The removal of geometrical distortion in airborne reconnaissance imagery caused by vertical instability of the sensor

K. J. Hermiston, D. M. Booth

Defence Evaluation and Research Agency
St Andrews Road
Malvern
WR14 3PS
England
kjherm@dera.gov.uk
dmbooth@dera.gov.uk

Abstract

The geometrical fidelity of airborne reconnaissance imagery acquired by state of the art "pushbroom" sensors has been found during trials to be susceptible to vibration of the sensor pod. The resulting imagery displays non-stationary vibration known amongst the reconnaissance community as "sensor jitter"; artefacts which must be removed prior to exploitation. Restoration of this unusual distortion has received little or no attention in the literature.

An algorithm is proposed for the removal of image jitter based on the global minimisation of a cost function consisting of two terms. The first seeks to maximise coherence of local image features, whilst the second maintains some consistency with the original image. Performances are compared using two optimisation strategies: simulated annealing; and a less computationally demanding application of Mean Field Approximation (MFA) to a mapping of the cost function onto a Hopfield Artificial Neural Network (ANN).

1 Introduction

An algorithm is described for removing geometrical distortions from optical and infra-red imagery obtained from a LOROP (LOng Range Oblique Photographic) sensor. The sensor is a pushbroom type, which comprises of a vertical column of sensor elements which are sampled periodically. The forward motion of the aircraft provides the horizontal shifts required to form a complete image. Geometrical distortion results from non-stationary, vertical vibration modes of the sensor which in turn results from a turbulent airflow around the sensor pod on a fast jet. The extent of the problem can be seen in Figures 1, 4, 7 and 10. Note that horizontal vibration manifests itself as local shortening and elongation of image features, and as such is not perceivable without reference to ground truth. The removal of the vertical component is an essential precursor to exploitation, particularly by automatic feature detection, image registration and sensor fusion techniques.

The "dewiggler" algorithm proposed here considers and exploits the vertical nature of the distortion. It makes no assumptions with regard to the periodicity of the jitter - it is time varying. The basis of the algorithm is that scenes change in a coherent way, and if the sensor is unstable then image coherence is weakened. The goal is therefore to generate a vertical shift for each column of pixels in the image, such that their overall effect is to minimise the difference in grey level between neighbouring columns and hence maximise image coherence.

2 Cost Function

The set of shifts is determined by minimising a cost function, C, which in some respects is similar to that of the weak elastic string proposed by Blake and Zizzerman [1]. Both have terms for measuring faithfulness to data but here we replace Blake and Zizzerman's minimum deformation term with one designed to maximise image coherence over local regions. The cost function, C, is given by

$$C = S + \alpha D \qquad \text{Equation 1}$$

where α is a weighting factor. The first term, S, is effectively a computationally cheap measure of the correlation between neighbouring columns, and is summed across the whole image. The similarity between two individual columns is given by the sum of the absolute grey level differences between corresponding pixels. In fact it is necessary to consider correlations between columns which are separated over a range of distances if a wide band of frequences are to be suppressed. However, due to scene changes, columns are likely to be less highly correlated when separated by greater distances. Hence, the cost function considers a range of separations, with their respective contributions weighted by a Gaussian centred at zero displacement.

Ignoring the problems at image boundaries, the cost associated with a given reconstruction f by the correlation term, S, is given by

$$S = \sum_{p=1}^{n} \left(\sum_{i=-w}^{w} \left(\frac{1}{\sqrt{2\pi\sigma^2}} e^{-(p-i)^2/2\sigma^2} \sum_{j=1}^{m} |f_{p+i,j} - f_{p,j}| \right) \right)$$

Equation 2

where M and N are the number of rows and columns in the image, and W represents the distance over which correlations will be considered; the assignment of weights is determined by σ. $f_{x,y}$ represents the grey level at image co-ordinates (x,y).

Other correlation measures have been investigated, notably Pearson's correlation coefficient and the effect of processing the gradient image, thereby emphasising the changes in the image. The former did not appear to improve performance significantly (but is more expensive to compute), whereas in some cases the latter resulted in errors caused by mismatched lines (an aliasing like effect).

The second term applies a penalty cost to the total displacement of the columns from their original positions.

$$D = \sum_{1}^{N} (u_i)^2 \qquad \text{Equation 3}$$

where u_i represents the vertical displacement that has been applied to column i. The D term effectively ties the image down, and in particular, it suppresses a tendency for the algorithm to transform sloping edge structures into steps or horizontal structures.

3 Optimisation using simulated annealing

Notionally, the algorithm operates by repeatedly executing the following set of operations: it applies a small shift to a given column, the cost function is evaluated, and if there is an improvement then the change is accepted. However, the nature of the problem is such that a seemingly bad move can prove advantageous in the long term. An appropriate optimisation approach is to apply simulated annealing which accepts poor changes, at least in the early stages. As time goes on the willingness to accept bad moves is gradually reduced (according to a specified cooling schedule) until eventually the system converges on a "best" solution. The reader is referred to Geman and Geman [2] for a more detailed description.

Occasionally, the system converges to a state in which the image contains step edges, particularly in parts of the image where the features have similar orientations. The step edges can be detected easily from the sequence of column displacements and then largely removed by interpolation. One way of avoiding the steps is to increase α, but that reduces the algorithm's ability to remove distortion. An alternative is to define a "break" (that is, a step in column displacement greater than some threshold) and only allow them to occur at cost. In general this has a smoothing effect but the underlying step edges remain. Smoothed steps cannot be identified readily and therefore correction of residual distortion by post-processing is not possible.

4 Results of simulated annealing optimisation

The result of applying the algorithm to Figure 1 is shown in Figure 2. This example is reasonably straightforward because the density of image features is high. The distorions at the targeted frequencies have been removed successfully. Some residual low frequency drift remains in this image and there is a suggestion that some high frequency noise may have been introduced. This must be expected of an imperfect system. Note that smoothing the column displacements will increase the noise rather than suppress it. Overall, the results show a significant improvement. Similar performance has been achieved on a range of imagery including the imagery shown in Figures 4 and 7. Problems arise when features are sparse (which suggests that the image contains little of interest) or when a sizable majority of the features are oriented similarly but other than horizontally, such as in Figure 10. This image is the worst case imaginable and

poses the added problem that it was scanned from a hardcopy and so may have suffered from missalignment and rebinning, that is, the perceived distortion may no longer be vertical. Figures 11 and 13 show the results after optimisation and interpolation, respectively. Figures 2,5 and 8 show how the interpretability of features of interest can be improved.

5 Algorithm implementation using Hopfield ANN and MFA

A Mean Field Approximation (MFA) applied to a mapping of the cost function onto a Hopfield Artificial Neural Network (ANN) was examined in an effort to overcome the computational demands of simulated annealing.

The approach is better understood when one considers that the jitter removal problem maps directly onto that of the travelling salesman. That is, image columns map onto cities, column shifts map onto tour-positions, and the grey level similarity between pixel columns equates to distances between cities.

Hopfield and Tank [3] proposed optimising the travelling salesman problem (TSP) by mapping it onto a Hopfield neural network. Later, Wilson and Pawley [4] reported problems achieving valid solutions, that is, a tour solution that fails to visit each city once. Even when a valid solution was calculated it was often non-optimal. However, the use of Mean Field Approximation (MFA), which assigns a single graded neuron to all the tour positions of a single city, and whose output is a unit vector in space, reduces the redundency in the neural system and has improved network convergence to valid and acceptable solutions. This approach is examined by Muller and Reinhardt [5] and has been investigated by Van der Bout and Miller [6], Peterson and Soderberg [7] and Bultan and Aykanat [8] and found to facilitate reliable optimisation of the TSP. The annealing effect implemented with the MFA adapts the Hopfield network from a local optimiser into a global optimiser.

The "dewiggling" measure was formulated onto a Hopfield ANN as

$$E=\frac{d_{\max}[X]}{2}\sum_{i=0}^{S}\sum_{\substack{Y=X-W\\Y\neq X}}^{X+W}\sum_{\substack{j=0\\j\neq i}}^{S}V_{X,i}V_{X,j}$$

$$+\frac{1}{2}\sum_{i=0}^{S}\sum_{\substack{Y=X-W\\Y\neq X}}^{X+W}\sum_{\substack{j=0\\j\neq i}}^{S}G_W d(X,Y,i-j)V_{X,i}V_{Y,j}$$

Equation 4

where d_{max} [X] is assigned a value slightly larger than the maximum difference of pixels between any image column within a coherence width (W) and the column (X) over the allowable column-shift interval S. N is the total number of image-columns in the frame. G_W is the Gaussian weight function centred on the randomly selected vector (X). d(X,Y,i-j) is the grey level dissimilarity between image columns X and Y, where column Y has undergone a shift of j-i relative to X. $V_{Y,j}$ is the neuron output associated with column Y when shifted by j-i. The use of the $j\neq i$ in the summations of Equation 4 encourages consistency with the original image and effects the elastic term in the algorithm cost function.

The first term in Equation 4 introduces penalties for invalid solutions, that is, there should be one shift value per image column. The second term seeks to minimise the grey level dissimilarity between column X and its neighbours. Localised correlation is enforced by Gaussian weighting.

The set of graded neurons associated with the image columns can be considered to be a physical spin system. And the thermodynamic expectation value of the spin vectors can be calculated using saddle-point approximation. The canonical partition function of the spin vector ensemble in thermal equilibrium can be considered to be

$$Z=\sum_{[s]}e^{-E(\mathbf{s_1},\ldots,\mathbf{s_n})/T} \qquad \text{Equation 5}$$

which can be expressed as an integration over a set of conjugate, auxilliary variables $\mathbf{u_i}$, $\mathbf{v_i}$, $i\in 1,N$ as

$$Z=\sum_{[s]}\int\left(\prod_i d\mathbf{v_i}\right)\int\left(\prod_i\frac{d\mathbf{u_i}}{2\pi}\right)e^{-E(\mathbf{v_1},\ldots,\mathbf{v_n})/T}\prod_i e^{i\mathbf{u_i}.(\mathbf{s_i}-\mathbf{v_i})}$$

Equation 6

An analytic continuation to imaginary values of $\mathbf{u_i}$ terms such that $i\mathbf{u_i}\rightarrow\mathbf{u_i}$ produces

$$Z=\int\prod_i d\mathbf{v_i}\int\left(\prod_i\frac{d\mathbf{u_i}}{2\pi}\right)e^{-E(\mathbf{v_1},\ldots,\mathbf{v_n})/T-\sum_i \mathbf{u_i}.\mathbf{v_i}}\sum_{[s]}\prod_i e^{\mathbf{u_i}.\mathbf{s_i}}$$

Equation 7

and using the summation property

$\sum_{[s]}\prod_i e^{\mathbf{u_i}.\mathbf{s_i}}=\prod_i\sum_{s_i}e^{\mathbf{u_i}.\mathbf{s_i}}$ gives

$$Z = \int\prod_i dv_i \int (\prod_i \frac{du_i}{2\pi}) e^{-E(v_1,\ldots,v_n)/T - \sum_i u_i.v_i + \sum_i \ln z(u_i)}$$

Equation 8

Since the number of columns in the image will usually be large N>>1, the integrand will only return appreciable contributions from the region in the vicinity of the minimum and the saddle-point approximation can be used. At the saddle-point, the first derivatives of the exponent with respect to all the integration variables vanish and it can be shown that

$$\mathbf{v_i} = \mathbf{F_N(u_i)} \qquad \text{Equation 9}$$

$$\mathbf{u_i} = -1/T\ \delta E/\delta \mathbf{v_i} \qquad \text{Equation 10}$$

where $\mathbf{F_N(u)}$ has components

$$F_N^{\alpha}(u) = \frac{e^{u\alpha}}{\sum_{\beta} e^{u\beta}} \text{ and } \sum F_N^{\alpha}(u) = 1$$

With reference to Equation 4, the mean field of a neuron $V_{x,i}$ can be evaluated as

$$U_{X,i} = d_{\max}[X] \sum_{\substack{Y=X-W \\ Y\neq X}}^{X+W} \sum_{j=0,\, j\neq i}^{S} V_{X,j} + \sum_{\substack{Y=X-W \\ Y\neq X}}^{X+W} \sum_{j=0,\, j\neq i}^{S} G_W d(X,Y,j) V_{Y,j}$$

Equation 11

The neuron outputs, can be interpreted as the probabilities that the image column X should be shifted to position i. They are given by

$$V_{X,i} = \frac{e^{-E_{X,i}/T}}{\sum_{j=0}^{S} e^{-E_{X,j}/T}} \qquad \text{Equation 12}$$

The energy and mean field approximations are derived from the image columns immediately surrounding the randomly selected column, as these exhert strongest coherence. These local evaluations produce a noisy movement towards optimisation. At high temperatures, all the neurons associated with a particular image column will have similar values. However, as the temperature cools, the value of the optimum neuron will increase while the value of all the others decrease. The summation of all the neuron outputs associated with a particular column will always remain unity throughout the cooling process.

The ANN based restoration algorithm is summarised below.

1 Normalise pixel grey levels within each column to sum to unity.
2 Calculate and store the pairwise dissimilarities between pixel columns under all possible shifts within coherence width (W).
3 Initialise the neuron weights and temperature
4 Select an image column at random (X)
5 Calculate the present energy using Equation 4
6 compute mean field using Equation 11
7 Update the Neuron (Vxi) using Equation 12
8 Calculate new energy using Equation 4
9 If new energy is greater than present energy decrement temperature.
10 If temperature is in cooling range then goto step 4, otherwise stop execution.

6 Comparitive numerical results of simulated annealing and ANN implementations

A quantitative assessment of algorithm performance is a problem in the absence of ground truth. The solution adopted here was to introduce jitter artificially into reconnaisance imagery collected from another source: an optical framing sensor onboard a Jaguar aircraft. The jitter shifts were modelled by superpositioning low amplitude sinusoids of differing frequencies. The original image, providing ground-truth, and the artificially distorted version can be seen in Figures 14 and 15, respectively. Imagery restored using the simulated annealing and ANN optimisation are shown in Figures 16 and 17 respectively. Multiple runs of both implementations, using different pseudo-random sequence seeds, produced the statistical results shown in Table 1.

Table 1

Optimisation	pixel MSE from ground truth
Simulated annealing	5.514
ANN with MFA	0.310

Visual examination of Figure 16 shows that the simulated annealing implementation of the algorithm has smoothed the vibrations well. However, owing to the static, non-adaptive elastic weight, significant localised

drifting becomes evident when overlayed onto the ground truth image. The visual result of the ANN implementation shows another good restoration but displays more residual ripples than the image restored by the original technique. However, reference to Table 1 reveals that the ANN implementation produces better restoration to the ground truth image when small and stationary distortions are applied. In real-world examples, where distortions are not so controlled, the drifting characteristic of the simulated annealing implementation has been observed to produce better visual restorations. Specifically, the smoothing effect of the simulated annealing approach sequentially correcting columns, appears to recover from the low frequency distortions which lie beyond the bandwidth capabilities of the Gaussian window width. This is exemplified by the ANN restoration image in Figure 9. When the neurons of the ANN system are expanded to encompass large amplitude distortions, it has been noticed that the restoration solution degrades. It is thought that the degradation in the quality of the ANN solution arises from the increased system size which has also been encountered with TSP applications.

7 Conclusions

The two techniques for "dewiggling" imagery have proven to be effective on a varied sample of optical and infra-red imagery. The simulated annealing implementation of the "dewiggling" algorithm has been seen to provide a visually good restoration of imagery which only suffers from small localised drifting. The drifting is thought to arise from the use of a static elastic weight which does not adapt to column pixel statistics and encounters regions where the correlation is not conterbalanced by the elastic term of the cost function. In an example "worst case" image there remains the problem of sparse features, containing inclined lines with respect to the image axes, producing discontinuities which can be partially removed through post-processing. However, in the majority of restored test images, this was not necessary and the algorithm performed robustly.

Optimisation through simulated annealing has been recognised to be computationally intensive, even using a crude approximation of the Geman and Geman cooling schedule. A second implementation of the dewiggling algorithm which uses a Hopfield ANN with MFA to optimise the cost function, has demonstrated an efficient approach which produces an acceptable solution quality with greatly improved solution times. Such an implementation may find most use in real-world applications such as tactical reconnaissance analysis where restoration is required in a matter of seconds rather than hours.

References

[1] Blake, A. and Zisserman, A. Visual Reconstruction, MIT Press, 1987.

[2] Geman, S. and Geman, D. "Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images", IEEE Trans. PAMI-5, pp.721-741, 1984.

[3] Hopfield J. and Tank D. "Neural computations of decisions in optimisation problems", Biological Cybernetics, 52, pp. 141-152, 1985.

[4] Wilson G. V. and Pawley G. S. "On the stability of the travelling salesman problem algorithm of Hopfield and Tank", Biological Cybernetics, 58, p63, 1988.

[5] Muller,B. and Reinhardt, J. Neural networks: An introduction, Springer-Verlag, 2nd edition, 1991.

[6] Van der Bout, D. E. and Miller, T. K. "A Travelling Salesman Objective Function That Works", IEEE Proc. International Conference on Neural Nets, Vol II pp.299-303, 1988.

[7] Peterson, C. and Soderberg, B. "A new method for mapping optimisation problems onto neural networks", International Journal of Neural Systems, Vol 1, No. 1 (1989) pp. 3-22, 1989.

[8] Bultan, T. and Aykanat, C. Artificial Neural Networks, Elsevier Science Publishing Company B.V., 1991.

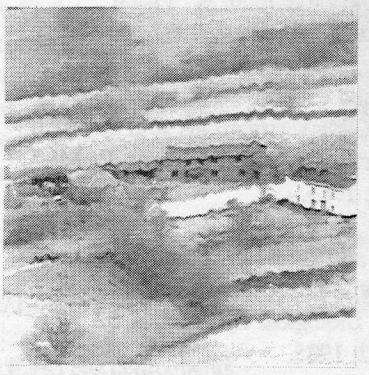

Figure 1 Original distorted image

Figure 2 Simulated annealing restored image

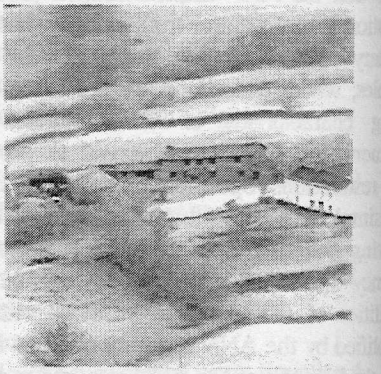

Figure 3 ANN restored image

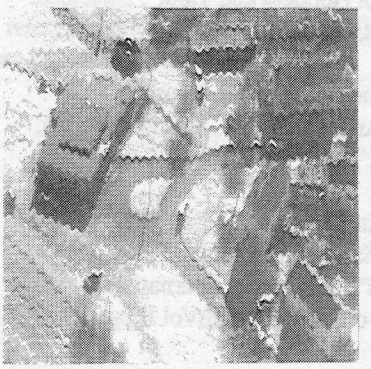

Figure 4 Original distorted image

Figure 5 Simulated annealing restored image

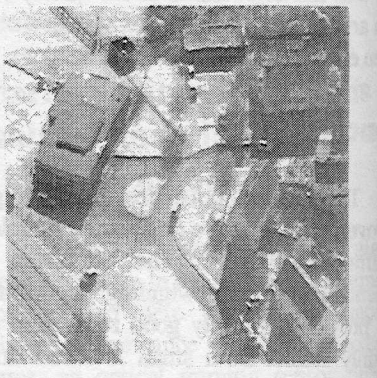

Figure 6 ANN algorithm image

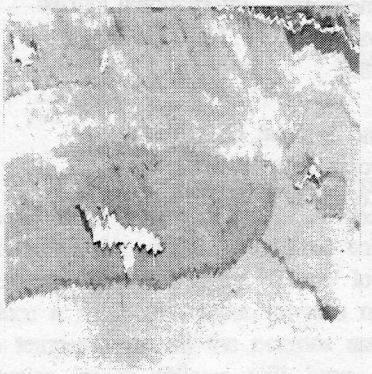

Figure 7 Original distorted image

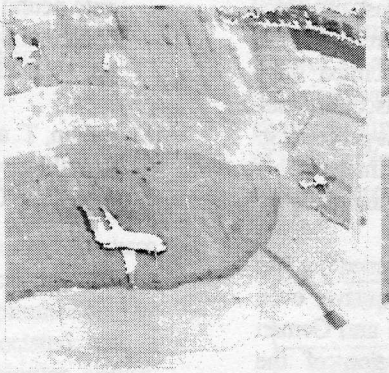

Figure 8 Simulated annealing restored image

Figure 9 ANN algorithm image

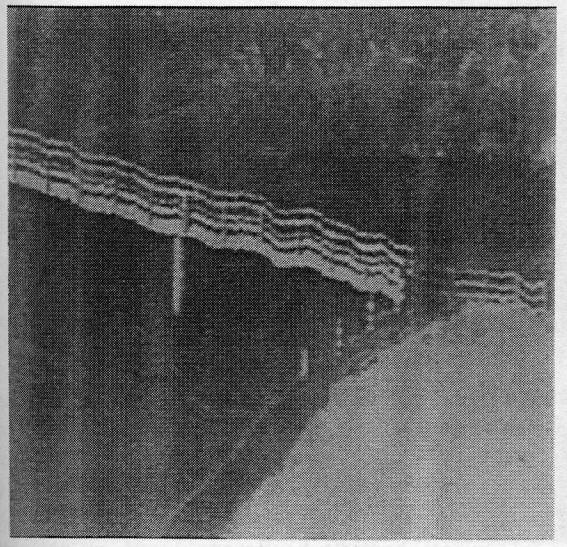

Figure 10 Worst case image

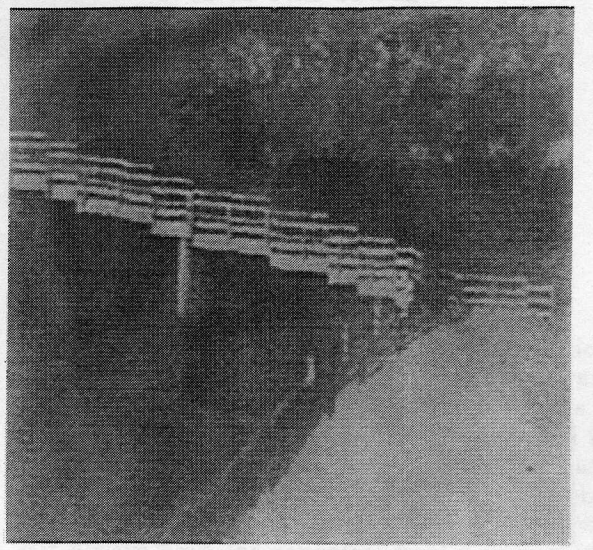

Figure 11 Simulated annealing restored image

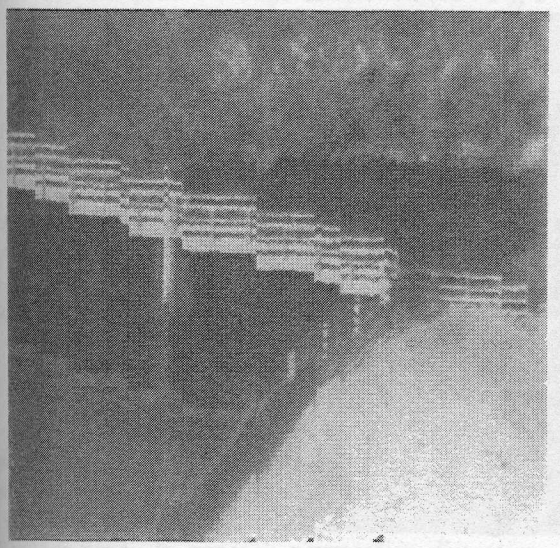

Figure 12 ANN restored image

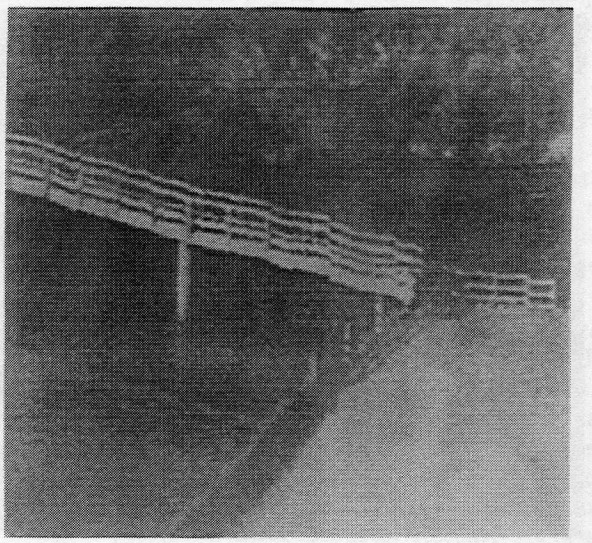

Figure 13 Image after step-edge post-processing

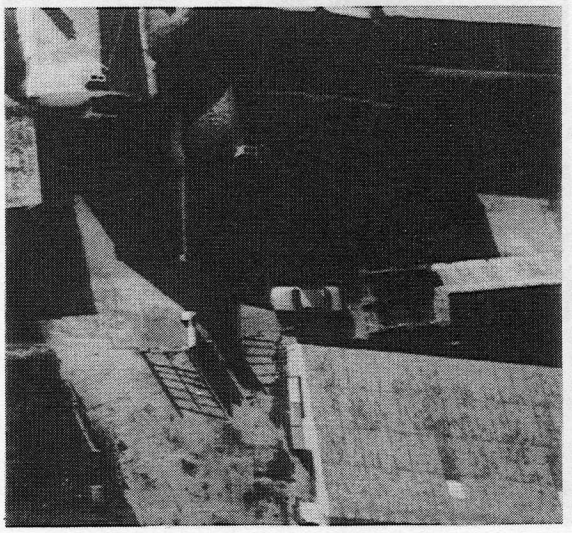

Figure 14 Original image

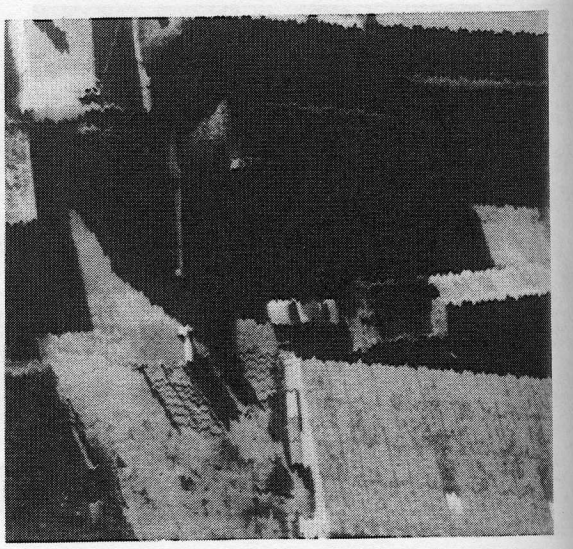

Figure 15 Image with artificial distortion applied

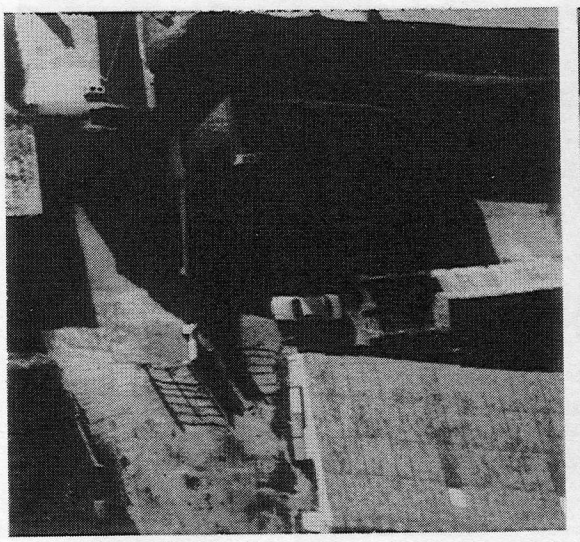

Figure 16 Simulated annealing restored

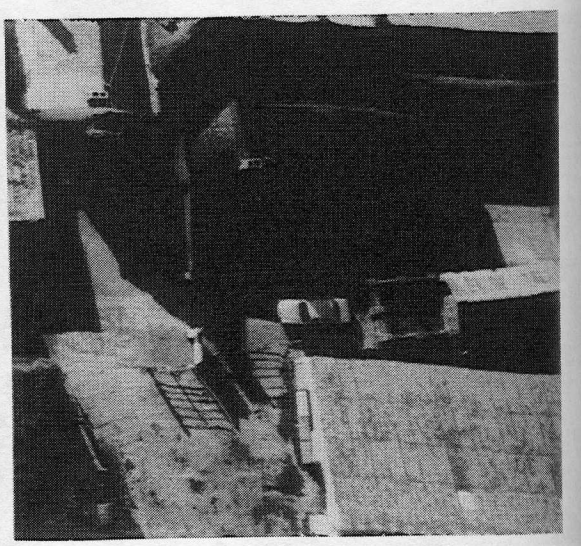

Figure 17 ANN restored image