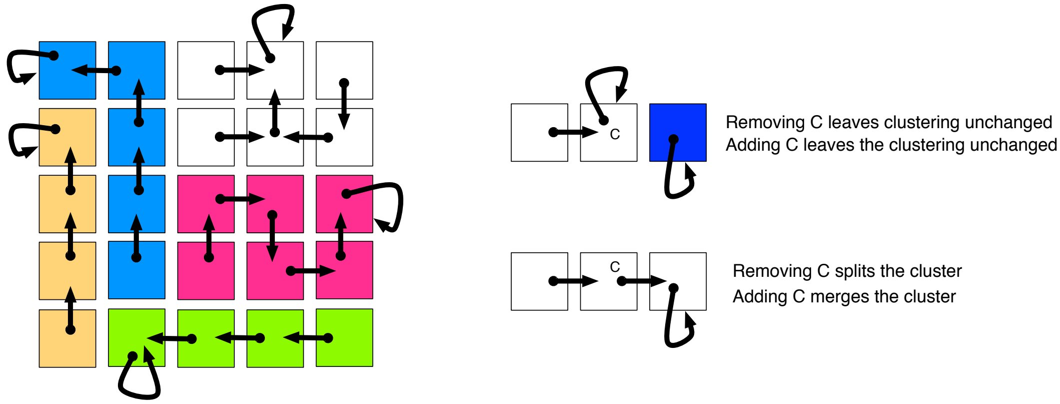


Figure 1: *Left:* An illustration of the relationship between the customer assignment representation and the table assignment representation. Each square is a data point (a pixel or superpixel) and each arrow is a customer assignment. Here, the distance window is of length 1. The corresponding table assignments, i.e., the clustering of these data, is shown by the color of the data points. *Right:* Intuitions behind the two cases considered by the Gibbs sampler. Consider the link from node $C$. When removed, it may leave the clustering unchanged or split a cluster. When added, it may leave the clustering unchanged or merge two clusters.

we will not attempt to survey here. Influential existing methods include approaches based on kernel density estimation [6], Markov random fields [3, 7], and the normalized cut spectral clustering algorithm [8, 9]. A recurring difficulty encountered by traditional methods is the need to determine an appropriate segment resolution for each image; even among images of similar scene types, the number of observed objects can vary widely. This has usually been dealt via heuristics with poorly understood biases, or by simplifying the problem (e.g., partially specifying each image's segmentation via manual user input [7]).

Recently, several promising segmentation algorithms have been proposed based on nonparametric Bayesian methods [10, 11, 12]. In particular, an approach which couples Pitman-Yor mixture models [13] via thresholded Gaussian processes [14] has lead to very promising initial results [10], and provides a baseline for our later experiments. Expanding on the experiments in [10], we analyze 800 images of different natural scene types, and show that the comparatively simpler ddCRP-based algorithms perform similarly to this work. Moreover, unlike previous nonparametric Bayesian approaches, the structure of the ddCRP allows spatial connectivity of the inferred segments to (optionally) be enforced. In some applications, this is a known property of all reasonable segmentations.

Our results demonstrate the practical utility of spatial ddCRP and hierarchical rddCRP models. We also provide the first rigorous comparison of nonparametric Bayesian image segmentation models.

## 2 Image Segmentation with Distance Dependent CRPs

Our goal is to develop a probabilistic method to segment images of complex scenes. Image segmentation is the problem of partitioning an image into self-similar groups of adjacent pixels. Segmentation is an important step towards other tasks in image understanding, such as object recognition, detection, or tracking. We model images as observed collections of "superpixels" [15], which are small blocks of spatially adjacent pixels. Our goal is to associate the features $x_i$ in the $i^{th}$ superpixel with some cluster $z_i$; these clusters form the segments of that image.

Image segmentation is thus a special kind of clustering problem where the desired solution has two properties. First, we hope to find contiguous regions of the image assigned to the same cluster. Due to physical processes such as occlusion, it may be appropriate to find segments that contain two or three contiguous image regions, but we do not want a cluster that is scattered across individual image pixels. Traditional clustering algorithms, such as $k$-means or probabilistic mixture models, do not account for external information such as pixel location and are not biased towards contiguous regions. Image locations have been heuristically incorporated into Gaussian mixture models by concatenating positions with appearance features in a vector [16], but the resulting bias towards elliptical regions often produces segmentation artifacts. Second, we would like a solution that determines the number