Representing Probability Distributions of Image Segments and Segmentations

Steven M. LaValle Seth A. Hutchinson
The Beckman Institute for Advanced Science and Technology
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801

Abstract **Segmentation is the problem of partitioning an image into a small number of regions, each of which is uniform or homogeneous in some sense. Although traditional approaches produce an optimal (or near-optimal) segmentation with respect to the chosen models, the problem is generally considered underconstrained. Consequently, the segmentation may not contain the best homogeneous regions needed by some higher-level process (i.e. a recognition system cannot exert complex model-based influences directly on the selection of an optimal segmentation).**

We develop a method for probabilistically maintaining sets of alternative homogeneous regions, and segmentations. Depending on the image size and complexity, and on the application, a probability distribution can be constructed over the entire image, or a distribution over partial segmentations can be formed. We develop an efficient representation structure, and a probabilistic mechanism for applying Bayesian, model-based evidence to guide the construction of the representation and influence the resulting posterior distribution over the space of alternatives. Our formalism is applied to range images using a piecewise-planar model with additive Gaussian noise.

1 INTRODUCTION

Image segmentation has been a topic of active research for over two decades. It is generally considered as the low-level vision task of determining a set of homogeneous regions (called segments) from an image, for the purpose of higher-level processing. Although considerable effort has yielded a number of approaches to the problem, segmentation remains a difficult problem in its general formulation. As Horn has pointed out, one of the primary difficulties in evaluating a segmentation method is the lack of a clear definition of the "correct" segmentation [7]. This definition usually depends on the intended application of the segmentation result. Szeliski has argued that low-level image models often underconstrain the solution, advocating the use of uncertainty estimation [14]. This type of difficulty in segmentation provides much of the impetus of our work: to represent multiple segments and segmentations, probabilistically, in a Bayesian framework.

Several statistically-based strategies for segmentation presently exist. Clustering approaches generate segments by iteratively grouping pixels or regions which have similar values in a feature metric space [8, 13]. Markov Random Fields (MRF's) encode local dependencies between image elements in a lattice of random variables, and yield a segmentation by optimizing a resulting energy function [3, 5, 14]. Probabilistic relaxation defines a compatibility measure between pairs of image elements, and yields a segmentation through iterative updating rules [11]. All of these methods define image models, perform estimation, and arrive at some optimal (or near-optimal) segmentation.

In the sections that follow we develop an organized and efficient method of building distributions of image segments and segmentations. Section 2 defines the probability spaces and representations, and provides an operation called *refinement* that iteratively constructs a distribution of segments. Section 3 describes the probability assignments that are made at each step in the construction. Section 4 presents an algorithm that guarantees finding the best n segments in a local portion of the image. Section 5 presents an experiment on a noisy synthetic range image, with a distribution of image segments, and also a distribution of full image segmentations. Conclusions and future work are presented in Section 6.

2 DEFINITIONS AND OPERATIONS

The input to the segmentation system is an array, D, of points. Associated with each element $D[i,j]$ is the representation of the point, and in this paper is a set of coordinates in $\Re^3$, denoted by $\mathbf{x} = [x_1\ x_2\ x_3]$. A given point $D[i,j]$ will have a set of points called *neighbors* which it is *adjacent* to. Using standard 4-neighbors, this set is: $D[i-1,j]$, $D[i+1,j]$, $D[i,j-1]$, $D[i,j+1]$.

Usually a segmentation is considered as a partition of the image D; however, in our work we will need to introduce some additional structure when defining a segmentation. We will introduce three terms that are used extensively in this paper: regions, segments, and segmentations.

A *region*, R, is some connected subset of D. By *connected* we mean that for any $D[i_1,j_1]$, $D[i_2,j_2] \in R$, there exists some sequence of elements of R, $\langle D[k_m,l_m]\rangle_{m=1,\ldots,n}$, in which each $D[k_{m-1},l_{m-1}]$ is

0-7803-0720-8/92 $3.00

adjacent to $D[k_m, l_m]$, and $D[k_1, l_1] = D[i_1, j_1]$, $D[k_n, l_n] = D[i_2, j_2]$. Two regions, R_1 and R_2, will be called *adjacent* if there exists some $D[i_1, j_1] \in R_1$ and $D[i_2, j_2] \in R_2$ that are adjacent.

A *segment*, T, is a connected set of pairwise-disjoint regions (e.g. $T = \{R_1, R_2, R_3\}$ is a segment consisting of three regions). A set of regions is *connected* if their union is connected, as defined above.

A *segmentation*, S, denotes a pairwise-disjoint set of segments (i.e. no region belongs to more than one segment) with the conditions that no pair of regions, throughout all the segments, may contain the same element of D, and that every point in D is an element of some region of some segment. For these conditions to be satisfied, the set of all regions must form a partition of D. Also, the set of segments define a partition of the regions.

Given a segmentation, S, and two adjacent segments[1] T_1, $T_2 \in S$, a new segmentation, S', can be formed by replacing T_1 and T_2 with $T_1 \cup T_2$, while keeping all other segments fixed. This corresponds to the familiar notions of region merging and segmentation, expressed in our set-theoretic terms.

It is usually profitable to begin with some initial partition of the image with small regions, and construct new segmentations by grouping regions, as is done in the region merging paradigm [12]. One justification for this is the savings in complexity from considering a smaller set of possible segmentations. Another reason is that often some minimal number of points is required in a region before the statistical models can be effectively employed. The initial segmentation represents the starting point in a region-merging algorithm. For instance, Silverman and Cooper begin with an initial image of *blocks*, which corresponds to an initial partition of the image into a grid of square regions [13]. Blocks are merged to yield *clusters*, which correspond to segments. This is also the approach taken in split-and-merge algorithms [6, 10].

In this paper, only the space of segmentations that can be generated by merging elements of some initial segmentation will be considered. This initial segmentation is called the *primitive segmentation*, S_p, with the regions being *primitive regions*. The set of primitive regions is denoted by $\mathcal{R}_p$. In this segmentation, every segment contains exactly one region. Let Π denote the set of all segmentations that can be constructed from the initial segmentation S_p. At one extreme, Π includes the partition induced by the original primitive regions, $S_p \in \Pi$. At another extreme, Π contains the partition corresponding to merging together all segments of S_p, giving $\{\mathcal{R}_p\} \in \Pi$.

2.1 The Segmentation Sample Space

We next define a sample space associated with Π, by considering subsets of Π and their associated probabilities. In real applications, the size of Π is quite large. Therefore in our approach, much of the sample space Π is represented implicitly by accessing only selected subsets of Π. This allows tremendous savings over representing every element $S \in \Pi$ and its corresponding probability.

[1] We mean that the union of the segments is connected.

The *segmentation sample space* (SSS) is represented by the probability triple $(\Pi, \mathcal{A}, P)$. Π is as previously defined. $\mathcal{A}$ represents the set of all subsets of Π (i.e. the power set of Π). P denotes a probability mapping, defined on $\mathcal{A}$, adhering to the standard probability axioms.

We have developed an organized structure for applying evidence to, and manipulating events on Π. The SSS is used to consider events directly on the space Π. A *segment sample space* (TSS) implicitly represents subsets of Π through the specification of groups of primitive regions, and is considered next.[2]

2.2 A Segment Sample Space

A segment sample space (TSS) essentially describes all ways to construct segments containing some specified, fixed region. It can be used to specify locally how to apply evidence, ultimately determining a probability mapping on Π, or to obtain a probability distribution of segments at some fixed image location. We next introduce the structure and operations on a TSS. Their relationship to the SSS is discussed in Section 2.5 and in [9].

Some $R_i \in \mathcal{R}_p$ is designated as the *initial region*, and Θ_i is the set of all possible segments that contain R_i. One such element of Θ_i is $\{R_i\}$. Another might contain R_i and several adjacent primitive regions. The entire set $\mathcal{R}_p$, which is connected, also belongs to Θ_i. By this definition, for any region $R_i \in \mathcal{R}_p$, there is a corresponding set of segments Θ_i.

A TSS is represented by the triple $(\Theta_i, \mathcal{B}_i, P)$. Θ_i is, as defined before, the space of all segments that include R_i. In the triple, $\mathcal{B}_i$ is the set of all subsets of Θ_i, and P is the probability mapping on $\mathcal{B}_i$, also adhering to the probability axioms. The events in $\mathcal{B}_i$ that contain exactly one segment will be termed *ground events*.

2.3 Compact representation of TSS events

We next define a method for implicitly representing elements of $\mathcal{B}_i$. For real image applications, the number of segments in Θ_i will be extremely large; the set $\mathcal{B}_i$ is exponentially larger. Any scheme that requires enumeration of either of these sets would be severely hindered by combinatoric explosion. It will be necessary to select certain elements of $\mathcal{B}_i$ in an organized manner, when evidence induced by the models is applied, and a representation of the TSS is built. This provides motivation for the definitions and representations that follow, which will be used to directly apply Bayesian evidence, affecting the probability distribution on Θ_i.

We can define an element $B \in \mathcal{B}_i$ by describing a set of primitive regions which must be included in every element of B, and a set of primitive regions which must *not* be included in any element of B. These two sets essentially contain all the information that is common to the segments in B. An *inclusion set*, I, is a set of primitive regions that includes R_i. An *exclusion set*, E, is a set of primitive regions in which each region in E is required to be adjacent to some region in I. Also, $I \cap E = \emptyset$.

[2] Note S denotes a segmentation, and hence the abbreviation SSS, and T denotes a segment T, hence TSS.

We define a function $\tau(I, E)$, which maps to some $B \in \mathcal{B}_i$, by

$$\tau(I, E) = \{T \in \Theta_i : I \subseteq T,\ E \cap T = \emptyset\}. \quad (1)$$

We state the following proposition[3]:

Proposition 1 *The mapping τ is well-defined and onto.*

The proposition implies that every event $B \in \mathcal{B}_i$ has a well-defined representation in terms of I and E sets. It is important to note that $\tau(I, E)$ provides a compact representation for potentially large subsets of Θ_i. The savings in representation are greater as the number of primitive regions in $\mathcal{R}_p$ increases.

2.4 Constructing approximate TSS representations

We now introduce two concepts, which enable us to incrementally build an approximate representation of the TSS: cover and refinement. A cover is a set of events that represent a partition of Θ_i. Refinement is an operation that takes a cover and produces a new cover, representing a better approximation of the TSS. The general strategy is to begin with an initial cover, and iteratively apply refinements until the desired representation is obtained. The probability of every event in a cover is known, and every ground event in Θ_i is represented by some element of the cover. The procedure is similar to impurity reduction through partitioning with classification and regression trees [2]. We are required to make probability assignments at each refinement, and this is detailed in Section 3.

Define a *cover*, C, of the TSS to be a set of pairwise-disjoint events in $\mathcal{B}_i$, forming a partition of Θ_i. Consider a cover $C = \{B_1, B_2, \ldots, B_k\}$ with specified probabilities on each of the events, $\{P(B_1), P(B_2), \ldots, P(B_k)\}$. If C is the set of all ground events in $\mathcal{B}_i$, then an exact representation of the TSS is obtained; all the elements of Θ_i are explicitly represented in C, and the probability is given for each of them. Suppose that $C = \{\Theta_i\}$. We know that $P(\Theta_i) = 1$, however, the probabilities of the other events in $\mathcal{B}_i$ cannot be directly determined. This corresponds to the poorest resolution, since no information is actually present about the TSS.

The goal is to obtain a cover that explicitly represents ground events that have high probability and events that have low probability only as subsets of Θ_i. The cover is a compact representation of the TSS since each event is represented by I and E sets. This representation is useful since it explicitly represents the most favorable segments and their corresponding probabilities.

The *refinement mapping*, ρ, takes a cover, C, an event, $B_\rho \in C$, and a primitive region ,R_ρ, and yields a *refined cover*, $C' = \rho(C, B_\rho, R_\rho)$. The region R_ρ, termed the *refinement region*, is adjacent to some region in I_ρ and $R_\rho \notin I_\rho \cup E_\rho$, in which $B_\rho = \tau(I_\rho, E_\rho)$. Also, B_ρ is termed the *refinement event*. The refinement mapping is formally defined as

$$C' = (C - B_\rho) \cup \{\tau(I_\rho \cup \{R_\rho\}, E_\rho), \tau(I_\rho, E_\rho \cup \{R_\rho\})\} \quad (2)$$

[3] This and subsequent propositions are proven in [9].

The only difference between C and C' is the replacement of B_ρ by $\tau(I_\rho \cup R_\rho, \{E_\rho\})$ and $\tau(I_\rho, E_\rho \cup \{R_\rho\})$. These new events are termed *refined events*. The refinement event B_ρ has been partitioned by using the refinement region R_ρ. All segments in B_ρ that include R_ρ are in $\tau(I_\rho \cup \{R_\rho\}, E_\rho)$. The remaining elements are in $\tau(I_\rho, E_\rho \cup \{R_\rho\})$.

The vast majority of all ground events will be implicitly represented as members of some event in the cover. Each of these events is a subset of Θ_i, and the cover can represent many segments that have very low probability, and are therefore uninteresting.

2.5 SSS to TSS relationships

The SSS describes the space of all possible segmentations. A TSS describes a distribution over segments, in which each segment represents a portion of each of the segmentations. The TSS can be used as a building block with which representations of events on the SSS can be constructed. It is necessary to relate the distribution over Θ_i to the distribution over Π since traditionally one is interested in full image segmentations. Also, it is useful to investigate the distribution arising from the composition of a few neighboring segments, since often only a few segments form a useful constraint for recognition [4].

A ground event, $\{T\}$, in a TSS corresponds to the SSS event containing all segmentations that contain $\{T\}$. When a TSS representation is built, a ground event can be selected, and a new TSS can be considered using the remaining primitive regions, $\mathcal{R}_p - \{T\}$. This type of iteration can be continued until every region in $\mathcal{R}_p$ is used, resulting in a segmentation. The probabilities of the segments in their respective TSS's are multiplied, yielding a probability value on the SSS. Further details of this relationship are discussed in [9].

3 PROBABILITY ASSIGNMENTS

Recall that each refinement removes one event in a cover of Θ_i, and replaces it with two disjoint events whose union is the original event. The basic strategy is to determine probability assignments of the new events when this step is performed. This requires deciding how to divide the probability of the original event between the two new events.

There are two basic mechanisms that exert influence on this probability assignment. As with any Bayesian framework, there is some prior distribution on the sample space. Also, after the application of evidence, some posterior distribution is obtained. Before constructing the representation, a prior distribution is implicitly defined on the TSS. Model-based evidence will be used, along with the prior distribution, to determine probability assignments at the refinement step.

Using the refinement mapping, successive partitions are constructed from Θ_i as prescribed by (2). Recall that in this operation, after selecting B_ρ and R_ρ, we partition $B_\rho = \tau(I_\rho, E_\rho)$ into $\tau(I_\rho \cup \{R_\rho\}, E_\rho)$ and $\tau(I_\rho, E_\rho \cup \{R_\rho\})$. For probabilistic consistency, it is necessary to have

$$P(B_\rho) = P(\tau(I_\rho \cup \{R_\rho\}, E_\rho)) + P(\tau(I_\rho, E_\rho \cup \{R_\rho\})). \quad (3)$$

A prior must be chosen to capture the notion of uniformity due to the lack of presented evidence. We split the probability assigned to B_ρ evenly, making

$$P(\tau(I_\rho \cup \{R_\rho\}, E_\rho)) = P(\tau(I_\rho, E_\rho \cup \{R_\rho\})). \quad (4)$$

This induces a prior probability mapping on the TSS.

3.1 Simplifications and the membership probability

An alternative way to represent the probability terms on the right side in (3) is by

$$P(\tau(I_\rho \cup \{R_\rho\}, E_\rho)) = P(\tau(I_\rho \cup \{R_\rho\}, E_\rho) \mid B_\rho)\ P(B_\rho) \quad (5)$$

and

$$P(\tau(I_\rho, E_\rho \cup \{R_\rho\})) = P(\tau(I_\rho, E_\rho \cup \{R_\rho\}) \mid B_\rho)\ P(B_\rho). \quad (6)$$

This is true since $\tau(I_\rho \cup \{R_\rho\}, E_\rho) \subseteq B_\rho$, and we have

$$P(\tau(I_\rho \cup \{R_\rho\}, E_\rho), B_\rho) = P(\tau(I_\rho \cup \{R_\rho\}, E_\rho)). \quad (7)$$

Similarly, for $\tau(I_\rho, E_\rho \cup \{R_\rho\}) \subseteq B_\rho$, we have

$$P(\tau(I_\rho, E_\rho \cup \{R_\rho\}), B_\rho) = P(\tau(I_\rho, E_\rho \cup \{R_\rho\})). \quad (8)$$

It is assumed that $P(B_\rho)$ is given by the previous iteration, and the conditionals of (5) and (6) are all that need to be determined. Recall that $B_\rho = \tau(I_\rho, E_\rho)$. The conditionals become

$$P_I \equiv P(\tau(I_\rho \cup \{R_\rho\}, E_\rho) \mid \tau(I_\rho, E_\rho)) \quad (9)$$

and

$$P_E \equiv P(\tau(I_\rho, E_\rho \cup \{R_\rho\}) \mid \tau(I_\rho, E_\rho)). \quad (10)$$

The probabilities, P_I and P_E, are called *membership probabilities*. P_I is the probability that R_ρ is a member of I_ρ, and P_E is the probability that R_ρ is a member of E_ρ. Since $P_I + P_E = 1$, only one of P_I and P_E needs to be computed. Hence we consider P_I, which can be rewritten by making use of the following proposition:

Proposition 2 *If I, I_1, I_2 are include sets with $I = I_1 \cup I_2$, and E, E_1, E_2 are exclude sets with $E = E_1 \cup E_2$ then $\tau(I, E) = \tau(I_1, E_1) \cap \tau(I_2, E_2)$.*

By letting $I_1 = I_\rho$, $I_2 = \{R_i, R_\rho\}$, $E_1 = E_\rho$, $E_2 = \emptyset$, (9) can be rewritten as

$$P_I = P(\tau(\{R_i, R_\rho\}, \emptyset), \tau(I_\rho, E_\rho) \mid \tau(I_\rho, E_\rho)) \quad (11)$$

This can be simplified by making use of the fact that for any two events B_1 and B_2, $P(B_1, B_2 | B_2) = P(B_1 | B_2)$. By using this directly, P_I can be reduced to

$$P(\tau(\{R_i, R_\rho\}, \emptyset) \mid \tau(I_\rho, E_\rho)). \quad (12)$$

Equation (12) is expressed in a form explicitly indicating the importance of adding R_ρ to I_ρ. This is the fundamental distinction between the event B_ρ and the two refined events. It is natural to expect that the probability due to evidence will depend directly on the new region that has been brought into consideration.

In general, for determining P_I (or P_E), an important issue between different models must be considered: the notion of IE-independence. Recall that R_i is an element of every segment in Θ_i. Intuitively the issue to consider is whether the probability of including (or excluding) R_ρ with R_i should depend on the other regions in I_ρ or E_ρ. Explicitly, this independence can be stated for P_I as

$$P(\tau(\{R_i, R_\rho\}, \emptyset) \mid \tau(I_\rho, E_\rho)) = P(\tau(\{R_i, R_\rho\}, \emptyset)). \quad (13)$$

When a model of evidence is being used under (13) it is termed *IE-independent* since the probability is independent of regions other than R_ρ and R_i. The original model is termed *IE-dependent*. The choice between these models depends on the application. In general the IE-independent model can be used when R_i contains sufficient information to make the effects of considering the rest of I_ρ, and E_ρ negligible. The IE-independent model will reduce the number of membership computations required because the same R_ρ can be chosen for numerous refinements. In this paper we consider only the IE-independent model. The IE-dependent model is discussed in [9].

3.2 Posterior evidence-based probabilities

A detailed derivation of posterior probabilities is beyond the scope of this paper. Consequently we state the resulting expressions and refer the reader to [9] for a complete treatment. The model treated here assumes that all objects in the scene are composed of planar patches. We use the Gaussian iid noise model, in which the noise component is normal to the plane [1].

We define a parameter space as $\mathbf{u} = [u_1\ u_2\ u_3]^T$ with $\|\mathbf{u}\| = 1$ and $u_3 \geq 0$, in which each point identifies a plane equation up to a scalar, d_k:

$$u_1 x_1 + u_2 x_2 + u_3 x_3 + d_k = 0. \quad (14)$$

For a fixed d_k, the space of all planar surfaces is represented by a hemisphere in the parameter space.

For a given parameter value $\mathbf{u}$ and value for d_k, we can use the region points to observe the sum-of-squares error. This is denoted by the random variable Y_k with values y_k. It can be shown that when the information in E is independent of R_i, the posterior membership probability can be expressed as:

$$P(\tau(\emptyset, \{R_i, R_\rho\}) | y_\rho, y_i) = \left[1 + \frac{1}{2\lambda_1(y_\rho, y_i)}\right]^{-1}. \quad (15)$$

in which $\lambda_1(y_\rho, y_i)$ is

$$\frac{\left[\int_{S^2} p(y_\rho | \mathbf{u}, \hat{d}_\rho) p(\mathbf{u}) d\mathbf{u}\right] \left[\int_{S^2} p(y_i | \mathbf{u}, \hat{d}_i) p(\mathbf{u}) d\mathbf{u}\right]}{\int_{S^2} p(y_\rho | \mathbf{u}, \hat{d}_{i\rho}) p(y_i | \mathbf{u}, \hat{d}_{i\rho}) p(\mathbf{u}) d\mathbf{u}}, \quad (16)$$

Above, $\hat{d}_\rho$, $\hat{d}_i$, and $\hat{d}_{i\rho}$ represent linear estimates of d_ρ, d_i, and $d_{i\rho}$ in R_ρ, R_i, and $R_i \cup R_\rho$, respectively. The functions of the form $p(y_k | u, \hat{d}_k)$ represent the χ^2 density, taken from considering the sum-of-squares error,

```
GET-TOP-SEGMENTS (R_p, R_i, n)
1    MEM-PROBS(R_i, FRONTIER(τ({R_i}, ∅)))
2    C_g ← ∅
3    C_h ← {τ({R_i}, ∅)}
4    repeat
5        B_ρ ←EXTRACT-MAX(C_h)
6        R_ρ ←SELECT-R_ρ(FRONTIER(B_ρ))
7        B_I ← τ(I_ρ ∪ {R_ρ}, E_ρ)
8        B_E ← τ(I_ρ, {E_ρ ∪ R_ρ})
9      if FRONTIER(B_I) ≠ ∅
10         then
11             INSERT(C_h, B_I)
12             MEM-PROBS(R_i, FRONTIER(B_I))
13         else
14             INSERT(C_g, B_I)
15     if FRONTIER(B_E) ≠ ∅
16         then
17             INSERT(C_h, B_E)
18         else
19             INSERT(C_g, B_E)
20   until [C_h = ∅] or
     [|C_g| ≥ n and P(NTH(C_g, n)) > P(MAX(C_h))]
21   return C_g
```

Figure 1: An algorithm returning the best n segments.

y_k, with respect to the plane given by u_k and $\hat{d}_k$. The function $p(\mathbf{u})$ represents a prior density on the parameter space, and is defined as $p(\mathbf{u}) \equiv 1/\pi$. Also, $\bar{S}^2$ represents the parameter hemisphere. These integrals are each evaluated by applying a spherical coordinate transformation and performing a 2-D, numerical volume integration.

4 ALGORITHM DETAILS

In this section we present an algorithm that generates an approximate TSS representation. We begin with the primitive regions, $\mathcal{R}_p$, the initial region, R_i, and the number of ground events with highest probability to represent, n. The adjacency structure for $\mathcal{R}_p$ is already provided. The formal description of the algorithm is given in Figure 1.

The cover at any given iteration is represented by two priority queues, C_g and C_h, sorted by probability. C_g contains the ground events, and C_h contains the remaining events. Although we use τ notation in Lines 3, 7, and 8, only the I and E sets, and the probability are represented in the implementation. The FRONTIER operation returns the set of all regions in $\mathcal{R}_p$ that can be chosen for refinement. The MEM-PROBS operation precomputes the membership probabilities for R_i and potential refinement regions. Also, NTH returns the n^{th} element in the queue.

The following proposition is used in Line 20 as a termination criterion to guarantee that the best n segments have been represented:

Proposition 3 *If* $P(\text{NTH}(C_g, n)) > P(\text{MAX}(C_h))$, *then for any* $\{T\} \in \Theta_i$ *such that* $\{T\} \notin C_g$, *then* $P(\{T\}) < P(\text{NTH}(C_g, n))$.

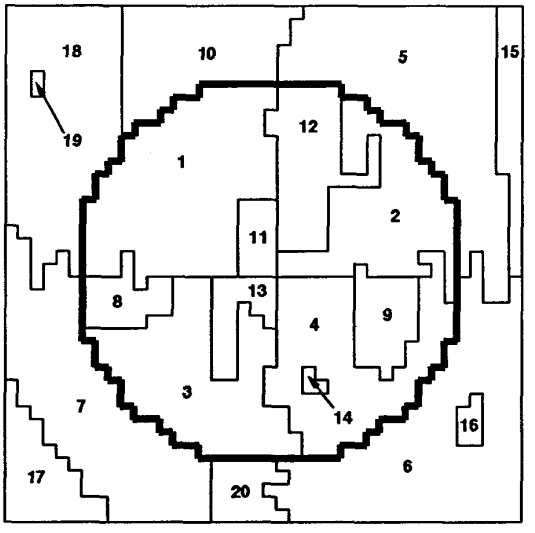

Figure 2: Range image with primitive segmentation indicated. Thick lines depict the true segment boundaries.

5 AN EXPERIMENTAL EXAMPLE

We have implemented the concepts presented here in Common Lisp and C, on a SPARC 1+ workstation. Experiments have been performed on several range images, both synthetic and real. We present one representative example here and refer the reader to [9] for more experiments.

The results that follow were obtained from a synthetic range image, consisting of 1600 points (40x40). There is unit spacing between the points and 0.707 standard deviation noise, corresponding to an extremely noisy image. The primitive segmentation is indicated in Figure 2. The thickened lines indicate where the true segment boundaries lie (in this image there are two segments). Without noise, the points inside the circle lie on the plane, $z = 0$, and the points outside the circle line on the plane, $z = 1$.

By selecting R_1 as the initial region, in ≈ 1 s we obtain the top twenty segments, shown in Table 1. Using an algorithm similar to Figure 1, that applies to the SSS, in ≈ 3 s we obtain a distribution of the top twenty segmentations, shown in Table 2. In each table, we observe that the leading choice is correct.

6 CONCLUSIONS AND FUTURE WORK

The cover and refinement concepts provide a useful way to consider distributions of image segments and segmentations. Together, the membership probability computation and the algorithm presented yield a computationally feasible approach to a problem with inherently difficult combinatorics.

We are presently developing a general framework for making the membership probability assignments, pertaining specifically to implicit polynomial surfaces with Gaussian noise, and in general to a much broader class of image/noise models.

Rank	Prob.	Primitive regions
1	.792622	(1 2 3 4 8 9 11 12 13 14)
2	.077962	(1 2 3 4 8 9 11 12 13 14 20)
3	.069580	(1 2 3 4 8 9 11 12 13)
4	.016383	(1 2 3 4 8 9 12 13 14)
5	.013355	(1 2 3 4 8 11 12 13 14)
6	.006844	(1 2 3 4 8 9 11 12 13 20)
7	.004409	(1 2 3 8 9 11 12 13)
8	.003684	(1 2 3 4 9 11 12 13 14)
9	.003679	(1 2 3 4 8 9 11 12 14)
10	.001611	(1 2 3 4 8 9 12 13 14 20)
11	.001438	(1 2 3 4 8 9 12 13)
12	.001314	(1 2 3 4 8 11 12 13 14 20)
13	.001196	(1 2 3 4 8 9 11 13 14)
14	.001172	(1 2 3 4 8 11 12 13)
15	.000828	(1 3 4 8 9 11 12 13 14)
16	.000434	(1 2 3 8 9 11 12 13 20)
17	.000362	(1 2 3 4 9 11 12 13 14 20)
18	.000362	(1 2 3 4 8 9 11 12 14 20)
19	.000323	(1 2 3 4 9 11 12 13)
20	.000323	(1 2 3 4 8 9 11 12)

Table 1: The TSS distribution for Θ_1.

Prob.	Segments
.382134	(A 9 11 14) (B 5 7 10 16 17 18 19 20)
.172566	(A 9 11 14) (B 5 7 10 16 17 18 20) (19)
.118369	(A 9 11 14) (B 5 7 10 16 17 20) (18 19)
.037907	(A 9 11 14 20) (B 5 7 10 16 17 18 19)
.033545	(A 9 11) (B 5 7 10 16 17 18 19 20) (14)
.026374	(A 9 11 14) (B 5 7 10 17 18 19 20) (16)
.017436	(A 9 11 14) (B 5 7 10 16 17 20) (18) (19)
.017118	(A 9 11 14 20) (B 5 7 10 16 17 18) (19)
.015149	(A 9 11) (B 5 7 10 16 17 18 20) (14)
.013153	(A 9 11 14) (5) (B 7 10 16 17 18 19 20)
.011976	(A 9 11 14 20) (B 5 10 16) (7 17 18 19)
.011910	(A 9 11 14) (B 5 7 10 17 18 20) (16)
.010391	(A 9 11) (B 5 7 10 16 17 20) (14) (18 19)
.008170	(A 9 11 14) (B 5 7 10 17 20) (16) (18 19)
.007899	(A 9 14) (B 5 7 10 16 17 18 19 20) (11)
.007671	(A 9 11 14) (B 5 7 16 17 18 19 20) (10)
.006439	(A 11 14) (B 5 7 10 16 17 18 19 20) (9)
.005940	(A 9 11 14) (5) (B 7 10 16 17 18 20)
.004207	(A 9 11 14) (5 10 18 19) (B 7 16 17 20)
.003567	(A 9 14) (B 5 7 10 16 17 18 20) (11)

Table 2: The SSS distribution. "A" represents regions 1, 2, 3, 4, 8, 12, and 13. "B" represents regions 6 and 15.

ACKNOWLEDGMENTS

This research was sponsored by NSF under grant number IRI-9110270. We thank Becky Anderson for helpful suggestions.

REFERENCES

[1] R. M. Bolle and D. B. Cooper. On optimally combining pieces of information, with application to estimating 3-D complex-object position from range data. *IEEE Trans. Pattern Anal. Machine Intell.*, 8(5):619–638, September 1986.

[2] P. A. Chou. Optimal partitioning for classification and regression trees. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(4):340–354, April 1991.

[3] P. B. Chou and C. M. Brown. The theory and practice of Bayesian image labeling. *International Journal of Computer Vision*, 4:185–210, 1990.

[4] O. D. Faugeras and M. Hebert. The representation, recognition, and locating of 3-D objects. *International Journal of Robotics Research*, 5(3):27–52, Fall 1986.

[5] D. Geman and S. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intell.*, 6(6):721–741, November 1984.

[6] R. M. Haralick and L. G. Shapiro. Image segmentation techniques. *Comp. Vision, Graphics, and Image Proc.*, 29:100–132, January 1985.

[7] B. K. P. Horn. *Robot Vision.* MIT Press, Cambridge, MA, 1986.

[8] J. Jolion, P. Meer, and S. Bataouche. Robust clustering with applications in computer vision. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(8):791–801, August 1991.

[9] S. M. LaValle. A Bayesian framework for considering a distribution of image segmentations. Master's thesis, Univ. of Illinois, Urbana/Champaign, August 1992.

[10] T. Pavlidas. *Structural Pattern Recognition.* Springer-Verlag, New York, NY, 1977.

[11] S. Peleg. A new probabilistic relaxation scheme. *IEEE Trans. Pattern Anal. Machine Intell.*, 2:362–369, 1980.

[12] A. Rosenfeld and A. C. Kak. *Digital Picture Processing.* Academic Press, New York, NY, 1982.

[13] J. F. Silverman and D. B. Cooper. Bayesian clustering for unsupervised estimation of surface and texture models. *IEEE Trans. Pattern Anal. Machine Intell.*, 10(4):482–496, July 1988.

[14] R. Szeliski. Bayesian modeling of uncertainty in low-level vision. *International Journal of Computer Vision*, 5(3):271–301, December 1990.