

Spatial Random Tree Grammars for Modeling Hierarchical Structure in Images with Regions of Arbitrary Shape

J.M. Siskind, J. Sherman, Jr., I. Pollak, M.P. Harper, and C.A. Bouman

Abstract—We present a novel probabilistic model for the hierarchical structure of an image and its regions. We call this model *spatial random tree grammars (SRTGs)*. We develop algorithms for exact computation of likelihoods and MAP estimates and exact EM updates for model-parameter estimation. We collectively call these algorithms the *center-surround algorithm*. We use the center-surround algorithm to automatically estimate the ML parameters of SRTGs, classify images based on their likelihood and based on the MAP estimate of the associated hierarchical structure. We apply our method to the task of classifying natural images and demonstrate that the addition of hierarchical structure significantly improves upon the performance of a baseline model that lacks such structure.

Index Terms—Bayesian methods for image understanding, multiscale analysis.

I. Introduction

Hierarchical organization can be very valuable in image analysis and classification. Recursive bipartitioning [14], [49], [53] and region merging [33], [37], using features such as intensity gradient, region shape, region color, and region texture, can yield a hierarchy of segmentations of an image. However, these methods are feature data driven and generally are not based on an overall explicit model of hierarchical structure in the image. An alternative to these approaches is the use of hierarchical Bayesian statistical models of image structure.

Researchers in the fields of speech and computational linguistics have long exploited the value of structural priors in language processing. Hidden Markov models (HMMs) are a non-hierarchical prior that has had enormous impact in speech processing [43]. Context-free grammars

(CFGs) have been used to model hierarchical sentence structure [10], [11], [12]. Probabilistic context-free grammars (PCFGs) [47] have been used to construct probabilistic hidden tree models for sentences and to robustly parse naturally-occurring text [31], [36]. An important property of these tree models is that the structure of the tree is random and can adapt to the hidden structure of the observations. As with most priors, HMMs and PCFGs require the estimation of numerous parameters. This can be done using the expectation-maximization (EM) algorithm [16] or, equivalently, the Baum-Welch reestimation procedure [2]. In both cases, the E and M steps can be computed exactly. However, in the case of HMMs, the E step is computed using the forward-backward algorithm [43], whereas in the case of PCFGs, it is computed using the more complex inside-outside algorithm [1], [26]. Variants of both algorithms can be used to compute the most probable state sequence, in the case of HMMs, or most probable parse, in the case of PCFGs [52].

Structural priors are also valuable for the analysis and processing of images and complex scenes. For example, HMMs have been used for document image recognition [22] and layout analysis [20]; and PCFGs have been applied to the event-recognition task, by processing a 1D aspect of a video stream [5]. However, the generalization of 1D priors to 2D has often proved to be very difficult. Markov random fields (MRFs) [3] are a non-hierarchical prior which is the most common extension of HMMs to 2D. For example, an MRF model of image regions was used in [32] for image interpretation. While effective methods have been developed to approximately compute maximum a posteriori (MAP) estimates [18] and maximum likelihood (ML) parameter estimates [4], [41], the exact computation of these estimates for MRFs is fundamentally intractable. HMMs have also been generalized to tree structures used for probabilistic reasoning [35] and multiscale image modeling [6], [8], [9], [15], [29], [28], [30], [45]. However, an important distinction between these models and PCFGs is that, in the former case, the tree structure is fixed and nonrandom.

Conditional random fields (CRFs) [25] model the posterior distribution directly, without imposing a prior. While the recent adaptations of CRFs to image analysis [19], [24], [50] offer interesting and promising alternatives to MRFs,

This work was supported in part by a National Science Foundation (NSF) CAREER award CCR-0093105, NSF grants 9987576, 9980054, and 0329156IIS, a Xerox Foundation grant, by ARDA under contract number MDA904-03-C-1788, and by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR0011-06-C-0023. Part of this work was carried out while the fourth author was on leave at NSF. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the view of NSF.

J.M. Siskind, I. Pollak, M.P. Harper, and C.A. Bouman are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47906-1285, e-mail e-mail qobi.ipollak, harper, bouman@ecn.purdue.edu. J. Sherman, Jr. is with the Department of Electrical and Computer Engineering, University of Maryland, College Park MD 20742 USA, e-mail shermanj@umd.edu.

they still result in intractable estimation problems which can only be solved approximately.

Various strategies have been used to apply grammar formalisms to images [17], [46], [48]. For example, K. S. Fu’s research in syntactic pattern recognition incorporated the use of PCFGs for the classification of 2D structures [17]. Deterministic context-free grammars called *X-Y trees* and associated algorithms were constructed in [23], [34] to model 2D layout of document images. Probabilistic grammars have more recently been applied to optical character recognition [42].

In this paper, we present a novel probabilistic model for the hierarchical structure of an image and its constituent components. Specifically, we present a novel method for reformulating PCFGs to model 2D images or other multidimensional data. We call this method *spatial random tree grammars* (SRTGs) and the associated algorithms the *center-surround algorithm*. The center-surround algorithm plays an analogous role to the inside-outside algorithm in that it allows for exact computation of likelihoods and MAP estimates and exact EM updates for model-parameter estimation. We use SRTGs to formulate priors on image-region hierarchies; and we use the associated center-surround algorithm to automatically estimate the ML parameters of such priors, compute the MAP estimate of the associated parse tree, and classify images based on their likelihood. We extend prior work [38], [39], [40], [54] in several important directions. The algorithms of [38], [39], [40], [54] are based on constructing probabilistic models for pixel intensities and grouping image pixels into non-overlapping rectangular regions. The models and algorithms proposed in this paper are more flexible, in that they are able to handle feature images where the features are extracted from possibly overlapping regions of arbitrary shapes. This results both in more modeling choices and in improved computational efficiency. In addition, there are important theoretical distinctions between the construction in [54] and the framework introduced in this paper. While the model of [54] hard-wires image regions to grammar symbols, we introduce a more flexible framework which is more similar to the traditional 1D PCFG models. In our current framework, the problem of uniquely associating a tree with an image is not straightforward. To solve this problem, we introduce in Section III a number of theoretical tools which insure that each tree has at most one associated image, much like in the standard 1D PCFGs. In addition, our methods of both [54] and the present paper extend PCFGs to continuous observations, much like the work in [27] extended HMMs.

In order to show the potential value of our methods, we apply them to the special case of classification of natural images and demonstrate a significant increase in average classification accuracy over baseline non-hierarchical models. We also show that substantial information is contained in the MAP image parses extracted by our methods.

The remainder of this paper is organized as follows.

Section II presents a nontechnical overview of our method. Section III presents the technical details of our method. Section IV presents experiments that illustrate our implementation classifying images. Section VI concludes with a summary of the contributions of this paper.

II. Overview

Context-free grammars (CFGs) can be used to impose hierarchical structure on sentences, i.e. 1D word strings. For example, the CFG in figure 1(a) imposes the parse in figure 1(b) on the sentence *The woman loves the man*. Nodes of such a parse tree correspond to *constituents*, where a constituent is a contiguous substring. Internal nodes might correspond to substrings that contain multiple words, while leaf nodes correspond to singleton constituents consisting of a single word.

For this 1D problem, the parse tree unambiguously specifies an ordering of the words from left to right. This implies that there is a unique mapping from the nodes of the parse tree to constituents of the sentence. We call this mapping the *constituency function*. Moreover, the unambiguous ordering of the nodes implies that each parse tree produces a unique sentence. This property allows a probability distribution over parse trees to induce a probability distribution over strings. We elaborate on this in section III. When a constituency function exists which maps the nodes of the parse tree to parts of a sentence (or, later, image), we say that the parse tree is a *parse* of the sentence (respectively, image).

The fundamental difficulty in extending CFGs to 2D lies in the unambiguous specification of the constituency functions so that each parse tree produces a unique image. To see that this property does not automatically hold in 2D, consider the example in figure 2. Figure 2(b) illustrates a parse tree generated by the CFG in figure 2(a). Figures 2(c) and 2(e) illustrate two different possible images, O and O' , parsed by this parse tree, with corresponding constituency functions illustrated in figures 2(d) and 2(f), respectively. Both images O and O' have the same set of regions, 1, 2, and 3. Note that the regions in the two images are not assumed to form a regular 2D lattice.

In this paper, we introduce a new class of grammars called *spatial tree grammars* (STGs) that ensures that no two distinct images have the same parse. STGs augment each branching production in a CFG with a *production class*. An STG generates parse trees whose branching nodes are *tagged* with production classes. The tags constrain the allowed relationships between the constituents that are associated with parse-tree nodes and those that are associated with their children. The list of all valid partitions of a constituent tagged with a specific production class into two subconstituents, is specified by a data structure called a *constituent hierarchy*. We formally define this data structure in definition 1 of section III-C.2. Figures 2(g) and 2(h) illustrate how production classes are used. In this example, we use two production classes: “h” and “v”. Intuitively,

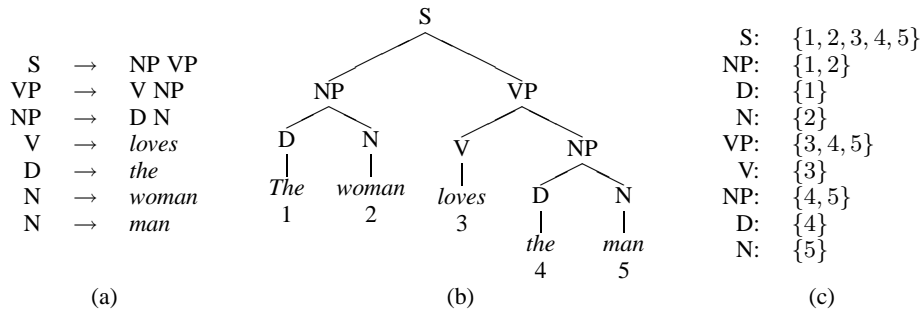


Fig. 1. The CFG (a) imposes the parse (b) on the sentence *The woman loves the man*. Note that the parse tree unambiguously specifies an ordering of the words from left to right which leads to the unique constituency function (c).




Fig. 2. (a) A CFG. (b) A parse tree generated by (a). (c) An image O with regions 1, 2, and 3, and region intensities $O(1) = u$, $O(2) = u$, and $O(3) = w$. (d) A constituency function that maps the nodes of (b) to the constituents of (c). (e) Another image, O' , with the same regions 1, 2, and 3 as image O , but with different intensities: $O'(1) = u$, $O'(2) = w$, and $O'(3) = u$. (f) A constituency function that maps the nodes of (b) to the constituents of (e). (g) An STG derived from (a) by augmenting the branching productions with production classes. (h) A parse tree generated by (g). This parse tree is the same as (b) except that the branching nodes have been tagged with production classes.

production class “h” (horizontal split) constrains a parse-tree node to be associated with a constituent that is split into upper and lower subconstituents that are associated with the node’s children, while production class “v” (vertical split) constrains a parse-tree node to be associated with a constituent that is split into left and right subconstituents that are associated with the node’s children. Figure 2(g) illustrates an STG derived from the CFG in figure 2(a) by augmenting the branching productions with these two production classes. Figure 2(h) illustrates a parse tree generated by this STG. This new parse tree is the same as that from figure 2(b) except that the branching nodes have been tagged with production classes. However, this parse tree does not parse the image in figure 2(e). Figure 2(f) is no longer a valid constituency function because the constituent $\{1, 3\}$ does not lie above the constituent $\{2\}$,

as is required by the tag “h” on the parse-tree node ρ_1 . It can be shown that, in fact, figure 2(c) is the only image parsed by figure 2(h).

Note that, while the example in figure 2 contains only rectangular constituents and partitions image constituents only horizontally or vertically, our current method, described in section III, supports nonrectangular constituents and nonhorizontal, nonvertical constituent partitionings as well. In section III-D, we provide specific conditions that ensure that no two distinct images have the same parse, and in section IV we demonstrate a useful application of this more general case.

In the case of CFGs, one can interpret a branching production like $S \rightarrow NP VP$ as saying that one can construct an S by concatenating an NP with a VP . CFGs are called such because they are *context free*. One can

concatenate *any* NP with *any* VP to construct an S. It has been shown that context-free languages can be parsed in polynomial time using the CKY algorithm [21], [56].

In the special case of STGs with rectangular constituents considered in this overview, one can interpret a branching production like $U \xrightarrow{\text{“h”}} V Z$ as concatenating matrices vertically, and a branching production like $V \xrightarrow{\text{“v”}} X Y$ as concatenating matrices horizontally. However, unlike the case of CFGs, where it is possible to concatenate any two strings, it is only possible to vertically concatenate matrices of the same width and horizontally concatenate matrices of the same height. Analogous restrictions hold for the more general case of arbitrary constituents and constituent partitionings that we consider in section III. This means that STGs are *context sensitive*. However, it can be easily shown that STGs are a special class of context-sensitive grammars—i.e. that there exist context-sensitive languages which cannot be generated by an STG. It moreover is shown below that the parsing algorithm for STGs is polynomial in the number of constituents. One way to ensure that our methods use algorithms with polynomial complexity is therefore to limit the number of possible constituents to be polynomial in the size of the input. This can be done using methods discussed in section IV-A.

One may induce a probability distribution over the parse trees generated by a CFG by imposing probabilities on the productions. Such an augmented CFG is called a probabilistic context-free grammar (PCFG). In a similar fashion, one may induce a probability distribution over the parse trees generated by an STG by imposing probabilities on the productions. We call such an augmented STG a *spatial random tree grammar* (SRTG).

III. The Technical Details of our Method

In this section, we formally introduce the concepts of *feature images* and *parse trees* and develop a framework for associating parse trees with a feature image, given its spatial organization. We construct a probability distribution over parse trees and show that our framework allows this probability distribution to induce a probability distribution on the feature images given their spatial organization.

A. Feature Images

We start with an image or any other multidimensional data. From this data, we determine a set of *locations* $D = \{v_1, \dots, v_{|D|}\}$, and, at each of these locations v_n , we extract a *feature vector* $O(v_n) \in \mathbb{R}^q$. We define the *feature image*, $\mathbf{O} \in \mathbb{R}^{q|D|}$, to be the vector consisting of all these feature vectors: $\mathbf{O} \triangleq (O(v_1), \dots, O(v_{|D|}))$.

For example, in our experiments reported in section IV, each location corresponds to an image region obtained through a segmentation algorithm, and the feature vector $O(v_n)$ for each location v_n is formed by extracting information from the corresponding region. An illustration of




Fig. 3. Illustration of a feature image defined on domain $D = \{v_1, v_2, v_3, v_4, v_5\}$ of size $|D| = 5$, for feature dimension $q = 1$. Here, the locations v_1, \dots, v_5 are the centroids of five constant-intensity regions of an image, and the feature $O(v_n)$ is the intensity at v_n .

this procedure, for a feature image with five regions ($|D| = 5$) and one feature dimension ($q = 1$), is shown in figure 3. The preprocessing segmentation step has the advantage of reducing the number of nodes to be processed and therefore the total computation. However, it also imposes practical limits on how much local information can be passed to the SRTG model.

B. Parse Trees

Let \mathcal{I} be a finite set of *production classes*, let \mathcal{J} be a finite set of *nonterminals*, and let the set \mathbb{R}^q be the set of *terminals*. We use i to denote production classes, u to denote terminals, and j, k , and l to denote nonterminals. We define a *parse tree* T to be a finite ordered tree where

- leaf nodes (*terminal nodes*) are *labeled* with terminals,
- nonleaf nodes (*nonterminal nodes*) are *labeled* with nonterminals,
- nonterminal nodes have either one child (*nonbranching nonterminal nodes*) or two children (*branching nonterminal nodes*),
- branching nonterminal nodes are *tagged* with production classes,
- the child of every nonbranching nonterminal node is a terminal node, and
- the children of every branching nonterminal node are nonterminal nodes.

Figure 2(h) shows an example of a parse tree with $\mathcal{I} = \{\text{“h”}, \text{“v”}\}$, $\mathcal{J} = \{j, k, l\}$, branching nonterminal nodes ρ_1, ρ_2 , nonbranching nonterminal nodes ρ_3, ρ_4, ρ_5 , and terminal nodes ρ_6, ρ_7, ρ_8 .

C. Constituents, Constituent Hierarchies, and Constituency Functions for SRTGs

1) *Motivation*: With PCFGs, one constructs a probability distribution over all trees generated by the PCFG and defines the probability¹ of each string to be the sum of probabilities of all its parses. This method of constructing probabilities is only correct if any two distinct strings have

¹If the terminals are continuous-valued, as in section III-D, then the probability *density* of each string is the sum of probability densities of all its parses.

disjoint sets of parses. This is ensured by the fact that, for PCFGs, two distinct strings cannot have the same parse. The reason that two distinct strings cannot have the same parse is because the unique string corresponding to any parse tree is simply obtained by ordering the leaves of the tree from left to right.

However, a potential problem with SRTGs is that, if improperly designed, a single parse tree can parse two distinct images. This problem can result from the fact that there is no unique natural ordering of the elements of a domain D . This can make it impossible for a probability distribution over trees to induce a probability distribution over feature images in a manner similar to PCFGs. This difficulty is avoided by imposing organization on the elements of D . Specifically, we introduce the novel concepts of *constituents* of D and *constituent hierarchies* which impose structure on the ways in which locations in a domain D can be combined. The following section introduces precise definitions of these concepts which ensure that a probability distribution on the feature images can be properly defined, and moreover allow for computationally tractable estimation algorithms.

2) *Constituents and Constituent Hierarchies*: We designate some subsets of a domain D to be *constituents*, and use \mathcal{C} to denote the set of all constituents of D . As we stated earlier, not all subsets $V \subset D$ will be constituents, so \mathcal{C} will generally be a proper subset of 2^D , the power set of D . Furthermore, for each production class $i \in \mathcal{I}$ and each constituent $V \in \mathcal{C}$, we will specify all valid two-way partitions of V into its subconstituents under production class i . To do this, we define the set $\mathcal{L}(i, V)$ of all *left i -subconstituents* of V , with the interpretation that a partition of V into V' and $V \setminus V'$ under the production class i is only valid if $V' \in \mathcal{L}(i, V)$. We call the quadruple $\mathcal{H} = \langle D, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$ a *constituent hierarchy*. We now formally define the concepts of constituents and constituent hierarchy.

Definition 1. Let D be a finite domain, and \mathcal{I} a finite set of production classes. Let \mathcal{C} be a set of subsets of D , and let \mathcal{L} be a function from $\mathcal{I} \times \mathcal{C}$ to $2^{\mathcal{C}}$. The quadruple $\mathcal{H} \triangleq \langle D, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$ is called a *constituent hierarchy* if it satisfies the four conditions C1–C4 given below. In this case, a subset $V \subset D$ is called a *constituent of D* iff $V \in \mathcal{C}$. A set V' is called a *subconstituent of V* if $V' \in \mathcal{L}(i, V)$ (in this case V' is called a *left i -subconstituent of V*) or if $V \setminus V' \in \mathcal{L}(i, V)$ (in this case V' is called a *right i -subconstituent of V*).

(C1) *The entire domain D is a constituent:*
 $D \in \mathcal{C}$

(C2) *Every subconstituent of every constituent V is a proper, nonempty subset of V :*
 $(\forall i \in \mathcal{I}, V \in \mathcal{C}, V' \in \mathcal{L}(i, V)) \Rightarrow$
 $(V' \subset V \text{ and } V' \neq \emptyset \text{ and } V' \neq V)$

(C3) *Every nonsingleton constituent can be partitioned in at least one way into two subconstituents:*
 $(\forall V \in \mathcal{C} \text{ s.t. } |V| > 1) \Rightarrow (\exists i \in \mathcal{I} \text{ s.t. } \mathcal{L}(i, V) \neq \emptyset)$

(C4) *Every constituent has at most one left i -subconstituent of given cardinality:*
 $(\forall i \in \mathcal{I}, V \in \mathcal{C}, V', V'' \in \mathcal{L}(i, V) \text{ s.t. } |V'| = |V''|) \Rightarrow (V' = V'')$

Given a D and an \mathcal{I} , it is typically easy to find such \mathcal{C} and \mathcal{L} that the quadruple $\mathcal{H} = \langle D, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$ meets C1-C3. Condition C4 is not as straightforward; however, we now describe a simple algorithm to modify any \mathcal{H} satisfying C1-C3 to obtain a valid constituent hierarchy. Specifically, suppose that you have a quadruple $\mathcal{H} = \langle D, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$ that meets C1-C3 but not C4. The following algorithm constructs a quadruple $\mathcal{H}' = \langle D, \mathcal{I}, \mathcal{C}', \mathcal{L}' \rangle$ that meets all of the conditions.

Algorithm 1. Let \prec be a total ordering on constituents. Initialize \mathcal{L}' to be the same as \mathcal{L} . Terminate, if there are no violations of C4. Let V be the first constituent, according to \prec , that violates C4. That means that there exists a production class $i \in \mathcal{I}$ and two constituents $V_1, V_2 \in \mathcal{L}'(i, V)$ such that $|V_1| = |V_2|$. Without loss of generality, assume $V_1 \prec V_2$. Remove V_2 from $\mathcal{L}'(i, V)$. Repeat this process until there are no violations of C4. \square

Section IV-A shows how to construct \mathcal{H} for a certain class of domains D , resulting in the number of constituents which is polynomial in $|D|$, and leading to polynomial-time inference algorithms.

3) *Constituency Functions*: To associate parse trees with feature images, we introduce the concept of constituency function. Suppose that $\mathcal{H} = \langle D, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$ is a constituent hierarchy, T is a parse tree generated using the set of production classes \mathcal{I} , and $\mathbf{O} \in \mathbb{R}^{q|D|}$ is a feature image. Consider a function F which maps the nodes of T to the constituents of D and which has the following properties.

- For the root node ρ , $F(\rho) = D$.
- For every terminal node ρ , labeled u , $F(\rho) = \{u\}$, where $O(v) = u$.
- For every nonbranching nonterminal node ρ , with child ρ_1 , $F(\rho) = F(\rho_1)$.
- For every branching nonterminal node ρ , tagged i , with left and right children ρ_1 and ρ_2 , $F(\rho_1) \in \mathcal{L}(i, F(\rho))$ and $F(\rho_2) = F(\rho) \setminus F(\rho_1)$.

In this case, we say that F is an \mathcal{H} -constituency function for T and \mathbf{O} ; and we say that T is an \mathcal{H} -parse of \mathbf{O} . We show that conditions C1-C4 guarantee that any \mathcal{H} -parse T of $\mathbf{O} \in \mathbb{R}^{q|D|}$ cannot also be an \mathcal{H} -parse of another feature image $\mathbf{O}' \in \mathbb{R}^{q|D|}$, and that moreover the \mathcal{H} -constituency function for T and \mathbf{O} is unique.

Theorem 1. Let $\mathcal{H} = \langle D, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$ be a constituent hierarchy. Let T be an \mathcal{H} -parse of $\mathbf{O} \in \mathbb{R}^{q|D|}$, and let F be an \mathcal{H} -constituency function for T and \mathbf{O} . Let T' also be an

\mathcal{H} -parse of $\mathbf{O}' \in \mathbb{R}^{q|D|}$, and let F' be an \mathcal{H} -constituency function for T and \mathbf{O}' . Then $F = F'$ and $\mathbf{O} = \mathbf{O}'$.

PROOF: See appendix. \square

D. Probabilistic Modeling with SRTGs

Using the framework described in the previous subsection, we can define a probability distribution over the feature images. We do this according to the following plan:

- Construct a probability distribution over trees, similar to PCFGs.
- Construct a quadruple \mathcal{H} which satisfies conditions C1-C3.
- If necessary, use algorithm 1 to enforce C4 and form a constituent hierarchy.
- Construct the probability (or probability density) of a feature image $\mathbf{O} \in \mathbb{R}^{q|D|}$ by summing the probabilities (or probability densities) of all \mathcal{H} -parses of \mathbf{O} .

We first specify the method for generating random parse trees. With probability r_k generate a root nonterminal node with label k . Next, repeatedly select a nonterminal node, ρ , that does not yet have children and denote its label by j . With conditional probability t_j , ρ will be a nonbranching node; and with conditional probability $1 - t_j$, it will be a branching node. If ρ is a nonbranching node, then its child must be a terminal node. In this case, the child's label, u , is chosen to be a conditionally normal random vector with mean μ_j and covariance Σ_j . If ρ is a branching nonterminal node, then it must have a production class tag which we call i , and left and right children whose labels we call k and l . In this case, choose $\langle i, k, l \rangle$ with conditional probability p_{jikl} . The distribution r and the conditional distributions p_j obey the following normalization constraints:

$$\sum_{j \in \mathcal{J}} r_j = 1 \quad (\text{N1})$$

$$(\forall j \in \mathcal{J}) \sum_{i \in \mathcal{I}, k, l \in \mathcal{J}} p_{jikl} = 1 \quad (\text{N2})$$

Note that while we take u to be conditionally normal, it is possible to use any conditional distribution, discrete or continuous or hybrid, so long as one replaces the reestimation procedures for μ_j and Σ_j that are presented in figure 8 with suitable variants for the alternate distribution. We take a *spatial random tree grammar* (SRTG) G to be an octuple $\langle \mathcal{I}, \mathcal{J}, \mathbb{R}^q, r, t, p, \mu, \Sigma \rangle$, where \mathcal{I} is the set of production classes, \mathcal{J} is the set of nonterminals, and \mathbb{R}^q is the set of terminals. The above process defines $P(T|G)$, the probability density over all trees generated by such a process, given an SRTG G .

Whenever we discuss an SRTG G and a constituent hierarchy \mathcal{H} , we assume that both use the same set \mathcal{I} of production classes. In this case, we define, for any feature

image $\mathbf{O} \in \mathbb{R}^{q|D|}$,

$$P(\mathbf{O}|G, \mathcal{H}) \triangleq \sum_{T \text{ is an } \mathcal{H}\text{-parse of } \mathbf{O}} P(T|G)$$

Note that, strictly speaking, $P(\mathbf{O}|G, \mathcal{H})$ is not a probability density since not every tree generated by the above stochastic process is necessarily an \mathcal{H} -parse of a feature image in $\mathbb{R}^{q|D|}$.² Therefore, it may happen that $\int_{\mathbb{R}^{q|D|}} P(\mathbf{O}|G, \mathcal{H}) d\mathbf{O} < 1$. A true probability density P' can then be obtained as follows:

$$P'(\mathbf{O}|G, \mathcal{H}) = \frac{P(\mathbf{O}|G, \mathcal{H})}{\int_{\mathbb{R}^{q|D|}} P(\mathbf{O}'|G, \mathcal{H}) d\mathbf{O}'} \quad (\text{N3})$$

The probability density $P'(\mathbf{O}|G, \mathcal{H})$ is, in general, very difficult to compute because of the denominator in equation N3. Fortunately, we will not need to compute it in any of our algorithms.

Note that a single SRTG G can be used to define $P(\mathbf{O}|G, \mathcal{H})$ for *any* arbitrary constituent hierarchy \mathcal{H} , as long as G and \mathcal{H} use the same set \mathcal{I} of production classes. This is an important property which is exploited in our algorithms.

E. Inference Algorithms

1) *Likelihood Calculation, Center Variables, and Surround Variables:* Given an SRTG G , a constituent hierarchy \mathcal{H} , and a feature image $\mathbf{O} \in \mathbb{R}^{q|D|}$, the likelihood $P(\mathbf{O}|G, \mathcal{H})$ can be computed using equations 1 and 2 of figure 4. For any nonterminal j and the domain D , we let the *center variable* $c(j, D)$ be the conditional probability density of the feature image \mathbf{O} given that the root label of its \mathcal{H} -parse is j : $c(j, D) = \sum_{T \text{ is an } \mathcal{H}\text{-parse of } \mathbf{O} \text{ with root label } j} P(T|G)$. The cen-

ter variable $c(j, V)$ for any other constituent V of D is defined similarly. Equation 2 of figure 4 relates the center variable $c(j, V)$ of any nonsingleton constituent V with the center variables of its subconstituents. Specifically, equation 2 takes into account all possible partitions of V into a left i -subconstituent V_1 and a right i -subconstituent $V \setminus V_1$, and sums over all such partitions, over all production classes i , and over all possible labels k and l of V_1 and $V \setminus V_1$. Equation 1 is the formula for the center variable of a singleton constituent. In this equation, $N(u|\mu, \Sigma)$ denotes the normal density. The center variables for all constituents and all nonterminals j can therefore be computed by first using equation 1 for each singleton constituent, and then using equation 2 for every nonsingleton constituent. The feature image likelihood is then obtained as follows:

$$P(\mathbf{O}|G, \mathcal{H}) = \prod_{j \in \mathcal{J}} r_j c(j, D).$$

²We expand upon this observation in section V where we discuss the relationships between our framework and traditional PCFGs.

$$\begin{aligned}
c(j, \{v\}) &= t_j N(O(v) | \mu_j, \Sigma_j) & (1) \\
c(j, V) &= \sum_{i,k,l} t_j \sum_{V_1 \in \mathcal{L}(i,V)} (1-t_j) p_{jikl} c(k, V_1) c(l, V \setminus V_1) & |V| > 1 \quad (2) \\
s(j, D) &= r_j & (3) \\
s(j, V) &= \sum_{i,k,l} r_j \sum_{V_1: V \in \mathcal{L}(i, V_1)} (1-t_k) p_{kijl} s(k, V_1) c(l, V_1 \setminus V) + & V \neq D \quad (4) \\
&\quad \sum_{i,k,l} \sum_{V_1: V_1 \setminus V \in \mathcal{L}(i, V_1)} (1-t_k) p_{kilj} s(k, V_1) c(l, V_1 \setminus V)
\end{aligned}$$

Fig. 4. The center and surround recursions for a single feature image \mathbf{O} . (1–2) The center recursions. (3–4) The surround recursions. Each summation over i, k, l is performed over the entire ranges of these variables, i.e. over $i \in \mathcal{I}$, $k, l \in \mathcal{J}$. Recall that $\mathcal{L}(i, V)$ is the set of left i -subconstituents of V , and that t_j is the probability that a node labeled j is nonbranching. We use $N(u|\mu, \Sigma)$ to denote the normal density.

The *surround variable* $s(j, V)$ for every nonterminal j and every constituent V is computed via the surround recursions shown in equations 3 and 4 of figure 4. Starting with the entire domain D , the surround variables for every constituent are recursively computed in terms of the surround variables of larger constituents and the center variables. These recursions, together with the center recursions, are used in the parameter estimation algorithm which is described below. Note that the center variables must be precomputed before calculating the surround variables. The center and surround recursions generalize the inside and outside recursions [1], [26].

2) *Estimation of the MAP Parse:* Given an SRTG G , a constituent hierarchy \mathcal{H} , and a feature image $\mathbf{O} \in \mathbb{R}^{q|D|}$, the MAP \mathcal{H} -parse of \mathbf{O} is $\operatorname{argmax}_{T \text{ is an } \mathcal{H}\text{-parse of } \mathbf{O}} P(T|G)$. It can be found using equations 5 through 7 of figure 7 in appendix. Equation 6 and equation 7 recursively find the most likely partition of a constituent V labeled j and tagged i into two subconstituents V_1 and $V \setminus V_1$, labeled k and l . Once the most likely quadruple $\langle i, k, l, V_1 \rangle(j, V)$ for each j and each V is determined and stored, the MAP parse T for \mathbf{O} can be constructed as follows:

- 1) Let ρ be the root node of T . Label ρ with $\operatorname{argmax}_{j \in \mathcal{J}} \hat{c}(j, D)$ and let $F(\rho) = D$.
- 2) For each node ρ in T , labeled j , where $|F(\rho)| > 1$, let $\langle i, k, l, V_1 \rangle = \langle i, k, l, V_1 \rangle(j, F(\rho))$, tag ρ with i , add left and right children ρ_1 and ρ_2 to ρ , label ρ_1 with k and ρ_2 with l , and let $F(\rho_1) = V_1$ and $F(\rho_2) = V \setminus V_1$.
- 3) For each node ρ in T , labeled j , where $F(\rho) = \{v\}$, add a single child ρ_1 to ρ , label ρ_1 with $O(v)$, and let $F(\rho_1) = \{v\}$.

This algorithm, based on equations 5 through 7 of figure 7 in appendix, generalizes the Viterbi [52] algorithm.

3) *Parameter Estimation:* Suppose we are given M training feature images, $\mathbf{O}_1 \in \mathbb{R}^{q|D_1|}, \dots, \mathbf{O}_M \in \mathbb{R}^{q|D_M|}$, and further suppose that each feature image has a corresponding constituent hierarchy, $\mathcal{H}_m = \langle D_m, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$,

for $m = 1, \dots, M$. We seek $\operatorname{argmax}_G \prod_{m=1}^M P(\mathbf{O}_m | G, \mathcal{H}_m)$.

We address this problem via the algorithm in equations 8 through 14 of figure 8 in appendix which use the center and surround recursions given in equations 1 through 4 of figure 4.

The center and surround recursions, together with equations 8 through 14 of figure 8, constitute the EM algorithm [2], [16]. Equations 8 and 9 constitute the E step while equations 10 through 14 constitute the M step, i.e. the reestimation formulas for r_j , t_j , p_{jikl} , μ_j and Σ_j . We collectively refer to the equations in figures 4 through 8 as the *center-surround algorithm*.

The standard implementation of the algorithm in figures 4 through 8 memoizes the center and surround recursions. This leads to algorithms for likelihood calculation, MAP estimation, and parameter reestimation that are polynomial in $|D|$, so long as the number of c and s values to be memoized is polynomial in $|D|$ and the summations and maximizations range over sets of indices whose size is polynomial in $|D|$. These are both true, so long as the number of constituents is polynomial in $|D|$.

IV. Experiments

A. Constituent Hierarchies: An Example

In our experiments, we obtain the domain D by segmenting an image. In this case, every location $v \in D$ corresponds to a region, i.e. a set of pixels in an image. We let $\bar{x}(v)$ and $\bar{y}(v)$ denote the mean of x and y coordinates of the pixels in region v respectively. We now present a method for constructing a constituent hierarchy \mathcal{H} for any domain D obtained in this way.

To define the set \mathcal{C} of constituents of D , we take a nonempty subset $V \subset D$ to be a constituent iff there exist four numbers x_1, x_2, y_1 , and y_2 such that for all $v \in V$, $x_1 \leq \bar{x}(v) \leq x_2$ and $y_1 \leq \bar{y}(v) \leq y_2$. Note that the number of constituents is $O(|D|^4)$. This results in time complexity $O(|\mathcal{I}| \cdot |\mathcal{J}|^3 \cdot |D|^5)$ for the likelihood calculation, one step of EM, and MAP tree estimation.

We take the set \mathcal{I} of production classes to be {"row", "column"}. In order to specify a constituent hierarchy, it remains to specify the sets $\mathcal{L}(i, V)$ of left i -subconstituents of V , for both production classes i and for every constituent V . Suppose V is a constituent. We take a proper nonempty subset V_1 of V to be a left "row"-subconstituent of V iff there exists a number x such that $\bar{x}(v) < x$ for all $v \in V_1$ and $\bar{x}(v) > x$ for all $v \in V \setminus V_1$. We take a proper nonempty subset V_1 of V to be a left "column"-subconstituent of V iff there exists a number y such that $\bar{y}(v) < y$ for all $v \in V_1$ and $\bar{y}(v) > y$ for all $v \in V \setminus V_1$. Formulating $\mathcal{H} = \langle D, \mathcal{I}, \mathcal{C}, \mathcal{L} \rangle$ in this fashion




Fig. 5. Illustration of a domain D consisting of five regions. Their centroids are marked with black dots. According to the constituent hierarchy construction in Section IV-A, the following are valid constituents: $1 \cup 2$, $1 \cup 4$, $2 \cup 3$, $2 \cup 5$, $3 \cup 4$, $3 \cup 5$, $4 \cup 5$, $1 \cup 2 \cup 3$, $2 \cup 3 \cup 5$, $3 \cup 4 \cup 5$, $1 \cup 3 \cup 4$, $2 \cup 3 \cup 4$, $1 \cup 2 \cup 3 \cup 5$, $2 \cup 3 \cup 4 \cup 5$, $1 \cup 2 \cup 3 \cup 4$, and $1 \cup 2 \cup 3 \cup 4 \cup 5$.

clearly meet conditions C1–C4.

Figure 5 illustrates these definitions. For example, $1 \cup 2$ is a valid constituent since the centroids of regions 1 and 2 are adjacent vertically and $3 \cup 4$ is a valid constituent since the centroids of regions 3 and 4 are adjacent horizontally. For $1 \cup 2$, the only left “row”-subconstituent is 1 and the only left “column”-subconstituent is 1 as well. For $3 \cup 4$, the unique left “row”-subconstituent is 4 and the unique left “column”-subconstituent is 3. Note that, despite the fact that regions 1 and 3 are neighbors (i.e. share a common boundary), their union is not a valid constituent since their centroids are separated by region 2’s centroid in the vertical direction and by region 4’s centroid in the horizontal direction. Thus, even though regions 1 and 4 are not neighbors, their union is a valid constituent.

B. Implementation of the Parameter Estimation Algorithm

When using SRTGs, it is sometimes convenient to designate a nonterminal j to only label nonbranching nonterminal nodes, i.e. to have $t_j = 1$. This can be accomplished during training by initializing t_j to 1. Indeed, it follows from the equations in figure 8 that if $t_j = 1$, t_j will never change by reestimation. In this case, we refer to j as a *nonbranching nonterminal*. Similarly, we can initialize $t_j = 0$, and obtain a nonterminal j which can only label branching nonterminal nodes. In this case, we refer to j as a *branching nonterminal*. It is also sometimes useful to have the possible labels of the root node comprise only a small subset of \mathcal{J} . If the root node can be labeled with a nonterminal j , we refer to j as a *root nonterminal*. Any nonterminal j can be prevented from being a root nonterminal by initializing $r_j = 0$ during training.

C. Experiments with a House-Car Dataset

1) *Data Collection*: To evaluate our methods, we first applied them to the task of distinguishing images of houses from images of cars. We took 100 photographs each of houses and cars in the University Farms subdivision of

West Lafayette IN. The houses were photographed from the front, the cars were photographed from the drivers side, and the target house or car was centered and sized to fill the field of view. The photographs were taken with an Intel Pocket Digital PC Camera at 640×480 resolution.³ The original images were converted to 256-level greyscale and subsampled to 160×120 without filtering. All subsequent processing was performed on the subsampled greyscale images. Some sample images from this data set are illustrated in figure 6(a). Note that the house images often have (partially occluded) cars parked in front and the car images often have (partially occluded) houses in the background. Also note that the images often have other occluding objects such as trees.

We segmented the images with Ratio Cut [53].⁴ All 200 images in our data set contained exactly 10 regions after segmentation. For each image, we define the domain $D = \{v_1, \dots, v_{10}\}$ where v_1, \dots, v_{10} are the ten extracted regions. The results of segmenting the images from figure 6(a) are shown in figure 6(b).

We constructed a 4-element feature vector as follows. Let $\Lambda(v)$ denote the covariance matrix of the coordinates of the pixels in region v . We associated each region v with the following features:

- the area of v , i.e. the number of pixels in v ,
- the average intensity of the pixels in v ,
- the orientation of the principle eigenvector of $\Lambda(v)$, and
- the ratio of the smallest to the largest eigenvalue of $\Lambda(v)$.

We normalized each component of the feature vector to be in $[0, 1]$ by dividing the area by the total image area, the average pixel intensity by the maximal possible pixel intensity, and the principle eigenvector orientation (whose value was shifted to lie in $[0, \pi)$) by π . The eigenvalue ratio is already in $[0, 1]$ since it is a ratio of the smaller to the larger eigenvalue.

2) *Classifiers*: We use two different methods for constructing a classifier with the algorithm of figures 4 through 8:

method A Train a distinct SRTG for each class on training feature images from that class and classify a test feature image \mathbf{O} by asking which SRTG G maximizes $P(\mathbf{O}|G, \mathcal{H})$.

method B Train a single SRTG G with distinct root nonterminals for each of the classes on all of

³This data set, as well as all of the source code and scripts used to perform the experiments reported in this section, are available from <ftp://ftp.ecn.purdue.edu/qobi/pami2007.tgz>.

⁴We used the iterated region-based segmentation technique with the blocking heuristic. We used a linear decreasing function g , a block size of 32×32 , and a homogeneity threshold $H_T = 720$ for the first iteration. We post-processed the results of the first iteration using the area-merging technique with $A_T = 100$. We then performed a second iteration with $H_T = 700$ and post-processed the results of this iteration by repeatedly merging the two adjacent regions with the largest cut ratio until the number of regions was 10.

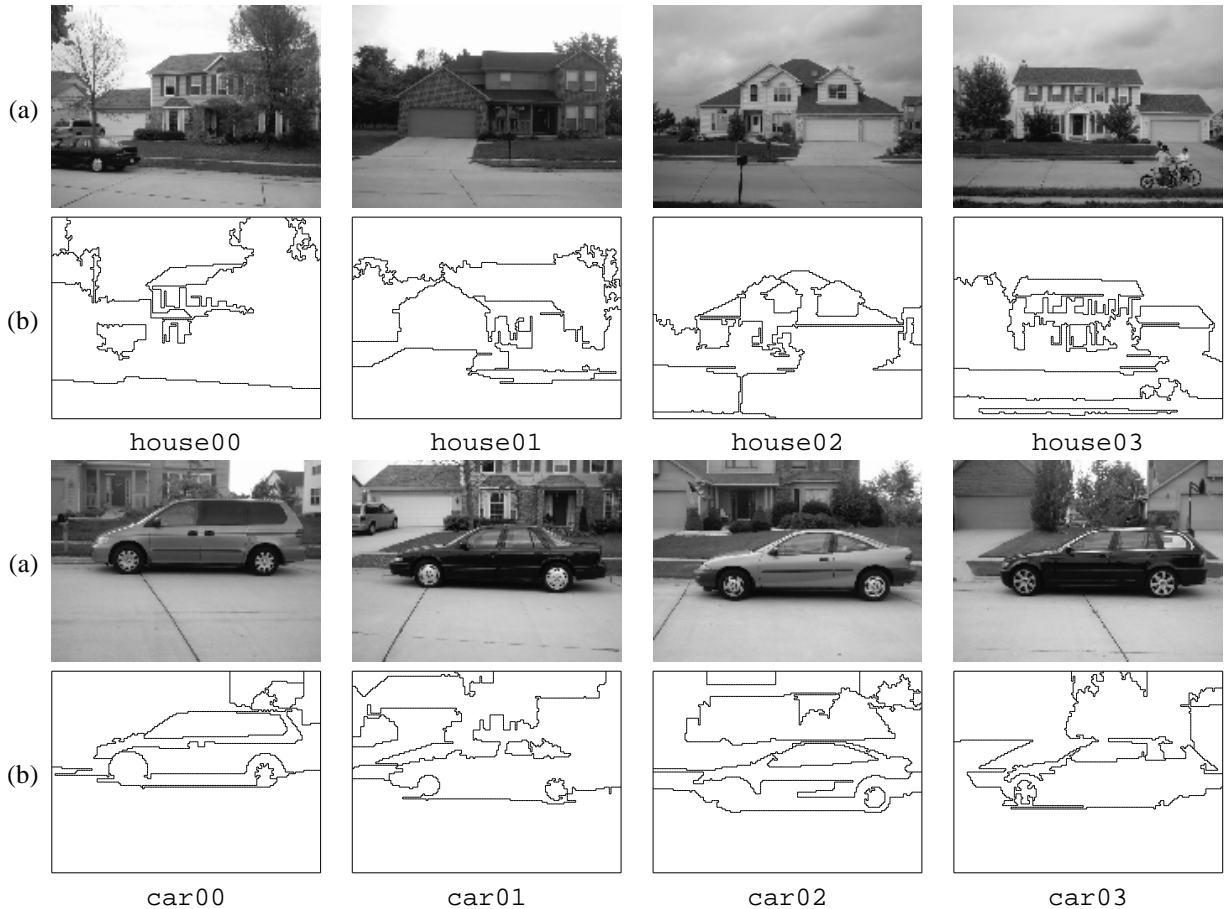


Fig. 6. (a) Sample images from our house-car data set. (b) The results of segmenting the images from (a) with Ratio Cut.

the training feature images, where those feature images are labeled with their class, and classify a test feature image \mathbf{O} by finding the MAP parse of \mathbf{O} and examining the root-node label.

When training with method B, r_j is taken to be one when j corresponds to the root nonterminal for the feature image being trained on and is taken to be zero otherwise. When classifying with method B, the single SRTG constructed jointly for all classes weights the r_j equally among the root nonterminals j for all classes. Method B allows a common vocabulary of terminals across different classes, reflected in the shared μ_j and Σ_j parameters, as well as a common vocabulary of nonterminals across different classes, reflected in the shared p_{jkl} parameters.⁵

⁵When reestimating r_j , t_j , and p_{jkl} , we clip them to zero when they are less than ϵ . Further, when reestimating t_j , we clip it to one when it is greater than $1 - \epsilon$. We renormalize the distributions r and p_j after such clipping to maintain the normalization equations N1 and N2. When estimating the parameters Σ_j for our baseline model, we perform an eigenvalue/eigenvector decomposition of Σ_j , clip the eigenvalues to ϵ_Σ when they are less than ϵ_Σ , and recompute Σ_j from the original eigenvectors and the clipped eigenvalues. In our experiments, we found that our algorithms are not excessively sensitive to the values of the parameters ϵ and ϵ_Σ . In particular, we found that the values $\epsilon = 10^{-5}$ and $\epsilon_\Sigma = 10^{-3}$, which we used in all the experiments in section IV, worked well.

3) *Experimental Results:* We conducted four experiments to compare the classification accuracy of SRTGs to that of a baseline. For each experiment, we performed a set of five round-robin training and classification runs. For each round-robin run, we partitioned our data set into a training set of eighty images for each class and a test set of twenty images for each class. The test sets partition the entire data set.

We used two different baselines, both of which were Gaussian mixture models. One, the *separate baseline*, trained separate mixtures of ten Gaussians, one on houses and one on cars. Each model contained distinct μ_j , Σ_j , and π_j parameters for each class. The other, the *joint baseline*, trained a single mixture of ten Gaussians on both houses and cars. Each model contained distinct π_j parameters for each class, but the models for both classes shared the same μ_j and Σ_j parameters.

Distinct separate and joint baseline models were trained for each round-robin run. The separate baseline model was trained in the traditional fashion. The joint baseline model was trained by first training a single set of μ_j , Σ_j , and π_j parameters on the combined set of house and car training images. The π_j mixture proportions were then discarded. New sets of π_j mixture proportions were then trained

| test set | SVM accuracy | |
|----------|-----------------|-------------|
| | on training set | on test set |
| 00–19 | 83.1% | 75.0% |
| 20–39 | 84.4% | 65.0% |
| 40–59 | 81.9% | 75.0% |
| 60–79 | 81.3% | 77.5% |
| 80–99 | 83.1% | 77.5% |
| mean | 82.8% | 74.0% |

TABLE II: THE RESULTS OF CLASSIFYING THE house-car DATA SET USING SUPPORT VECTOR MACHINES.

separately on just houses or just cars using the same μ_j and Σ_j parameters for both houses and cars. Images were classified with the baseline models by selecting the model that best fit that image.

To focus the comparison between SRTGs and the baseline on the advantages of the hierarchal structure provided by SRTGs, we used the same μ_j and Σ_j parameters for both the baseline models and the SRTGs in each round-robin run in each experiment. To train the SRTGs for each round-robin run in each experiment, we took the μ_j and Σ_j parameters for that round-robin run from the baseline model, discarded the π_j mixture proportions, and trained just the p_{jkl} parameters of an SRTG. All SRTGs were trained with twenty nonterminals, ten of which were constrained to be branching nonterminals and ten of which were constrained to be nonbranching nonterminals with the fixed μ_j and Σ_j parameters from the baseline model. To train each SRTG, we started reestimation with two different random values for p_{jkl} , performed 300 iterations of reestimation on each using equations 1–4 and 8–12, and then selected the one that yielded the highest likelihood on the training set. Each initial p contained 8000 nonzero entries. The number of nonzero entries in the final values of p is indicated in tables I through IV. Note that typically less than 2% of the entries are nonzero after training.

a) Experiment 1: In our first experiment, we compared the classification accuracy of SRTGs to the separate baseline models using method A from section IV-C.2. For each round-robin run, we trained both a separate baseline model and an SRTG for each class. We then classified both the training images and test images using both the baseline models and the SRTGs. The results of this classification are shown in table I. SRTGs consistently out-performed the baseline, with only 9 misclassifications on 200 test images, compared to 32 baseline misclassifications. It is also interesting to note that our training procedure drastically reduced the model order of SRTGs: while the initial number of nonzero p_{jkl} ’s was 8000, the final number was usually about 100.

To provide another benchmark for our experiments, we list in table II the classification performance of a

support vector machine⁶ (SVM) [51]. The SVM classification experiments used the same feature vectors as the baseline and SRTG models in table I. Comparing tables I and II, we see that SVMs performed substantially worse than both SRTGs and the Gaussian mixture baseline. In addition, SVMs performed substantially worse than SRTGs in experiments 2 and 3 described below.

b) Experiment 2: Our second experiment differed from experiment 1 in that it used the joint baseline models instead of the separate baseline models. We again compared the classification accuracy of SRTGs to the baseline models using method A from section IV-C.2. The results of this classification are shown in table III. Again, SRTGs consistently out-performed the baseline, this time misclassifying only 6 test images out of 200, compared to 54 baseline misclassifications.

c) Experiment 3: Experiments 1 and 2 used method A from section IV-C.2. Our third experiment differed from experiment 2 in that it used method B instead of method A. We again compared the classification accuracy of SRTGs to the joint baseline models. We used the same joint baseline model for each class and round-robin run as the one from the corresponding class and round-robin run from experiment 2. (This means that the baseline numbers in tables III and IV are identical.) For each round-robin run, we trained a single SRTG on both the houses and cars in the corresponding training set. We constrained each SRTG to contain two root branching nonterminals corresponding to the training labels of the two image classes. We then classified both the training images and the test images using both the baseline models and the SRTGs. The results of this classification are shown in table IV. SRTGs again significantly out-performed the baseline, with 6 misclassifications on 200 test images, compared to 54 baseline misclassifications.

d) Experiment 4: We evaluated the parses produced by SRTGs with a further experiment. For each round-robin run in this experiment, we used the same joint baseline model and SRTG as the corresponding round-robin run of experiment 3. We attempted to classify houses and cars using just the shape of their parses without any image data, feature vectors, parse-tree node tags, or parse-tree node labels. For each round-robin run, we computed the set of all best parses for each image in our data set with the SRTG trained on the training set for that round-robin run. Note that there can be—and often are—multiple distinct parses with the same maximal probability. We computed the set of all such parses for each image in our data set. We then took the training set in each round-robin run as the exemplars and classified the test set in each round-robin against these exemplars with a k -nearest-neighbor classifier using pivot distance as the distance metric. Our definition of pivot distance is given in appendix. Since

⁶We used the SVM software of <http://svmlight.joachims.org/>

| test set | nonzero p entries | | accuracy on training set | | | accuracy on test set | | |
|----------|---------------------|-----|--------------------------|--------|----------|----------------------|--------|----------|
| | house | car | baseline | SRTG | Δ | baseline | SRTG | Δ |
| 00–19 | 108 | 95 | 95.6% | 100.0% | 4.4% | 80.0% | 92.5% | 12.5% |
| 20–39 | 114 | 110 | 95.0% | 100.0% | 5.0% | 85.0% | 95.0% | 10.0% |
| 40–59 | 113 | 108 | 93.8% | 100.0% | 6.3% | 90.0% | 100.0% | 10.0% |
| 60–79 | 108 | 107 | 95.6% | 100.0% | 4.4% | 82.5% | 100.0% | 17.5% |
| 80–99 | 107 | 102 | 95.0% | 100.0% | 5.0% | 82.5% | 90.0% | 7.5% |
| mean | | | 95.0% | 100.0% | 5.0% | 84.0% | 95.5% | 11.5% |

TABLE I: THE RESULTS OF EXPERIMENT 1, CLASSIFYING THE house–car DATA SET, USING BOTH THE BASELINE MODELS AND SRTGS, FOR VARIOUS ROUND-ROBIN RUNS. THIS EXPERIMENT USED THE SEPARATE BASELINE MODELS AND METHOD A FROM SECTION IV-C.2.

| test set | nonzero p entries | | accuracy on training set | | | accuracy on test set | | |
|----------|---------------------|-----|--------------------------|--------|----------|----------------------|--------|----------|
| | house | car | baseline | SRTG | Δ | baseline | SRTG | Δ |
| 00–19 | 106 | 106 | 75.6% | 100.0% | 24.4% | 75.0% | 95.0% | 20.0% |
| 20–39 | 117 | 109 | 75.6% | 100.0% | 24.4% | 77.5% | 95.0% | 17.5% |
| 40–59 | 108 | 112 | 75.6% | 100.0% | 24.4% | 75.0% | 100.0% | 25.0% |
| 60–79 | 113 | 118 | 80.0% | 100.0% | 20.0% | 62.5% | 95.0% | 32.5% |
| 80–99 | 98 | 105 | 72.5% | 100.0% | 27.5% | 75.0% | 100.0% | 25.0% |
| mean | | | 75.9% | 100.0% | 24.1% | 73.0% | 97.0% | 24.0% |

TABLE III: THE RESULTS OF EXPERIMENT 2, CLASSIFYING THE house–car DATA SET, USING BOTH THE BASELINE MODELS AND SRTGS, FOR VARIOUS ROUND-ROBIN RUNS. THIS EXPERIMENT USED THE JOINT BASELINE MODELS AND METHOD A FROM SECTION IV-C.2.

| test set | nonzero p entries | accuracy on training set | | | accuracy on test set | | |
|----------|---------------------|--------------------------|--------|----------|----------------------|--------|----------|
| | | baseline | SRTG | Δ | baseline | SRTG | Δ |
| 00–19 | 158 | 75.6% | 99.4% | 23.8% | 75.0% | 92.5% | 17.5% |
| 20–39 | 158 | 75.6% | 100.0% | 24.4% | 77.5% | 97.5% | 20.0% |
| 40–59 | 153 | 75.6% | 98.1% | 22.5% | 75.0% | 95.0% | 20.0% |
| 60–79 | 161 | 80.0% | 100.0% | 20.0% | 62.5% | 100.0% | 37.5% |
| 80–99 | 141 | 72.5% | 98.8% | 26.3% | 75.0% | 100.0% | 25.0% |
| mean | | 75.9% | 99.3% | 23.4% | 73.0% | 97.0% | 24.0% |

TABLE IV: THE RESULTS OF EXPERIMENT 3, CLASSIFYING THE house–car DATA SET, USING BOTH THE BASELINE MODELS AND SRTGS, FOR VARIOUS ROUND-ROBIN RUNS. THIS EXPERIMENT USED THE JOINT BASELINE MODELS AND METHOD B FROM SECTION IV-C.2.

both exemplars and test images could have multiple best parses, we computed the pivot distance between all pairs consisting of an exemplar best parse and a test image best parse and selected the minimal pivot distance as the distance metric between that exemplar and that test image. Note that pivot distance measured only the shape of the parses and was not sensitive to any parse-tree node tags or labels. We repeated this for all odd values of $1 \leq k \leq 13$ and compared the classification accuracy against the baseline models. The results of this classification are shown in table V. Note that, on average, the SRTG-based method which used only the extracted hierarchal structure, out-performed the baseline which directly used the region feature vectors.

4) *Discussion of the Experiments:* The experiments reported in section IV-C.3 used the same μ_j and Σ_j parameters from the baseline models when constructing SRTGs

and did not reestimate those parameters. This was done to focus the experiments on measuring the added leverage afforded by hierarchal structure. SRTGs are strictly more powerful than the baseline models because the baseline models are equivalent to degenerate SRTGs with a single branching nonterminal. Our methods can obviously be used to reestimate the μ_j and Σ_j parameters to yield SRTGs with an even better fit to the training data.

Experiments 1 through 3 demonstrate the added leverage afforded by hierarchal structure. The baseline models and SRTGs in each round-robin run of each experiment differ only in that the baselines use nonhierarchal mixture proportions, reflected in the π_j parameters, while the SRTGs use hierarchal structure, reflected in the p_{jikl} parameters. In all three experiments, SRTGs consistently out-perform the corresponding baseline model, often significantly. This indicates that the presence of hierarchal structure signifi-

| | baseline | SRTG | | | | | | |
|-------|----------|---------|---------|---------|---------|---------|----------|----------|
| | | $k = 1$ | $k = 3$ | $k = 5$ | $k = 7$ | $k = 9$ | $k = 11$ | $k = 13$ |
| 00–19 | 75.0% | 72.5% | 77.5% | 75.0% | 72.5% | 75.0% | 82.5% | 80.0% |
| | Δ | -2.5% | 2.5% | 0.0% | -2.5% | 0.0% | 7.5% | 5.0% |
| 20–39 | 77.5% | 77.5% | 70.0% | 70.0% | 65.0% | 70.0% | 67.5% | 72.5% |
| | Δ | 0.0% | -7.5% | -7.5% | -12.5% | -7.5% | -10.0% | -5.0% |
| 40–59 | 75.0% | 82.5% | 85.0% | 85.0% | 82.5% | 82.5% | 82.5% | 82.5% |
| | Δ | 7.5% | 10.0% | 10.0% | 7.5% | 7.5% | 7.5% | 7.5% |
| 60–79 | 62.5% | 75.0% | 72.5% | 85.0% | 82.5% | 82.5% | 82.5% | 82.5% |
| | Δ | 12.5% | 10.0% | 22.5% | 20.0% | 20.0% | 20.0% | 20.0% |
| 80–99 | 75.0% | 82.5% | 75.0% | 82.5% | 72.5% | 80.0% | 80.0% | 75.0% |
| | Δ | 7.5% | 0.0% | 7.5% | -2.5% | 5.0% | 5.0% | 0.0% |
| mean | 73.0% | 78.0% | 76.0% | 79.5% | 75.0% | 78.0% | 79.0% | 78.5% |
| | Δ | 5.0% | 3.0% | 6.5% | 2.0% | 5.0% | 6.0% | 5.5% |

TABLE V: THE RESULTS OF EXPERIMENT 4, CLASSIFYING THE house-car DATA SET, USING BOTH THE BASELINE MODELS AND SRTGs, FOR VARIOUS ROUND-ROBIN RUNS AND VALUES OF k . FOR EACH ROUND-ROBIN RUN IN THIS EXPERIMENT, WE USED THE SAME JOINT BASELINE MODEL AND SRTG AS THE CORRESPONDING ROUND-ROBIN RUN OF EXPERIMENT 3. CLASSIFICATION WITH SRTGs WAS PERFORMED USING A k -NEAREST-NEIGHBOR CLASSIFIER BETWEEN THE ENTIRE TRAINING SET TAKEN AS EXEMPLARS AND EACH ELEMENT OF THE TEST SET. THE MINIMAL PIVOT DISTANCE BETWEEN ALL BEST PARSES OF AN EXEMPLAR AND ALL BEST PARSES OF A TEST IMAGE WAS TAKEN AS THE DISTANCE METRIC. NOTE THAT PIVOT DISTANCE MEASURED ONLY THE SHAPE OF THE PARSES AND WAS NOT SENSITIVE TO ANY PARSE-TREE NODE TAGS OR LABELS.

| accuracy on training set | | | accuracy on test set | | |
|--------------------------|-------|----------|----------------------|-------|----------|
| baseline | SRTG | Δ | baseline | SRTG | Δ |
| 95.0% | 97.8% | 2.8% | 90.2% | 95.4% | 5.2% |

TABLE VI: THE RESULTS OF EXPERIMENT 5, CLASSIFYING THE CALTECH house-car DATA SET, USING BOTH THE BASELINE MODELS AND SRTGs.

cantly improves classification accuracy.

Experiment 4 demonstrates that the trained SRTGs allow derivation of image parses that contain sufficient information to support classification. In other words, it is possible to classify images solely on the basis of their hierarchal structure reflected in the shape of their parses. One can imagine using this for image database retrieval. In fact, it is interesting to note that, in experiment 4, classification based solely on the extracted hierarchal structure outperforms, on average, the baseline classification method which directly uses the region feature vectors.

D. Experiments with Additional Datasets

To further illustrate our methods, we conducted two more experiments where we used different datasets. In addition, in order to illustrate the applicability of our classification methods in conjunction with a variety of preprocessing segmentation algorithms, we employed two segmentation algorithms which are different from the segmentation method used in experiments 1–4: Normalized Cuts [49] and Mean Shift [13].

a) *Experiment 5:* In our first additional experiment, we used images from the Caltech house dataset⁷ and the Caltech car dataset⁸. We used the Normalized Cuts algorithm [49] to segment each image, followed by a single pass of Ratio Cut in order to get 10 regions per image. Our basic setup was similar to experiment 1. Specifically, we trained both a separate baseline mixture of 10 Gaussians and an SRTG for each class; the same μ_j and Σ_j parameters were used for the baseline and SRTG models. Out of the total of 441 car images and 237 house images used in the experiment, we randomly selected 160 images of each class as the training set and used the remaining images as the test set. The results of this experiment, summarized in table VI, show that SRTGs significantly outperform the baseline, achieving misclassification rates which are more than twice as low as the baseline misclassification rates for both the training data and the test data.

b) *Experiment 6:* In our second additional experiment, we combined the house and car images used in Experiments 1–4 with another set of 200 images taken around the University Farms subdivision of West Lafayette, IN: 100 images of mailboxes and 100 images of basketball hoops. We segmented the whole dataset with the Mean Shift algorithm [13], followed by three passes of Ratio Cut in order to get 20 or fewer regions per image. We again used an experimental setup similar to that of experiment 1. We conducted five round-robin runs where 80 images in

⁷http://www.vision.caltech.edu/Image_Datasets/Pasadena-Houses-2000.tar

⁸http://www.vision.caltech.edu/Image_Datasets/cars_brad/cars_brad.tar

| | accuracy on training set | | | accuracy on test set | | |
|---------------|--------------------------|-------|----------|----------------------|-------|----------|
| | baseline | SRTG | Δ | baseline | SRTG | Δ |
| hoop-mailbox | 85.3% | 95.5% | 10.2% | 83.0% | 88.0% | 5.0% |
| hoop-car | 97.2% | 99.1% | 1.9% | 97.5% | 98.5% | 1.0% |
| hoop-house | 94.8% | 96.8% | 2.0% | 93.5% | 93.0% | -0.5% |
| mailbox-car | 94.2% | 98.7% | 4.5% | 92.0% | 96.5% | 4.5% |
| mailbox-house | 95.1% | 97.8% | 2.7% | 92.0% | 95.0% | 3.0% |
| car-house | 95.1% | 99.3% | 4.2% | 92.5% | 96.0% | 3.5% |
| mean | 93.6% | 97.9% | 4.3% | 91.75% | 94.5% | 2.75% |

TABLE VII: RESULTS FOR ALL TWO-WAY CLASSIFICATIONS FOR THE hoop-mailbox-house-car DATA SET IN EXPERIMENT 6.

| class | accuracy on training set | | | accuracy on test set | | |
|---------|--------------------------|-------|----------|----------------------|-------|----------|
| | baseline | SRTG | Δ | baseline | SRTG | Δ |
| mailbox | 80.0% | 97.7% | 17.7% | 73.0% | 89.0% | 16.0% |
| hoop | 80.0% | 89.2% | 9.2% | 78.0% | 79.0% | 1.0% |
| house | 89.2% | 96.2% | 7.0% | 82.0% | 92.0% | 10.0% |
| car | 92.5% | 97.2% | 4.7% | 91.0% | 92.0% | 1.0% |
| mean | 85.4% | 95.1% | 9.7% | 81.0% | 88.0% | 7.0% |

TABLE VIII: RESULTS FOR FOUR-WAY CLASSIFICATIONS FOR THE hoop-mailbox-house-car DATA SET IN EXPERIMENT 6.

each class were used as a training set, and the remaining 20 images were used as a test set. The test sets for the five runs were disjoint. The correct classification rates for each pairwise classification task, averaged over the five round-robins, are given in table VII. On average, SRTGs reduce the misclassification rate, as compared to the baseline, by about a factor of three on the training set and a factor of 1.5 on the test set. Note that these average numbers are very similar to the numbers for the two newly added classes: the SRTG misclassification rate for hoops vs mailboxes is about three times smaller than baseline’s misclassification rate on the training set, and about 1.5 times smaller than baseline’s misclassification rate on the test set. The correct classification rates for one-of-four classification are given in table VIII, again showing that SRTGs achieve approximately three and 1.5 times smaller misclassification rates than the baseline, on the training and test sets, respectively.

V. SRTGs, PCFGs, and our Prior Work

Our notation for the probability distribution of images is different from the traditional notation for PCFGs in that, in our case, the probability distribution for \mathbf{O} is conditioned not only on the grammar G but also on the constituent hierarchy \mathcal{H} . However, the traditional PCFGs implicitly assume the constituent hierarchy that corresponds to concatenating 1D strings, and therefore the probabilities induced by a PCFG are implicitly conditioned on this constituent hierarchy.

More precisely, let $D_M = \{1, \dots, M\}$, and let $\mathcal{I} = \{i\}$ consist of a single production class. Let us form a constituent hierarchy $\mathcal{H}_M = \langle D_M, \mathcal{I}, \mathcal{C}_M, \mathcal{L}_M \rangle$ as follows.

We form the set \mathcal{C}_M by taking every contiguous subset of D_M to be a constituent. We take any proper nonempty subset V_1 of V to be a left i -subconstituent of V iff there exists an integer n such that $v \leq n$ for all $v \in V_1$ and $v > n$ for all $v \in V \setminus V_1$. Formulating \mathcal{H}_M in this fashion clearly meets conditions C1–C4. Let $\mathcal{H}^{\mathbb{Z}^+} = \{\mathcal{H}_1, \mathcal{H}_2, \dots\}$ be the collection of these constituent hierarchies \mathcal{H}_M , for all M .

Now let G be any SRTG which uses the same set $\mathcal{I} = \{i\}$ consisting of a single production class. Let G_{PCFG} be the PCFG which has the same sets of terminals and nonterminals as G , and which has the same parameters r_j , p_{jkl} , t_j , μ_j , and Σ_j . For simplicity, we assume that $q = 1$ (i.e. the terminals are scalars) and that the PCFG G_{PCFG} is proper.⁹ It can then be shown that the pair $\langle G, \mathcal{H}^{\mathbb{Z}^+} \rangle$ is equivalent to the PCFG G_{PCFG} , in the following sense. For any $\mathbf{O} \in \mathbb{R}^M$, the quantity $P(\mathbf{O}|G, \mathcal{H}_M)$, as defined in section III-D, is equal to the probability of the string \mathbf{O} induced by the PCFG G_{PCFG} . In this sense, PCFGs can be viewed as a special case of SRTGs which is obtained by using a singleton set of production classes and the collection $\mathcal{H}^{\mathbb{Z}^+}$ of constituent hierarchies.

A similar relationship can be shown to exist between SRTGs as formulated in the present paper, and our prior work [38], [39], [40] where a specialized version of the center-surround algorithm was developed for the case when every constituent is a rectangular set of image pixels. Another version of spatial random tree models and the center-surround algorithm was developed in [54]. In that work, the problem of uniquely associating the leaves of

⁹A PCFG is called *proper* [7] if the total probability of all infinite trees generated by the PCFG is zero.

a parse tree with image pixels was solved by hard-wiring image regions to nonterminals. This resulted in a context-free grammar which models images specified on a fixed finite domain, which is a departure both from [38], [39], [40] and from the framework developed in the present paper.

We finally point out an interesting distinction between PCFGs and our general formulation of SRTGs. Given an SRTG G and a constituent hierarchy \mathcal{H} for a domain D , it is easy to see that a parse tree generated by G will not necessarily be an \mathcal{H} -parse of some feature image $\mathbf{O} \in \mathbb{R}^{q|D|}$, since the number of leaves of the parse tree may not necessarily be equal to $|D|$. This is similar to PCFGs. However, even if the number of leaves of a parse tree T is equal to $|D|$, it may happen that T is not an \mathcal{H} -parse of any feature image. This is in contrast to PCFGs where every parse tree generated by a grammar must necessarily be a parse for some string. This distinction, however, does not influence any of our algorithms.

VI. Conclusion

We have presented a novel method for formulating priors on the hierarchal organization of a feature image and its constituents using SRTGs. The center-surround algorithm can be used to estimate the parameters of SRTGs from sets of training feature images and to classify images based on their likelihood and based on the MAP estimate of the associated hierarchal structure. We have demonstrated the efficacy of these methods by training on and classifying natural images.

There are several open problems associated with our framework that we are currently investigating. We are developing alternative methodologies both for constructing constituent hierarchies [54] and for feature selection to improve the overall classification performance. We are also investigating methods to improve parameter estimation—i.e. to avoid convergence to poor local maxima.

Acknowledgments

We would like to thank Shawn Brownfield, Bingrui Foo, Jacob Harris, Yan Huang, Eduardo Maia, William Nagel, and Eric Théa for experimental work leading to this paper as well as for stimulating and helpful discussion. The experiments reported in this paper utilized the significant computational resources made available to us by the ITaP Research Computing Services Group.

APPENDIX PROOF OF THEOREM 1

The proof is by contradiction. Let ρ be the first node in T , in a preorder traversal from left to right, for which these two functions differ, i.e. for which $F(\rho) \neq F'(\rho)$. Since \mathcal{H} -constituency functions map the root node to D , ρ cannot be the root node.

case 1: ρ is an only child. Since \mathcal{H} -constituency functions map every nonbranching nonterminal node to the

same constituent as its child, F and F' must map the parent of ρ to different constituents. This contradicts the premise, since parents are traversed before their children.

case 2: ρ is a right child. Let V_2 , V , and V_1 denote the constituents to which F maps ρ , its parent, and its sibling and let V'_2 , V' , and V'_1 denote the constituents to which F' maps them respectively. By traversal, $V = V'$ and $V_1 = V'_1$. Since \mathcal{H} -constituency functions map the children of every branching nonterminal node to constituents that partition the constituent mapped to by their parent, $V_2 = V'_2$, which contradicts the premise.

case 3: ρ is a left child. Let i be the tag of the parent of ρ . Let V_1 and V denote the constituents to which F maps ρ and its parent, and let V'_1 and V' denote the constituents to which F' maps them respectively. By traversal, $V = V'$. Furthermore, $|V_1| = |V'_1|$, since both equal the number of leaf nodes dominated by ρ . By condition C4, $V_1 = V'_1$, which contradicts the premise.

This proves that $F = F'$. It is easy to see that conditions C1-C4, together with our definition of a parse tree, imply that, for every $v \in D$, there exists a leaf node ρ of T such that $F(\rho) = \{v\}$. For this location v , we have, from the definition of an \mathcal{H} -constituency function, $O(v) = u$ where u is the label of ρ . Since $F = F'$, it follows that $F'(\rho) = \{v\}$, and therefore $O'(v) = u = O(v)$. Since this is true for every $v \in D$, it follows that $\mathbf{O} = \mathbf{O}'$. ■

MAP PARSE AND PARAMETER ESTIMATION FORMULAS

In figure 7, we give the recursive formulas for computing the MAP \mathcal{H} -parse of an image, as discussed in section III-E. Note that equation 5 is identical to equation 1 of figure 4, and equation 6 is obtained by replacing each “ \sum ” with a “max” in equation 2.

Equations 8–14 of figure 8 are the formulas for the parameter updates, as discussed in section III-E. In equations 8–14, the parameters corresponding to the m -th training feature image \mathbf{O}_m are indexed by m . In particular, c_m stands for the center variables computed by applying the center recursions to \mathbf{O}_m ; and s_m stands for the surround variables computed by applying the surround recursions to \mathbf{O}_m . If we take G' to be $\langle \mathcal{I}, \mathcal{J}, \mathbb{R}^q, r', t', p', \mu', \Sigma' \rangle$, it can be shown that
$$\prod_{m=1}^M P(\mathbf{O}_m | G', \mathcal{H}_m) \geq \prod_{m=1}^M P(\mathbf{O}_m | G, \mathcal{H}_m).$$
 Repeating this process converges to a local maximum [44], [55].

PIVOT DISTANCE

Let us define the *pivot distance* between pairs of parse trees as follows. Let $\text{BRANCHING}(\rho)$ be true if ρ is a branching nonterminal node and false if it is a nonbranching nonterminal node. If ρ is a branching nonterminal node, let $\text{LEFT}(\rho)$ and $\text{RIGHT}(\rho)$ denote the left and

$$\begin{aligned}\hat{c}(j, \{v\}) &= t_j N(O(v) | \mu_j, \Sigma_j) \\ \hat{c}(j, V) &= \max_{i,k,l} \max_{V_1 \in \mathcal{L}(i,V)} (1 - t_j) p_{jikl} \hat{c}(k, V_1) \hat{c}(l, V \setminus V_1) \quad |V| > 1 \quad (5) \\ \langle i, \widehat{k, l, V_1} \rangle(j, V) &= \arg \max_{i,k,l} \max_{V_1 \in \mathcal{L}(i,V)} (1 - t_j) p_{jikl} \hat{c}(k, V_1) \hat{c}(l, V \setminus V_1) \quad |V| > 1 \quad (7)\end{aligned}$$

Fig. 7. The Viterbi variant of the center recursions, used for computing the MAP \mathcal{H} -parse of a feature image \mathbf{O} defined on a domain D .

$$f_m(j, i, k, l, V) = \frac{s_m(j, V)(1 - t_j) p_{jikl} \sum_{V_1 \in \mathcal{L}(i,V)} c_m(k, V_1) c_m(l, V \setminus V_1)}{\sum_{j \in \mathcal{J}} c_m(j, D_m) r_j} \quad (8)$$

$$g_m(j, V) = \frac{s_m(j, V) c_m(j, V)}{\sum_{j \in \mathcal{J}} c_m(j, D) r_j} \quad (9) \quad p'_{jikl} = \frac{\sum_{m=1}^M \sum_{V \in \mathcal{C}_m, |V| > 1} f_m(j, i, k, l, V)}{\sum_{m=1}^M \sum_{V \in \mathcal{C}_m, |V| > 1} g_m(j, V)} \quad (12)$$

$$r'_j = \frac{1}{M} \sum_{m=1}^M g_m(j, D_m) \quad (10) \quad \mu'_j = \frac{\sum_{m=1}^M \sum_{v \in D_m} g_m(j, \{v\}) O_m(v)}{\sum_{m=1}^M \sum_{v \in D_m} g_m(j, \{v\})} \quad (13)$$

$$t'_j = \frac{\sum_{m=1}^M \sum_{v \in D_m} g_m(j, \{v\})}{\sum_{m=1}^M \sum_{V \in \mathcal{C}_m} g_m(j, V)} \quad (11) \quad \Sigma'_j = \frac{\sum_{m=1}^M \sum_{v \in D_m} g_m(j, \{v\}) [O_m(v) - \mu'_j] [O_m(v) - \mu'_j]^t}{\sum_{m=1}^M \sum_{v \in D_m} g_m(j, \{v\})} \quad (14)$$

Fig. 8. Parameter reestimation: (8–9) The E step of the reestimation procedure, (10–14) The M step of the reestimation procedure. We use the superscript t to denote the transpose of a vector.

right child of ρ respectively. Define a *left pivot* to be the transformation from the right to left below and a *right pivot* to be the reverse transformation.



Note that it is only possible to left pivot a branching node whose right child is a branching node. Similarly note that it is only possible to right pivot a branching node whose left child is a branching node. Let $\text{LEFTPIVOT}(\rho)$ and $\text{RIGHTPIVOT}(\rho)$ denote the left and right pivots of ρ respectively, when they exist, and ρ , when they do not exist.

The pivot distance $\|\rho_1, \rho_2\|$ between ρ_1 and ρ_2 is defined as the minimal number of pivots that must be applied to ρ_1 and ρ_2 so that the resulting trees have the same shape (i.e. are the same ignoring labels and tags). The pivot distance $\|\rho_1, \rho_2\|$ can be computed as shown in the equation at the top of the next page.

This can be computed in polynomial time by memoizing $\|\rho_1, \rho_2\|$. Note that the pivot distance between trees with different numbers of nodes is infinite and that the pivot distance between trees with the same number of nodes is finite.

REFERENCES

[1] J. K. Baker. Trainable grammars for speech recognition. In *Speech Communications Papers for the 97th Meeting of the Acoustical Society of America*, pages 547–550, 1979.

[2] L. E. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, 41(1):164–171, 1970.

[3] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society B*, 36(2):192–236, 1974.

[4] J. Besag. Efficiency of pseudolikelihood estimation for simple Gaussian fields. *Biometrika*, 64(3):616–618, 1977.

[5] Aaron F. Bobick and Yuri A. Ivanov. Action recognition using probabilistic parsing. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 196–202, Santa Barbara, CA, June 1998.

[6] C. A. Bouman and M. Shapiro. A multiscale random field model for Bayesian image segmentation. *IEEE Transactions on Image Processing*, 3(2):162–177, March 1994.

[7] Zhiyi Chi and Stuart Geman. Estimation of probabilistic context-free grammars. *Computational Linguistics*, 24(2):299–305, June 1998.

[8] Hyeokho Choi and Richard G. Baraniuk. Multiscale document segmentation using wavelet-domain hidden Markov models. In *Proceedings of IST/SPIE's 12th Annual International Symposium—Electronic Imaging 2000, Science & Technology, San Jose, CA, January 2000*, January 2000.

[9] Hyeokho Choi and Richard G. Baraniuk. Multiscale image segmentation using wavelet-domain hidden Markov models. *IEEE Transactions on Image Processing*, 10(9):1309–1321, September 2001.

[10] Noam Chomsky. *The Logical Structure of Linguistic Theory*. Plenum, New York, NY, 1955.

[11] Noam Chomsky. *Syntactic Structures*. Mouton, The Hague, 1957.

[12] Noam Chomsky. On certain formal properties of grammars. *Information and Control*, 2:137–167, 1959.

[13] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, May 2002.

[14] Ingemar J. Cox, Satish B. Rao, and Yu Zhong. Ratio regions: A technique for image segmentation. In *Proceedings of the International Conference on Pattern Recognition*, pages 557–564, Santa Barbara, CA, August 1996.

[15] Matthew S. Crouse, Robert D. Nowak, and Richard G. Baraniuk. Wavelet-based statistical signal processing using hidden Markov

$$\|\rho_1, \rho_2\| = \begin{cases} 0 \\ \infty \\ \infty \\ \min \left(\begin{array}{l} 1 + \min \left(\begin{array}{l} \|\text{LEFTPIVOT}(\rho_1), \rho_2\|, \\ \|\text{RIGHTPIVOT}(\rho_1), \rho_2\|, \\ \|\rho_1, \text{LEFTPIVOT}(\rho_2)\|, \\ \|\rho_1, \text{RIGHTPIVOT}(\rho_2)\| \end{array} \right), \\ \|\text{LEFT}(\rho_1), \text{LEFT}(\rho_2)\| + \\ \|\text{RIGHT}(\rho_1), \text{RIGHT}(\rho_2)\| \end{array} \right), \end{cases} \begin{pmatrix} \neg\text{BRANCHING}(\rho_1), \\ \neg\text{BRANCHING}(\rho_2) \\ \text{BRANCHING}(\rho_1), \\ \neg\text{BRANCHING}(\rho_2) \\ \neg\text{BRANCHING}(\rho_1), \\ \text{BRANCHING}(\rho_2) \\ \text{BRANCHING}(\rho_1), \\ \text{BRANCHING}(\rho_2) \end{pmatrix}$$

models. *IEEE Transactions on Signal Processing*, 46(4):886–902, April 1998.

- [16] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society B*, 39:1–38, 1977.
- [17] K. S. Fu. *Syntactic Pattern Recognition and Applications*. Prentice-Hall, Englewood Cliffs, NJ, 1982.
- [18] Stuart Geman and Donald Geman. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, November 1984.
- [19] Xuming He, Richard S. Zemel, and Miguel Á. Carreira-Perpián. Multiscale conditional random fields for image labeling. In *Proc. CVPR*, 2004.
- [20] Tapas Kanungo and Song Mao. Stochastic language models for style-directed layout analysis of document images. *IEEE Transactions on Image Processing*, 12(5):583–596, May 2003.
- [21] T. Kasami. An efficient recognition and syntax algorithm for context-free languages. Scientific Report AFCRL-65-758, Air Force Cambridge Research Laboratory, Bedford MA, 1965.
- [22] Gary E. Kopec and Philip A. Chou. Document image decoding using Markov source models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6):602–617, June 1994.
- [23] Mukkai Krishnamoorthy, George Nagy, Sharad C. Seth, and Mahesh Viswanathan. Syntactic segmentation and labeling of digitized pages from technical journals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(7):737–747, July 1993.
- [24] Sanjiv Kumar and Martial Hebert. A hierarchical field framework for unified context-based classification. In *Proc. ICCV2005*, October 2005.
- [25] John Lafferty, Andrew McCallum, and Fernando Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proc. International Conference on Machine Learning*, pages 282–289, 2001.
- [26] K. Lari and S. J. Young. The estimation of stochastic context-free grammars using the inside-outside algorithm. *Computer Speech and Language*, 4(1):35–56, 1990.
- [27] Stephen E. Levinson. Continuously variable duration hidden Markov models for speech analysis. In *Proceedings of the International Conference on Acoustic and Speech Signal Processing*, pages 1241–1244, 1986.
- [28] Jia Li and Robert M. Gray. Context-based multiscale classification of document images using wavelet coefficient distributions. *IEEE Transactions on Image Processing*, 9(9):1604–1616, September 2000.
- [29] Jia Li, Robert M. Gray, and Richard A. Olshen. Multiresolution image classification by hierarchical modeling with two-dimensional hidden Markov models. *IEEE Transactions on Information Theory*, 46(5):1826–1841, August 2000.
- [30] M. R. Luetzgen, W. C. Karl, A. S. Willsky, and Robert R. Tenney. Multiscale representations of Markov random fields. *IEEE Transactions on Signal Processing*, 41(12), December 1993.
- [31] C. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA, 1999.
- [32] James W. Modestino and Jun Zhang. A Markov random field model-based approach to image interpretation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(6):606–615, June 1992.
- [33] J.-M. Morel and S. Solimini. *Variational Methods in Image Segmentation*. Birkhauser, 1995.
- [34] George Nagy and Sharad C. Seth. Hierarchical image representation with application to optically scanned documents. In *Proc. ICPR84*, pages 347–349, Montreal, Canada, July 1984.
- [35] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [36] Fernando Pereira and Yves Schabes. Inside-outside reestimation from partially bracketed corpora. In *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, pages 128–135, Newark, DE, 1992.
- [37] I. Pollak, A. S. Willsky, and H. Krim. Image segmentation and edge enhancement with stabilized inverse diffusion equations. *IEEE Transactions on Image Processing*, 9(2), February 2000.
- [38] Ilya Pollak, Jeffrey Mark Siskind, Mary P. Harper, and Charles A. Bouman. Modeling and estimation of spatial random trees with application to image classification. In *Proceedings of ICASSP*, April 2003.
- [39] Ilya Pollak, Jeffrey Mark Siskind, Mary P. Harper, and Charles A. Bouman. Parameter estimation for spatial random trees using the EM algorithm. In *Proceedings of ICIP*, September 2003.
- [40] Ilya Pollak, Jeffrey Mark Siskind, Mary P. Harper, and Charles A. Bouman. Spatial random trees and the center-surround algorithm. Technical Report TR-ECE-03-03, Purdue University, School of Electrical and Computer Engineering, January 2003.
- [41] Gerasimos G. Potamianos and John K. Goutsias. Partition function estimation of Gibbs random filed images using Monte Carlo simulation. *IEEE Transactions on Information Theory*, 39(4):1322–1332, July 1993.
- [42] D. Potter. *Compositional Pattern Recognition*. PhD thesis, Brown University, 1999. <http://www.dam.brown.edu/people/dfp>.
- [43] L.R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proc. of the IEEE*, 77(2):257–286, 1989.
- [44] Richard A. Redner and Homer F. Walker. Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review*, 26(2):195–239, April 1984.
- [45] Justin K. Romberg, Hyeokho Choi, and Richard G. Baraniuk. Bayesian tree-structured image modeling using wavelet-domain hidden Markov models. *IEEE Transactions on Image Processing*, 10(7):1056–1068, July 2001.
- [46] A. Rosenfeld. *Picture Languages*. Academic Press, New York, NY, 1979.
- [47] D. Sankoff. Branching processes with terminal types: application to context-free grammars. *Journal of Applied Probability*, 8:233–240, 1971.
- [48] A. C. Shaw. A formal picture description scheme as a basis for picture processing systems. *Information and Control*, 14:9–52, 1969.
- [49] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, August 2000.
- [50] Antonio Torralba, Kevin P. Murphy, and William T. Freeman.

Contextual models for object detection using boosted random fields. In *NIPS*, 2004.

- [51] Vladimir N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.
- [52] Andrew J. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13:260–267, 1967.
- [53] Song Wang and Jeffrey Mark Siskind. Image segmentation with ratio cut. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(6):675–690, June 2003.
- [54] Wiley Wang, Ilya Pollak, Tak-Shing Wong, Charles A. Bouman, Mary P. Harper, and Jeffrey Mark Siskind. Hierarchical stochastic image grammars for classification and segmentation. *IEEE Transactions on Image Processing*, 15(10):3033–3052, October 2006.
- [55] C.F. Jeff Wu. On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11(1):95–103, 1983.
- [56] D. H. Younger. Recognition and parsing of context-free languages in time $O(n^3)$. *Information and Control*, 10(2):189–208, 1967.

PLACE
PHOTO
HERE

Dr. Ilya Pollak received the B.S. and M.Eng. degrees in 1995 and Ph.D. in 1999, all from the Massachusetts Institute of Technology, Cambridge, all in electrical engineering. In 1999–2000, he was a post-doctoral researcher at the Division of Applied Mathematics, Brown University, Providence, RI. Since 2000, he has been with Purdue University, West Lafayette, IN, where he is currently an Associate Professor of Electrical and Computer Engineering. He has held short-term visiting positions at the Institut National de Recherche en Informatique et en Automatique in Sophia Antipolis, France, and at Tampere University of Technology, Finland. His research interests are in image and signal processing, specifically, hierarchical statistical models, fast estimation algorithms, nonlinear scale-spaces, and adaptive representations.

Dr. Pollak received a CAREER award from the National Science Foundation in 2001. He is an Associate Editor of the *IEEE Transactions on Signal Processing*, a member of the *IEEE Signal Processing Society's* Technical Committee on Signal Processing Theory and Methods, and the Chair of the Signal Processing Chapter of the Central Indiana Section of the IEEE. He is a Co-Chair of the SPIE/IS&T Conference on Computational Imaging.

PLACE
PHOTO
HERE

Dr. Jeffrey M. Siskind received the B.A. degree in computer science from the Technion, Israel Institute of Technology, Haifa, in 1979, the S.M. degree in computer science from the Massachusetts Institute of Technology (M.I.T.), Cambridge, in 1989, and the Ph.D. degree in computer science from M.I.T. in 1992. He did a postdoctoral fellowship at the University of Pennsylvania Institute for Research in Cognitive Science from 1992 to 1993. He was an assistant professor at the University of Toronto

Department of Computer Science from 1993 to 1995, a senior lecturer at the Technion Department of Electrical Engineering in 1996, a visiting assistant professor at the University of Vermont Department of Computer Science and Electrical Engineering from 1996 to 1997, and a research scientist at NEC Research Institute, Inc. from 1997 to 2001. He joined the Purdue University School of Electrical and Computer Engineering in 2002 where he is currently an associate professor. His research interests include machine vision, artificial intelligence, cognitive science, computational linguistics, child language acquisition, and programming languages and compilers.

PLACE
PHOTO
HERE

Dr. Mary P. Harper (M'89–SM'02) received the Sc.M. and Ph.D. degrees in computer science from Brown University, Providence, RI in 1986 and 1990, respectively. In 1989, she joined the faculty of Purdue University, West Lafayette, IN, where she currently holds the rank of Professor in the School of Electrical and Computer Engineering. Dr. Harper's research focuses on computer modeling of human communication with a focus on methods for incorporating multiple types of knowledge sources,

including lexical, syntactic, prosodic, and most recently visual sources. Recent research involves the integration of speech and natural language processing systems, the integration of speech, gesture, and gaze, and the utilization of hierarchical structure to improve the classification accuracy of documents and images.

Dr. Harper was a speaker for the IEEE Computer Society Distinguished Visitors Program and Chapter Tutorial Program from 1997 through 2000. She is currently an associate editor for *IEEE Transactions on Speech, Audio, and Language Processing* and sits on the IEEE Speech Technical Committee of the IEEE Signal Processing Society. She also sits on the Board for the North American Chapter of the Association for Computational Linguistics (NAACL).

James J. Sherman

PLACE
PHOTO
HERE

PLACE
PHOTO
HERE

Dr. Charles A. Bouman (S'86-M'89-SM'97-F'01) received the B.S.E.E. degree from the University of Pennsylvania, Philadelphia, in 1981, the MS degree in electrical engineering from the University of California, Berkeley, in 1982, and the Ph.D. degree in electrical engineering from Princeton University, Princeton, NJ. From 1982 to 1985, he was a full staff member at MIT Lincoln Laboratory, Lexington, MA. In 1989, he joined the faculty of Purdue University, West Lafayette, IN, where he holds

the rank of Professor with a primary appointment in the School of Electrical and Computer Engineering and a secondary appointment in the School of Biomedical Engineering. Dr. Bouman's research focuses on the use of statistical image models, multiscale techniques, and fast algorithms in applications including medical and electronic imaging.

Dr. Bouman is a Fellow of the IEEE, a Fellow of the American Institute for Medical and Biological Engineering (AIMBE), a Fellow of the society for Imaging Science and Technology (IS&T), a member of the SPIE professional society, a recipient of IS&T's Raymond C. Bowman Award for outstanding contributions to digital imaging education and research, and a University Faculty Scholar of Purdue University. He is currently the Editor-in-Chief Elect for the IEEE Transactions on Image Processing and a member of the Steering Committee for the IEEE Transactions on Medical Imaging. He has been an associate editor for the IEEE Transactions on Image Processing and the IEEE Transactions on Pattern Analysis and Machine Intelligence. He has also been Co-Chair of the 2006 SPIE/IS&T Symposium on Electronic Imaging, Co-Chair of the SPIE/IS&T conferences on Visual Communications and Image Processing 2000 (VCIP), a Vice President of Publications and a member of the Board of Directors for the IS&T Society, and he is the founder and co-chair of the SPIE/IS&T conference on Computational Imaging.