious application for the algorithm presented (i.e., spatial clustering of a *binary* field) is target detection in remotely sensed images. To test the algorithm further, a more 'realistic' irregular polygon shape was created using the `locator` command in Splus™ (Figure 4a). This shape does not correspond to any real feature to be found in remotely sensed images, but it does represent a more complex geometry than the circle and linear feature. First, since it was created by hand it does not correspond to any simple geometric shape. Second, it has two concave sides that make it more difficult to recreate using the optimization algorithm.

The image of the irregular polygon (Figure 4a) was degraded to a coarser spatial resolution to create an image of proportions (hypothetically of land cover) with pixels of 7 sub-pixels by 7 sub-pixels (Figure 4b). The optimization algorithm was applied to the image of

proportions shown in Figure 4b. Initially, the proportions per-pixel were used to allocate hard binary classes to the sub-pixels within each pixel (Figure 4c). Then, the optimization algorithm proceeded iteratively to a solution. Convergence was reached in around 15 iterations. This time the differences between the predicted image and the target image are more obvious. This is reflected in the RMSE error which does not decrease to as small a proportion of the initial RMSE. Nevertheless, even for this more complex shape the algorithm does perform reasonably well (compare the RMSE of the predicted image with that of the random allocation).

5. Discussion

5.1 Comparison to SSA

The optimization algorithm presented in this paper is simple and efficient. However, there are several important differences between its implementation here and the implementation of spatial simulated annealing algorithms (e.g., van Groenigen, 1999).

First, the SSA-type algroithm is based on the *random* selection of two sites and their swapping if the value of some objective function is increased as a result of the swap. In the present algorithm the entire image of some distance weighted function is calculated and the two sites (per-pixel) most in need of swapping (i.e., that will increase the value of the objective function the most) are swapped, as long as the objective function is increased as a result. This difference makes the present algorithm very fast.

Another important difference between the present algorithm and SSA is that the present algorithm has no stochastic element (the random initialization aside). A consequence of this may be that for larger, more complex shapes, there may be a greater likelihood of falling into local minima. For the features shown in Figures 2-4, several runs (with different initial sub-pixel allocations) resulted in the same predicted images, but this may not be the case for more complex shapes.

In the future, SSA-type algorithms will be implemented and compared to the present algorithm.

5.2 Choice of parameters

Several parameters (including the choice of distance weighted function) needed to be set by the investigator, and of these the most important were the number of neighbours and the non-linear parameter of the exponential function. The number of neighbours was initially set to 1 to increase the speed of the algorithm, but this was found to produce unsatisfactory results. Numbers of neighbours larger than 2 did actually lead to slightly more accurate prediction in general. However, the use of Splus™ to implement the optimization algorithm meant that run times were longer than would be expected for, say, the equivalent C or Fortran code. Therefore, the number of neighbours was set to 2 in the present case to ensure speedy convergence. Where lower-level code is impemented, use of a larger number of neighbours may be profitable. However, an important consideration is that the number of neighbours should not be so large that a given sub-pixel is attracted to other sub-pixels that it cannot neighbour (e.g., from a separate feature or from a pixel that its own pixel does not neighbour).

The non-linear parameter of the exponential model was chosen through experimentation. The parameter was set to be quite large (5 pixels) relative to initial personal expectation (remember, in geostatistical terms a non-linear parameter of 5 pixels relates to an 'effective' range of about 15 pixels).

Also important was the zoom factor, the increase in spatial resolution from the pixel-level image of proportions to the sub-pixel level image of hard land cover classes. In the examples presented in this paper, the zoom factor was equal to 7. This number was chosen because it has been used previously by Tatem *et al*. (2001a) and because it allowed rapid development of the algorithm. However, any zoom factor, including larger factors, could have been chosen. The results of Tatem *et al*. (2001a) suggest that larger zoom factors increase the precision of prediction.

5.3 Application to real images?

The optimization algorithm developed in this paper was implemented in Splus™. The time taken to converge for the linear feature (9 iterations) was approximately 6 minutes on a P-233 processor. The circle took approximately 10 minutes and the polygon took approximately 14 minutes. To apply the algorithm to real imagery it will be beneficial to write the code in a lower-level programming language such as C.

The algorithm implemented here could be applied to detect targets in real remotely sensed imagery. For example, targets could include floods (in particular, delineation of flood boundaries) and ice flows in coarser spatial resolution imagery, and vehicles and buildings in finer spatial resolution imagery. In the future, the algorithm will be extended to handle multiple land cover classes simultaneously to allow super-resolution land cover classification.

6. Conclusion

A simple, efficient algorithm has been presented as an alternative to the HNN algorithm for super-resolution target mapping in remotely sensed imagery. In its present form it allows the mapping of hard binary land

cover classes at a super-resolution (finer spatial resolution) from soft land cover proportions at an original spatial resolution. The algorithm was demonstrated to produce reasonable results for three simple images. Additional research is necessary to explore further the properties of the current algorithm, to extend the current algorithm to handle multiple land cover classes simultaneously and to find more efficient, but equally simple algorithms (e.g., SSA-type algrotihms) for super-resolution mapping.

7. Acknowledgements

The author thanks the Department of Geography, University of Toronto and, in particular, Dr. Ferenc Csillag and Kent Todd, for the use of computing and general office facilities during the summer of 2001.

8. References

Adams, J.B., Smith, M.O. and Johnson, P.E. (1985) "Spectral mixture modelling: a new analysis of rock and soil types at the Viking Lander 1 site," *Journal of Geophysical Research*, vol. 91, pp. 8098-8112.

Aplin, P. and Atkinson, P.M. (2001) "Per-field classification using sub-pixel allocation of soft land cover proportions," *International Journal of Remote Sensing* (in press).

Atkinson, P.M., Cutler, M.E.J. and Lewis, H. (1997) "Mapping sub-pixel proportional land cover with AVHRR imagery," *International Journal of Remote Sensing*, vol.18, pp. 917-935.

Atkinson, P.M. and Tate, N.J. (2000), "Spatial scale problems and geostatistical solutions: a review," *Professional Geographer*, vol. 52, pp. 607-623.

Atkinson, P.M. and Tatnall, A.R. (1997) "Introduction: neural networks in remote sensing," *International Journal of Remote Sensing*, vol. 18, pp. 699-709.

Bezdek, J.C. (1981) "*Pattern Recognition with Fuzzy Objective Function Algorithms*" . Plenum: New York.

Bezdek, J.C., Ehrlich, R. and Full, W. (1984) "FCM: The fuzzy *c*-means clustering algorithm," *Computers and Geosciences*, vol. 10, pp. 191-203.

Brown, M., Gunn, S.R. and Lewis, H.G. (1999) "Support vector machines for optimal classification and spectral unmixing," *Ecological Modelling*, vol.120, pp. 167-179.

Curran, P.J. and Atkinson, P.M. (1998) "Geostatistics and remote sensing," *Progress in Physical Geography*, vol.22, pp. 61-78.

Flack, J., Gahegan, M. and West, G. (1994) "The use of sub-pixel measures to improve the classification of remotely sensed imagery of agricultural land," *Proceedings of the 7th Australasian Remote Sensing Conference*, Melbourne, pp. 531-541.

Foody, G.M. (1998) "Sharpening fuzzy classification output to refine the representation of sub-pixel land cover distribution," *International Journal of Remote Sensing*, vol.19, pp. 2593-2599.

Garcia-Haro, F.J., Gilabert, M.A. and Melia, J. (1996) "Linear spectral mixture modelling to estimate vegetation amount from optical spectral data," *International Journal of Remote Sensing*, vol.17, pp. 3373-3400.

Goovaerts, P. (1997) "*Geostatistics for Natural Resources Evaluation*". Oxford University Press: New York.

Hopfield, J. and Tank, D.W. (1985) "Neural computation of decisions in optimization problems," *Biological Cybernetics*, vol.52, pp. 141-152.

Journel, A.G. (1996) "Modelling uncertainty and spatial dependence: sochastic imaging," *International Journal of Geographical Information Systems*, vol. 10, pp. 517-522.

Jupp, D.L.B., Strahler, A.H. and Woodcock, C.E. (1988) "Autocorrelation and regularization in digital images I. Basic theory," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 26, pp. 463-473.

Jupp, D.L.B., Strahler, A.H. and Woodcock, C.E. (1989) "Autocorrelation and regularization in digital images II. Simple image models," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 27, pp. 247-258.

Justice, C.O., and Townshend, J.R.G. (1981) "Integrating ground data with remote sensing," in *Terrain Analysis and Remote Sensing* (J.R.G. Townshend, Ed.), Allen and Unwin: London, pp. 38-101.

Matheron, G. (1965) "*Les variables régionalisées et leur estimation*". Masson: Paris.

Schneider, W. (1993) "Land use mapping with subpixel accuracy from landsat TM image data," *Proceedings of the 25th International Symposium on Remote Sensing and Global Environmental Change*, pp. 155-161.

Schneider, W. (1999) "Land cover mapping from optical satellite images employing subpixel segmentation and radiometric calibration," in I. Kanellopoulos, G. Wilkinson and T. Moons, *Machine Vision and Advanced Image Processing in Remote Sensing*. Springer: London.

Steinwendner, J., Schneider, W. and Suppan, F. (1998) “Vector segmentation using spatial subpixel analysis for object extraction,” *International Archives of Photogrammetry and Remote Sensing*, vol.32, pp. 265-271.

Tatem, A.J., Lewis, H.G., Atkinson, P.M. and Nixon, M.S. (2001a) "Super-resolution target identification from remotely sensed images using a Hopfield neural network," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, pp. 781-796.

Tatem, A.J., Lewis, H.G., Atkinson, P.M. and Nixon, M.S. (2001b) "Multiple class land cover mapping at the sub-pixel scale using a Hopfield neural network," *International Journal of Applied Earth Observation and Geoinformation* (in press).

Tatem, A.J., Lewis, H.G., Atkinson, P.M. and Nixon, M.S. (2001c) "Land cover simulation and estimation at the sub-pixel scale using a Hopfield neural network," *Remote Sensing of Environment* (in press).

Verhoeye *et al.* (2000) *IGARSS2001 Scanning The Present and Resolving the Future*, IEEE, Sydney, Australia, CD-ROM.

Van Groenigen, J.-W. (1999) "*Constrained optimisation of spatial sampling. A Geostatistical Approach*," ITC Publication Series, No. 65, ITC: Enschede, the Netherlands.

Woodcock, C.E., and Strahler, A.H. (1987) "The factor of scale in remote sensing," *Remote Sensing of Environment*, vol. 21, pp. 311-322.

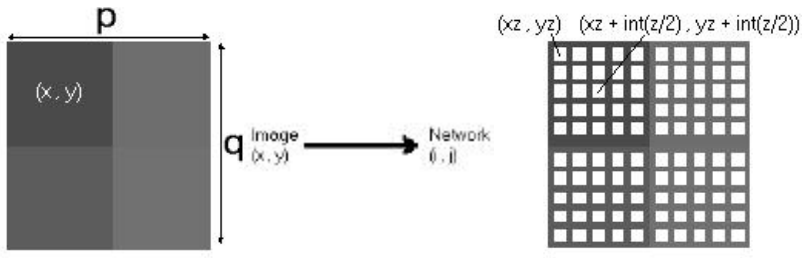

(a) (b)

Figure 1. (a) 2x2 pixel image, *p* and *q* represent the image dimensions, (*x*,*y*) represents the image co-ordinates; (b) Representation of a Hopfield network for the image in (a), *i* and *j* represent the neuron co-ordinates (int = integer value).

Figure 2. Super-resolution mapping of a circle: (a) test image, (b) image of proportions input to the optimization algorithm, (c) random initial allocation to sub-pixels, (d) solution after 4 iterations, (e) solution after 8 iterations, (f) solution after 12 iterations, (g) solution after 16 iterations and (h) plot of RMSE against number of iterations.

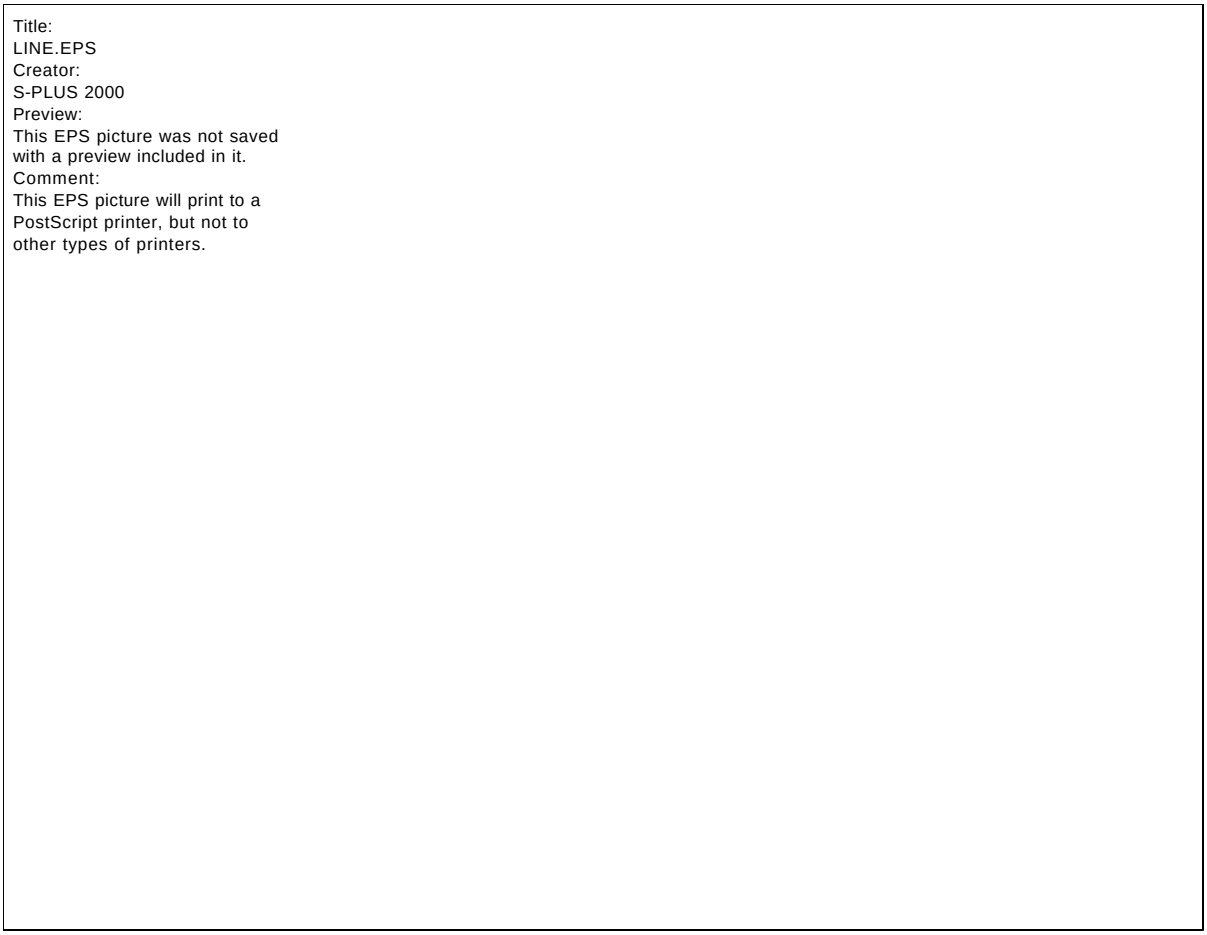

Figure 3. Super-resolution mapping of a linear feature: (a) test image, (b) image of proportions input to the optimization algorithm, (c) random initial allocation to sub-pixels, (d) solution after 3 iterations, (e) solution after 6 iterations, (f) solution after 9 iterations, (g) solution after 12 iterations and (h) plot of RMSE against number of iterations.

Figure 4. Super-resolution mapping of an irregular polygon: (a) test image, (b) image of proportions input to the optimization algorithm, (c) random initial allocation to sub-pixels, (d) solution after 5 iterations, (e) solution after 10 iterations, (f) solution after 15 iterations, (g) solution after 20 iterations and (h) plot of RMSE against number of iterations.