



Figure 2: *Left:* A reference mesh in which links (yellow arrows) currently define three parts (connected components). *Right:* Each part undergoes a distinct affine transformation, generated as in Equation (2).

The final seating arrangement gives a partition of the data, where each occupied table corresponds to a part in the final segmentation.

Although described sequentially, the CRP induces an exchangeable distribution on partitions, for which the segmentation probability is invariant to the order in which triangle allocations are sampled. This is inappropriate for mesh data, in which nearby triangles are far more likely to lie in the same part. The ddCRP alters the CRP by modeling customer links not to tables, but to other customers. The link c_m for customer m is sampled according to the distribution

$$p(c_m = n \mid D, f, \alpha) \propto \begin{cases} f(d_{mn}) & m \neq n, \\ \alpha & m = n. \end{cases} \quad (1)$$

Here, d_{mn} is an externally specified distance between data points m and n , and α determines the probability that a customer links to themselves rather than another customer. The monotonically decreasing decay function $f(d)$ mediates how the distance between two data points affects their probability of connecting to each other. The overall link structure specifies a partition: two customers are clustered together if and only if one can reach the other by traversing the link edges.

We define the distance between two triangles as the minimal number of hops, between adjacent faces, required to reach one triangle from the other. A “window” decay function of width 1, $f(d) = \mathbf{1}[d \leq 1]$, then restricts triangles to link only to immediately adjacent faces. Note that this doesn’t limit the size of parts, since all pairs of faces are potentially reachable via a sequence of adjacent links. However, it does guarantee that only spatially contiguous parts have non-zero probability under the prior. This constraint is preserved by our MCMC inference algorithm.

2.2 Modeling Part Deformation via Affine Transformations

Articulated object deformation is naturally described via the spatial transformations of its constituent parts. We expect the triangular faces within a part to deform according to a coherent part-specific transformation, up to independent face-specific noise. The near-rigid motions of interest are reasonably modeled as affine transformations, a family of co-linearity preserving linear transformations. We concisely denote the transformation from a reference triangle to an observed triangle via a matrix $A \in \mathbb{R}^{3 \times 4}$. The fourth column of A encodes translation of the corresponding reference triangle via homogeneous coordinates x_{bn} , and the other entries encode rotation, scaling, and shearing.

Previous approaches have treated such transformations as parameters to be estimated during inference [8, 9]. Here, we instead define a prior distribution over affine transformations. Our construction allows transformations to be analytically marginalized when learning our part-based segmentation, but retains the flexibility to later estimate transformations if desired. Explicitly modeling transformation uncertainty makes our MCMC inference more robust and rapidly mixing [7], and also allows data-driven determination of an appropriate number of parts.

The matrix of numbers encoding an affine transformation is naturally modeled via multivariate Gaussian distributions. We place a conjugate, matrix normal-inverse-Wishart [10, 11] prior on the affine transformation A and residual noise covariance matrix Σ :

$$\begin{aligned} \Sigma &\sim \mathcal{IW}(n_0, S_0) \\ A \mid \Sigma &\sim \mathcal{MN}(M, \Sigma, K) \end{aligned} \quad (2)$$