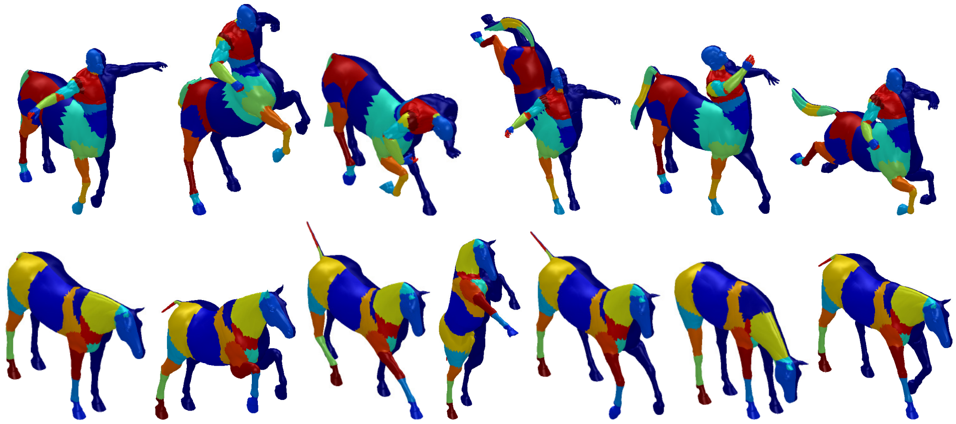

Figure 3: Segmentations produced by mesh-ddcrp on synthetic Tosca meshes [20]. The first mesh in each row displays the chosen reference mesh. For illustration, we have only segmented the right half of each mesh.

Here, A^*_{tk} is the least squares estimate of the single affine transformation responsible for mapping $X_{b_t k}$ to Y_{tk}. Note that Equation (12) is trivially zero for a degenerate solution wherein each mesh face is assigned to its own part. However, segmentations of similar resolution may safely be compared using Equation (12), with lower errors corresponding to better segmentations.

On our test set of human meshes, the mesh-ddcrp model produces an error of $\mathcal{E} = 1.39$ meters, which corresponds to sub-millimeter accuracy when normalized by the number of faces. Figure 4 displays a plot comparing the errors achieved by the different methods. Mesh-ddcrp is significantly better than all other methods, including for settings of K which allocate 50% more parts to competing approaches, according to a Wilcoxon's signed rank test (5% significance level).

Next, we demonstrate the benefits of sharing information among differently shaped bodies. We selected an illustrative articulated pose for each of the two training subjects in addition to their respective reference poses (Figure 4). The chosen poses either exhibit upper or lower body deformations, but not both. The meshes were then segmented both independently for the two subjects and jointly sharing information across subjects. Figure 5 demonstrates that the independent segmentations exhibit both undersegmented (legs in the first set) and oversegmented (head in the second) parts. However, sharing information among subjects results in parts which correspond well with physical human bodies. Note that with only two articulated poses, we are able to generate meaningful segmentations in about an hour of computation. This data-limited scenario also demonstrates the benefits of the ddCRP prior: as shown in Figure 5, the parts extracted by mesh-crp are "patchy", spatially disconnected, and physically implausible.

5 Discussion

Adapting the ddCRP to collections of 3D meshes, we have developed an effective approach for the discovery an unknown number of parts underlying articulated object motion. Unlike previous methods, our model guarantees that parts are spatially connected, and uses transformations to model instances with potentially varying body shapes. Via a novel application of matrix normal-inverse-Wishart priors, our sampler analytically marginalizes transformations for improved efficiency. While we have modeled part motion via affine transformations, future work should explore more accurate Lie algebra characterizations of deformation manifolds [21].

Experiments with dozens of real human body poses provide strong quantitative evidence that our approach produces state-of-the-art segmentations with many potential applications. We are currently exploring methods for using multiple samples from the ddCRP posterior to characterize part uncertainty, and scaling our Monte Carlo learning algorithms to datasets containing thousands of meshes.

Acknowledgments This work was supported in part by the Office of Naval Research under contract W911QY-10-C-0172. We thank Eric Rachlin, Alex Weiss, and David Hirshberg for acquiring and aligning the human meshes, and Aggeliki Tsoli for her helpful comments.