Scene Partitioning via Statistic-based Region Growing

Joseph Weber

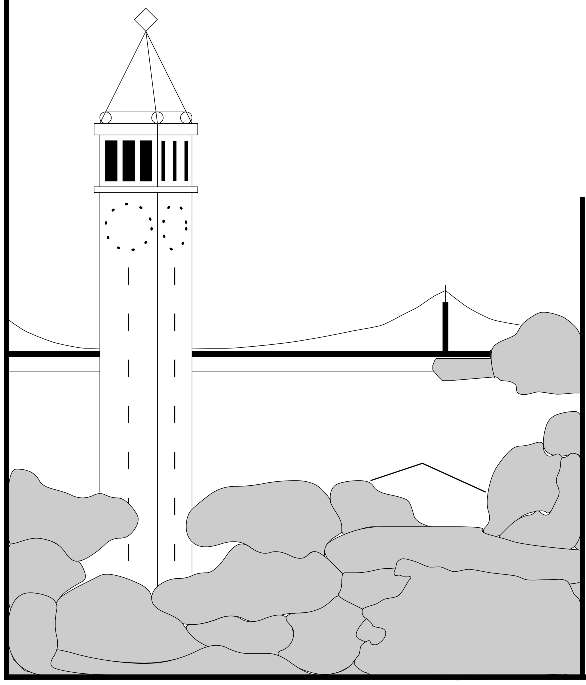

Report No. UCB/CSD 94-817

Computer Science Division (EECS)
University of California
Berkeley, California 94720

Scene Partitioning via Statistic-based Region Growing

Joseph Weber

Abstract

The scene partitioning problem *is to delineate regions in an image that correspond to the same object according to some underlying object model. Examples include partitioning an intensity image into piecewise constant intensities or identifying separate planar regions in a disparity map. A general algorithm for solving the partitioning problem in cases of linear models is presented. The algorithm uses a statistical test in a region growing formalism. The algorithm relies on the assumption that the correct image partition is connected in image coordinates. Experiments are performed on a series of models with a range of state dimensions.*

1 Introduction

Often in computational vision we make the assumption that an image can be divided into continuous regions which conform to some simplifying model. Examples include segmenting an image into piecewise constant intensities, segmenting disparity into distinct planar regions and modeling optical flow as locally affine.

The underlying model for each distinct region is usually associated with a set of model parameter values, $\boldsymbol{\mu}$. The model provides a relationship between the model parameters and the measured image function which is defined on a lattice of pixel values, $\mathbf{x}_i$. Usually the model is chosen such that a linear relationship exists between the parameters and the measurements, $\mathbf{x}_i = A_i\boldsymbol{\mu}$.

The *scene partitioning problem* is to correctly assign each pixel location with the correct set of model parameters under the simultaneous constraints of (1) minimizing error between model prediction and measurement and (2) keeping the number of partitions small.

A number of techniques have been developed to solve the partition problem. When the relationship A_i is invertible, simple thresholding or the Hough transform [DH73] can be used to distinguish groups. Region based techniques [BF70, HP74] begin with the parameters determined from subsets of the image and combine or split these subsets depending on some uniformity criterion. Energy approaches transform the two constraints into an energy functional which is minimized [MS85]. Markov interactions between neighboring pixel sets are used to introduce local priors [GG84]. More recent continuation methods begin with a smooth approximate energy functional which has a simple solution and slowly transform this solution toward a final piecewise constant solution [BZ87, Lec89, GY91].

Our approach uses region growing in which the uniformity criterion is based on statistics calculated from each region. This statistic is used to test whether adjacent regions come from the same distribution and thus can be merged into a single region. The validity of the hypothesis test for merging depends on the model used and the assumptions about the noise distribution.

In Section 2 we define the problem and enumerate the assumptions made about the scene, model and noise. In Section 3 we develop the statistic used in the merging hypothesis. Section 4 develops the algorithm used to merge initial sets into a final partition set. In Section 6 the performance of the algorithm on various partition problems with a range of state and measurement dimensions is examined.

2 Problem Statement

We are given a collection of measurements, $\mathbf{x}_i, 1 \leq i \leq n, \mathbf{x}_i \in \mathbb{R}^m$ which we assume to be the linear projection of an element of the set of possible states, $\boldsymbol{\mu}^j \in \{\boldsymbol{\mu}\}, \boldsymbol{\mu}^j \in \mathbb{R}^n$, plus Gaussian noise with known variance:

$$\mathbf{x}_i = A_i\boldsymbol{\mu}^j + \boldsymbol{\nu}_i, \boldsymbol{\mu}^j \in \{\boldsymbol{\mu}\}, \boldsymbol{\nu}_i \sim N(0, I\sigma_i^2) \tag{1}$$

A_i is a known $m \times n$ matrix relating the state to the measurements.

The measurements $\mathbf{x}_i$ are defined on a lattice with each location i having a set of neighbor locations, N_i. We also assume that each $\boldsymbol{\mu}^j$ forms a connected set in image coordinates. The state set $\{\boldsymbol{\mu}\}$ is assumed to be finite. The elements of $\{\boldsymbol{\mu}\}$ are not known, but the minimal difference between set elements, Δ, is assumed known, i.e. $\forall \boldsymbol{\mu}^i, \boldsymbol{\mu}^j \in \{\boldsymbol{\mu}\} : |\boldsymbol{\mu}^i - \boldsymbol{\mu}^j| \geq \Delta$.

We wish to associate the correct state $\boldsymbol{\mu}^j$ to each measurement $\mathbf{x}_i \rightarrow \boldsymbol{\mu}^j$ while simultaneously minimizing the difference between measurement and model prediction, and the number of distinct connected regions. In addition we would like to estimate the state set $\{\boldsymbol{\mu}\}$. We could form a cost functional which incorporates the two terms to be minimized as a function of state assignment. The cost of assigning states $\boldsymbol{\mu}_i$ to each pixel location i could be written:

$$\sum_i \left[(\mathbf{x}_i - A_i \boldsymbol{\mu}_i)^2/\sigma_i^2 + \alpha \sum_{j \in N_i} (1 - \delta(|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j|))/2 \right] \tag{2}$$

The second summation is a penalty term for introducing a discontinuity in the state assignment. It is equal to the length of the boundaries between state regions. Equation (2) would be minimized over the set of assignments $\boldsymbol{\mu}_i$. The resulting field $\boldsymbol{\mu}_i$ is piecewise constant on the lattice. This is the same cost functional used by Leclerc [Lec89]. Such cost functions are the basis for continuation methods.

Unfortunately we are not given the membership identities of the $\mathbf{x}_i$. That is, we do not know which $\boldsymbol{\mu}^j$ each $\mathbf{x}_i$ comes from. We will exploit the assumption that the regions in which $\boldsymbol{\mu}_i$ is constant are connected in using a region-growing technique to find the connected sets.

3 Merging Statistic

Suppose we have correctly partitioned two adjacent sets of measurements, ω_1 and ω_2. The sets have n_1 and n_2 elements respectively. The least squares estimates for $\boldsymbol{\mu}^i$ in each group are $\hat{\boldsymbol{\mu}}_1$ and $\hat{\boldsymbol{\mu}}_2$ where $\hat{\boldsymbol{\mu}}_j = (H_j^T H_j)^{-1} H_j^T X_j$, H_j is a $(n_j m) \times n$ matrix composed of the normalized individual A_i/σ_i matrices stacked onto each other and X_j is a column vector comprised of the normalized individual $\mathbf{x}_i/\sigma_i$ vectors stacked on top of each other.

$$H_j^T = (A_1^T/\sigma_1 \; A_2^T/\sigma_2 \ldots A_{n_j}^T/\sigma_{n_j}), \; X_j^T = (\mathbf{x}_1^T/\sigma_1 \; \mathbf{x}_2^T/\sigma_2 \ldots \mathbf{x}_{n_j}^T/\sigma_{n_j}) \tag{3}$$

Under the assumption that these groups were *correctly* partitioned, i.e. they contain pixels which come from a single model, these estimates have distributions $N(\boldsymbol{\mu}^1, (H_1^T H_1)^{-1})$ and $N(\boldsymbol{\mu}^2, (H_2^T H_2)^{-1})$ respectively. The residual error for each group is the first term in the cost functional. For group j this is:

$$S_j^2 = \sum_{i \in \omega_j} (\mathbf{x}_i - A_i \hat{\boldsymbol{\mu}}_j)^2/\sigma_i^2 = ||X_j - H_j \hat{\boldsymbol{\mu}}_j||^2 \tag{4}$$

The quantity S_j^2 has distribution $\chi^2(n_j - 1)$ since it is the sum of n_j squared $N(0,1)$ variables.

If we were to merge sets w_1 and w_2 into a single set, we would have a new residual, S_{1+2}, and least squares estimate, $\hat{\boldsymbol{\mu}}_{1+2}$. The least-squares estimate $\hat{\boldsymbol{\mu}}_{1+2}$ is a linear combination of measurements from both sets. It can also be written as a linear combination of the two least-squares estimates from the separate sets.

$$
\begin{aligned}
\hat{\boldsymbol{\mu}}_{1+2} &= (H_{1+2}^T H_{1+2})^{-1} H_{1+2}^T X_{1+2} \\
&= \hat{\boldsymbol{\mu}}_1 - (H_1^T H_1 + H_2^T H_2)^{-1} H_2^T H_2 (\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2) \\
&= \hat{\boldsymbol{\mu}}_2 - (H_1^T H_1 + H_2^T H_2)^{-1} H_1^T H_1 (\hat{\boldsymbol{\mu}}_2 - \hat{\boldsymbol{\mu}}_1) \qquad (5)
\end{aligned}
$$

$$(6)$$

After some matrix algebra, S_{1+2}^2 can be written as

$$
\begin{aligned}
S_{1+2}^2 &= \sum_{i \in \omega_1} (\mathbf{x}_i - A_i \hat{\boldsymbol{\mu}}_{1+2})^2 / \sigma_i^2 + \sum_{i \in \omega_2} (\mathbf{x}_i - A_i \hat{\boldsymbol{\mu}}_{1+2})^2 / \sigma_i^2 \\
&= S_1^2 + S_2^2 + (\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2)^T H_1^T H_1 (H_1^T H_1 + H_2^T H_2)^{-1} H_2^T H_2 (\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2) \qquad (7)
\end{aligned}
$$

We can see that the combined residual, S_{1+2}^2, is the sum of the separate residuals plus a positive term. This positive term has an expectation of zero if $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2$ since the expectation of the two least squares estimates is the same: $< \hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2 >= 0$.

We define the following statistic based on this positive term

$$
F = \frac{n_1 + n_2 - n}{n} \frac{S_{1+2}^2 - S_1^2 - S_2^2}{S_1^2 + S_2^2} \qquad (8)
$$

where again n is the dimension of the state variable. This statistic is the ratio of the residual when the two regions are merged and the residual when they are separate. From [Rao73] we know that the statistic F has the distribution $F(n, n_1 + n_2 - n)$ under the hypothesis that the sets are from the same model, i.e. $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2$.

We can use the value of F to decide whether two sets can be merged. Large values of F indicate that the hypothesis that the two sets have the same state variable is probably false. Choosing an α-significance level such that $P(F > F_\alpha) = \alpha$, we reject the hypothesis if $F > F_\alpha$.

The variables mentioned so far and and their distributions are enumerated below.

$$
\begin{aligned}
S_j^2 &= \sum_{i \in A_i \omega_j} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_j)^2 / \sigma_i^2 = ||X_j - H_j \hat{\boldsymbol{\mu}}_j||^2 \\
S_j^2 &\sim \chi^2(n_j - 1) \\
\hat{\boldsymbol{\mu}}_j &= (H_j^T H_j)^{-1} H_j^T X_j \\
\hat{\boldsymbol{\mu}}_j &\sim N(\boldsymbol{\mu}_j, (H_j^T H_j)^{-1}) \\
S_{1+2}^2 &= \sum_{i \in \omega_1} (\mathbf{x}_i - A_i \hat{\boldsymbol{\mu}}_{1+2})^2 / \sigma_i^2 + \sum_{i \in \omega_2} (\mathbf{x}_i - A_i \hat{\boldsymbol{\mu}}_{1+2})^2 / \sigma_i^2 \\
&= S_1^2 + S_2^2 + (\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2)^T H_1^T H_1 (H_1^T H_1 + H_2^T H_2)^{-1} H_2^T H_2 (\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2) \\
\hat{\boldsymbol{\mu}}_{1+2} &= (H_{1+2}^T H_{1+2})^{-1} H_{1+2}^T X_{1+2} \\
&= \hat{\boldsymbol{\mu}}_1 - (H_1^T H_1 + H_2^T H_2)^{-1} H_2^T H_2 (\hat{\boldsymbol{\mu}}_1 - \hat{\boldsymbol{\mu}}_2) \\
&= \hat{\boldsymbol{\mu}}_2 - (H_1^T H_1 + H_2^T H_2)^{-1} H_1^T H_1 (\hat{\boldsymbol{\mu}}_2 - \hat{\boldsymbol{\mu}}_1)
\end{aligned}
$$

Hypothesis	Statistic	Distribution
$\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2$	$F = \frac{n_1+n_2-n}{n} \frac{S^2_{1+2}-S^2_1-S^2_2}{S^2_1+S^2_2}$	$F(n, n_1+n_2-n)$
$\lvert\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\rvert < \Delta$	$F' = \frac{n_1+n_2-n}{n} \frac{S^2_{1+2}-D-S^2_1-S^2_2}{S^2_1+S^2_2}$	$\simeq F(n, n_1+n_2-n)$

Table 1: The hypotheses and statistics used to test them. $D = \Delta^T H_1^T H_1 (H_1^T H_1 + H_2^T H_2)^{-1} H_2^T H_2 \Delta$.

$$\begin{aligned} F &= \frac{n_1+n_2-n}{n} \frac{S^2_{1+2}-S^2_1-S^2_2}{S^2_1+S^2_2} \\ &\sim F(n, n_1+n_2-n) \end{aligned}$$

The statistic F is used to decide between two hypotheses: the distributions are from the same state or they are not. However the image model which assumes that the image can be segmented discretely into regions fitting the linear model is usually only an approximation. The image function is expected to deviate from this model. We are actually concerned not with regions with different model parameters but with *significantly* different model parameters.

We therefore would like to test the hypothesis that $|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2| < \Delta$. The fit error would be

$$\min_{|\boldsymbol{\mu}'_1 - \boldsymbol{\mu}'_2| < \Delta} S'^2_{1+2} = \sum_{i \in \omega_1} (\mathbf{x}_i - A_i \hat{\boldsymbol{\mu}}'_1)^2/\sigma_i^2 + \sum_{i \in \omega_2} (\mathbf{x}_i - A_i \hat{\boldsymbol{\mu}}'_2)^2/\sigma_i^2 \tag{9}$$

The minimal value of S'_{1+2} is difficult to calculate. We notice from equation (7) however that S'^2_{1+2} is bounded from below by $S^2_{1+2} - \Delta^T H_1^T H_1 (H_1^T H_1 + H_2^T H_2)^{-1} H_2^T H_2 \Delta = S^2_{1+2} - D$. That is, the less restrictive constraint can reduce the residual error by as much as D. We introduce a new statistic which includes this term:

$$F' = \frac{n_1+n_2-n}{n} \frac{S^2_{1+2}-D-S^2_1-S^2_2}{S^2_1+S^2_2} \tag{10}$$

and assume it also has a $F(n, n_1+n_2-n)$ distribution. We can then use the F_α confidence threshold to decide if two regions can be merged under the less restrictive hypothesis. The two hypotheses and their distributions are enumerated in Table 1.

Previous uses of the F-test in computer vision include Laprade and Doherty's work [LD87] where still images were segmented into piecewise constant intensity regions. This is the same as our first experiment in Section 6. Bouthemy and Rivero [BR87] and [HHNO94] segment optical flow into affine regions using the difference in residuals. This case is also treated in the experimental section of this paper.

4 Algorithm

Our region-growing algorithm uses the statistic F' to merge connected regions into larger regions of uniform state, assuming that these regions contain points from a single state. Using the statistic is valid only if the initial regions are correctly partitioned. We form initial regions by grouping neighboring pixels into small initial sets and assuming that these small sets contain elements from a single state. These initial sets can also be as small as a single pixel and thus are guaranteed to contain measurements from only one state.

Small initial sets from neighboring pixels are formed.
Δ is set to zero.
While $\Delta < \Delta_{final}$
    While sets can merge
        Each set determines via the F' statistic if any of the adjacent sets come from the same distribution. If they do, the sets are merged.
    Increase Δ
The final sets are tested for possible non-adjacent sets being from the same distribution.

Table 2: Region growing by statistical measure algorithm.

Each set then determines via the F' statistic if it has the same distribution as any of its neighbor sets. Neighbors having a value of the statistic below the α-confidence level are merged into a single region. This continues until no more adjacent regions can be merged.

When the sets are very small, it is difficult to have a high confidence in the merging statistic since it is based on statistics calculated from small samples. As the sets grow in size, the statistics become more reliable. Thus, we begin with $\Delta = 0$, merging sets which have very similar state values. When all sets have merged under this more restrictive hypothesis, the value of Δ is increased. This repeats until Δ reaches a final value, Δ_f. This final value is the maximum difference between states that can be considered insignificant. The algorithm is outlined in Table 2.

The linearity of the least-squares estimate makes implementation of the algorithm very simple. Each set needs to keep track only of the summations $X_j^T X_j, H_j^T H_j, H_j^T X_j$ each of which have fixed dimension. The values of the error, S^2, and least-squared estimate, $\hat{\boldsymbol{\mu}}$, when combining two sets are simple linear calculations given the summation terms. Thus merge and hypothesis testing operations are computationally inexpensive.

Certain data structures can be used to implement the algorithm efficiently. The current list of neighbors for each set could for example be stored in an adjacency graph [BB82]. The nodes of the graph would store the summations necessary for merging.

5 Connection to Continuation Methods

A connection can be made between the algorithm outlined above and continuation methods. In continuation methods [BZ87, Lec89, GY91], the cost functional is slowly transformed from one with an easy solution, to the final desired functional. The solution is tracked, usually by local gradient descent methods, during the transformation from the easy solution to what is hoped to be the true solution at the final cost functional.

The cost functional (2) in Section 2 can be written as the sum of an error term plus the cost of line boundaries. The cost for a given segmentation is

$$E = \sum_i \left[(\mathbf{x}_i - A_i \boldsymbol{\mu}_i)^2 / \sigma_i^2 + \alpha \sum_{j \in N_i} (1 - \delta(|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j|))/2 \right] = \sum_j \left[S_j + \alpha \sum_{i \in N_j} L(i,j)/2 \right] \tag{11}$$

where $L(i,j)$ is the length of the boundary between regions i and j, and N_j is the set of regions adjacent to region j. The first summation is over pixel locations while the second summation is over regions. When $\alpha = 0$, any partition of the image into sets of fewer than n points, (recall that n is the dimension of $\boldsymbol{\mu}$), is a minimal (zero) energy solution. This corresponds to the initial partition of the image in our algorithm. As the value of α is slowly increased towards its final value, regions will merge when the cost of the boundary between them outweighs the error caused by merging them, i.e. if

$$S_{1+2} - S_1 - S_2 < \alpha L(1,2) \tag{12}$$

The merging step in this case decreases the total energy, just as in gradient descent algorithms. Comparing (12) to the definition of F', we see that increasing the value of Δ is analogous to increasing α, the relative weight given to boundaries in the cost function. As Δ increases, the D term in the numerator of F' increases. Thus larger increases in fit error can be tolerated when merging regions since the cost of a boundary between them has increased.

6 Experiments

We applied the algorithm to a number of segmentation problems. The state spaces are of a range of dimensions. We begin with the one-dimensional state space representing constant brightness in section 6.1. In this model we are attempting to segment an image into piecewise constant intensities. In section 6.2 we look at higher dimensional state spaces. In the first experiment we model range data as being piecewise planar. The state space then has dimension 3. In the second experiment, we model optical flow as being piecewise affine. Thus the state space is of dimension 6 representing affine parameters, and the measurements are of dimension 2, representing the optical flow.

In all experiments, the α-confidence level was set to 0.9. The value of Δ was increased in eleven steps beginning with a value of zero and increasing by $\Delta_f/10$ at each step where Δ_f was the final value. Thus there were a total of eleven iterations of the merging step.

6.1 One dimensional experiments

Figure 1 shows a piecewise constant image (wedding cake) with unit difference between intensities corrupted by Gaussian noise with variance 0.2 units. This figure also shows a histogram of the intensity values and the final boundaries between constant intensity regions found by the algorithm. Note from the histogram that the distributions overlap considerably. The initial sets consisted of 2×2 pixel neighborhoods. The final value of Δ was 0.8.

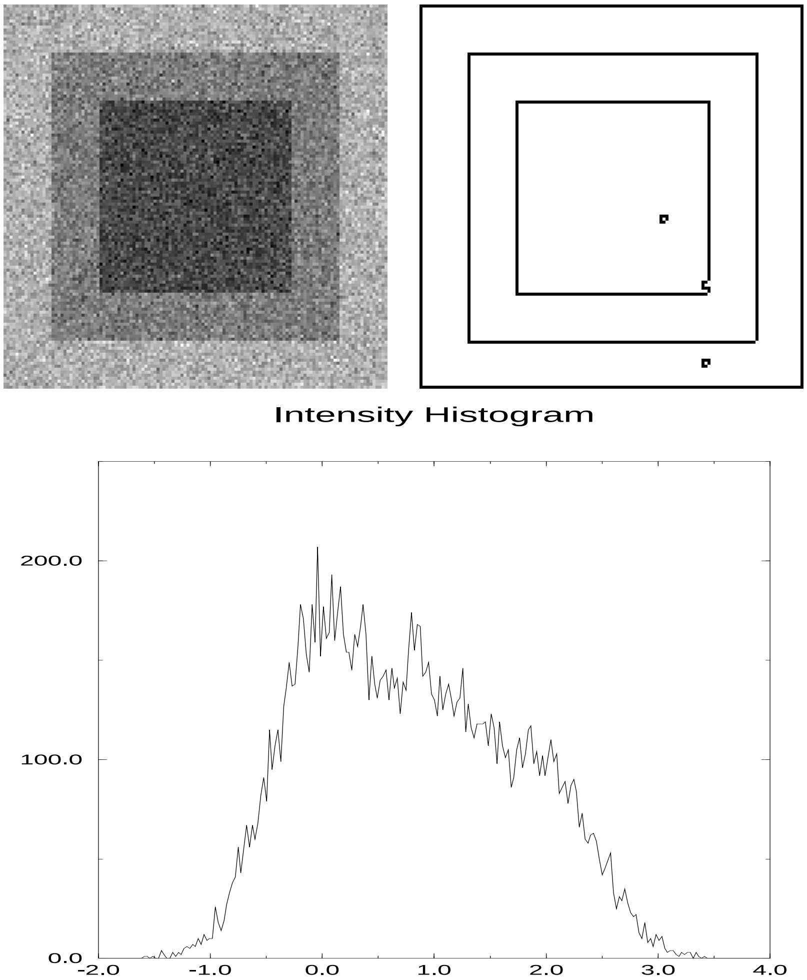

Figure 1: Piecewise constant image corrupted with Gaussian noise, the boundaries between constant intensities found by the algorithm, and the input image's histogram. Original intensity values were $0, 1$ and 2 units. Noise variance was 0.2 units.

We ran the algorithm on an aerial image of an earthquake damaged region near Los Angles. In this example, the value of Δ_f was set to 60. The image and the region boundaries are shown in Figure 2.

One thing to notice about both images is how the lack of any prior on boundaries affects the results. In both experiments we see isolated patches of only 4 pixels. The least-squares estimate in these patches is usually more than Δ units different than the surround. These isolated patches are consistent with our algorithm. However, many previous algorithms include priors to make the results conform more closely to expected human interpretations. These priors would favor straight over jagged boundaries, attempt to connect broken but aligned line segments, and prevent isolated sets such as the ones that occur in both simulations. Presently, our algorithm does not include any priors on boundaries.

6.2 Multi-dimensional experiments

6.2.1 3-dimensional state

The range map for simple objects is often modeled as being piecewise planar. Thus the range as a function of image position is modeled as $I(x, y) = a(x - x_0) + b(y - y_0) + c$. In this case we have a three-dimensional state, $(a, b, c)^T$. Surface normals can be recovered directly from the state parameters.

Figure 3 shows range data of a scene containing some simple objects, and the recovered boundaries between piecewise planar regions along with the direction of the depth gradient. We can see that for the curved object, the model fails and the curved surface is approximated by a series of planes. The minimal difference between states was set to $\Delta_f = (0.5, 0.5, 5)$.

Since the dimension of the state is 3 we need to start with initial sets of at least 3 pixels. We started with 2×2 pixel sets. This sometimes created a set which contained two different distributions, thus violating the assumption that all sets are correctly distinguished. These initial sets combine only with other sets containing the same two distributions. This creates the thin sets along boundaries that can be seen in Figure 3.

6.2.2 Affine Optical Flow

In this experiment we assume that optical flow can be modeled as piecewise affine. This has been used to segment a scene based on motion since the optical flow produced from a planar object is affine in retinal coordinates [Adi85]. Equating the measurement to the optical flow, $X_i = \begin{pmatrix} u \\ v \end{pmatrix}_i$, and the state to affine

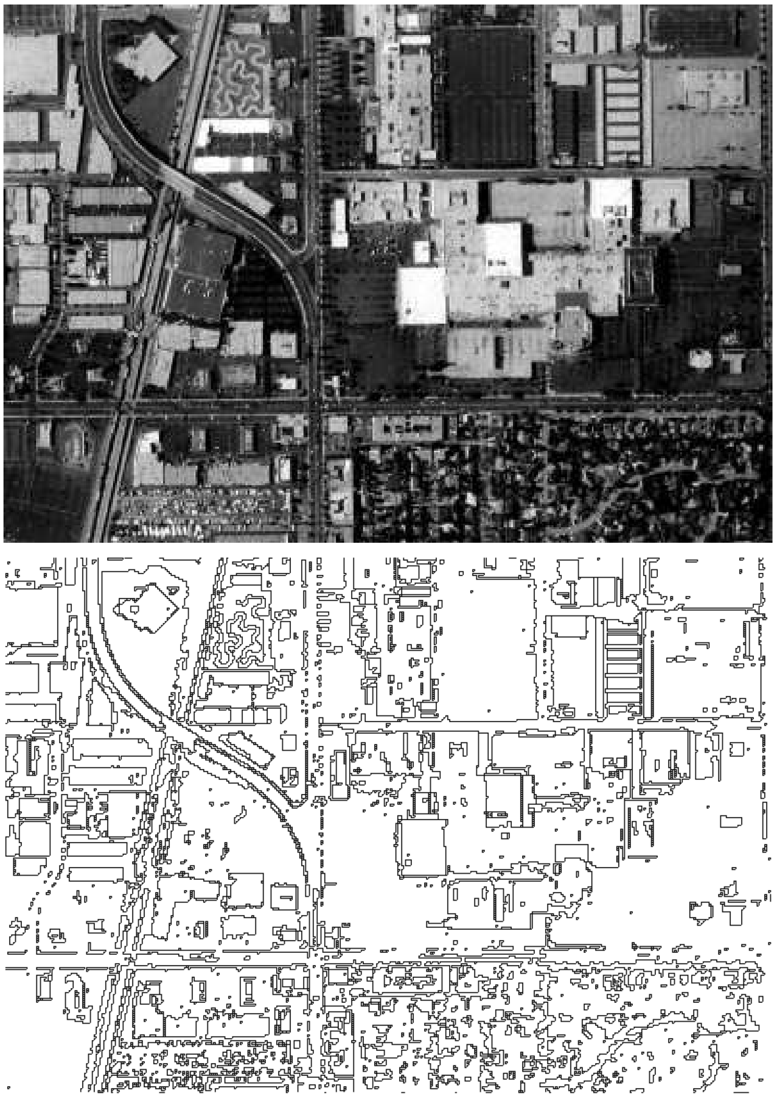

Figure 2: Aerial view of an earthquake damage in Northridge, CA. and the boundaries between constant intensities found by the algorithm.

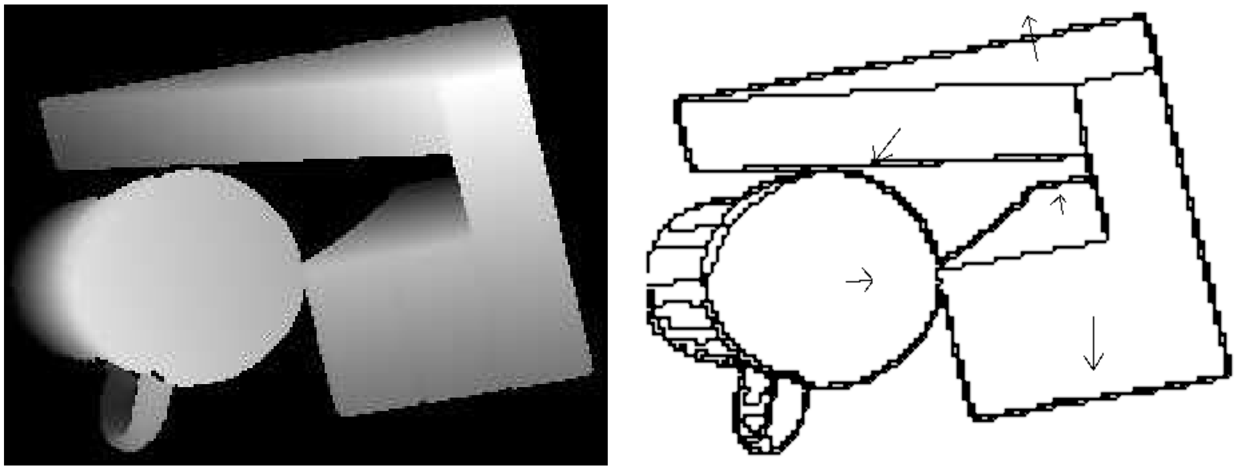

Figure 3: On the left, range data from some simple objects. On the right are the set boundaries between piecewise planar regions and the depth gradient directions found from the algorithm.

parameters $\boldsymbol{\mu} = (a, b, c, d, e, f)^T$, we have

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x - x_0 \\ y - y_0 \end{pmatrix} + \begin{pmatrix} e \\ f \end{pmatrix} = \begin{pmatrix} x - x_0 & y - y_0 & 0 & 0 & 1 & 0 \\ 0 & 0 & x - x_0 & y - y_0 & 0 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \\ d \\ e \\ f \end{pmatrix} = A_i \boldsymbol{\mu}^j \tag{13}$$

where (x_0, y_0) is the centroid of the (x, y) position of the set of points.

We used equation (13) to model the optical flow field for the "yosemite" sequence corrupted by white noise. Figure 4 shows the input field overlayed on one frame from the sequence and the resulting field with the set boundaries overlayed. The initial sets were formed from 4×4 pixel regions. A minimum of 6 pixels was needed because of the 6 dimensional state. The value of Δ_f was set to $(0.01, 0.01, 0.01, 0.01, 1, 1)/2$. Again there appears a thin region of initial sets which overlap two different model regions. These regions fail to combine with any neighbors and produce the double border we see in the Figure.

7 Conclusions

A general algorithm for solving the partitioning problem in cases of linear models is presented and tested on a series of experiments with a range of state dimensions. The algorithm use a statistical test in a region growing formalism. The algorithm relies on the assumption that the image partition is connected in image coordinates. For non-connected partitions, the algorithm would need to be extended to include non-local operations. There would also be a combinatorial explosion of hypothesis testing steps in the algorithm.

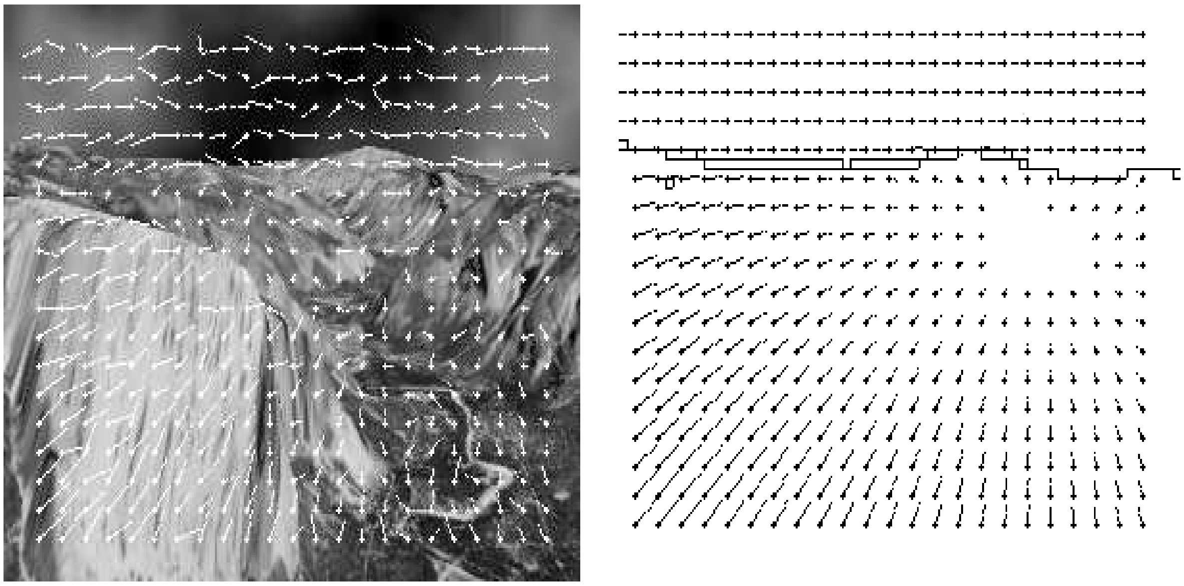

Figure 4: Noisy optical flow overlayed with single frame from sequence and the recovered affine flow region boundaries with recovered affine flow.

The region growing algorithm includes only merging operations. Beginning with initial sets, a finite number of merging and hypothesis tests are required. Due to the linearity of the merging operation, only image measurements and coordinate summations need be stored. Hypothesis testing and statistic formulation are simple linear operations on these summations. Efficient algorithms exist for implementing the model.

We find from the experiments that the lack of any prior on segmentation boundaries causes some segmentation results which may be non-intuitive. Future work will focus on incorporating priors while maintaining the algorithmic simplicity of the model.

Acknowledgements

I wish to thank Peter Belhumeur and Davi Geiger for useful discussions. This work was supported by the California Department of Transportation through the PATH project grant No. MOU-83.

References

[Adi85] Gilad Adiv. Determining three-dimensional motion and structure from optical flow generated by several moving objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(4):384–401, 1985.

[BB82] Dana H. Ballard and Christopher M. Brown. *Computer Vision*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1982.

[BF70] C. Brice and C. Fennema. Scene analysis using regions. *Artificial Intelligence*, 1(3):205–226, 1970.

[BR87] P. Bouthemy and J. Santillana Rivero. A hierarchical likelihood approach for region segmentation according to motion-based criteria. In *ICCV*, pages 463–467, London, 1987.

[BZ87] Andrew Blake and Andrew Zisserman. *Visual reconstruction.* MIT press, 1987.

[DH73] R. Duda and P. Hart. *Pattern Classification and Scene Analysis.* John Wiley & Sons, New York, Chichester, Brisbane, Toronto, Singapore, 1973.

[GG84] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, November 1984.

[GY91] Davi Geiger and Alan Yuille. A common framework for image segmentation. *International Journal of Computer Vision*, 6(3):227–243, 1991.

[HHNO94] H. Kollnig H.-H. Nagel, G. Socher and M. Otte. Motion boundary detection in image sequences by local stochastic tests. In *Proceedings of the third EECV*, volume II, pages 305–316, Stockholm, 1994.

[HP74] S.L. Horowitz and T. Pavlidis. Picture segmentation by a directed split-and-merge procedure. In *Proceedings of the 2nd International Conference on Pattern Recognition*, pages 424–433, 1974.

[LD87] Robert H. Laprade and Mark F. Doherty. Split-and-merge segmentation using an f test criterion. In *Proceedings of the SPIE, Volume 758: Image Understanding and the Man-Machine Interface*, pages 74–79, 1987.

[Lec89] Yvan G. Leclerc. Constructing simple stable descriptions for image partitioning. *International Journal of Computer Vision*, 3:72–102, 1989.

[MS85] David Mumford and J. Shah. Boundary detection by minimizing functionals. In *IEEE Computer Society Computer Vision and Pattern Recognition conference proceedings*, volume 22, 1985.

[Rao73] C. R. Rao. *Linear Statistical Inference and its Appliations.* Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, New York, London, Sydney,Toronto, 1973.