

---

# Stochastic Segmentation Trees for Multiple Ground Truths

---

**Jake Snell**

University of Toronto  
Vector Institute

**Richard S. Zemel**

University of Toronto  
Vector Institute  
Canadian Institute for Advanced Research

## Abstract

A strong machine learning system should aim to produce the full range of valid interpretations rather than a single mode in tasks involving inherent ambiguity. This is particularly true for image segmentation, in which there are many sensible ways to partition an image into regions. We formulate a tree-structured probabilistic model, the *stochastic segmentation tree*, that represents a distribution over segmentations of a given image. We train this model by optimizing a novel objective that quantifies the degree of match between statistics of the model and ground truth segmentations. Our method allows learning of both the parameters in the tree and the structure itself. We demonstrate on two datasets, including the challenging Berkeley Segmentation Dataset, that our model is able to successfully capture the range of ground truths and to produce novel plausible segmentations beyond those found in the data.

## 1 INTRODUCTION

Humans have a remarkable ability to resolve ambiguity in real-world situations. For example, a skilled interpreter knows that many valid translations of a sentence to a target language exist, and which is the most appropriate depends on context. Yet even in the absence of disambiguating information, humans are able to internally represent a wide range of plausible interpretations when presented with a stimulus. As the central aim of AI is to build machines that can reason about the world in a human-like way, strong AI systems should also be able to maintain a distribution over plausible outputs when confronted with an ambiguous task. Many current machine learning approaches are unable to do this despite




Figure 1: Multiple ground truths from the Berkeley Segmentation Dataset and samples from our model. Stochastic segmentation trees produce segmentations that generalize through novel combinations of the level of detail in different areas of the image.

the prevalence of ambiguity in domains such as natural language parsing, machine translation, image caption generation, and visual scene understanding.

Tasks involving ambiguity can often naturally be cast as structured output learning problems, in which outputs have internal dependencies. Examples of structured outputs are natural language sentences or paths in a graph. It is often infeasible to independently model each element of the exponentially large structured output spaces typically considered in such problems. Fortunately, graphical models provide a powerful and flexible framework for building rich distributions over structured

output spaces. This is a boon for tasks involving ambiguity as they tend to naturally have multiple ground truths, each reflecting a different yet correct interpretation of the input. For example, the Berkeley Segmentation Dataset (BSD) (Martin et al., 2001; Arbelaez et al., 2011) contains segmentations of natural images collected from multiple human annotators. Due to variability in interpretation or attention, some annotators provided more detailed segmentations or emphasized different contour types than others (Hou et al., 2013). See Figure 1 for an example of such variation.

In this work we introduce the *stochastic segmentation tree* (SST) model, which treats multiple ground truths as samples from a target distribution and learns to predict a distribution specific to the given image. Motivated by the observation that biological visual systems parse scenes according to a hierarchy (Sharon et al., 2006), the model first builds a tree-structured latent region hierarchy of the image. It then predicts a distribution over binary variables corresponding to nodes in the tree that indicate which regions in the tree should be grouped in the same segment.

We make a number of contributions in this paper:

- We introduce an aim that to our knowledge is novel, of producing for a single test input multiple valid structured outputs that directly capture variations in ground truths (GTs).
- We formulate a probabilistic model, a *stochastic segmentation tree*, that represents a distribution over segmentations of a given image.
- We develop a method of forming a tree specific to a given image, and assigning probabilities to its nodes, from a base set of pixels or superpixels (groupings of pixels).
- We define a novel objective that quantifies the degree of match between the model and ground truth segmentations, and formulate an optimization method that allows learning of both the parameters in a tree and the tree structure itself.

## 2 RELATED WORK

The first line of research related to our SSTs concerns learning systems that produce multiple structured outputs. Many of these approaches can be viewed as structured output generalizations of common multi-label prediction, in which an input can be associated with any number of outputs. One way this can be done is to redefine the output space to be over sets rather than single outputs (Lampert, 2011). Another strategy is to train an ensemble of models to produce multiple hypotheses

at test time, either independently (Guzman-Rivera et al., 2012) or in a cascade (Dey et al., 2015). Multiple outputs can also be produced by applying a specialized inference procedure, such as Diverse M-Best MAP (Batra et al., 2012; Yadollahpour et al., 2013; Kirillov et al., 2015) or Diverse Beam Search (Vijayakumar et al., 2016). These approaches have some overlap with our SSTs but differ in aims and formulation. None of these methods produce alternative valid interpretations given an input but instead predict a set of outputs such that one of them is good. Others have attempted to model diverse (and thereby multiple) outputs more directly via Structured DPPs (Kulesza et al., 2012; Gillenwater et al., 2012).

Our model is closely related to approaches that build a graphical model whose structure is conditioned on input data. For example, Wong & Mooney (2006) learn a set of context-free grammar rules which are used to produce a parse tree for a given natural language sentence. They also learn model parameters that define a distribution over meaning representations for the predicted parse tree. Their graphical model parameters are globally learned whereas ours are predicted by the model given the input. The machine translation model of Tromble et al. (2008) forms a distribution over sentences via paths through a translation lattice given a source sentence.

Our model also relates to structured output models, but differ in the training and evaluation criteria. Structured output models are conventionally trained with maximum likelihood given the ground truths for each input. Models are typically evaluated, however, via a task-specific loss comparing a single model output against the ground truths. Examples of task losses include mean average precision for rankings and BLEU (Papineni et al., 2002) for text generation. This disconnect between training and test has spurred many attempts to minimize task loss at training time (Hazan et al., 2010; Song et al., 2015; Shen et al., 2015; Ranzato et al., 2015). These training procedures often improve task loss performance at test but require additional loss-augmented inference or dynamic programming steps during training. In contrast, our method constructs a target distribution from the ground truths and learns to minimize the distance between the model and target distributions. When the model is amenable to such distance computation, this is a more efficient way to learn than task loss minimization since no inference is required. Moreover, the evaluation of model outputs against multiple ground truths via a task loss depends on specific characteristics of the (often few in number) ground truths. BLEU, for example, compares a model generated sentence against n-grams found in ground truth sentences. Yet replacing words in the ground truth sentences with direct synonyms would likely be equally plausible outputs. Using

domain knowledge to construct a target distribution invariant to such low-level variations is a more direct way to train and evaluate models, and is the approach we take in this work.

We now review approaches specific to image segmentation. The goal of image segmentation is to partition an image into meaningful regions. Several task losses are used to measure segmentation performance; one of the most popular is the Probabilistic Rand Index (Unnikrishnan & Hebert, 2005), a generalization of Rand Index (Rand, 1971) to the single-output multiple ground-truth setting. A common approach to image segmentation is to form a hierarchical region grouping, in which regions of finer-scale segmentations are subregions of those at coarser scales, and then cut the tree to obtain an output segmentation. The gPb-OwT-UCM (Arbelaez et al., 2011) algorithm (called UCM, hereafter) utilizes the oriented watershed transform on gPb (Maire et al., 2008) boundaries, followed by greedy merging of regions. It forms a weighted boundary map such that thresholding it at any level produces a valid segmentation, and the threshold value controls the scale. Jain et al. (2011) propose a method that learns to agglomerate superpixels into hierarchies, by learning a similarity function based on ground-truth clustered data. The agglomerative region clustering method (Ren & Shakhnarovich, 2013) uses a learned boundary probability model to merge regions until the estimated probability of merging is below a threshold. Arbelaez et al. (2014) propose a multiscale combinatorial grouping that uses a fast normalized cuts algorithm to produce a hierarchical segmentation that leverages information across different scales.

In general, these methods provide alternative approaches to construct the visual hierarchy from the image; our primary aim is to utilize this tree structure to generate alternative segmentations. The two approaches that most resemble our model are first the work by Ion et al. (2011), which composes multiple figure-ground hypotheses obtained by applying constraints at different locations and scales, into multiple larger interpretations of the entire image. More recently, Hu et al. (2015) propose a probabilistic generative model for segmentations, which can efficiently generate segmentations of varying granularity; unlike our SSTs this work does not propose a novel tree construction process, and more importantly, is trained via maximum likelihood as opposed to matching the set of human annotations.

### 3 SST MODEL

Our SST model produces multiple segmentation hypotheses from a single image in three steps: (1) a tree structure is constructed; (2) probabilities are associated

with nodes in the tree; (3) sample segmentations are generated from the tree based on the node probabilities. Here we detail these steps, initially assuming that the tree structure is given. We first describe how nodes in a tree relate to regions of an image, and how a binary labeling of nodes in such a tree corresponds to a segmentation of the image. We then present the probabilistic model over segmentations based on these binary labelings, and a metric comparing the model samples to ground-truth segmentations of an image that can be used to evaluate a multiple segmentation system. This metric is then the objective utilized in learning the network that predicts the node probabilities from an image at test time. Finally at the end of this section we describe our own procedure for generating trees, which is also based on this metric.

#### 3.1 PROBABILISTIC TREE MODEL

A segmentation can be represented as a vector of numbers, where each pixel is assigned a particular region index  $r \in \mathcal{R}$ . In such a representation, pixels assigned the same index  $r$  are considered to be grouped together. An alternative representation of the same information is as a symmetric binary matrix  $\mathbf{S}$ , where  $S_{ij} = 1$  if pixel  $i$  and  $j$  are assigned to the same region, and  $S_{ij} = 0$  otherwise. We adopt this representation here.

We also utilize a third representation of a segmentation, with respect to a *segmentation tree*. A segmentation tree has  $M$  leaves, one for each of  $M$  non-overlapping and complete image superpixels. A superpixel  $m \in \{1, \dots, M\}$  represents a grouping of contiguous pixels in the image; we denote the set of pixels assigned to superpixel  $m$  as  $C(m)$ . As a slight abuse of notation, we use  $m$  to denote both a superpixel and its corresponding leaf in the tree. Each internal node  $k$  represents a grouping of pixels that is the union of pixels assigned to its children:  $C(k) = C(\ell(k)) \cup C(r(k))$ , where  $\ell(k)$  and  $r(k)$  are the left and right child of  $k$ , respectively. This assignment of nodes to left and right children of parent nodes define the tree structure.

Given a particular tree structure, a segmentation is a set of binary values  $\mathbf{t}$ , where each leaf  $m$  has  $t_m = 1$ , and the binary value  $t_k$  associated with each internal tree node  $k$  represents whether or not the set of pixels in its left child and the pixels of the right child belong to the same image region:

$$t_k = 1 \Rightarrow S_{ij} = 1, \quad \forall i \in C(\ell(k)), j \in C(r(k)) \quad (1)$$

These binary values are constrained to ensure that the values in the pairwise pixel matrix  $\mathbf{S}$  correspond to a valid segmentation. This is achieved by forcing  $t_k = 0$  if either  $t_{\ell(k)} = 0$  or  $t_{r(k)} = 0$ . This constraint results in tree labelings in which any path that starts at a leaf and




Figure 2: Left: image and color-coded superpixels. Middle: two sample segmentations, with corresponding binary values ( $\mathbf{t}$ ) of the segmentation tree. Right: each node  $k$  in the tree has an associated marginal probability  $\pi_k$ , which governs a specific region of the pairwise segmentation matrix.

ends at the root has an initial value of 1, and if the value switches to 0 along the path then it stays at 0 through the root.

The last node along this path with a value of 1 defines a sub-tree of superpixels that all belong to the same image region, which corresponds to a block in the segmentation matrix. The relationship between the tree labeling  $\mathbf{t}$  and the matrix  $\mathbf{S}$  is illustrated in Figure 2.

In order to represent the variety of segmentations of a given image, with different levels of detail in different regions, we formulate a probabilistic model based on a leaf-to-root sampling scheme. The probabilities are assigned to nodes in the tree, which directly correspond to probabilities in the segmentation matrix. These equations define the SST probabilistic model:

$$P(t_k = 1 | \mathbf{t}_{-k}) = \begin{cases} p_k & \text{if } t_{\ell(k)} = t_{r(k)} = 1 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where  $\mathbf{t}_{-k}$  is the vector of binary values associated with the nodes below  $k$  in the tree. Given this definition, and a sampling scheme that samples binary values  $t_k$  starting at the leaves and moving up the tree, then it can be shown that the marginal probability  $\pi_k$  that node  $k$  takes on value 1 is a product of the local probability of that node and the marginals of its children:

$$\pi_k \equiv P(t_k = 1) = \prod_{d \in \text{subTree}(k)} p_d = p_k \pi_{\ell(k)} \pi_{r(k)} \quad (3)$$

where  $\text{subTree}(k)$  are all the nodes in the sub-tree rooted at node  $k$ .

The probabilistic model plays an important role in our system, specifying a distribution over segmentations of an image, which can be sampled to generate hypothesized segmentations. An example probability vector

$P(\mathbf{t})$  for a given tree, and some corresponding samples are illustrated in Figure 2.

### 3.2 EVALUATING AN SST

The aim of the stochastic segmentation tree is to produce segmentations such that their statistics match the statistics of the reference segmentations for a given image. One metric that is well-suited to this goal is the *Hellinger distance*, which is a metric for comparing two distributions.

In our setting we utilize the marginal pairwise probabilities as the underlying probabilities to be compared via the Hellinger distance. For the model, the marginal probability that pixel  $i$  is in the same segment as pixel  $j$  is captured by  $\pi_k$ , where  $i \in C(\ell(k))$  and  $j \in C(r(k))$ , or vice versa. Given a set of ground-truth segmentations  $\{\mathbf{S}^1, \dots, \mathbf{S}^U\}$ , the target marginals  $q_{ij}$  can be computed easily:  $q_{ij} = \frac{1}{U} \sum_u S_{ij}^u$ . The squared-Hellinger distance  $D(\pi, q)$  is defined to be:

$$1 - \frac{1}{\binom{N}{2}} \left( D_0 + \sum_{\substack{k, i \in C(\ell(k)) \\ j \in C(r(k))}} \left[ \sqrt{\pi_k q_{ij}} + \sqrt{(1 - \pi_k)(1 - q_{ij})} \right] \right) \quad (4)$$

where  $N$  is the number of pixels. Here  $D_0$  is the summed Hellinger distance of each of the super-pixels, which sums over the SST leaves of any errors in the super-pixelation; these errors are constant and cannot be repaired by the SST.

We propose that matching pairwise statistics is an appropriate goal for the multiple-segmentation scenario. Intuitively, the metric considers all pairs of pixels in the image, and compares how often that pair are grouped together by humans versus the model. An alternative is

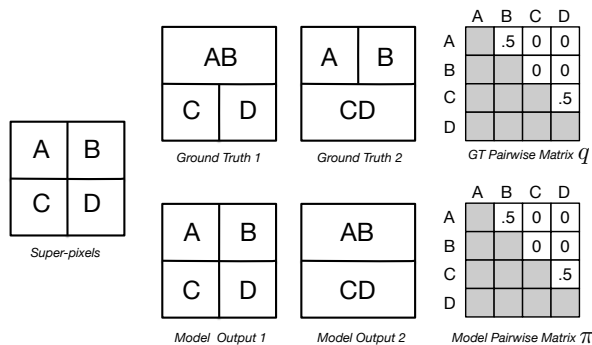


Figure 3: Illustration of the properties of the Hellinger distance as it relates to segmentations. In this scenario, there are four superpixels and two ground truth segmentations. When viewed in the pairwise matrix form, both the ground truth and model distributions are identical despite the fact that the model outputs segmentations not found in the ground truths. We hypothesize that these are sensible generalizations.

to attempt to match higher order statistics between the model and ground-truth segmentations, e.g., triplets of pixels, but given that typically only a small reference set of segmentations are available, these higher-order statistics are not as reliable. A visualization of the properties of the Hellinger distance is shown in Figure 3.

### 3.3 NETWORK TO ESTIMATE NODE MARGINALS

If the reference segmentations are available, the node probabilities and their corresponding marginals for a given tree structure can be optimized (e.g., via gradient descent) to minimize the Hellinger objective (Equation 4). At test time, however, with no reference segmentation these node marginals must be estimated from the image. We formulate a neural network that estimates node marginals, and optimize it to minimize the Hellinger objective on the training set. Then given a tree for a test image, we can estimate marginals, and sample segmentations.

Our network to estimate the marginals  $\pi_k$  for any node  $k$  is a three layer feed-forward neural net. The inputs to the network consist of visual features extracted from regions corresponding to  $\ell(k)$  and  $r(k)$ , as well as a small number of geometric features of the regions, including centroid location and size. The first layer of the network applies the same weights to these features of the left and right regions. The second layer takes as input the concatenation of the hidden left and right representations with their element-wise squared difference. The third and final layer connects to the network output, a single sigmoid unit, which is the estimated marginal prob-

ability for the  $\ell(k)$ ,  $r(k)$  merge. The same network is used to predict merge marginal probabilities throughout the entire tree.

Learning is performed via SGD with backpropagation, to minimize the Hellinger distance with respect to the parameters of our model. This is made simpler by the fact that  $D(\pi, q)$  decomposes over nodes of the tree. The learned network predicts marginal probabilities independently for each node and therefore it is possible that the network’s predictions are not consistent with our tree assumptions from Section 3.1. For example, if the marginal probability predicted at a certain node is greater than that predicted at either of its children, there will be no setting of the conditional node probabilities  $\mathbf{p}$  consistent with the network outputs that also satisfies Equation 3. Thus as a final step, we perform a separate constrained optimization using L-BFGS (Liu & Nocedal, 1989) to find the setting of  $\mathbf{p}$  that has minimal Hellinger distance to the marginals outputted by the network while still obeying our tree assumptions. We then recompute the marginal probabilities corresponding to this setting of  $\mathbf{p}$  using Equation 3 and treat these reconciled marginals as the final outputs of the model.

Sample segmentations can be generated from the model using a simple bottom-up procedure. The non-terminal nodes in the tree are traversed in bottom-up order. For each node  $k$ , draw  $x_k \sim \text{Uniform}(0, 1)$ . If  $x_k < p_k$  and  $t_{\ell(k)} = t_{r(k)} = 1$ , then  $t_k \leftarrow 1$ . Otherwise,  $t_k \leftarrow 0$ . This setting of  $\mathbf{t}$  defines a segmentation that is a sample from the model (see Figure 2 for a example  $\mathbf{t}$  and the corresponding segmentation).

### 3.4 TREE CONSTRUCTION

The discussion above assumed the tree structure was given, as could be provided by some existing segmentation tree algorithm such as UCM (Arbelaez et al., 2011) or IS CRA (Ren & Shakhnarovich, 2013). Here we extend our formulation to also construct a tree for a given image. Our tree construction algorithm, called SST-SWAPTREE, begins with an arbitrary tree structure and seeks to iteratively improve it by making moves within the space of tree structures. We consider swaps, in which the subtree rooted at a node in the tree is exchanged with the subtree rooted at another node. Refer to Figure 4 for an illustration of a sample swap. Alternate strategies for constructing trees, such as agglomerative clustering, are possible but we opt to iteratively improve tree structures rather than building them from scratch for two main reasons. First, maintaining valid tree structures makes it easier to track performance throughout the tree building process. Second, it is simpler to estimate the effect of a move on the entire tree structure because the global con-




Figure 4: An example swap. **Left:** The tree structure before swapping B and C. **Right:** The structure after the swap. Note that when non-leaf nodes are swapped, their descendants also accompany them in the swap.

text is known, as opposed to agglomerative clustering in which part of the tree structure does not yet exist.

At each iteration of SST-SWAPTREE, the first step involves enumerating allowable swaps. A swap between nodes is allowable if neither is a descendant of the other, as otherwise the resulting structure is ill-defined. It also must preserve the property that children of any non-terminal node in the tree are neighboring. Both of these properties can be checked efficiently (see the supplementary material for details).

The next step involves scoring a swap. The key idea in this step is that any tree can be evaluated based on the minimum achievable Hellinger distance between the tree and the ground-truth segmentations. Hence a swap can be scored based on how much it lowers this minimum distance. At training time we have access to the ground truth segmentations, and so the minimum achievable Hellinger distance for a given structure can be easily computed via numerical minimization of  $D(\pi, q)$  with respect to  $\pi$ . At test time however this is not possible, so instead we utilize our learned network that estimates node marginals to approximate the change in achievable Hellinger distance effected by a swap, which determines its score. Every iteration of the algorithm greedily selects the highest scoring swap and produces a new tree structure. This procedure is repeated for a fixed number of iterations chosen based on validation performance. Space does not permit a detailed description here of our SST-SWAPTREE algorithm for enumerating and scoring swaps, and so further details can be found in the supplementary material.

## 4 EXPERIMENTS

In order to assess the effectiveness of SSTs in producing multiple plausible segmentations, we learned models on two segmentation datasets: the Berkeley Segmentation Dataset 500 (BSDS500) (Arbelaez et al., 2011) and the Penn-Fudan pedestrian parsing dataset (Bo & Fowlkes, 2011; Wang et al., 2007). In this section, we first describe the datasets we used. We then provide details of

our model training procedure and show experimental results.

### 4.1 DATASETS

The BSDS500 dataset is a natural fit for our task because it contains multiple human-annotated ground truth segmentations for each image. The Penn-Fudan pedestrian dataset is a semantic image dataset that we adapt for our purposes by generating four synthetic GT segmentations per example, each of which represents a different prototypical parse of the image. We generate these GTs based on merging different body parts in different segmentations. Details of the merging algorithm, along with visualizations of several ground truths and corresponding images are shown in the supplementary material.

The BSDS500 dataset contains predefined training, validation, and test splits of 200, 100, and 200 examples, respectively. The Penn-Fudan dataset contains only 169 examples, so we randomly created five splits with 20% of the examples as test and the remainder divided into training and validation in an 80%-20% ratio. The reported evaluation results for this dataset are averaged over test predictions across the five splits.

### 4.2 SST TRAINING DETAILS

We took the regions output by the UCM algorithm, prior to thresholding, as the leaves of our trees. For BSDS500 this yielded a large number of leaves per image, so we thresholded the weighted contour map for each example independently such that 100 regions remained. All methods, including baselines, used the same base-level regions.

Our network for estimating marginal probabilities used two sets of features as input. The first are visual features computed by average-pooling activations from the fourth convolutional layer of the VGG-19 network (Simonyan & Zisserman, 2014). The second set were geometric features regarding the regions. Both sets of features have the benefit that the features of a parent region can efficiently be computed given features of the child regions.

The marginal probability estimation network was trained with SGD using the Adam optimizer (Kingma & Ba, 2014). Early stopping was done on the Hellinger distance of the model merge probability assigned trees on the validation set. Dropout was used on the first two layers of the network for regularization. Further training and architectural details are contained in the supplementary material.

### 4.3 RESULTS

We evaluate the learned SST models along with several baselines, which we describe here. We distinguish between methods that produce a single output and those that generate multiple outputs.

Of the single output baselines, the simplest is SUPERPIXELS, the segmentation consisting of the starting base-level regions themselves. BOYKOV-JOLLY is a baseline inspired by Boykov & Jolly (2001) in which strengths between regions are computed as a monotonically increasing function of the squared difference in mean intensity between them. This induces a tree structure, which can be thresholded at any strength in order to form a segmentation. UCM SINGLE is the standard single segmentation produced by UCM, based on thresholding the hierarchy at the threshold yielding the best performance on the validation set.

For multiple-output baselines, we consider an extension of UCM SINGLE, in which a Gaussian distribution over thresholds is constructed. Multiple segmentations can be generated by first sampling a threshold from this Gaussian, and then generating a segmentation by merging regions until this threshold is reached. The mean of this Gaussian is set to the corresponding UCM SINGLE threshold and the variance is tuned on the validation set. We call this extension UCM MULTIPLE.

We compare these baselines to two methods utilizing SSTs. SST-UTREE is our model trained using the UCM tree structures as described in Section 4.2. SST-SWAPTREE is a modified version of the procedure described in Section 3.4 to iteratively improve the trees generated by UCM. The only modification we made was to use the true node marginals when choosing swaps to make. As such these tree structures are oracle trees in that they can only be computed if the optimal  $\pi_k^*$  is known at test time. This tests whether our tree construction procedure produces better trees, provided the  $\hat{\pi}$  estimator is sufficiently accurate. The final Hellinger distance computed on the tree structures does however use the network marginal outputs to assign probabilities to each node in the procedure described in Section 3.3.

We also evaluate the methods on a multiple output variant of the Rand Index that we call *matching Rand Index (MRI)*. For each method a set of candidate segmentations is first generated by sampling from the model. Then an optimal matching is found between these and the ground-truth segmentations; this can easily be computed via a max-weighted bipartite matching algorithm. Each ground truth matches the segmentation from the candidate set with the greatest Rand Index with respect to itself, subject to the constraint that no two ground truths

Table 1: Test Hellinger distances (lower is better). The first methods all output a single segmentation, while the others output several.

Method	BSDS 500	Penn-Fudan
SUPERPIXELS	0.1451	0.3485
BOYKOV-JOLLY	0.1447	0.3172
UCM SINGLE	0.1315	0.2974
UCM MULTIPLE	0.1315	0.2971
SST-UTREE	0.0864	0.2070
SST-SWAPTREE	<b>0.0689</b>	<b>0.1901</b>

Table 2: Test matching Rand Index (higher is better).

Method	Samples	BSDS 500	Penn-Fudan
SUPERPIXELS	1	0.8107	0.6366
BOYKOV-JOLLY	1	0.8111	0.6612
UCM SINGLE	1	0.8247	0.6813
UCM MULTIPLE	50	0.8687	0.7037
SST-UTREE	50	0.8786	0.7160
SST-SWAPTREE	50	<b>0.9116</b>	<b>0.7726</b>
UCM MULTIPLE	100	0.8718	0.7072
SST-UTREE	100	0.8813	0.7180
SST-SWAPTREE	100	<b>0.9148</b>	<b>0.7792</b>

match the same model output. The MRI is then the mean Rand Index of the matched segmentations. The MRI results are shown in Table 2. For methods that output a single segmentation we report its average Rand Index to the ground-truths; MRI simplifies to probabilistic Rand Index (PRI) (Unnikrishnan & Hebert, 2005) in this setting. For the multiple output methods we include results with both 50 and 100 samples drawn from the respective models. SST significantly outperforms the baselines here even though it was trained to optimize Hellinger distance, demonstrating that Hellinger is a sensible objective for training models in a multiple output setting. Also, the SST-SWAPTREE MRI shows the gains obtained by improving the tree structure.

Segmentations sampled from the SST-UTREE model are visualized in Figure 5 for BSDS500. We observe that the samples vary across regions of the image, capturing the idea of generating segmentations at different levels of granularity.

### 4.4 APPLICATION OF SST: SEMANTIC SEGMENTATION

The stochastic segmentation tree provides an informative representation of the image structure. The tree, with each node assigned its own well-calibrated probability of the associated merge, describes a rich distribution over




Figure 5: Samples from the SST-UTREE model on BSDS500. From top to bottom: image, samples 1-4 from the model. All images are from the test set.

segmentations. In addition, there are a number of potential applications of SSTs. Any algorithm  $\mathcal{A}$  that uses superpixels as inputs could benefit from this rich representation. If  $\mathcal{A}$  wants  $k$  superpixels, the SST can produce multiple instances with approximately the same number, but with varying granularity in different regions of the image. Other existing tree-based methods do not have the same ability. For example, UCM generates alternative segmentations, but these obey a fixed ordering, so they cannot produce higher granularity in one region and lower in another in one segmentation, and vice versa in a second segmentation. Running  $\mathcal{A}$  on each of these superpixel inputs provides a natural measure of uncertainty in the algorithm’s output.

Here we explore one such algorithm, semantic segmentation, that utilizes superpixels as inputs. In semantic segmentation, the goal is to assign each pixel in an image one of a fixed set of labels. We consider a scenario in which an SST is used to generate multiple segmentations, which are used as superpixel input to a semantic segmentation algorithm. In order to decouple the effects of the specific segmentation algorithm from the performance of the SST as a superpixel generator, we consider an oracle setting in which each region is assigned a single label that gives the highest segmentation accuracy as measured by Hamming distance or intersection-over-union (IOU) between the semantic segmentation and the

ground truth.

Our experiments use the Penn-Fudan dataset, as they already contain semantic segmentation labels. We varied the sigmoid output bias of our SST-UTREE model in order to encourage it to generate superpixelations with various numbers of regions per image. We chose the biases to produce roughly 30, 50, 80, and 100 superpixels per image, and computed both the oracle segmentation accuracy and number of segments of 1,000 samples from our model on each test example. The results are shown in Figure 6 for Hamming distance and in Figure 7 for mean intersection-over-union (IOU). IOU was computed for each of the 12 non-background classes in the Penn-Fudan dataset and then averaged across classes to produce mean IOU. For comparison, we also display the oracle segmentation accuracy of the UCM SINGLE algorithm over a range of thresholds and the BOYKOV-JOLLY baseline. The oracle accuracy of SST-UTREE is higher for each of these scenarios than either UCM SINGLE or BOYKOV-JOLLY. In order to illustrate how the sampled segmentations vary, Figure 8 shows sampled oracle segmentations from the SST-UTREE model that produces roughly 50 superpixels per image.






Figure 6: Oracle segmentation accuracy (Hamming) for SST-UTREE, UCM-SINGLE, and BOYKOV-JOLLY on the Penn-Fudan dataset.

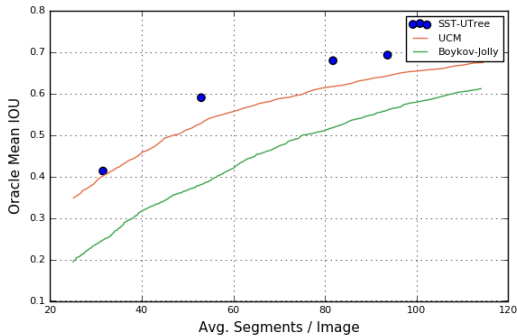


Figure 7: Oracle segmentation accuracy (mean IOU) for SST-UTREE, UCM-SINGLE, and BOYKOV-JOLLY on the Penn-Fudan dataset.

## 5 CONCLUSION & FUTURE WORK

We have presented a model for producing multiple segmentations. The model extends current hierarchical segmentation approaches by formulating a probabilistic model using its structure. The probabilistic objective, based on the Hellinger distance, provides a fairly simple metric for evaluating and training this model. We formulate a learning approach to estimate merge probabilities for the internal nodes of a given tree. We then utilize this same estimator in order to optimize the tree structure, based on a sequence of sub-tree swaps, which effectively improves our loss function. We tested the model against two strong baselines, the standard Boykov-Jolly as well as a top current segmentation approach. The model provides a significant win on two datasets, including the challenging BSDS500. We also showed its utility in semantic segmentation, as a method of forming superpixels at different levels of granularity.

We are currently investigating a number of directions, including further improvements in the tree construction




Figure 8: Oracle segmentation visualizations. From left to right: image, ground truth segmentation, UCM-SINGLE oracle segmentation, and oracle segmentations from three samples from SST-UTREE. Superpixels were generated from the  $\sim 50$  superpixels per image setting for both UCM-SINGLE and SST-UTREE. Superpixels are outlined in red.

procedure. If the complexity of using multiple tree structures can be managed, this should provide a performance boost as well. Multiple trees would allow the model to handle diversity in the segmentations that go beyond that captured in a single tree, such as alternative groupings of ambiguous regions. In addition, allowing the merge probability estimator to take into account more information about regions, such as the level of the tree, and features at multiple levels, may also be beneficial.

The proper treatment of ambiguity is important to the future progress in AI. We believe that predicting the structure and parameters of a graphical model given an input, as we have done here for image segmentation, is applicable to other domains as well. Such models facilitate the minimization of distance to a target distribution constructed from multiple ground truths, which when appropriately built according to the task is a sensible objective. We hypothesize that this approach can be used to improve both training and evaluation of models in a wide range of important domains, particularly in tasks involving ambiguity in computer vision and natural language processing.

## Acknowledgements

We thank Chris Williams and Raquel Urtasun for their helpful comments. This research was supported by Samsung and the Natural Sciences and Engineering Research Council of Canada.

## References

- Arbelaez, P., Maire, M., Fowlkes, C., and Malik, J. Contour Detection and Hierarchical Image Segmentation. *IEEE PAMI*, 2011.
- Arbelaez, P., Pong-Tuset, J., Barron, J., Marques, F., and Malik, J. Multiscale Combinatorial Grouping. In *CVPR*, 2014.
- Batra, Dhruv, Yadollahpour, Payman, Guzman-Rivera, Abner, and Shakhnarovich, Gregory. Diverse M-Best Solutions in Markov Random Fields. In *ECCV*, 2012.
- Bo, Y. and Fowlkes, C. Shape-based Pedestrian Parsing. In *CVPR*, 2011.
- Boykov, Y. and Jolly, M.-P. Interactive Graph Cuts for Optimal Boundary & Region Segmentation of Objects in N-D Images. In *ICCV*, 2001.
- Dey, Debadepta, Ramakrishna, Varun, Hebert, Martial, and Bagnell, J Andrew. Predicting Multiple Structured Visual Interpretations. In *ICCV*, 2015.
- Gillenwater, Jennifer, Kulesza, Alex, and Taskar, Ben. Discovering Diverse and Salient Threads in Document Collections. In *EMNLP*, 2012.
- Guzman-Rivera, A., Batra, D., and Kohli, P. Multiple Choice Learning: Learning to Produce Multiple Structured Outputs. In *NIPS*, 2012.
- Hazan, Tamir, Keshet, Joseph, and McAllester, David A. Direct Loss Minimization for Structured Prediction. In *NIPS*, 2010.
- Hou, X., Yuille, A., and Koch, C. Boundary Detection Benchmarking: Beyond F-Measures. In *CVPR*, 2013.
- Hu, Shell X, Williams, Christopher KI, and Todorovic, Sina. Tree-Cut for Probabilistic Image Segmentation. *arXiv preprint arXiv:1506.03852*, 2015.
- Ion, A, Carreira, J, and Sminchisescu, C. Image Segmentation by Figure-Ground Composition into Maximal Cliques. *ICCV*, 2011.
- Jain, V., Turaga, S., Briggman, K., Helmstaedter, M., Denk, W., and Seung, H. S. Learning to Agglomerate Superpixel Hierarchies. In *NIPS*, 2011.
- Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kirillov, Alexander, Savchynskyy, Bogdan, Schlesinger, Dmitriy, Vetrov, Dmitry, and Rother, Carsten. Inferring M-Best Diverse Labelings in a Single One. In *ICCV*, 2015.
- Kulesza, Alex, Taskar, Ben, et al. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*, 5(2-3):123-286, 2012.
- Lampert, C. Maximum Margin Multi-Label Structured Prediction. In *NIPS*, 2011.
- Liu, Dong C and Nocedal, Jorge. On the Limited Memory BFGS Method for Large Scale Optimization. *Mathematical programming*, 45(1-3):503-528, 1989.
- Maire, Michael, Arbeláez, Pablo, Fowlkes, Charless, and Malik, Jitendra. Using Contours to Detect and Localize Junctions in Natural Images. In *CVPR*, 2008.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In *ICCV*, 2001.
- Papineni, Kishore, Roukos, Salim, Ward, Todd, and Zhu, Weijng. BLEU: a Method for Automatic Evaluation of Machine Translation. In *ACL*, 2002.
- Rand, W. Objective Criteria for the Evaluation of Clustering Methods. *JASA*, 12 1971.
- Ranzato, Marc'Aurelio, Chopra, Sumit, Auli, Michael, and Zaremba, Wojciech. Sequence Level Training with Recurrent Neural Networks. *arXiv preprint arXiv:1511.06732*, 2015.
- Ren, Z. and Shakhnarovich, G. Image Segmentation by Cascaded Region Agglomeration. In *CVPR*, 2013.
- Sharon, E., Galun, M., Sharon, D., Basri, R., and Brandt, A. Hierarchy and Adaptivity in Segmenting Visual Scenes. *Nature*, 442(7104):810-813, August 2006.
- Shen, Shiqi, Cheng, Yong, He, Zhongjun, He, Wei, Wu, Hua, Sun, Maosong, and Liu, Yang. Minimum Risk Training for Neural Machine Translation. *arXiv preprint arXiv:1512.02433*, 2015.
- Simonyan, Karen and Zisserman, Andrew. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Song, Yang, Schwing, Alexander G, Zemel, Richard S, and Urtasun, Raquel. Direct Loss Minimization for Training Deep Neural Nets. *arXiv preprint arXiv:1511.06411*, 2015.
- Tromble, Roy W, Kumar, Shankar, Och, Franz, and Macherey, Wolfgang. Lattice Minimum Bayes-Risk Decoding for Statistical Machine Translation. In *EMNLP*, 2008.
- Unnikrishnan, Ranjith and Hebert, Martial. Measures of Similarity. In *IEEE ACV*, 2005.
- Vijayakumar, Ashwin K, Cogswell, Michael, Selvaraju, Ramprasad R, Sun, Qing, Lee, Stefan, Crandall, David, and Batra, Dhruv. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*, 2016.
- Wang, L., Shi, J., Song, G., and Shen, I. Object Detection Combining Recognition and Segmentation. In *ACCV 2007*, 2007.
- Wong, Yuk Wah and Mooney, Raymond J. Learning for Semantic Parsing with Statistical Machine Translation. In *ACL*, 2006.
- Yadollahpour, Payman, Batra, Dhruv, and Shakhnarovich, Gregory. Discriminative Re-ranking of Diverse Segmentations. In *CVPR*, 2013.