These results indicate that shadow detection is difficult prior to segment noise cleaning, but may be possible for the nonnoise segments. Shadows can be used to verify the recognition of segments belonging to buildings and roads; the former should have shadows if they are appropriately oriented, but the latter should not.

VII. Concluding Remarks

The approach used here is quite straightforward. It proceeds in an essentially bottom-up fashion, with no provision for feedback between levels, and it makes no use of higher-level information, e.g., that buildings are alongside roads, or that roads form a connected network. Its relative success in spite of these restrictions illustrates the possibility of achieving reasonable performance with a simple bottom-up approach.

The programs described in this correspondence made use of a number of empirically chosen constants. In some cases, these simply represented liberal thresholds defined by round numbers (e.g., 9, 0.2, 8, 5 percent, and the rate of linear falloff in Section IIff; two-thirds and 155° in Section III-A; 25° and 10 percent in Section III-B; etc.) In other cases, they were based on information about scale (i.e., the sizes (in pixels) of the buildings and roads that were to be detected; cf. the four-pixel strip width used in Section II) or grayscale (i.e., their contrasts; cf. the ten gray level range in Section III-A), and would have to be adjusted for different types of imagery. In any case these parameters worked well for all five of the examples on which the program was tested [8], two of which are given here, as well as for two other examples taken from an aerial photograph of a different part of the country.

These programs were not designed to be computationally efficient; their running time was 10–20 min on a time-shared Univac 1108. It is evident, however, that the method used here could be implemented very efficiently using suitable parallel hardware, since the processing of segments is largely local.

Many specific improvements in the approach are possible at each of its stages. The process of fitting straight line segments to connected components of edge pixels is somewhat order-dependent, and tends to produce overshoots; it might be better to use a Hough-like approach to detect clusters of collinear edge pixels. Rather than making a succession of decisions about segments, pairs of segments, and groups of segments, it might be better to design a hierarchical relaxation scheme in which segments are assigned probabilities of being parts of roads, buildings, etc., and iteratively adjusting the probabilities based on their compatibilities with those of related segments. Of course, one should not expect that the probability adjustments in this relaxation scheme could be based on a simple algebraic formula; more likely, they would be computed by a probabilistic "decision tree" associated with each segment. An approach to defining quantitative compatibilities for pairs of segments that continue one another or are antiparallel is investigated in [9]–[10]. These modifications of the approach should yield improved performance and still further improvement should be obtained using higher-level information and a more flexible control structure. However, much work will have to be done before human-level performance at tasks of this type can be achieved.

Acknowledgment

References

[1] R. Nevatia and K. Babu, "Linear feature extraction and description," *Computer Graphics Image Processing*, vol. 13, pp. 257–269, 1980.

[2] M. J. Fischler, G. J. Agin, H. G. Barrow, R. C. Bolles, L. H. Quam, J. M. Tenenbaum, and H. C. Wolf, "The SRI road expert: an overview," in *Proc. Image Understanding Workshop*, Nov. 1978, pp. 13–19.

[3] M. Nagao, T. Matsuyama, and Y. Ikeda, "Region extraction and shape analysis in aerial photographs," *Computer Graphics Image Processing*, vol. 10, pp. 195–223, 1979.

[4] M. Nagao and T. Matsuyama, *A Structural Analysis of Complex Aerial Photographs*. New York: Plenum, 1980.

[5] S. Peleg and A. Rosenfeld, "Straight edge enhancement and mapping," Univ. of Maryland, College Park, MD, Comput. Sci. TR-694, Sept. 1978.

[6] G. J. VanderBrug, "Experiments in iterative enhancement of linear features," *Computer Graphics Image Processing*, vol. 6, pp. 25–42, 1977.

[7] B. J. Schachter, A. Lev, S. W. Zucker, and A. Rosenfeld, "An application of relaxation methods to edge reinforcement," *IEEE Trans. Systems, Man, Cybernetics*, vol. 7, pp. 813–816, 1977.

[8] M. Tavakoli and A. Rosenfeld, "Toward the recognition of buildings and roads on aerial photographs," Univ. of Maryland, College Park, MD, Comput. Sci. TR-913, July 1980.

[9] M. Tavakoli and A. Rosenfeld, "A method for linking pairs of compatible linear features," Univ. of Maryland, College Park, MD, Comput. Sci. TR-930, Aug. 1980.

[10] M. Tavakoli and A. Rosenfeld, "A method for finding pairs of antiparallel features," Univ. of Maryland, College Park, MD, Comput. Sci. TR-943, Sept. 1980.

Image Smoothing and Segmentation by Cost Minimization

K. A. NARAYANAN, DIANNE P. O'LEARY, AND AZRIEL ROSENFELD, FELLOW, IEEE

Abstract — Many types of images are composed of regions in which the gray level is approximately constant. Such images can be smoothed and segmented by constructing piecewise constant functions whose constant parts correspond to the regions. The desired function must be (piecewise) smooth and must also be close to the original image; thus we can regard it as minimizing a two-part cost measure, in which one component measures roughness and the other measures distance, e.g., from the original image. The functions obtained by minimizing the cost with respect to various measures of this type are compared. The method of steepest descent is used for cost minimization. The relationship of this approach to other methods of image smoothing, including relaxation methods, is also discussed.

I. Introduction

In 1971 Martelli and Montanari [1] formulated the problem of digital image smoothing as a problem of minimizing a two-part cost measure summed over the image, where one component of the cost measures "roughness" (e.g., squared gradient magnitude) and the other measures discrepancy from the original noisy image (e.g., squared difference). They applied this idea to smoothing an array of local slope values obtained from a fingerprint image, rather than to smoothing image gray level. More recently, Weszka and Rosenfeld [2] discussed the role of roughness and discrepancy measures in image segmentation by thresholding, and studied methods of choosing thresholds so as to minimize the roughness of the thresholded image.

This correspondence investigates the use of roughness plus discrepancy cost measures in gray level smoothing. A steepest descent method is used to minimize the cost. When applied to an image that is approximately piecewise constant, this tends to yield a piecewise constant result, thus defining a segmentation of the image. The relationship of this approach to other iterative methods of image smoothing and segmentation, in particular relaxation methods, is also discussed. The approach turns out to

be computationally expensive, but it is of interest because of its conceptual significance.

II. THE MINIMIZATION PROCESS

Let $F(x, y)$ be the given (noisy) image. We want to find an image $G(x, y)$ that minimizes

$$C \equiv \sum_x \sum_y \left[R^2 + \alpha(F - G)^2 \right]$$

where R^2 is a roughness measure computed for a neighborhood of each point in G, and α is a factor which determines the relative weight given to roughness and discrepancy. The sum is computed pointwise over the entire image.

The minimization technique used in this correspondence is the method of steepest descent. We construct a sequence of images $F \equiv F^{(0)}, F^{(1)}, F^{(2)}, \cdots$ by adjusting the current $F^{(k)}$ at each point in a direction that reduces

$$C^{(k)} \equiv \sum_x \sum_y \left[(R^{(k)})^2 + \alpha(F^{(0)} - F^{(k)})^2 \right]$$

subject to the restriction that the values of $F^{(k)}$ remain within the allowed gray level range. Specifically, let $g^{(k)}(x, y) = -\partial C^{(k)}/\partial F^{(k)}(x, y)$; then we set

$$F^{(k+1)}(x, y) = F^{(k)}(x, y) + \lambda^{(k)} g^{(k)}(x, y)$$

where $\lambda^{(k)}$ is chosen to minimize $C^{(k+1)}$ and insure that $F^{(k+1)}$ remains within the range.

Note that $\partial C^{(k)}/\partial F^{(k)}(x, y) = 2\sum_x \sum_y R^{(k)} \partial R^{(k)}/\partial F^{(k)}(x, y) - 2\alpha(F^{(0)} - F^{(k)})(x, y)$, where the sum has only a few nonzero terms, since $R^{(k)}$ depends on $F^{(k)}(x, y)$ only at a few points near (x, y). In Section III, we give the $\partial R/\partial F$ terms for several roughness measures R, and in Section IV we discuss the results of the process when these roughness measures are used and when discrepancy is measured between $F^{(0)}$ and $F^{(k)}$. In Section V we discuss other ways of measuring discrepancy and their possible advantages.

III. ROUGHNESS MEASURES

The following roughness measures R were used in our experiments (superscript (k) omitted).

a) The Digital Laplacian:

$$R_L(x, y) \equiv 4F(x, y) - [F(x - 1, y) + F(x, y - 1) + F(x + 1, y) + F(x, y + 1)].$$

In this case it can be verified that

$$\frac{1}{2} \frac{\partial \sum_u \sum_v R_L^2(u, v)}{\partial F(x, y)} = 20F(x, y) - 8[F(x - 1, y)$$
$$+ F(x, y - 1) + F(x + 1, y)$$
$$+ F(x, y + 1)] + 2[F(x - 1, y - 1)$$
$$+ F(x - 1, y + 1) + F(x + 1, y - 1)$$
$$+ F(x + 1, y + 1)]$$
$$+ [F(x - 2, y) + F(x, y - 2)]$$
$$+ [F(x + 2, y) + F(x, y + 2)].$$

Equivalently, it is equal to

$$\sum_{u=x-1}^{x+1} \sum_{v=y-1}^{y+1} R_L(u, v) [\partial R_L(u, v)/\partial F(x, y)].$$

b) The Digital Gradient Magnitude:

$$R_G(x, y) \equiv \left[(F(x, y) - F(x - 1, y))^2 + (F(x, y) - F(x, y - 1))^2 \right]^{1/2}.$$

Here

$$\frac{1}{2} \frac{\partial \sum_u \sum_v R_G^2(u, v)}{\partial F(x, y)} = \frac{1}{2} \sum_{u=x}^{x+1} \sum_{v=y}^{y+1} \frac{\partial R_G^2(u, v)}{\partial F(x, y)}$$
$$= 4F(x, y) - [F(x - 1, y)$$
$$+ F(x, y - 1) + F(x + 1, y)$$
$$+ F(x, y + 1)]$$

c) The "Cornerity" Measure [3]:

$$R_C(x, y) \equiv \frac{1}{9} U(x, y) V(x, y) - \frac{1}{16} W^2(x, y)$$

where

$$U(x, y) \equiv F(x - 1, y - 1) + F(x - 1, y) + F(x - 1, y + 1)$$
$$+ F(x + 1, y - 1) + F(x + 1, y)$$
$$+ F(x + 1, y + 1) - 2[F(x, y - 1) + F(x, y)$$
$$+ F(x, y + 1)]$$
$$V(x, y) \equiv F(x - 1, y - 1) + F(x, y - 1) + F(x + 1, y - 1)$$
$$+ F(x - 1, y + 1) + F(x, y + 1) + F(x + 1, y + 1)$$
$$- 2[F(x - 1, y) + F(x, y) + F(x + 1, y)]$$
$$W(x, y) \equiv F(x - 1, y + 1) + F(x + 1, y - 1)$$
$$- F(x - 1, y - 1) - F(x + 1, y + 1).$$

This measure, like others studied in [4], is high when both the gradient magnitude and the rate of change of gradient direction are high—in other words, at strong edges which are turning sharply. Note that the gradient and Laplacian magnitudes are high at strong edges even if they are not turning.

For this measure we have

$$\frac{1}{2} \frac{\partial \sum_u \sum_v R_C^2(u, v)}{\partial F(x, y)} = \sum_{u=x-1}^{x+1} \sum_{v=y-1}^{y+1} R_C(u, v) \frac{\partial R_C(u, v)}{\partial F(x, y)}$$
$$= \sum_{u=x}^{x+1} \sum_{v=y-1}^{y+1} R_C(u, v)$$
$$\cdot \left[\frac{1}{9} U(u, v) \frac{\partial V(u, v)}{\partial F(x, y)} + \frac{1}{9} V(u, v) \right.$$
$$\cdot \frac{\partial U(u, v)}{\partial F(x, y)} - \frac{1}{8} W(u, v) \frac{\partial W(u, v)}{\partial F(x, y)} \right].$$

(The closed form in terms of the values of F in the neighborhood of (x, y) is quite complicated and will not be given here.) The advantage of this measure is that, unlike the Laplacian and gradient, it does not penalize straight or gently curved edges in the image; this seems more appropriate if we are trying to approximate the image by a set of constant regions of simple shapes.

IV. RESULTS

In this section we present the results obtained when the roughness measures of Section III are used, and when discrepancy is measured between $F^{(0)}$ and $F^{(k)}$, as in Section II. The initial images are shown in Fig. 1; the first is an infrared image of a tank, and the second is a portion of a Landsat image of the Monterey, CA area to which approximately Gaussian noise has been added.

Fig. 2(a)–(c) shows results for the tank image using the roughness measures R_L, R_G, and R_C, respectively. In Fig. 2(a) and (b) results are shown for $\alpha = 0.1, 1, 5, 10$ after 10 iterations, and in
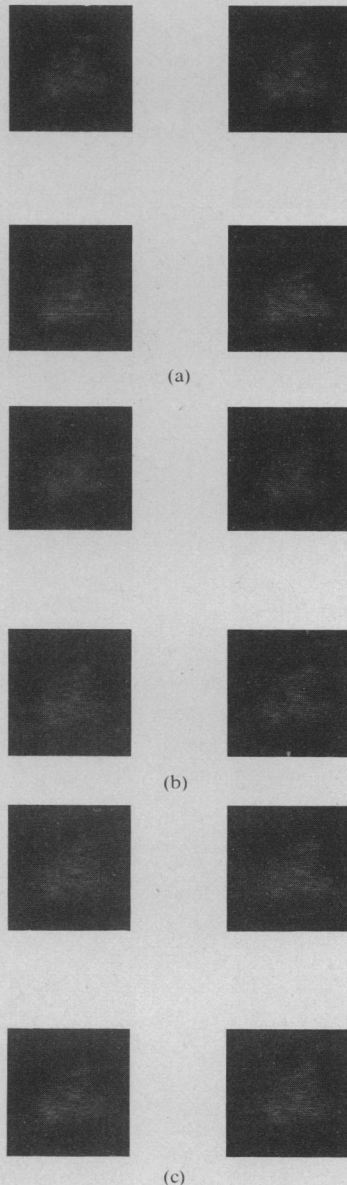
Fig. 1.   Original images.

(a)

(b)

(c)

Fig. 2.   Results for tank images using discrepancy from original image in cost.
(a) Roughness measure R_L. (b) Roughness measure R_G. (c) Roughness measure R_C.

Fig. 2(c) for $1, 10, 20, 30$ after 30 iterations. For the first two measures, low values of α (i.e., giving little weight to the discrepancy) yield blurred results. On the other hand, high values of α for the first two measures, and all values for the third measure, yield results in which the noise cleaning is very slight.

The difference in performance between the first two measures and the third measure is not hard to understand. When roughness is measured by the squared magnitude of the gradient or Laplacian, a piecewise constant image does not have very low cost, since these measures are high at the edges. When the cornerity measure is used, on the other hand, a piecewise constant image will have low cost provided the edges bounding the pieces are not sharply curved. (Note that this conforms to the intuitive concept of "piecewise constant" as "composed of constant pieces of appreciable sizes and simple shapes".) Thus when little weight is given to discrepancy, the gradient and Laplacian measures can only be reduced by blurring the image, whereas the cornerity measure can be reduced even while retaining sharp contrasts in the image. Note also that it takes more iterations for the cornerity results to converge, since its derivative varies over a wide range, and only a few points are corrected at each iteration.[1]

The results obtained in this section are not very satisfactory; they are not nearly as good, for example, as the results obtained in [5], [6] using relaxation methods. One reason for this seems to be that in the relaxation approach, we are specifying that there are two classes—i.e., we are trying to obtain a piecewise constant function that takes on two values. In the next section we consider possible ways of incorporating the knowledge that there are two classes into the smoothing process.

V.   SPECIFYING TWO CLASSES

One way of specifying that two classes are desired is to measure discrepancy not from the original image, but from a thresholded version of that image. Since such a version is two-valued, this introduces a tendency toward a two-valued solution. Fig. 3 shows results for the tank image when discrepancy is measured from $F^{(0)}$ thresholded at its mean rather than from $F^{(0)}$ itself. In Fig. 3(a) and (b), $\alpha = 0.1, 1, 5, 10$ and there were 10 iterations, while in Fig. 3(c), $\alpha = 1, 10, 20, 30$, and there were 30 iterations. These results appear to be much more strongly two-valued, but they are rather noisy. (Fig. 3(a), using the Laplacian roughness measure, is the least noisy, presumably because the Laplacian is very sensitive to noise.) This is probably because the thresholded image itself (Fig. 4) is noisy, and the discrepancy measure penalizes any attempt to remove this noise. In general, it will not be possible to avoid noise in the thresholded image; in fact, if it were easy to threshold so as to produce a nonnoisy result, the smoothing process would not be necessary, since the thresholded image would be smooth.

An alternative approach which makes use of thresholding is to define the cost function to favor values that are far away from the threshold, rather than penalizing discrepancy from the thresholded image. In other words, our cost function now depends on roughness and on closeness to the threshold, rather than on roughness and discrepancy. Since we are no longer using discrepancy, the noisiness of the thresholded image no longer causes a problem. For threshold t, the new cost function is

$$C'^{(k)} \equiv \sum_x \sum_y \left[\left(R^{(k)2} - \alpha(t - F^{(k)})^2 \right) \right].$$

Results using t equal to the mean of $F^{(0)}$ are shown in Figs. 5–7 for both images. In Figs. 5 and 6 $\alpha = 0.1, 1, 5, 10$, and there were 10 iterations; in Fig. 7 $\alpha = 1, 10, 20, 30$, and there were 30 iterations. The results are still somewhat noisy, but there is more smoothing than in Fig. 3, and the images still tend toward values at the extremes of the gray level range, especially for large α.

The first few iterations produce the greatest changes, and the image then remains relatively stable, even though the process has not completely converged. Figs. 8 and 9 show R_L and R_G results after 30 iterations for $\alpha = 1$ and 5, corresponding to the second and third parts of Figs. 5(a) and 6(a), and Fig. 10 shows R_C results after 80 iterations for $\alpha = 20$ and 30, corresponding to the third and fourth parts of Fig. 7(a). Fig. 11 shows R_L results for $\alpha = 10$ after iterations 2, 4, 6, and 8.

---

[1] The process converges when $C^{(K+1)} = C^{(K)}$, i.e., when all the components of the gradient are zero. In our experiments, however, we stopped after a relatively small number of iterations, by which time the cost had become significantly lower than at the start of the process. An example illustrating how the cost changes during the first few iterations will be given in the next section.
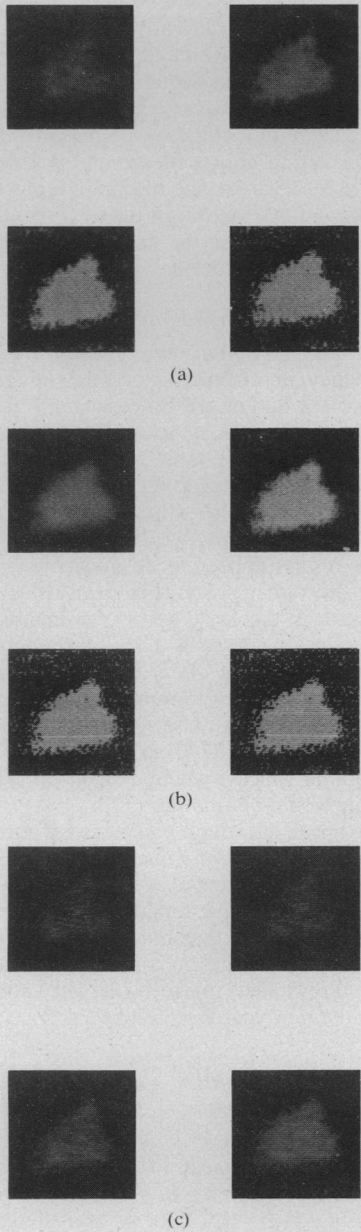
(a)

(b)

(c)

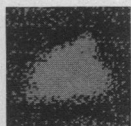Fig. 3. Analogous to Fig. 2 using discrepancy from thresholded image in cost.

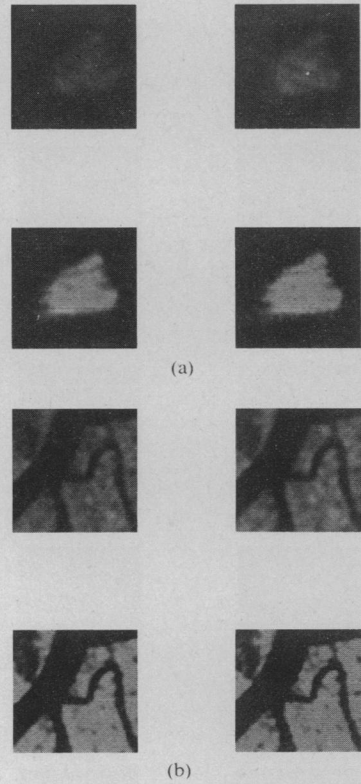Fig. 4. Tank image thresholded at its mean.

(a)

(b)

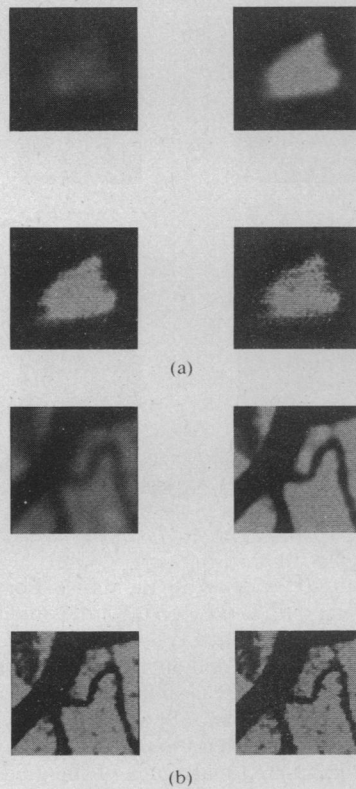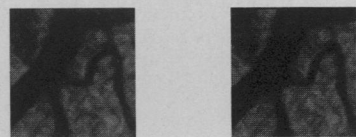Fig. 5. Results using closeness to gray level mean in cost for roughness measure R_L.

(a)

(b)

Fig. 6. Analogous to Fig. 5 for roughness measure R_G.

(a)

(b)

Fig. 7. Analogous to Fig. 5 for roughness measure R_C.

Fig. 8. R_L results after 30 iterations for $\alpha = 1, 5$.
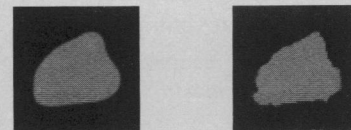

Fig. 9. R_G results after 30 iterations for $\alpha = 1, 5$.
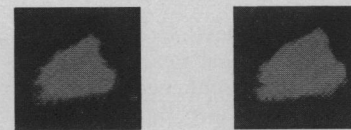

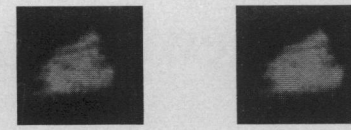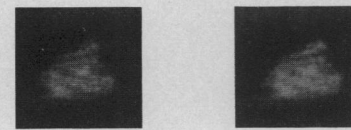Fig. 10. R_C results after 80 iterations for $\alpha = 20, 30$.

Fig. 11. R_L results after 2, 4, 6, and 8 iterations for $\alpha = 10$.

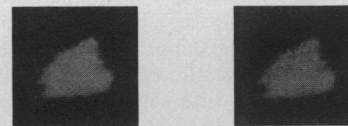Fig. 12. Results using Faugeras relaxation scheme.

Fig. 13. Results using Hummel–Zucker and Peleg relaxation schemes.

VI. COMPARISON WITH RELAXATION RESULTS

A "relaxation" method of image smoothing based on iterative local probability adjustment was investigated by Faugeras and Berthod in [5]. Initially, the probability that each point is white is taken to be proportional to its gray level, i.e., if B and W represent the ends of the gray level range, the white probability of F is $p = (F - B)/(W - B)$, while its black probability is $1 - p$. The local inconsistency of the probability assignment $p(x, y)$ can be measured by the distance between $p(x, y)$ and $q(x, y)$, where $q(x, y)$ is a function of the neighboring probabilities (see [5] for the details). The ambiguity of $p(x, y)$ can be locally measured by $p(1 - p)$, which is a maximum for $p = 0.5$. Faugeras' relaxation process uses a variant of steepest descent to iteratively adjust the p's so as to minimize a linear combination of inconsistency and ambiguity, summed over all x and y. Note that the inconsistency term corresponds to the first part of our cost function, since it is high when the neighboring values are very different; and the ambiguity term corresponds to the second part, since it is minimized when the values are far from the middle of the range. Thus Faugeras' process, as applied to the smoothing problem, is analogous to the approach used here, but using cost function components derived from white and black probability estimates rather than from the gray levels themselves. In [6] Faugeras points out that (for the case of L classes) minimizing $\Sigma_L p_i (1 - p_i) = 1 - \Sigma_L p_i^2$ and minimizing $\Sigma_L (p_i - q_i)^2$ have opposite effects, since minimizing $\Sigma_L (p_i - q_i)^2$ tends to minimize $\Sigma_L p_i^2$ and $\Sigma_L q_i^2$. He therefore proposes as an alternative simply minimizing $-\Sigma_L p_i q_i$; this tends to make $\Sigma (p_i - q_i)^2$ small without minimizing Σp_i^2 and Σq_i^2, and yields good results without the need for a $\Sigma p_i (1 - p_i)$ component.

The results of ten iterations for Faugeras' cost function $-\Sigma pq$, using our steepest descent algorithm, are shown in Fig. 12; they are very similar to our results (e.g., the last pictures in Fig. 5). For comparison, results obtained using two other relaxation processes based on probability estimation [7] are shown, for the tank image (eight iterations), in Fig. 13; they are somewhat cleaner but otherwise comparable.

The relaxation approach has a straightforward generalization to more than two classes. However, initializing the probabilities of these classes requires some knowledge of the ideal classes (e.g., obtained by fitting a mixture of Gaussians to the image's histogram). Given such knowledge, cost functions could be defined based on distance from the class means, or on closeness to the midpoints between the classes, so that our approach could also be applied to more than two classes.

The steepest descent procedure is computationally more expensive than the probabilistic relaxation processes, since it requires

several subiterations at each iteration in order to find the maximum feasible $\lambda(x, y)$: starting from a small initial λ, computing $F(x, y)$ and the associated cost, incrementing λ, and repeating the process as long as the cost keeps decreasing. Other descent techniques such as the conjugate gradient method would require less computation but more storage. Other optimization methods, such as nonlinear Gauss–Seidel iteration, might be more effective.

VII. Concluding Remarks

Image smoothing and segmentation can often be achieved by constructing a piecewise constant function whose pieces correspond to the desired regions. We have used the method of steepest descent to construct such a function having minimum roughness and minimum discrepancy from the original image. If we know that there should be only two types of regions, light and dark, we can use discrepancy from a thresholded image, or better, closeness to the mean gray level, as a cost component in place of discrepancy from the original image. The latter method yields results very similar to those obtained by relaxation methods, but without the need to introduce "probabilities" that the pixels belong to the two types of regions. It would be of interest to apply a similar approach based on cost function minimization to other problems, such as edge and curve enhancement, that have been successfully handled using relaxation methods.

In summary, we have investigated the possibility of deriving simple piecewise constant approximations to an image by cost function minimization, using two types of images (tank and terrain). This approach yields results similar to those obtainable by various other methods. It is of conceptual interest, but it does not seem to be of great practical value because of its relatively high computational cost.

References

[1] A. Martelli and U. Montanari, "Optimal smoothing in picture processing: an application to fingerprints," in *Proc. IFIP Congr.*, 1971, Booklet TA-2, pp. 86–90.
[2] J. S. Weszka and A. Rosenfeld, "Threshold evaluation techniques," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-8, pp. 622–629, 1978.
[3] P. R. Beaudet, "Rotationally invariant image operators," in *Proc. 4th Int. Joint Conf. on Pattern Recognition*, 1976, pp. 579–583.
[4] L. Kitchen and A. Rosenfeld, "Gray level corner detection," Comput. Vision Lab., Comput. Sci. Center, Univ. of Maryland, College Park, MD, TR-887, Apr. 1980.
[5] O. Faugeras and M. Berthod, "Scene labeling: An optimization approach," *Pattern Recognition*, vol. 12, pp. 339–347, 1980.
[6] M. Berthod and O. Faugeras, "Using context in the global recognition of a set of objects: An optimization approach," in *Information Processing 80*, S. H. Lavington, Ed. Amsterdam, The Netherlands: 1980, pp. 695–698.
[7] R. Smith and A. Rosenfeld, "Thresholding using relaxation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-3, 1981, pp. 598–606.

Book Reviews

Modern Power System Analysis—I. J. Nagrath and D. P. Kothari (New Delhi, India: Tata McGraw-Hill Ltd., 1980, 375 pp.). *Reviewed by R. C. Desai, Department of Electrical Engineering, M.S. University of Baroda, Vadodara 390001, India.*

Modern power systems, owing to their interconnected, complex, and integrated natures, are assuming the proportions of large-scale systems. As such, the solutions of their problems have to be obtained by using the latest computer-oriented techniques devised for large-scale systems and the undergraduate electrical engineering students and the practicing power-system engineers must be trained accordingly to keep them abreast of rapid developments in the field. Most of the books written before the last decade lack the modern approach to power-system analysis. The few books written by power engineers during the last decade discuss only a few aspects of modern techniques. However a person with a background in control system engineering is better suited to write a book on "Power System Analysis," with the large-scale systems approach.

Thus Professor Nagrath (coauthor of a book titled *Control Systems Engineering*, published by Wiley Eastern Limited) in association with Professor Kothari has done a very thorough job of writing this book which satisfies the stated need by integrating the basic principles of power-system analysis (illustrated through simple system structures) with analysis techniques for the large-scale systems found in practice.

The book contains the following topics: Inductance, resistance and capacitance of transmission lines, representation of power-system components, characteristics and performance of power transmission lines, load flow studies, optimal system operation, automatic generation and voltage control, economic dispatch, symmetrical fault analysis and symmetrical components, unsymmetrical fault analysis, and power-system stability.

This book includes digital-computer algorithms for various system studies such as load flow, fault level analysis, stability, etc. As a special

feature it also covers the latest and practically useful topics such as unit commitment, generation reliability, optimal thermal scheduling, optimal hydro–thermal scheduling, and decoupled load flow.

In essence the book is highly comprehensive, well-organized, up-to-date and (above all) lucid and easy to follow for self-study. The book is amply illustrated with solved examples for every concept and technique employing two-, three-, or four-bus structure, as necessary and the numerous examples given at the end of each chapter makes it a very useful text for teaching purposes. The reviewer rates this book highly and recommends it to students and practicing engineers in power systems. He congratulates the authors for carrying out a splendid job.

Simulation of Control Systems—I. Troch, Ed. (Amsterdam, Netherlands and New York: North-Holland, 1978, pp. 311). *Reviewed by S. G. Tzafestas, Control Systems Laboratory, University of Patras, Greece.*

This book contains the papers presented at the International Association for Mathematics and Computers in Simulation Symposium on "Simulation of Control Systems with Emphasis on Modelling and Redundancy," held at the Technical University of Vienna (Sept. 27–29, 1978).

The contributed papers have been arranged according to their main subject in three sections. A separate section involves the invited papers by R. Tomovic, H. Rzehak, E. Pavlik, and R. Vichnevetsky. These papers give important surveys on the control of large systems, redundancy in hardware and software of process computers, interdependence of process model and simulation tool, and the difficulties of computing optimal control problems.