OPTIMAL PROCESSING TECHNIQUES FOR SAR

David Stewart, Rod Cook, Ian McConnell

NA Software, Roscoe House, 62 Roscoe Street, Liverpool, UK
Tel: +44 151 709 4738, Fax: +44 151 709 5645
email: stos@nasoftware.co.uk

Chris Oliver

DRA, Malvern, Worcestershire, UK
Tel: +44 1684 895165, Fax: +44 1684 894481
email: chris@sar.dra.hmg.gb

KEY WORDS: SAR, Global Optimisation, Segmentation, Classification.

ABSTRACT: In the history of SAR image processing, many algorithms have been proposed to tackle the problems of segmentation, classification and edge detection. They are typically heuristic in basis, and more successful on some types of imagery than others. With the development of global optimisation methods it has now become possible to produce optimal techniques; that is, those which can genuinely achieve the optimal solution of the posed problem. The problem is characterised by an objective function and the chosen optimisation technique. The most successful and wide-spread method has been simulated annealing and we detail its application in the fields of segmentation and classification. The performance of the resulting algorithms on various SAR imagery is given.

INTRODUCTION

Simulated annealing is alone amongst other global optimisers, such as taboo search and genetic algorithms, in that it guarantees convergence to a global minimum. Geman and Geman (1984) were among the first to adapt the technique to image processing and there have been many other applications since. Typically, however, a less than complete understanding of the nature of the problem has led to the introduction of heuristics into both the objective function and the algorithmic implementation. After a brief resume of the basics of simulated annealing, we will describe a cost function for evaluating image segmentations and classifications that is free from heuristics. We will then give a brief description of an algorithm incorporating this cost function and demonstrate its applicability to all sorts of SAR imagery.

SIMULATED ANNEALING

Simulated annealing is a technique for finding the global minimum of a given function, called the objective or cost function, over a configuration space $\mathcal{S}$. The cost function is such that the lower the value of $C(s \in \mathcal{S})$, the better the corresponding configuration. To be successful it is necessary that the implementation allows the algorithm access to all areas of the configuration space.

In this paper we will illustrate the application of this method in the fields of segmentation and classification. $\mathcal{S}$ will thus represent the space of all possible segmentations or all possible classifications into a fixed number of regions or classes respectively.

Simulated annealing algorithms proceed by randomly changing from one state in the configuration space, say s_i, to another, s_j. By evaluating the change in cost, $\Delta C_{ij} = C(s_j) - C(s_i)$, the algorithm decides via a probabilistic acceptance criterion whether to accept the new configuration or keep the current one. The acceptance criterion is dependent on the value of a parameter T, known as the temperature, which itself is updated after a fixed number of configuration changes have been considered. The criterion is such that a configuration change entailing an increase in cost is more likely to be accepted at a higher temperature than at a lower one. By slowly decreasing the value of T, we will avoid being trapped in local minima and will reach the optimal global solution.

The parameters within the simulated annealing algorithm need careful tuning for the particular application, in particular the behaviour of the temperature T. Aarts et al. (1986) give a methodology for approximating the parameters which works reasonably well in practise.

THE COST FUNCTION

The definition of the objective function is dependent on the intended application and on what particular image features are desired in the final image. For general purpose segmentation and classification, the objective function is commonly composed of two competing terms. The first is derived from a single point statistical model of SAR imagery whilst the second limits the behaviour of the boundaries between segments or classes. The addition of the shape term is necessary for otherwise the segments in the optimal segmentation would simply follow the lines of speckle in the original image whilst the optimal classification would just classify the speckle. Minimising this objective function is then equivalent to maximising the likelihood of the data fitting the segmentation or classification with the appropriate degree of qualification upon region shape. It is often the lack of knowledge about the relationship between these two competing terms that leads to the introduction of heuristics into the cost function. We can write the cost function as

$$C = C_\Lambda + sC_S$$

where C_Λ denotes the likelihood term and C_S the shape term. We define examples of these for the instance of intensity segmentation and classification below. The constant s quantifies the relationship between the two terms and we indicate how this value can be rigorously defined later.

The Individual Terms

For the Likelihood term, C_Λ, we use the well known fact that the Gamma distribution provides a good model for SAR intensity data. We suppose, therefore, that we have an image containing r homogeneous regions or classes of gamma distributed L-look data. Let the data be represented by $\{x_i | i = 1, \cdots, N\}$ and I_j and μ_j, $j = 1, \cdots, r$, denote the indexing sets and means of the appropriate regions or classes. Then the probability for the data fitting the distribution is given by

$$\prod_{j=1}^{r} \prod_{i \in I_j} \frac{L^L}{\Gamma(L)\mu_j^L} x_i^{L-1} \exp(-Lx_i/\mu_j)$$

Let $N_j = |I_j|$ and let $\overline{m_j}$ denote the sample mean, i.e. $\overline{m_j} = \frac{1}{N_j}\sum_{i \in I_j} x_i$. The log-likelihood is then calculated to be

$$\Lambda = a - L\sum_{j=1}^{r} N_j(\log(\mu_j) - \overline{m_j}/\mu_j)$$

where the constant terms have been grouped into a. When the parameters μ_j are not known *a priori*, the estimates $\overline{\mu_j} = \overline{m_j}$ provide the maximum likelihood solution for the log-likelihood. Since we want to maximise the likelihood, and simulated annealing is designed to minimise a cost function, we multiply the log likelihood by -1 to give the cost function:

$$C_\Lambda = \sum_{j=1}^{r} N_j \log(\overline{m_j})$$

Note that any constants can be removed since the annealing is driven solely by the change of cost.

There have been many proposed models to describe the type of shape we expect segments to have and, analogously, the connectivity we expect of classes. Most of these are expressions of the size of the boundary or edge set. Here we choose one of the most simple and define the shape term to be given by

$$C_S = \frac{1}{2}\sum_{i}\sum_{j \in \mathrm{Nhbd(i)}} (1 - \delta_{L_i L_j})$$

where L_i is the region label for a pixel i and Nhbd(i) a defined neighbourhood.

The Relationship Between the Terms

To define the shape penalty constant s we need to understand how the terms of the cost function interact. This is likely to change with respect to the looks of the image and the number of regions or classes.
Since the annealing is driven by the change in cost

$$\Delta C = \Delta C_{\Lambda} + s\Delta C_{S}$$

it is necessary to determine the distributions for the likelihood difference and the shape cost difference. We want to find these distributions for uniform background speckle since in that way we can quantify the extent to which the change in shape term dominates the change in likelihood.
To be more explicit, we wish to calculate a probability of false alarm, i.e. the probability, given there is no real edge, that the edge position is at a certain pixel on the strength of the change in likelihood rather than the change in shape. Mathematically we define

$$\begin{aligned} P_{fa} &= \mathrm{Prob}(\Delta C_{\Lambda} < -s\Delta C_{S} | \Delta C_{S} > 0) \\ &\quad + \mathrm{Prob}(\Delta C_{\Lambda} > -s\Delta C_{S} | \Delta C_{S} < 0) \\ &\quad + \mathrm{Prob}(|\Delta C_{\Lambda}| > 0 | \Delta C_{S} = 0) \end{aligned}$$

Having a universal value for P_{fa} will then give the same degree of shape effect for whatever SAR imagery fits our models and for whatever the value of looks and region or class size.
To calculate the above we thus need the distributions for the log-likelihood difference and the change in shape. Both are analytically intractable but can be approximated via simulation. The variation of the distributions with respect to region sizes, number of classes and looks of the data can also be approximated in this fashion. Typical distributions are given in figure 1.

THE IMPLEMENTATION

There are two main considerations when implementing an algorithm to minimise the cost function of the previous section. Namely

1. In order to ensure the success of the annealing it is necessary to provide the capability to visit, in theory, all points of the configuration space. As such, for the results that follow, we employ a free topological model and alter the segmentation on a pixel by pixel basis. By making the smallest admissible changes, we also expect to achieve results of the highest possible resolution.
2. To avoid the use of heuristics as far as possible.

An outline of the segmentation and classification algorithm is as follows:

1. Initialise the program. An initial tessellation into a fixed number of regions or classes is required along with the determination of the annealing parameters and the shape penalty constant s using the analysis described above.
2. For a number of outer iterations.
   (a) Set T. This is usually a function of the number of outer iterations.
   (b) For a number of inner iterations.
      i. Perturb the current segmentation or classification. In the segmentation case this involves moving a pixel between neighbouring segments and as long as this maintains segment connectivity and existence. In the classification case this is simply a random change of the pixel's class label.
      ii. Evaluate the change in cost, ΔC.
      iii. Refer to the acceptance criteria. The acceptance probability most widely used equals 1 if $\Delta C < 0$ and $\exp(-\Delta C/T)$ if $\Delta C > 0$.
      iv. If the new configuration is rejected, reinstate the old one.
3. Tidy up. In the case of segmentation it is feasible to run a post-process merging stage to remove spurious regions.

RESULTS

We apply the algorithms detailed above to various SAR imagery in figures 2 to 5. The ratio image of figure 2 indicates how successfully the segmentation fits the original data. The segmentation and classification of the AIRSAR image, figure 4, are both visually impressive

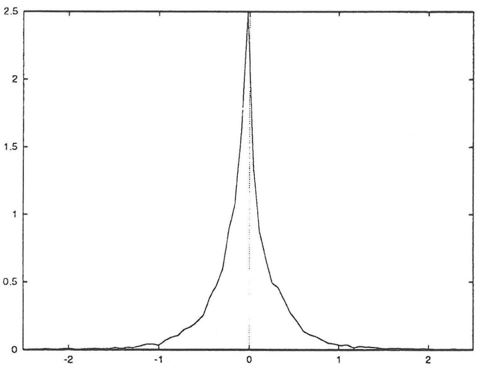

(a) Example of the distribution for ΔC_Λ derived from simulation.

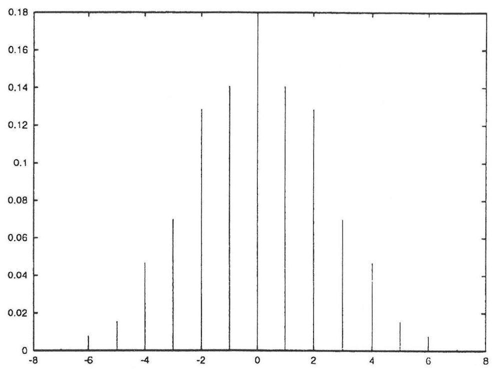

(b) Example of the distribution for ΔC_S derived from simulation. It is a mixture of several multinomial distributions.

Figure 1: Typical distributions for the likelihood difference and change in shape

whilst differing in some small details. Figure 5 gives the result of using the segmentation scheme in the important application of differentiating between virgin rainforest and clearing.

The shape penalties used for each segmentation and classification are widely different, yet all correspond to the same probability of false alarm. It can be seen that the resultant shape effect is roughly equivalent in all the images. Future work involves finding improvements in the approximations for the log-likelihood difference and the change in shape distributions. It is hoped that this will make this effect even more apparent.

REFERENCES

1. Aarts, E.H.L., de Bont, F.M.J., Habers, E.H.A., van Laarhoven, P.J.M., 1986. Parallel implementations of the Statistical Cooling Algorithm. *INTEGRATION, the VLSI journal*, 4, pp. 209-238.

2. Cook, R., McConnell, I., Stewart, D., and Oliver, C.J., 1996. Segmentation and Simulated Annealing. *Proc. EUROPTO (Sicily) 1996 SPIE*, 2958, pp. 134-154.

3. Geman, S., and Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6), pp. 721-741.

4. Oliver, C.J., Blacknell, D., and White, R.G., 1994. Optimum edge detection in SAR. *IEE Proc. Radar, Sonar and Navig*, 143(1), pp. 31 - 40.

5. Oliver, C.J., Stewart, D. and McConnell, I., 1996. Optimum Texture Segmentation in SAR. *Proc. EUSAR (Konigswinter)* pp. 234-245.

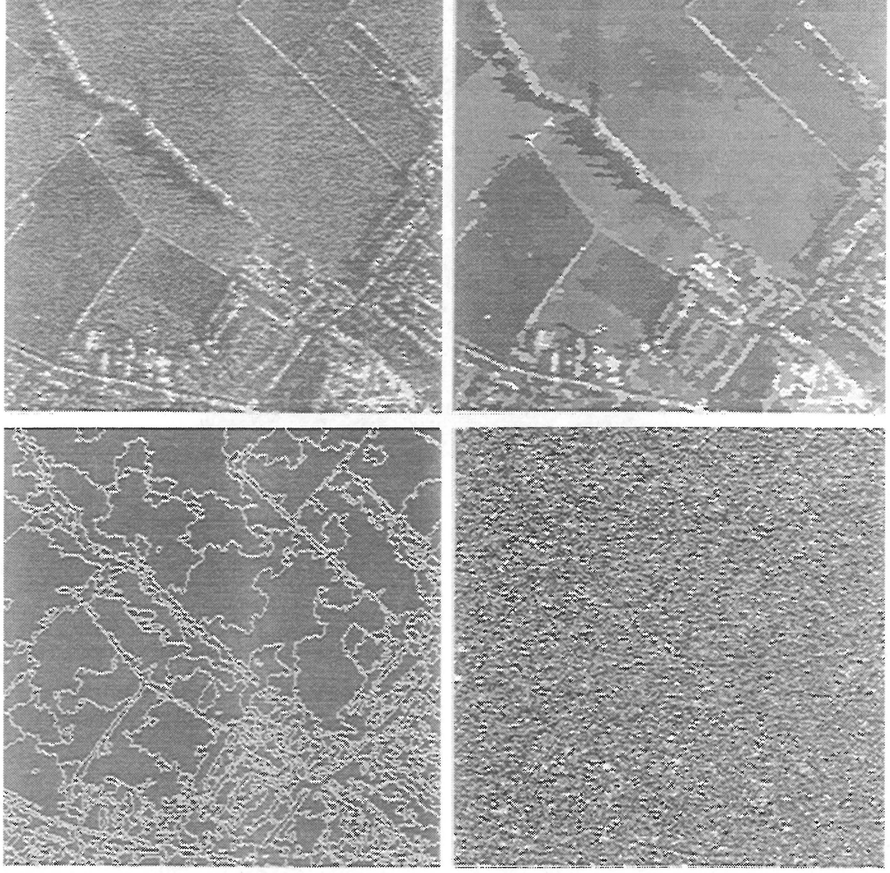

Figure 2: Top: X-band SAR and annealing segmentation. Bottom: Edge map and ratio image.

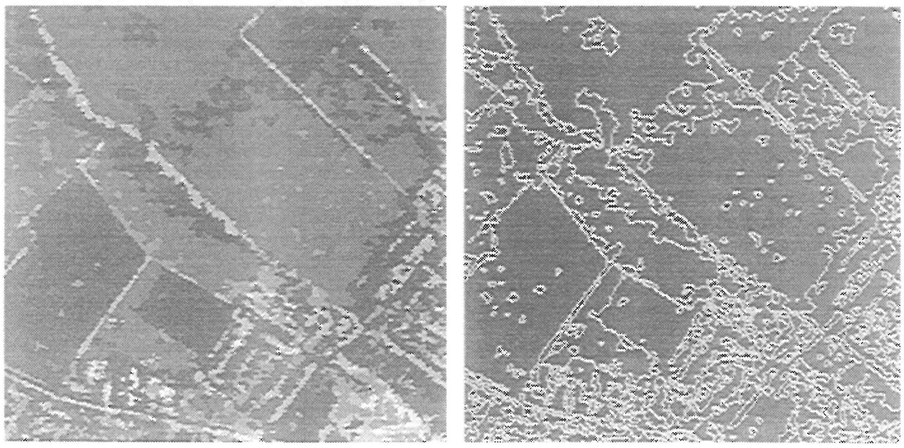

Figure 3: Annealing classification and edge map for the SAR image of figure 2

Figure 4: Top: AIRSAR image. Middle: Annealing segmentation and edge map. Bottom: Annealing classification and edge map.

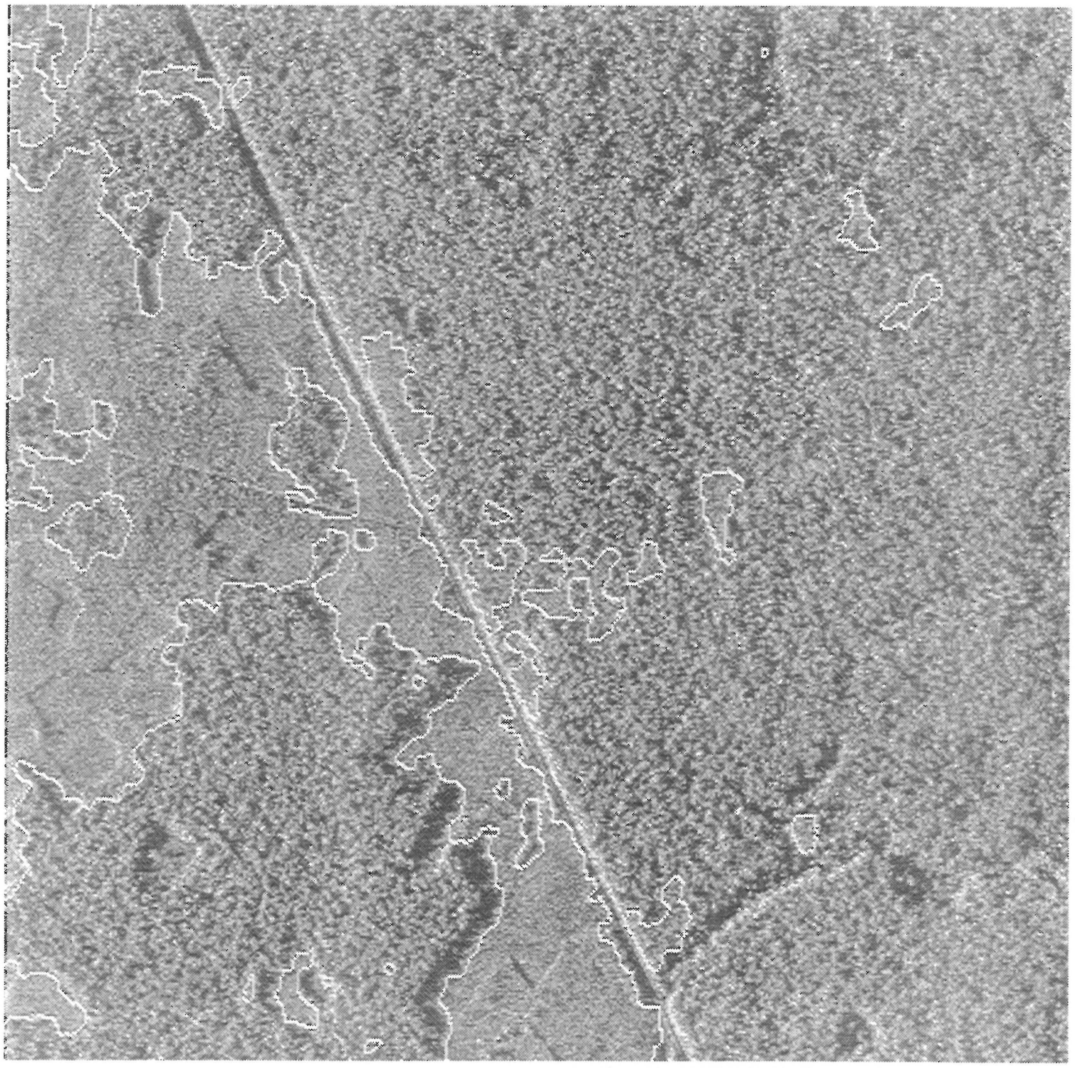

Figure 5: Identification of clearing in the Brazilian rainforest. The original image was processed by a 7 by 7 mask to provide an image of normalised log estimates. This image was processed by the annealing segmenter and then thresholded. Finally, the resulting edge map was overlaid on the original data.