N O T I C E

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

80-10267

AgRISTARS

SR-PO-00454 NASA CR-
NAS9-15466 *160697*

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

Supporting Research

February 1980

Technical Report

Pixel Labeling by Supervised Probabilistic Relaxation

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

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

NASA

TECHNICAL REPORT

PIXEL LABELING BY SUPERVISED

PROBABILISTIC RELAXATION

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.

FEBRUARY 1980

PIXEL LABELING BY SUPERVISED
PROBABILISTIC RELAXATION[*]

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

Abstract

A simple modification to existing probabilistic relaxation procedures is suggested which allows the information contained in initial labels to exert an influence on the direction of relaxation throughout the process. In this manner, the initial labels assume more importance than with conventional algorithms and are used in combination with the outcome of relaxation at each iteration to produce a cooperative estimate of the correct label for a particular object. Pixel labeling examples are presented which show the performance that can be obtained with the modified algorithm. The procedure is readily generalized to allow other data to influence the process.

INTRODUCTION

Probabilistic relaxation procedures iteratively modify an initial estimate of the labeling of a scene element by reference to spatial context. Several algorithms have been proposed for this purpose; Rosenfeld et al [1] have devised a technique that introduces context by means of correlations of labels between objects and their neighbors. Zucker and Mohammed [2] have suggested schemes that depend instead upon the conditional probability of occurrence of a particular label on an object in view of the labeling on neighbors. More recently an algorithm, also based upon a probabilistic interpretation of context, and which has a probabilistic rather than heuristic basis, has been proposed [3] as have variations on Rosenfeld's algorithm [4,5] and an algorithm derived from a constrained optimization of a mixed consistency and ambiguity criterion [6].

It is an essential ingredient of the above schemes that the initial scene labeling is used only once, viz., when the algorithm is initialized, and thereafter the success of the final labeling is dependent upon both the attributes of the algorithm and the accuracy of the contextual data, both of which can be envisaged as assuming more significance relative to the initial labeling as relaxation proceeds. This may not be a difficulty in picture-labeling problems such as the "toy triangle" example often used [1,2] since the initial labeling is seen mainly as an initialization procedure and the context information is often known with certainty. The situation can be quite different, however, in pixel labeling exercises such as those undertaken in the interpretation of Landsat images. For example, when it is desired to determine a label for every pixel in an image, the contextual information would generally not be known exactly and indeed may only be an estimate based upon "typical" image data of a similar type. Further, the initial labeling, by

and large, would represent "the best one could do" based upon all information
at hand, apart from context. In such a situation, the information is therefore
contained very much in <u>both</u> the context <u>and</u> the initial labels. As relaxation
is applied, it is desirable that both of these sources be used to produce
final labels which are. as far as possible, consistent with both the context
and the initial labels. The <u>complexity</u> of pixel labeling also makes this
desirable. Unlike many simple object labeling exercises in picture processing,
where there exists only a small number of possible (final) label distributions,
ambiguity in pixel labeling can be enormous. Clearly the reduction of
ambiguity already undertaken in the production of the initial labels is there-
fore worth maintaining to some extent during relaxation to steer the process
towards a reasonable and narrow subset of all possible label distributions.
In the following, a procedure is presented for achieving this by allowing the
algorithm to keep sight of the initial labels while creating contextual con-
sistency.

PIXEL RELAXATION LABELING

As a vehicle for the discussion, consider the arithmetic-averaging-rule
procedure of Zucker and Mohammed [2], which can be written

$$p_i^{k+1}(\lambda) = p_i^k(\lambda)q_i^k(\lambda)/\sum_\lambda p_i^k(\lambda)q_i^k(\lambda) \tag{1}$$

$$\text{with } q_i^k(\lambda) = \sum_j d_{ij} \sum_{\lambda'} p_{ij}(\lambda|\lambda')p_j^k(\lambda')$$

in which $p_i^{k+1}(\lambda)$ is the $k+1^{th}$ estimate of the probability that pixel i
has label λ given that pixel j has label λ', and the d_{ij} are a set of weighting
constants which satisfy $\sum_j d_{ij} = 1$. Often, as is the case here, all d_{ij} are
considered equal. When this is not the case, they allow some neighbors to be
more influential than others in modifying the label probability estimates.

In applying an algorithm of the type (1) to pixel labeling, it is first necessary to establish the "neighborhood" that is considered significant. In this note the neighborhood is taken to include only those pixels immediately above, below, and to each side of the pixel of interest (the four "nearest neighbors"). It is also important to establish a set of probabilities on the image boundaries that can be used when relaxing pixels adjacent to boundaries. In the absence of any indication to the contrary, it seems reasonable to assume that all labels are equiprobable on the boundary, an assumption that has been adopted throughout the remainder of this paper.

Figures 1a and 1b show the results of two simple relaxation exercises where either of two labels W or $\underline{b}$ (representing "blank") has to be assigned to each of a set of pixels. In each case, the context conditional probabilities were determined by counting joint and individual occurrences in the "true" data. However, rather than computing four different sets of these corresponding to each different neighbor type (left, right, above, and below), a single set was calculated by counting joint occurrences both vertically and horizontally.

These examples illustrate a difficulty with simple application of the scene relaxation algorithm of (1); the same will be true of other techniques also. In Figure 1a, it is seen that final labeling has apparently occurred in 7 iterations. Inspection of the set of label probabilities, however, reveals that although most of the pixel probabilities at this stage are near a fixed point (i.e., $p_i(\lambda) = 0$ or 1)*, those near the corners of the W field are not and are thus still susceptible to modification. As relaxation proceeds beyond 7 iteration, corners begin to disappear and all corner regions are substantially weakened in the final labeling as seen in the Figure.

The reason for the disappearance of the W labels can be appreciated by examining the relaxation mechanism in the vicinity of a corner pixel. The

---

*For a discussion of conditions regarding fixed points, see Zucker et al [7].

probability that the corner pixel has label W is modified by its neighbors via the conditional probabilities $p(W|W)$ and $p(W|\underline{b})$. Similarly the probability that its label is blank is modified via the context conditionals $p(\underline{b}|\underline{b})$ and $p(\underline{b}|W)$. However, once the label probabilities of the neighbors have approached the fixed points of 0 or 1, only one context conditional probability in each of the above pairs is significant, depending upon the label at a neighbor which has the higher probability. This is depicted in Figure 2. Suppose $p(\underline{b}|\underline{b}) > p(W|W)$; then for a two-label problem $p(\underline{b}|W) > p(W|\underline{b})$ also. Consequently the situation in Figure 2b is dominant, i.e., the context probabilities acting to enhance the probability that the corner pixel is blank are stronger than those which act to enhance the probability that the label is W. As a result, the W label weakens as observed. However, should $p(W|W) > p(\underline{b}|\underline{b})$, the corner pixel labeling obviously would not have weakened in this example, although such an arrangement of conditionals would lead to a loss of the blank label on the corner pixel of a blank region (i.e., an internal corner within a W field). This effect is evident in the example of Figure 1b.

Label degradation effects similar to those mentioned above could be expected to occur when segment boundaries intersect a boundary of the image (on which a set of boundary condition probabilities has been established).

As a result of the above effects, labeling accuracy, although improving in the early stages of a relaxation exercise, will undergo a subsequent gradual degradation; this trend can also be seen in results presented by other investigators [3].

III. SUPERVISED RELAXATION LABELING

It is proposed now that the degradation described above could be reduced by giving the algorithm access to the initial data, thus enabling it to form an overall impression of the correct label for a particular pixel. Clearly

not too much weight should be placed on the initial labels since they may also contain errors. However, some judicious combination of both the context data and the initial labeling would seem worthwhile.

The initial labeling can be made to exert an influence on the direction of the relaxation process in several ways. One would be to permit the pixel currently under modification to contribute to its own neighborhood function in an appropriate manner. Another, which has some useful generalizations, is presented below.

At the k^{th} iteration, the probability that pixel i has the particular label λ is given, from an application of (1), as $p_i^k(\lambda)$. If the corresponding initial labeling probability is $p_i^0(\lambda)$, then it would be desirable that $p_i^k(\lambda)$ be increased relative to the other label estimates at the i^{th} pixel if $p_i^0(\lambda)$ is the largest initial estimate. This can be achieved by modifying each of the label probabilities at the k^{th} iteration according to

$$p^i(\lambda)^+ = p^k(\lambda) \ [1 + \beta(Np_i^0(\lambda) - 1)] \tag{2}$$

followed by a normalization, to ensure the results are also properly probabilities. In (2), $p^k(\lambda)^+$ is the label probability estimate modified by reference to the supervising data, N is the number of possible labels, and β is a factor that can be used to adjust the degree of influence the initial labeling probability has in the modification procedure. (In particular $\beta = 0$ corresponds to no modification, leading to the relaxation procedure given simply by (1).) The new label estimates are re-entered in (1) to proceed with the next iteration of relaxation. For a two-label problem, a variation on (2) -- which is easily implemented computationally and which leads to probabilities which require no further normalization -- is illustrated in Figure 3. The parameter α shown in that diagram controls the amount by which $p_i^k(\lambda)$ is changed, and is given by

$$\alpha = \tfrac{1}{2} + \beta(p_i^0(\lambda) - \tfrac{1}{2}). \qquad (3)$$

Owing to the manner in which the iterates are enhanced or weakened by reference to external (initial) data, the authors have referred to this modified procedure as <u>supervised</u> relaxation labeling.

IV. EXPERIMENTAL RESULTS

The example of Figure 1a was chosen to test the supervised relaxation algorithm; the final labeling achieved is shown in Figure 4. Comparing Figures 4 and 1a, it is seen that incorporating supervision into the relaxation procedure has circumvented label weakening at the corners which would otherwise occur. Inspection of the label probabilities confirms that the results in Figure 4, including the corner region pixels, are at fixed points of 0 or 1 after about 40 iterations and thus cannot change further, i.e., cannot be weakened. By comparison, the label probabilities for the corner pixels without supervision weaken to zero, i.e., $W \to \underline{b}$. Had the initial labeling been erroneous on a corner, that error of course would also have appeared in the final labeling, even with supervision.

Figure 5 shows the example of Figure 1b redone with supervision incorporated into the relaxation algorithm. As seen, there is again a substantial improvement offered by supervising with the initial label estimates owing to preservation of the internal corners.

The model data sets of the previous examples were used to enable individual pixels to be examined in determination of labeling errors. As an indication of the performance of supervised relaxation on more extensive, real data, the results shown in Figure 6 are presented. This shows labeling error versus number of iterations, using supervised and unsupervised relaxation, obtained in a wheat/nonwheat labeling exercise. In this, an area in Kansas was classified

from multitemporal Landsat data, using a minimum distance to means classifier to provide the initial labeling. Using the neighborhood assumption chosen earlier and determining β empirically, the results shown were obtained. As seen, without supervision the labeling error passes through a minimum and then degrades; with supervision, the error falls in a monotonic fashion to a value not too different from the minimum in the unsupervised curve.

In general, there is an optimum value of β that should be used. A large β imposes a large degree of modification on each iterate and thus presumably reflects a high degree of confidence in the accuracy of the initial labels. Obviously, errors in the initial labels are more likely to influence the final labeling for large β. As yet no theoretical guidelines have been derived that permit a value for β to be chosen on analytical grounds. Consequently, in practice it would be desirable to have training data available to enable this parameter to be determined just as classifier parameters are established using prototype information.

V. CONCLUSIONS

In regions of an image where the initial labeling is in error, the relaxation mechanism and the direction provided by supervision are often in conflict. Although this conflict is ultimately resolved, it does tend to slow the initial convergence of the process, as can be noted in a comparison of Figures 1, 4 and 5. This slower start is, however, more than offset in the examples investigated by the authors, since the supervised procedure settles into a fairly stable, steady state whereas the weakening processes, referred to in Section 2 above, persist for quite some time in the absence of supervision, leading to gradual changes in labeling.

The structure of the algorithm proposed in Section III above does not, in principle, restrict the supervising information to the initial labeling.

Rather it would be perfectly possible to use any form of information about the area to be labeled to supervise the process. A situation could be envisaged, for example, where ancillary data becomes available after an extensive initial labeling exercise has been undertaken. These data could then be used separately or together with the initial labeling to guide the relaxation process. In such a situation the $p_i^0(\lambda)$ in (2) would be altered to reflect the influence of the ancillary information.

BIBLIOGRAPHY

1. A. Rosenfeld, R. Hummel & S. Zucker, "Scene Labeling by Relaxation Algorithms," _IEEE Trans. Systems, Man, and Cybernetics_, vol. SMC-6, June 1976, pp. 420-433.

2. S. Zucker & J. Mohammed, "Analysis of Probabilistic Relaxation Labeling Processes," _Proc. 1978 IEEE Conf. Pattern Recognition and Image Processing_, Chicago, IL, May 1978, pp. 307-312.

3. S. Peleg, "A New Probabilistic Relaxation Scheme," _Proc. 1979 IEEE Conf. Pattern Recognition and Image Processing_, Chicago, IL, Aug. 1979, pp. 337-434.

4. S. Peleg & A. Rosenfeld, "Determining Compatibility Coefficients for Curve Enhancement Relaxation Processes," _IEEE Trans. Systems, Man, and Cybernetics_, vol. SMC-8, July 1978, pp. 548-555.

5. H. Yamamoto, "A Method of Deriving Compatibility Coefficients for Relaxation Operators," _Comp. Graph. Image Processing_, vol. 10, 1979, pp. 256-271.

6. O. Faugeras & M. Berthod, "Scene Labeling: An Optimization Approach," _Proc. 1979 IEEE Conf. Pattern Recognition and Image Processing_, Chicago, IL, Aug. 1979, pp. 318-326.

7. S. Zucker, E. Krishnamurthy & R. Haar, "Relaxation Processes for Scene Labeling: Convergence, Speed and Stability," _IEEE Trans. Systems, Man, and Cybernetics_, vol. SMC-8, Jan. 1978, pp. 41-48.
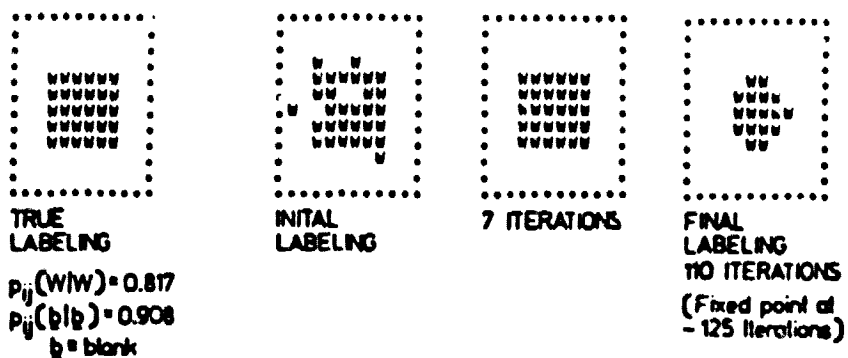
a
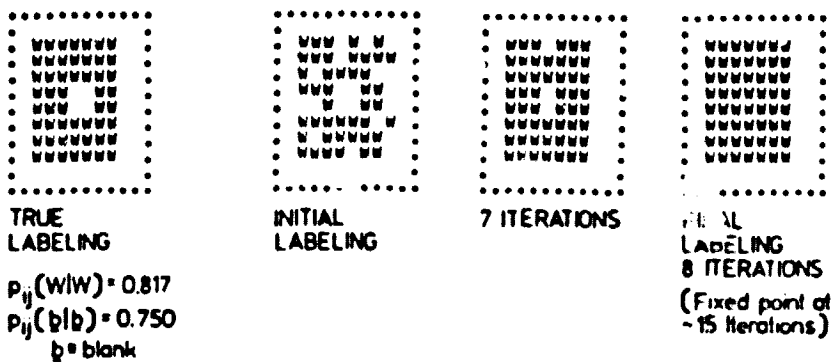
b

Figure 1.   Pixel labeling examples using the algorithm of (1).
   (a)   Demonstration of "W" label weakening
        at corners.
   (b)   Demonstration of "b" (blank) label weakening
        at corners.
   In the initial labeling, pixels shown as W were
   initialized with p(W) = 0.9, whereas those shown
   as blank were initialized with p(b) = 0.9.   This
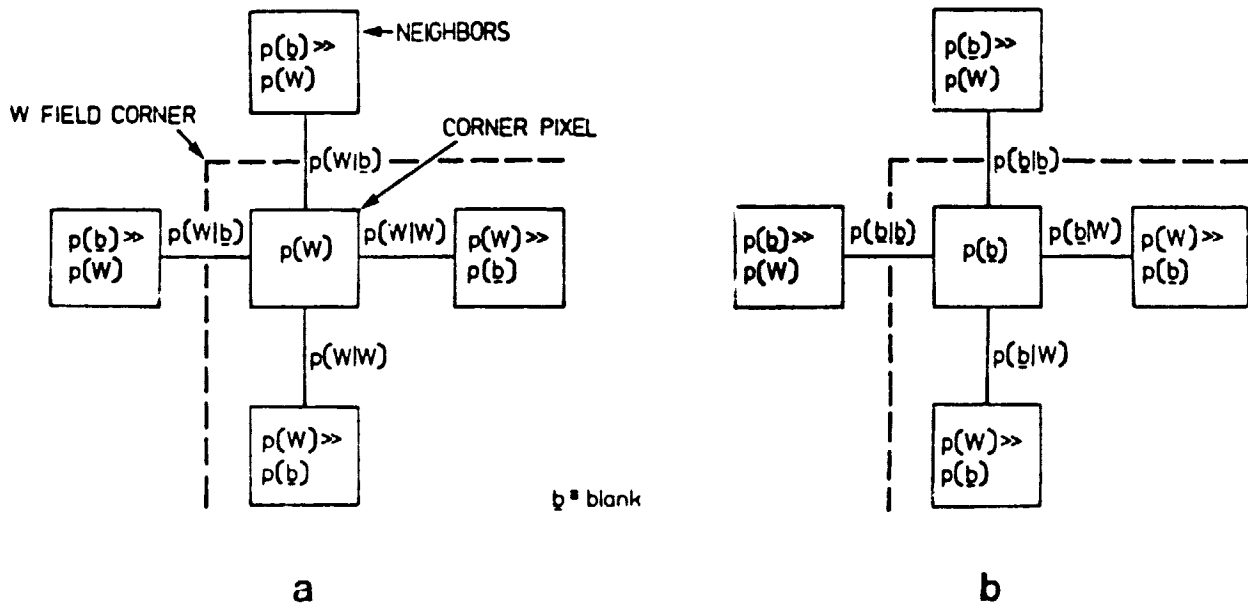   choice was made in all model examples.

Figure 2. The significant context conditional probabilities acting to establish a label of (a)W and (b) b (blank) on a corner pixel, at a stage in the relaxation procedure where all pixel probabilities are near 0 or 1.

$0 \leqslant p_i^K(\lambda) \leqslant \frac{1}{2}$:

$p_i^K(\lambda)^* = 2a \, p_i^K(\lambda)$

$\frac{1}{2} \leqslant p_i^K(\lambda) \leqslant 1$:

$p_i^K(\lambda)^* = 2(1-a) p_i^K(\lambda)$
$\qquad -1 + 2a$

$p_i^K(\lambda)$ — Label probability estimate given from an application of (1).

$p_i^K(\lambda)^*$ — Label probability estimate modified by reference to supervising data, implicit in a as defined in (3). This estimate is fed back to equation (1).

Figure 3. Prescription for modifying iterates by reference to supervising information for a two-label problem.

TRUE
LABELING

INITIAL
LABELING

FINAL
LABELING
20 ITERATIONS

(Fixed point at
~50 Iterations)

$p_{ij}(W|W) = 0.817$

$p_{ij}(b|b) = 0.908$

$\underline{b} = blank$

$\beta = 0.30$

Figure 4.   Demonstration of the usefulness of
supervision during relaxation in
avoiding loss of "W" labels at field
corners.   Compare with Figure 1a.

TRUE
LABELING

INITAL
LABELING

FINAL
LABELING
26 ITERATIONS

(Fixed point at
~40 Iterations)

$p_{ij}(W|W) = 0.817$

$p_{ij}(b|b) = 0.750$

$\beta = 0.25$

Figure 5.   Demonstration of the usefulness of
supervision during relaxation in
avoiding loss of blank labels at
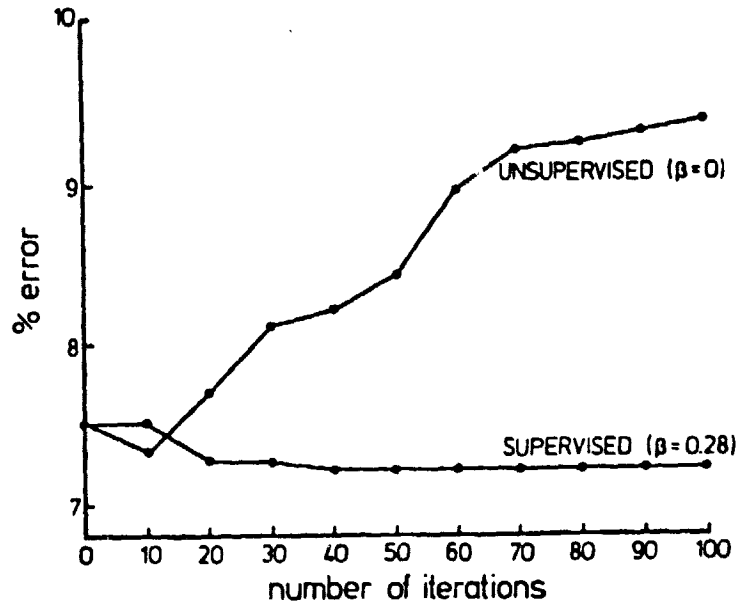internal corners.   Compare with Figure
1b.

Figure 6. Label error with and without supervision for
a wheat/nonwheat classification exercise.
The image consisted of 4000 pixels which
were labeled initially as wheat or nonwheat
by using a minimum distance to means classifier
on multitemporal Landsat acquisitions over
Kansas.