Here, $n_0 \in \mathbb{R}$ and $S_0 \in \mathbb{R}^{3\times 3}$ control the variance and mean of the Wishart prior on Σ^{-1}. The mean affine transformation is $M \in \mathbb{R}^{3\times 4}$, and $K \in \mathbb{R}^{4\times 4}$ and Σ determine the variance of the prior on A. Applied to mesh data, these parameters have physical interpretations and can be estimated from the data collection process. While such priors are common in Bayesian regression models, our application to the modeling of geometric affine transformations appears novel.

Allocating a different affine transformation for the motion of each part in each pose (Figure 2), the overall generative model can be summarized as follows:

1. For each triangle n, sample an associated link $c_n \sim$ ddCRP (α, f, D). The part assignments z are a deterministic function of the sampled links $c = [c_1, \dots, c_N]$.
2. For each pose j of each part k, sample an affine transformation A_{jk} and residual noise covariance Σ_{jk} from the matrix normal-inverse-Wishart prior of Equation (2).
3. Given these pose-specific affine transformations and assignments of mesh faces to parts, independently sample the observed location of each pose triangle relative to its corresponding reference triangle, $y_{jn} \sim \mathcal{N}(A_{jz_n} x_{b_j n}, \Sigma_{jz_n})$.

Note that Σ_{jk} governs the degree of non-rigid deformation of part k in pose j. It also indirectly influences the number of inferred parts: a large S_0 makes large Σ_{jk} more probable, which allows more non-rigid deformation and permits models which utilize fewer parts. The overall model is

$$p(\mathbf{Y}, c, A, \Sigma \mid \mathbf{X}, b, D, \alpha, f, \eta) = p(c \mid D, f, \alpha) \prod_{j=1}^{J} \left[\prod_{k=1}^{K(c)} p(A_{jk}, \Sigma_{jk} \mid \eta) \right] \left[\prod_{n=1}^{N} \mathcal{N}(y_{jn} \mid A_{jz_n} x_{b_j n}, \Sigma_{jz_n}) \right] \quad (3)$$

where $\mathbf{Y} = \{Y_1, \dots, Y_J\}$, $\mathbf{X} = \{X_1, \dots, X_B\}$, $b = [b_1, \dots, b_J]$, the ddCRP links c define assignments z to $K(c)$ parts, and $\eta = \{n_0, S_0, M, K\}$ are likelihood hyperparameters. There is a single reference mesh X_b for each object instance b, and Y_j captures a single deformed pose of X_{b_j}.

2.3 Previous Work

Previous work has also sought to segment a mesh into parts based on observed articulations [8, 12, 13, 14]. The two-stage procedure of Rosman et al. [13] first minimizes a variational functional regularized to favor piecewise constant transformations, and then clusters the transformations into parts. Several other segmentation procedures [12, 14] lack coherent probabilistic models, and thus have difficulty quantifying uncertainty and determining appropriate segmentation resolutions.

Anguelov et al. [8] define a global probabilistic model, and use the EM algorithm to jointly estimate parts and their transformations. They explicitly model spatial dependencies among mesh faces, but their Markov random field cannot ensure that parts are spatially connected; a separate connected components process is required. Heuristics are used to determine an appropriate number of parts.

Ambitious recent work has considered a model for joint mesh alignment and segmentation [9]. However, this approach suffers from many of the issues noted above: the number of parts must be specified *a priori*, parts may not be contiguous, and their EM inference appears prone to local optima.

3 Inference

We seek the constituent parts of an articulated model, given observed data ($\mathbf{X}$, $\mathbf{Y}$, and b). These parts are characterized by the posterior distribution of the customer links c. We approximate this posterior using a collapsed Gibbs sampler, which iteratively draws c_n from the conditional distribution

$$p(c_n \mid c_{-n}, \mathbf{X}, \mathbf{Y}, b, D, f, \alpha, \eta) \propto p(c_n \mid D, f, \alpha) p(\mathbf{Y} \mid z(c), \mathbf{X}, b, \eta). \quad (4)$$

Here, $z(c)$ is the clustering into parts defined by the customer links c. The ddCRP prior is given by Equation (1), while the likelihood term in the above equation further factorizes as

$$p(\mathbf{Y} \mid z(c), \mathbf{X}, b, \eta) = \prod_{k=1}^{K(c)} \prod_{j=1}^{J} p(Y_{jk} \mid X_{b_j k}, \eta) \quad (5)$$