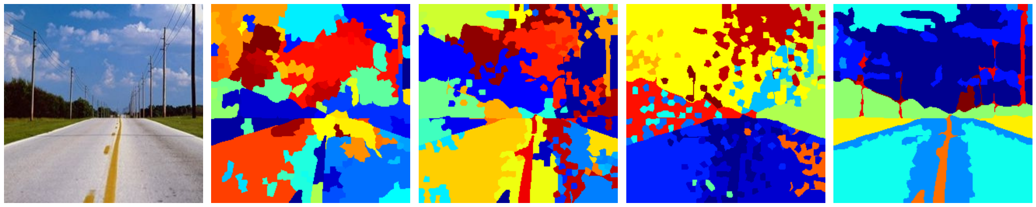

Figure 2: Comparison of distance-dependent segmentation priors. From left to right, we show segmentations produced by the ddCRP with $a = 1$,the ddCRP with $a = 2$, the ddCRP with $a = 5$, and the rddCRP with $a = 1$.

superpixel from another, with hops being allowed only amongst spatially neighboring superpixels. A "window" decay function of width a, $f(d) = \mathbb{1}[d \leq a]$, determines link probabilities. If $a = 1$, superpixels can only directly connect to adjacent superpixels. Note this does not explicitly restrict the size of segments, because any pair of pixels for which one is *reachable* from the other (i.e., in the same connected component of the customer assignment graph) are in the same image segment. For this special case segments are *guaranteed* with probability one to form spatially connected subsets of the image, a property not enforced by other Bayesian nonparametric models [10, 11, 12].

The full generative process for the observed features $x_{1:N}$ within a N-superpixel image is as follows:

1. For each table, sample parameters $\phi_k \sim G_0$.
2. For each customer, sample a customer assignment $c_i \sim \text{ddCRP}(\alpha, f, D)$. This indirectly determines the cluster assignments $z_{1:N}$, and thus the segmentation.
3. For each superpixel, independently sample observed data $x_i \sim P(\cdot \mid \phi_{z_i})$.

The customer assignments are sampled using the spatial distance between pixels. The partition structure, derived from the customer assignments, is used to sample the observed image features. Given an image, the posterior distribution of the customer assignments induces a posterior over the cluster structure; this provides the segmentation. See Figure 1 for an illustration of the customer assignments and their derived table assignments in a segmentation setting.

As in [10], the data generating distribution for the observed features studied in Section 4 is multinomial, with separate distributions for color and texture. We place conjugate Dirichlet priors on these cluster parameters.

2.3 Region-based hierarchical distance dependent CRPs

The ddCRP model, when applied to an image with window size $a = 1$, produces a collection of contiguous patches (tables) homogeneous in color and texture features (Figure 2). While such segmentations are useful for various applications [16], they do not reflect the statistics of manual human segmentations, which contain larger regions [17]. We could bias our model to produce such regions by either increasing the window size a, or by introducing a hierarchy wherein the produced patches are grouped into a small number of regions. This region level model has each patch (table) associated with a region k from a set of potentially infinite regions. Each region in turn is associated with an appearance model ϕ_k. The corresponding generative process is described as follows:

1. For each customer, sample customer assignments $c_i \sim \text{ddCRP}\,(\alpha, f, D)$. This determines the table assignments $t_{1:N}$.
2. For each table t, sample region assignments $k_t \sim \text{CRP}\,(\gamma)$.
3. For each region, sample parameters $\phi_k \sim G_0$.
4. For each superpixel, independently sample observed data $x_i \sim P\,(\cdot \mid \phi_{z_i})$, where $z_i = k_{t_i}$.

Note that this region level rddCRP model is a direct extension of the Chinese restaurant franchise (CRF) representation of the HDP [5], with the image partition being drawn from a ddCRP instead of a CRP. In contrast to prior applications of the HDP, our region parameters are not shared amongst images, although it would be simple to generalize to this case. Figure 3 plots samples from the rddCRP and ddCRP priors with increasing a. The rddCRP produces larger partitions than the ddCRP with $a = 1$, while avoiding the noisy boundaries produced by a ddCRP with large a (see Figure 2).