

On the Accuracy of Pixel Relaxation Labeling

by J. A. Richards, D. A. Landgrebe, P. H. Swain

Laboratory for Applications of Remote Sensing
Purdue University West Lafayette, Indiana 47906 USA
1980

AgRISTARS

SR-PO-004550-10185
NAS9-15466
NASA CR-
160691

"Made available under NASA sponsorship
in the interest of early and wide dis-
semination of Earth Resources Survey
Program information and without liability
for any use made thereof."

A Joint Program for
Agriculture and
Resources Inventory
Surveys Through
Aerospace
Remote Sensing

Supporting Research

March 1980

Technical Report

On the Accuracy of Pixel Relaxation Labeling

by J.A. Richards, D.A. Landgrebe,
and P.H. Swain

(E80-10185) ON THE ACCURACY OF PIXEL
RELAXATION LABELING (Purdue Univ.) 34 p
HC A03/MF A01 CSCI 05B

N80-27765

Unclas
G3/43 00185

Laboratory for Applications of Remote Sensing
Purdue University
West Lafayette, Indiana 47907



NASA



NOTICE

THIS DOCUMENT HAS BEEN REPRODUCED FROM
MICROFICHE. ALTHOUGH IT IS RECOGNIZED THAT
CERTAIN PORTIONS ARE ILLEGIBLE, IT IS BEING RELEASED
IN THE INTEREST OF MAKING AVAILABLE AS MUCH
INFORMATION AS POSSIBLE

Star Information Form

1. Report No SR-PO-00455	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle On the Accuracy of Pixel Relaxation Labeling		5. Report Date March 1980	
		6. Performing Organization Code	
7. Author(s) J. A. Richards, D. A. Landgrebe, P. H. Swain		8. Performing Organization Report No. 030180	
		10. Work Unit No.	
9. Performing Organization Name and Address Laboratory for Applications of Remote Sensing Purdue University 1220 Potter Drive West Lafayette, IN 47906		11. Contract or Grant No. NAS9-15466	
		13. Type of Report and Period Covered	
12. Sponsoring Agency Name and Address J. D. Erickson NASA/Johnson Space Center Houston, TX 77058		14. Sponsoring Agency Code	
		15. Supplementary Notes	
16. Abstract <p>An analysis of pixel labeling by probabilistic relaxation techniques is presented to demonstrate that these labeling procedures degenerate to weighted averages in the vicinity of fixed points. A consequence of this is that undesired label conversions can occur, leading to a deterioration of labeling accuracy at a stage after an improvement has already been achieved. Means for overcoming the accuracy deterioration are suggested and are used as the basis for a possible design strategy for using probabilistic relaxation procedures.</p> <p>The results obtained are illustrated using simple data sets in which labeling on individual pixels can be examined and also using Landsat imagery to show application to data typical of that encountered in remote sensing applications.</p>			
17. Key Words (Suggested by Author(s)) Pixel Labeling		18. Distribution Statement	
19. Security Classif. (of this report) U	20. Security Classif. (of this page) U	21. No. of Pages 31	22. Price*

SR-P0-00455

NAS9-15466

TECHNICAL REPORT
ON THE ACCURACY OF PIXEL RELAXATION LABELING

By

J. A. Richards

D. A. Landgrebe

P. H. Swain

This report describes activity carried
out in the Supporting Research Project.

Purdue University
Laboratory for Applications of Remote Sensing
West Lafayette, Indiana 47907, U.S.A.

MARCH 1980

"ON THE ACCURACY OF PIXEL RELAXATION LABELING"*

J. A. Richards
D. A. Landgrebe
P. H. Swain**

ABSTRACT

An analysis of pixel labeling by probabilistic relaxation techniques is presented to demonstrate that these labeling procedures degenerate to weighted averages in the vicinity of fixed points. A consequence of this is that undesired label conversions can occur, leading to a deterioration of labeling accuracy at a stage after an improvement has already been achieved. Means for overcoming the accuracy deterioration are suggested and are used as the basis for a possible design strategy for using probabilistic relaxation procedures.

The results obtained are illustrated using simple data sets in which labeling on individual pixels can be examined and also using Landsat imagery to show application to data typical of that encountered in remote sensing applications.

*This work was funded in part by NASA Contract No. NAS9-15466.

**The authors are with the School of Electrical Engineering and the Laboratory for Applications of Remote Sensing, Purdue University, West Lafayette, IN 47907. J. A. Richards is on leave from the School of Electrical Engineering, University of New South Wales, Kensington NSW 2033, Australia.

I. INTRODUCTION

Probabilistic relaxation procedures, which employ the information embedded in spatial context, appear to be attractive techniques for reducing labeling errors in various types of image data. In the results of some simple exercises such as the labeling of the sides of a triangle [1,2], this has indeed been the case with perfect labeling shown to be possible. However, in more complex labeling tasks such as line and edge enhancement [3,4,5] and pixel labeling [6,7], the results obtained to date detract somewhat from the appeal of relaxation since labeling accuracy has been observed to improve during the early iterations of the process only to be followed by a subsequent degradation. In pixel labeling, for example, the labeling error exhibits a turning point at a specific iteration and the final error, in some situations, can be worse than that initially; similarly in line enhancement applications, line broadening is observed to occur late in the process, degrading an otherwise acceptable labeling. From a practical viewpoint, this suggests that the relaxation process in these sorts of applications should be stopped at some particular point to avoid incipient deterioration of the results. However, since the iteration of minimum error will not be known, so that the optimum stopping point will not be evident, it is likely that the final labeling error will always be larger than necessary. To avoid this situation, it is clearly important that the degradation mechanism be understood so that, at worst, a stopping criterion can be devised or, better still, the deterioration of labeling accuracy can be minimized or avoided. Eklundh and Rosenfeld [8] and Peleg [9] have addressed the task of determining suitable stopping rules. In particular, Eklundh and Rosenfeld observe that the convergence of relaxation is such that labeling error changes most in the earlier

iterations and only slightly in the latter stages of the process. As a result, they recommend that the average absolute difference between label estimates in pairs of sequential iterations be computed and that relaxation be terminated when this measure is an order of magnitude smaller than it was after the first iteration. Peleg bases his stopping criterion on a measure of the likelihood that the labeling at each iteration is the correct one. By establishing a likelihood measure that incorporates both the influence of the current label estimates and the effect of the current joint probabilities, he demonstrates (for two examples) that the most probable labeling occurs at some iteration before that where the minimum labeling error is observed. A second stopping rule proposed by Peleg, using information channel concepts, exhibits similar behavior. With both of these measures, the process is stopped significantly short of the iteration of minimum error and thus they must be regarded as sub-optimal criteria. Moreover, since the reason for the turning point in the error curve has not been determined, there is no theoretical reason to suppose that stopping rules of these types will circumvent deterioration of labeling accuracy.

The present treatment is directed towards understanding the mechanism during relaxation that causes labeling error to increase again after having reached a minimum. It is demonstrated that this is a process of local averaging once relaxation has approached a fixed point. As a consequence, it is shown that if the parameters in the relaxation algorithm are suitably chosen, the error versus number of iterations curve can be made to decrease monotonically to a fixed error.

2. THE RELAXATION ALGORITHM AND THE DEFINITION OF NEIGHBORHOOD

Consider the probabilistic relaxation algorithm of Rosenfeld, Hummel and Zucker [1]:

$$p_i^{k+1}(\lambda) = p_i^k(\lambda) Q_i^k(\lambda) / \sum_{\lambda} p_i^k(\lambda) Q_i^k(\lambda) \quad (1)$$

where $p_i^k(\lambda)$ is the k^{th} estimate of the probability that λ is the proper label for the i^{th} pixel, and $Q_i^k(\lambda)$ is the k^{th} estimate of the neighborhood function, given by

$$Q_i^k(\lambda) = 1 + \sum_{j \in J} d_j \sum_{\lambda'} r_{ij}(\lambda|\lambda') p_j^k(\lambda'). \quad (2)$$

In this expression the $r_{ij}(\lambda|\lambda')$ are the compatibility coefficients, the d_j are a set of neighbor weights that can be used to give different neighbors differing degrees of influence in the neighborhood function, and J defines the neighborhood about the particular pixel being considered. Owing to practical considerations, this neighborhood in pixel labeling is chosen either as the 3x3 set of pixels about the pixel under consideration or, more simply, as the pixels above, below, and to the sides of that pixel. Within these choices two variations appear to have been used. One includes the central pixel (i.e., that under consideration) as a member of that neighborhood [1] and the other excludes that pixel [10,11]. These will be referred to here as inclusive and exclusive neighborhoods respectively. The following analysis is based upon inclusive neighborhoods with results for exclusive neighborhoods, as required, being given as special cases.

3. LOCAL AVERAGING IN THE VICINITY OF FIXED POINTS AND ITS EFFECT ON GEOMETRIC FEATURES

Suppose a particular relaxation exercise has progressed to a point where the label estimates have all approached 0 or 1. (The stage where

the label estimates are at 0 or 1 is called a fixed point in the process. Fixed points with $p_i^k(\lambda)$ other than 0 or 1 can occur; however, they are infrequent in pixel labeling and will not be considered here.) Within homogeneous regions -- i.e., where all pixels in a neighborhood have the same predominant label -- the mutual support offered among neighbors will not allow the label estimate on any particular pixel to alter by any significant amount. In fact, those estimates will simply move closer to their fixed points. However, the situation at boundaries such as corners of one region within another, the ends of lines (single pixel wide), and isolated pixels can be quite different, as the following discussion reveals.

Consider a λ_1 pixel on the boundary between λ_1 and λ_2 regions. Evidently $p_i^k(\lambda_1)$ is the largest label estimate for that pixel and it is reasonable to assume for such a λ_1, λ_2 neighborhood that $p_i^k(\lambda_1) > p_i^k(\lambda_2) \gg p_i^k(\lambda_n)$, $\forall n \neq 1, 2$. Now consider whether the label estimate $p_i^k(\lambda_1)$ will be strengthened or weakened as relaxation proceeds. To do this, it is sufficient to consider the relative strengths of the neighborhood functions as defined in (2). In particular, if

$$Q_i^k(\lambda_1) > Q_i^k(\lambda_2)$$

the λ_1 label will be strengthened at the next iteration; otherwise it will weaken. This will continue with subsequent iterations (since the label estimates at neighbors will not change by any significant amount). Should $Q_i^k(\lambda_2) > Q_i^k(\lambda_1)$, the repeated application of relaxation will ultimately lead to λ_2 being the favored label at the pixel -- i.e., the λ_1 label will be removed by further iterations. Consequently even though labeling error could have been reduced in establishing the λ_1 label on that pixel, it

will now (gradually) increase owing to the loss of that label. To avoid this, it is necessary, therefore, to ensure (from (2)) that

$$1 + \sum_{j \in J} d_j \sum_{\lambda'} r_{ij}(\lambda_1 | \lambda') p_j^{k(\lambda')} > 1 + \sum_{j \in J} d_j \sum_{\lambda'} r_{ij}(\lambda_2 | \lambda') p_j^{k(\lambda')}$$

i.e. $\sum_{j \in J} d_j \sum_{\lambda'} \{r_{ij}(\lambda_1 | \lambda') - r_{ij}(\lambda_2 | \lambda')\} p_j^{k(\lambda')} > 0$ (3)

Note that the additive "1" in (2) has been of no significance in determining (3), so that (3) is a result general to all present relaxation algorithms which employ arithmetic averaging over the neighborhood, including particularly that where the $r_{ij}(\lambda | \lambda')$ are mapped to conditional probabilities in which case the "1" does not appear in (2). (See [10].)

The probability that the pixel's label could alter to that of a third class λ_3 has been ignored owing to the earlier assumptions regarding the relative strengths of the label estimates on that pixel.

Since it has been assumed that all the probability estimates are close to 0 or 1, (3) can be modified to

$$\sum_{j \in J} d_j \{r_{ij}(\lambda_1 | \lambda_j) - r_{ij}(\lambda_2 | \lambda_j)\} > 0$$
 (4)

where λ_j is the preferred label on the j^{th} neighbor.

Now consider the neighborhood definition explicitly. Let J' be the exclusive neighborhood so that $J : \{J', i\}$ where i is the pixel whose label is "currently" under consideration. Then (4) can be recast as

$$d_i \{r_{ii}(\lambda_1|\lambda_1) - r_{ii}(\lambda_2|\lambda_1)\} + \sum_{j \in J'} d_j \{r_{ij}(\lambda_1|\lambda_j) - r_{ij}(\lambda_2|\lambda_j)\} \geq 0,$$

$$\text{giving } d_i > \frac{\sum_{j \in J'} d_j \{r_{ij}(\lambda_1|\lambda_j) - r_{ij}(\lambda_2|\lambda_j)\}}{r_{ii}(\lambda_2|\lambda_1) - r_{ii}(\lambda_1|\lambda_1)} \quad (5)$$

as the condition that λ_1 be retained as the label for the i^{th} pixel.

To simplify further discussion, now consider some special cases of (5). First suppose the compatibilities $r_{ij}(\lambda|\lambda')$ have been chosen as conditional probabilities, and secondly consider only a two-label problem so that

$$r_{ij}(\lambda_2|\lambda_j) = p_{ij}(\lambda_2|\lambda_j) = 1 - p_{ij}(\lambda_1|\lambda_j).$$

Thus (5) becomes

$$d_i > \frac{\sum_{j \in J'} d_j \{1 - 2p_{ij}(\lambda_1|\lambda_j)\}}{-1 + 2p_{ii}(\lambda_1|\lambda_1)}$$

Within the conditional probability compatibility definition, it is logical that $p_{ii}(\lambda_1|\lambda_1) = 1$, although in (5) $r_{ii}(\lambda_1|\lambda_1)$ based upon other compatibility definitions need not necessarily be unity. To avoid loss of a λ_1 label on a border between λ_1 and λ_2 regions, we therefore have the condition

$$d_i > \sum_{j \in J'} d_j \{1 - 2p_{ij}(\lambda_1|\lambda_j)\} \quad (6)$$

Now consider the particular choice of neighborhood shown in Fig. i, and let the pixel under consideration be a corner pixel, as depicted. Suppose

$d_j = d \forall j$, and further assume the compatibility coefficients $p_{ij}(\lambda_1 | \lambda_j)$ are the same for each neighbor j of the corner pixel.* In view of these, (6) simplifies to

$$d_i > 4d \{p_{ij}(\lambda_2 | \lambda_2) - p_{ij}(\lambda_1 | \lambda_1)\}$$

Further suppose the d_j have been chosen such that $\sum_j d_j = 1$. Such a choice is strictly only required when the $r_{ij}(\lambda | \lambda')$ are chosen as correlations. However, it is a useful choice in general and here leads to $4d + d_i = 1$ so that we have

$$d_i > \eta(1+\eta)^{-1}, \quad \eta = p_{ij}(\lambda_2 | \lambda_2) - p_{ij}(\lambda_1 | \lambda_1) \quad (7)$$

as the required condition that λ_1 corner labels not be lost. This condition also applies to the preservation of single-pixel-wide λ_1 lines that pass through a λ_2 neighborhood. For the simple neighborhood chosen, the only other geometries that are subject to label conversion (deterioration) by the mechanism described are the ends of lines a single pixel wide, and single isolated pixels. From (6) it can be shown that the condition for the preservation of labels at the ends of lines of λ_1 within λ_2 regions is

$$d_i > \frac{3p_{ij}(\lambda_2 | \lambda_2) - p_{ij}(\lambda_1 | \lambda_1) - 1}{3p_{ij}(\lambda_2 | \lambda_2) - p_{ij}(\lambda_1 | \lambda_1) + 1} \quad (8)$$

Likewise, to preserve individual λ_1 labeled pixels in λ_2 regions, it is necessary that

$$d_i > \frac{2p_{ij}(\lambda_2 | \lambda_2) - 1}{2p_{ij}(\lambda_2 | \lambda_2)} \quad (9)$$

*This ignores, for example, systematic biases such as the unequal vertical and horizontal sampling rates present in Landsat imagery.

The predictions of (7-9) were checked using the data chosen in Figure 2. This is assumed to be a portion of an image for which the compatibilities are $p_{ij}(W|W) = 0.700, p_{ij}(\underline{b}|\underline{b}) = 0.800$, where \underline{b} implies blank. Using (7-9), the following conditions can be determined.

1. To avoid loss of a W corner in a \underline{b} region $d_i > 0.091$
2. To avoid loss of a \underline{b} corner in a W region $d_i > 0.111$
3. To avoid loss of a W line end in a \underline{b} region $d_i > 0.259$
4. To avoid loss of a b line end in a W region $d_i > 0.130$
5. To avoid loss of a W pixel in a \underline{b} region $d_i > 0.375$
6. To avoid loss of a \underline{b} pixel in a W region $d_i > 0.286$

Consequently we would expect that if

$d_i = 0$, only \underline{b} corners would be retained

$d_i = 0.100$ both \underline{b} and W corners would be retained

$d_i = 0.150$ the above plus \underline{b} lines would be retained

$d_i = 0.270$ the above plus W lines would be retained

$d_i = 0.400$ all corners, lines and isolated pixels would be retained

As seen in Figure 2, these predictions are accurate. The image was initialized very close to a fixed point by choosing the initial label estimates as $p_i^0(\underline{b} \text{ or } W) = 0.99$, and thus could be regarded as an image which has approached that condition by some preceding iterations of relaxation; moreover, it is useful to suppose the initial labeling represents the true labeling since then the label conversions observed in Figure 2 would represent the introduction of labeling errors.

An example with assumed initial labeling errors is shown in Figure 3. Again, initial label estimates of 0.99 were chosen to allow the prediction of (7) to be checked. As seen when d_i is chosen to avoid loss of corners, the relaxation process converges to the true labeling and remains there.

However, when d_j is less than the prescribed value, the corners are lost shortly after the true labeling has been achieved.

The predictions of (5) through (9), of course, only hold exactly for an image that has approached a 0,1 fixed point and thus tacitly assumes that the local averaging that gives rise to the conversion of border labels takes place when the label estimates are all near 0 or 1. While this is indeed the case, averaging also takes place earlier when the label probability estimates are not quite so extreme. By initializing label probabilities further from a fixed point, the predictions from equations such as (7) through (9) will be modified. As an indication of this, Figure 4 illustrates how the value of d_j predicted from (7) for the example of Figure 3 is modified for a range of initial label estimates. This graph was produced empirically and implies that (7) is a lower bound.

Should an exclusive neighborhood definition be used, then $d_j = 0$ in (5) through (9). Thus λ_1 label deterioration of the types considered will occur unless the right-hand sides of those equations are less than zero. A little thought reveals that these equations can never be satisfied for all complementary pairs of neighbor geometry (i.e., λ_1 corners in λ_2 regions and λ_2 corners in λ_1 regions) so that label degradation leading to an increase in labeling error would always be expected to occur with conventional probabilistic relaxation algorithms applied to real imagery when used with exclusive neighborhood definitions. The supervised algorithm proposed recently, however, can be adjusted to avoid the degradation problem on exclusive neighborhoods since it bears similarity to using an inclusive neighborhood definition with the conventional algorithm [12].

4. LABELING IMPROVEMENT DURING RELAXATION

The intention of applying relaxation to an image is to improve upon a labeling which has been generated beforehand by some "imperfect" process. In endeavoring to examine the improvement, it is useful to view the situation in the following manner. The relaxation algorithm does not know, of course, which are the correct and which are the incorrect labels. It only "knows" which labels are consistent and which are inconsistent with their neighbors. Consequently, an image with initial labeling errors will be treated by the relaxation algorithm as though it were correctly labeled and the "improvement" which it creates is a conversion of locally inconsistent labels. This conversion will take place by mechanisms such as those described in the previous section and, in particular, for pixels that are close to fixed points, equations such as (7) through (9) can be used to describe labeling improvement in addition to likely degradation. Indeed, in the special case when an image is intentionally initialized close to a fixed point, those expressions can be used very accurately to describe the labeling improvement phase as well as any deterioration in the labeling that might occur. In such a situation, the predictions of (7) through (9) (for a two-label example) allow the value of d_j to be chosen relative to the compatibilities and other neighbor weighting coefficients to ensure that some labels are intentionally converted (i.e., those in error), while others are retained. Clearly the requirements for improvement and for avoiding degradation will often conflict in real image segments and, in order to obtain clean-up during relaxation, some correct labels may have to be sacrificed. As an illustration of these comments, consider the results of Figure 5 and suppose any one of the isolated pixels happened to be correctly labeled initially. If corner W labels are not to be lost, (7)

demands $d_j > 0.274$, whereas (8) and (9) require $d_j < 0.36$, 0.429 respectively if line end pixels and individual pixels have to be removed during the relaxation "improvement" phase. Choosing $d_j = 0.280$ shows that all erroneous labels are modified as expected, except that shown by the arrow which forms a corner with the corner W region. However, if one of the isolated pixels was correctly labeled initially, as supposed above, then this is also now in error.

The discussions above and the supporting results presented have been based upon initial label probabilities being chosen close to 0 or 1. The graph of Figure 4 supports that these comments will apply also for initial label estimates different to 0, 1 but all being the same for the same class. Should the initial label estimates within a class be all different (as would happen, for example, should they be determined on the basis of Mahalanobis distance considerations in a classification [7]), some correctly (and weakly) labeled pixels will be removed preferentially during the early improvement phase in the relaxation. However, all label estimates will then move toward 0 or 1 and the remarks of Section 3 regarding deterioration still apply in principle.

5. RELEVANCE OF ACCURATE COMPATIBILITIES

In view of the comments of the previous two sections, it is clear that control of a relaxation process lies significantly in equations of the type (7) through (9) for a two-label problem and similar (albeit more numerous) manifestations of (5) for a multi-label exercise. Consequently, in the removal of initial labeling errors and in avoiding label degradation, the actual values of the compatibility coefficients ($r_{ij}(\lambda|\lambda')$ or $p_{ij}(\lambda|\lambda')$) appear not to be important in pixel labeling so much as their values relative

to each other and to d_i as described in (7) through (9). As a demonstration of this, consider again the example of Figure 5, and arbitrarily choose the compatibilities as $p_{ij}(W|W) = 0.600$ and $p_{ij}(b|b) = 0.700$ (compared with the true values of 0.500 and 0.875). Eqn (7) shows $d_i > 0.091$ for corner retention but loss of other geometries. Choosing $d_i = 0.096$, the results in Figure 6 are obtained showing the expected label improvement without subsequent degradation, notwithstanding the arbitrary choice of the compatibilities.

6. DESIGN OF POSSIBLE RELAXATION STRATEGIES FOR PIXEL LABELING

Since equations such as (7) through (9) describe the effect of relaxation (within the comments of the previous sections), it should be possible to specify an appropriate set of compatibilities and the weighting constant d_i to achieve certain desired results. In so doing, the following guidelines are significant in a two-label situation, with the neighborhood definition chosen:

- (i) If $p_{ij}(\lambda_2|\lambda_2) > p_{ij}(\lambda_1|\lambda_1)$ corner pixels labeled λ_2 protruding into λ_1 regions will never degrade.
- (ii) Practical lower bounds on the compatibilities $p_{ij}(\lambda_1|\lambda_1)$ and $p_{ij}(\lambda_2|\lambda_2)$ are 0.5. Otherwise the image must have consisted of isolated pixels or lines of pixels, depending upon the manner in which those compatibilities were calculated.
- (iii) for both $p_{ij}(\lambda_1|\lambda_1)$ and $p_{ij}(\lambda_2|\lambda_2)$ close to 0.5 (7) through (9) reveal that all conditions on d_i are approximately zero -- i.e., there would be no degradation and no improvement (as expected).

(iv) Convergence is faster for larger values of the compatibilities. As a result of the above comments, the following is proposed as one possible design strategy for pixel labeling relaxation procedures, which obviates the need to obtain reliable estimates for the compatibility coefficients. It turns out to be a sub-optimal procedure for practical image data, since its success depends upon foreknowledge of, or a feeling for, the prevailing geometry in a particular image; however, it is a significant improvement on choosing neighbor weights (d_j) in an arbitrary manner.

1. Choose all initial label estimates as 0.99 (or 0.01 as appropriate), unless there is good reason for doing otherwise. Such a choice (close to a fixed point) allows moderate accuracies in predictions made from (7) through (9).
2. Choose the compatibility for the most prevalent class to be the strongest since this automatically preserves corners in that class and will preempt lower overall final error.
3. Form an impression of the label geometries in which most label errors seem to lie (such as isolated pixels) and also of the label geometries which should not be allowed to degrade (such as corners) and choose d_j in order to remove only suspected errors.
4. If speed is a consideration, choose the magnitudes of the compatibilities to be as large as possible within the restraints imposed by the above considerations.

This procedure is now illustrated using two data sets. One consists of a multitemporal Landsat image acquired over a region in Kansas and contains an array of 117 x 196 pixels. The other is a 40 x 100 pixel portion of that same image. The latter was chosen to enable the results

to be inspected on a pixel-by-pixel basis, whereas the former is used to illustrate performance on imagery of the size typically encountered in remote sensing applications. Figure 7a represents two-category (wheat and non-wheat) ground truth for the smaller image and as such can be regarded as true labeling. Figure 7b shows the result of a crude classification of that portion. This classification was obtained using a minimum distance to means classifier on pattern vectors consisting of three only of the 16 possible spectral response features. These features were chosen beforehand on the basis of a separability measure computed over the full image. Similarly, training areas were selected from the full segment and Figure 7b represents only a small portion of the resulting classification map.

Inspection of Figure 7a suggests that it would be desirable to remove isolated W (wheat) pixels from any classification but that W line ends and corners of W fields should not be allowed to alter. Also, it would seem desirable to preserve blank (non-wheat) corners. On the basis of these observations, the relaxation parameters

$$p(W|W) = 0.600 \quad p(\underline{b}|\underline{b}) = 0.700 \quad d_i = 0.200$$

could be suggested as a possible choice which would remove scattered W pixels but retain all other geometries. This prediction can be checked on Figure 7c which shows labeling error, after 100 iterations of relaxation, as a function of d_i ; the image was initialized close to a fixed point. It is evident that $d_i = 0.200$ is a good choice for these particular data. The figure also displays discernible improvements in labeling at values of d_i corresponding to the preservation of the various geometric features noted on the diagram. Inspection of Figure 7d, which shows the final

labeling achieved with $d_i = 0.200$, reveals that only the isolated W pixels have been relabeled as required.

In passing, it is of interest to note that only when $d_i > 0.15$, relaxation reduces the labeling error below that in the initial labeling. For $d_i < 0.14$, the label degradation mechanism leads to worse error after relaxation than before for these particular data. This effect is so severe here because of the "large" number of geometries that need to be preserved.

As a basis for comparison of the results of Figure 7, Figure 8 shows label error as a function of d_i using "true" compatibilities calculated from the full image. Shown in that figure also are the predicted values of d_i from (7) through (9) relating to the preservation of particular geometries. Comparison of Figures 7c and 8 shows no essential difference in shape, supporting the notion that exact compatibilities are not required.

Figure 9 shows the result of relaxation over the complete image using the "true" values of the compatibilities. Again, the significant values of d_i are noted. Examination of this figure reveals that preservation of (too many) wheat corners is detrimental. The fact that this behavior is different from that of Figures 7 and 8 is indicative of the fact that the portion of the image used in these previous figures has a geometric character that is not representative of the complete segment. This is evident from an inspection of the full ground-truth map.

Figure 10 shows labeling error versus number of iterations for selected values of d_i in Figure 9. Note that for d_i less than optimum, labeling error initially decreases, passes through a turning point, and increases again before settling down to a pessimistic final value. For values of d_i near 0.15, the error curve does not exhibit the deterioration phase and has a final value which is almost as low as the minimum in the

previous curve. For larger d_j , while the curve is monotonically decreasing the final error is larger than necessary. Ultimately, for large d_j the curve will remain constant at the initial labeling error.

7. SUPERVISION AS A MEANS FOR SEGMENTING LABEL CONVERSION EFFECTS

An unfortunate observation on the results of the previous section is that once a choice of algorithm parameters has been made to retain some and remove other border geometries, this effect takes place over the complete image, except at those pixels that were so strongly labeled initially that their probability estimates reach 0 or 1 well ahead of others. In a practical pixel labeling exercise, particularly of the remote sensing variety, this is undesirable since there could be segments of an image where, for example, single pixels would wish to be preserved, whereas in other segments single pixels would want to be removed (e.g., urban versus agricultural regions). The only way this effect can be implemented with existing algorithms is to attempt to condition the initial probability estimates by, say, strengthening those corresponding to regions where single pixels are desirable. At best, this would be a time-consuming procedure that would also override any information implicit in the initial labeling. An alternative, and potentially attractive, technique for suitably segmenting the image is to make use of the supervised approach to relaxation proposed recently [12]. In that technique, the label estimates at each iteration of relaxation are modified by reference to some other data. Using the initial label estimates for that other data has given rise to a relaxation procedure that can be used to overcome the detrimental label conversion effects on boundaries, as discussed earlier. However, it should also be possible to derive the

supervising data from an array which overlays the image and which contains data on the likely geometries which should desirably be preserved in various image segments. As an early indication of the viability of such a scheme, Figure 11 shows an image which is supervised with the overlay array of probabilities indicated. As observed, it is possible to adjust the degree of influence of the supervision (via the parameter β) to retain certain geometric features in one image segment while allowing those same features to be relaxed out in other segments. Means for establishing the appropriate values of β have not yet been determined; however, it is believed that the results of Figure 11 demonstrate the usefulness of this approach.

8. DISCUSSION AND CONCLUSIONS

From the results presented in the previous sections, it is evident that the compatibilities should not need to be accurately characteristic of a particular image. Rather, as noted, it is biases in the compatibilities, along with the value of d_i (relative to the weights on the other neighbors -- here all taken to be the same) that substantially determine how relaxation will behave on particular image data. This is demonstrated in the fact that the compatibilities in Figures 5 and 7 are the same (by choice) and yet clearly the images are quite different. In those cases it was only necessary to choose the appropriate value for d_i in view of the compatibilities given. A little thought also reveals that for image data (of the Landsat type especially) the true compatibilities cannot be particularly significant since these are statistically averaged measures computed over the whole or even a part of an image where in fact some regions of an image may

bear no geometric or statistical resemblance to other areas of that same image.

The examples presented above have shown that it is possible with image data to choose compatibilities and specific values of the neighbor weights d_j such that the relaxation process will converge to a near optimum error which will not subsequently increase owing to label conversion (degradation) mechanisms. Owing to the greater degree of homogeneity present in the data typically encountered in picture processing tasks (such as noise removal), it is likely that the predictions concerning retention of label geometries presented herein may be more useful in those applications than with pixel labeling. For example, inspection of the noisy scene of a house used in [7] reveals that the most important features to be retained are the corners of one label type within another. For example, it is important to preserve corners of "sky" which protrude into regions of "brick" (sky is visible through the end of a veranda). With the compatibilities chosen by those authors, (5) can be used to specify a value of d_i beyond which sky-in-brick corners will be preserved. Those authors use a full 3x3 neighborhood and choose $d_j = 1, \forall j$ (the compatibility coefficients were defined by mutual information). Choosing $d_j = 1, \forall j \neq 1$, eqn (5) shows that for sky-in-brick corners to be preserved, it is necessary that $d_1 \geq 5.22$ (the actual value of d_1 depends upon the initial label estimates). Clearly with $d_1 = 1$ degradation will occur and it is to be expected, with regard to this feature at least, that label accuracy would improve early on during relaxation and then become poorer owing to sky label degradation. Indeed those authors report a degradation phase during this relaxation exercise and it is probably that a (major) component of it is a result of sky label loss at corners.

The theme of this paper has been to develop a model for the relaxation process that would permit the label degradation mechanism to be understood and thus avoided. Consequently attention has not been given to attempting to achieve the smallest possible labeling error. For example, the exercises presented have been initialized with label probabilities all close to a fixed point and all the same within the same class initially. Distributing the initial label estimates, however, according to some measure of confidence, would probably lead to overall lower error since erroneous labels that were weak initially would be removed before the relaxation mechanism fixed them by one of the preservation measures discussed earlier. Notwithstanding this, the predictions of equations such as (5) through (9) are important guidelines for controlling the relaxation improvement and degradation mechanisms and consequently for algorithm design as discussed.

REFERENCES

1. A. Rosenfeld, R. Hummel & S. Zucker, "Scene Labeling by Relaxation Algorithms," IEEE Trans. Sys. Man, Cyber, vol. SMC-6, June 1976, pp. 420-433.
2. S. Zucker, E. Krishnamurthy & R. Haar, "Relaxation Processes for Scene Labeling: Convergence, Speed and Stability," IEEE Trans. Sys. Man, Cyber, vol. SMC-8, Jan. 1978, pp. 41-48.
3. S. Zucker, R. Hummel & A. Rosenfeld, "An Application of Relaxation Labeling to Line and Curve Enhancement," IEEE Trans. Comp., C-26, Apr. 1977, pp. 394-403, plus correction, Sept. 1977, pp. 900-929.
4. B. Schachter, A. Lev, S. Zucker and A. Rosenfeld, "An Application of Relaxation to Edge Reinforcement," IEEE Trans. Sys. Man, Cyber, SMC-7, Nov. 1977, pp. 813-816.
5. S. Peleg & A. Rosenfeld, "Determining Compatibility Coefficients for Curve Enhancement Relaxation Processes," IEEE Trans. Sys. Man, Cyber, SMC-8, July 1978, pp. 548-555.
6. A. Lev, S. Zucker and A. Rosenfeld, "Iterative Enhancement of Noisy Images," IEEE Trans. Sys. Man, Cyber, SMC-7, June 1977, pp. 435-442.
7. J. Eklundh, H. Yamamoto & A. Rosenfeld, "Relaxation Methods in Multispectral Pixel Classifications," Technical Report 662, Computer Science Center, University of Maryland, College Park, July 1978.
8. J. Eklundh and A. Rosenfeld, "Convergence Properties of Relaxation," Technical Report 701, Computer Science Center, University of Maryland, College Park, Oct. 1978.
9. S. Peleg, "Monitoring Relaxation Algorithms Using Labeling Evaluation," Technical Report 842, Computer Science Center, University of Maryland, College Park, Dec. 1979.
10. S. Zucker and J. Mohammed, "Analysis of Probabilistic Relaxation Labeling Processes," Proc. 1978 IEEE Conf. Pattern Recognition and Image Processing, Chicago, IL, May 1978, pp. 307-312.
11. S. Peleg, "A New Probabilistic Relaxation Scheme," Proc. 1979 IEEE Conf. Pattern Recognition and Image Processing, Chicago, Aug. 1979, pp. 337-343.
12. J. A. Richards, D. Landgrebe and P. H. Swain, "Pixel Labeling by Supervised Probabilistic Relaxation," LARS Information Note 102979. To appear in IEEE Transactions on Pattern Analysis and Machine Intelligence.




Figure 1. Neighborhood definition used herein.

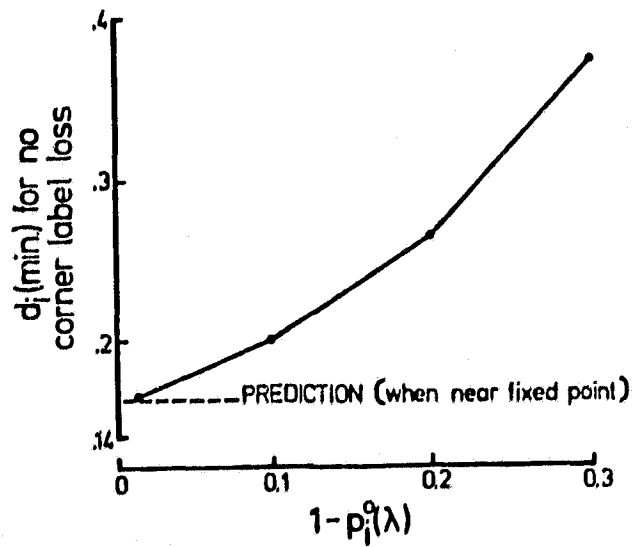


Figure 4. Modification of the prediction of equation (7) for label estimates further from a 0,1 fixed point.

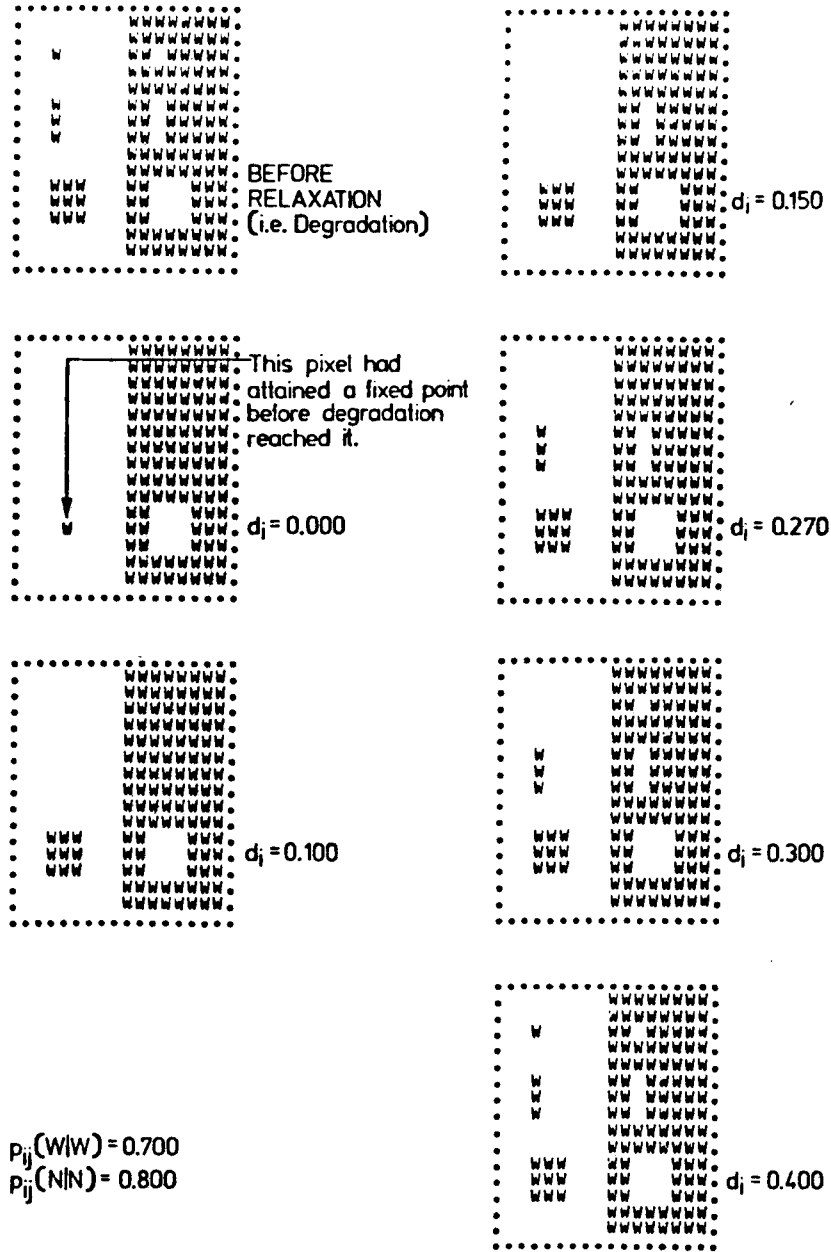


Figure 2. Verification of predictions made using equations (7) through (9).




Figure 3. Using the prediction of equation (7) to avoid loss of corner W labels

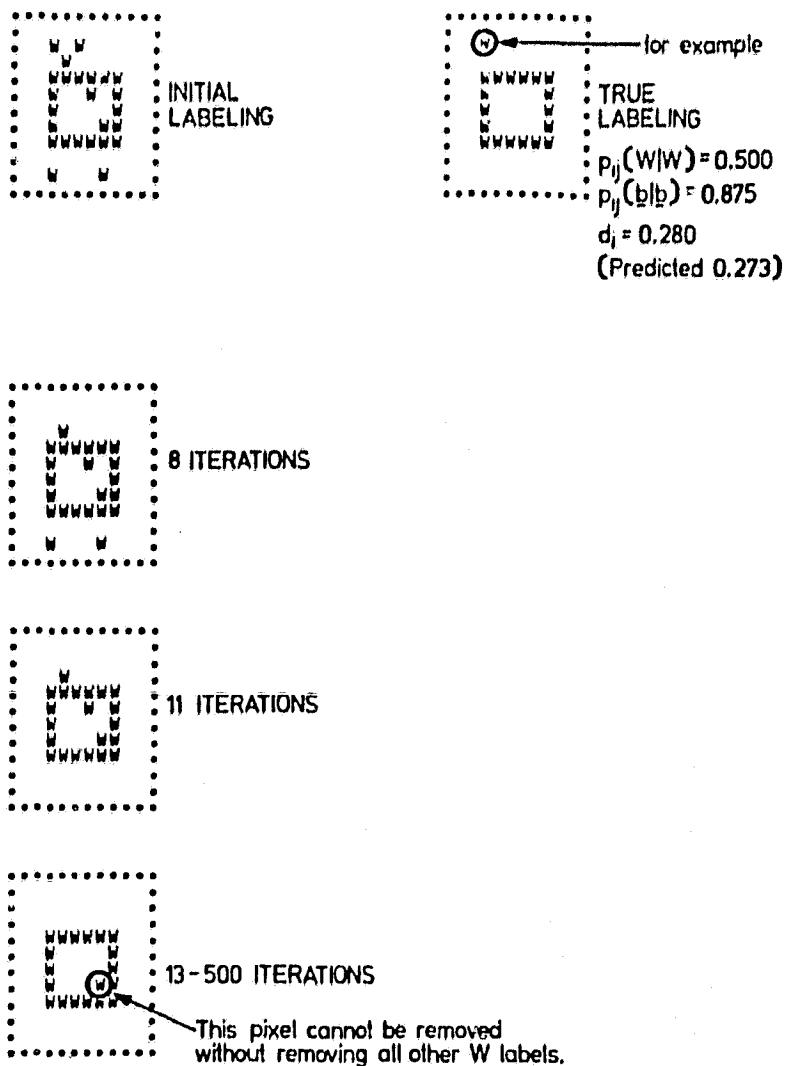



Figure 5. Illustration of a choice of the neighbor weight (d_i), associated with the central pixel of the neighborhood, to remove individual pixels and line-ends but to preserve corners. If an individual pixel was correctly labeled, as indicated, then it becomes erroneous.



TRUE LABELING
as for Figure 5
 $p_{ij}(W|W) = 0.600$
 $p_{ij}(b|b) = 0.700$
 $d_i = 0.096$
(Predicted $d_i > 0.091$)







Figure 6. Repetition of the example of Figure 5 but with "arbitrarily" chosen compatibilities and with d_i adjusted accordingly.



ORIGINAL PAGE IS
OF POOR QUALITY

Figure 7(a). Two-category (wheat-W, nonwheat-blank) ground truth for the 40x100 pixel Landsat multitemporal small image. This is considered as "true" labeling.




Figure 7(b). Initial labeling for the 40x100 pixel image, obtained from a simple minimum distance classifier.




Figure 7 (c). Remaining labeling error as a function of the central pixel neighbor weight d_i , after 100 iterations of relaxation on the 40x100 pixel small image. (100 iterations is sufficient to achieve the final labeling.) The compatibilities have been chosen as discussed in the text.




Figure 7(d). Final labeling of the small image after 100 iterations with $d_i = 0.200$. Note that only isolated W pixels have been removed but that all other geometric features both in the W and blank labels have been retained.




Figure 8. Effect of relaxation on the small image when the "true" values of the compatibilities are used.




Figure 9. Remaining labeling error after 100 iterations of relaxation on the full 117x196 pixel Landsat image, using the "true" compatibilities.




Figure 10. Labeling error versus number of iterations when relaxation is applied to the 117x196 pixel image, using several values of d_i from Figure 9.




Figure 11. Illustration of the usefulness of supervising the relaxation process with a map of probabilities in order to segment the label conversion process. The image was initialized with $p^0(W \text{ or } b) = 0.99$ and 300 iterations of relaxation were applied.