



Figure 1: Human body segmentation. *Left*: Reference poses for two female bodies, and those bodies captured in five other poses. *Right*: A manual segmentation used to align these meshes [6], and the segmentation inferred by our ddCRP model from 56 poses. The ddCRP segmentation discovers parts whose motion is nearly rigid, and includes small parts such as elbows and knees absent from the manual segmentation.

tion. To our knowledge, no previous methods for segmenting meshes combine information about deformation from multiple bodies to address this *corpus segmentation* problem.

In this paper, we develop a statistical model which addresses all of these issues. We adapt the *distance dependent Chinese restaurant process* (ddCRP) [4] to model spatial dependencies among mesh triangles, and enforce spatial contiguity of the inferred parts [5]. Unlike most previous mesh segmentation methods, our Bayesian nonparametric approach allows data-driven inference of an appropriate number of parts, and uses an affine transformation-based likelihood to accommodate object instances of varying shape. After developing our model in Section 2, Section 3 develops a Gibbs sampler which efficiently marginalizes the latent affine transformations defining part deformation. We conclude in Section 4 with results examining meshes of humans and other articulated objects, where we introduce a metric for quantitative evaluation of deformation-based segmentations.

2 A Part-Based Model for Mesh Deformation

Consider a collection of J meshes, each with N triangles. For some input mesh j , we let $y_{jn} \in \mathbb{R}^3$ denote the 3D location of the center of triangular face n , and $Y_j = [y_{j1}, \dots, y_{jN}] \in \mathbb{R}^{3 \times N}$ the full mesh configuration. Each mesh j has an associated N -triangle reference mesh, indexed by b_j . We let $x_{bn} \in \mathbb{R}^4$ denote the location of triangle n in reference mesh b , expressed in homogeneous coordinates ($x_{bn}(4) = 1$). A full reference mesh $X_b = [x_{b1}, \dots, x_{bN}]$. In our later experiments, Y_j encodes the 3D mesh for a person in pose j , and X_{b_j} is the reference pose for the same individual.

We estimate aligned correspondences between the triangular faces of the input pose meshes Y_j , and the reference meshes X_b , using a recently developed method [6]. This approach robustly handles 3D data capturing varying shapes and poses, and outputs meshes which have equal numbers of faces in one-to-one alignment. Our segmentation model does not depend on the details of this alignment method, and could be applied to data produced by other correspondence algorithms.

2.1 Nonparametric Spatial Priors for Mesh Partitions

The recently proposed distance dependent Chinese restaurant process (ddCRP) [4], a generalization of the CRP underlying Dirichlet process mixture models [7], has a number of attractive properties which make it particularly well suited for modeling segmentations of articulated objects. By placing prior probability mass on partitions with arbitrary numbers of parts, it allows data-driven inference of the true number of mostly-rigid parts underlying the observed data. In addition, by choosing an appropriate distance function we can encourage spatially adjacent triangles to lie in the same part, and *guarantee* that all inferred parts are spatially contiguous [5].

The Chinese restaurant process (CRP) is a distribution on all possible partitions of a set of objects (in our case, mesh triangles). The generative process can be described via a restaurant with an infinite number of tables (in our case, parts). Customers (triangles) i enter the restaurant in sequence and select a table z_i to join. They pick an occupied table with probability proportional to the number of customers already sitting there, or a new table with probability proportional to a scaling parameter α .