where $Y_{jk} \in \mathbb{R}^{3 \times N_k}$ is the set of triangular faces in part $k$ of pose $j$, and $X_{b_j k}$ are the corresponding reference faces. Exploiting the conjugacy of the normal likelihood to the prior over affine transformations in Equation (2), we marginalize the part-specific latent variables $A_{jk}$ and $\Sigma_{jk}$ to compute the marginal likelihood in closed form (see the supplement for a derivation):

$$p(Y_{jk} \mid X_{b_j k}, \eta) = \frac{|K|^{3/2}|S_0|^{(n_0/2)}\Gamma_3\left(\frac{N_k+n_0}{2}\right)}{\pi^{(3N_k/2)}|S_{xx}|^{(3/2)}|S_0+S_{y|x}|^{((N_k+n_0)/2)}\Gamma_3(\frac{n_0}{2})}, \tag{6}$$

$$S_{xx} = X_{b_j k}{X_{b_j k}}^T + K, \quad S_{yx} = Y_{jk}{X_{b_j k}}^T + MK, \tag{7}$$

$$S_{y|x} = Y_{jk}{Y_{jk}}^T + MKM^T - S_{yx}(S_{xx})^{-1}S_{yx}^T. \tag{8}$$

Instead of explicitly sampling from Equation (4), a more efficient sampler [4] can be derived by observing that different realizations of the link $c_n$ only make a small change to the partition structure. First, note that removing a link $c_n$ generates a partition $z(c_{-n})$ which is either identical to the old partition $z(c)$ or contains one extra part, created by splitting some existing part. Sampling new realizations of $c_n$ will give rise to new partitions $z(c_{-n} \cup c_n^{(new)})$, which may either be identical to $z(c_{-n})$ or contain one less part, due to a merge of two existing parts. We thus sample $c_n$ from the following distribution which only tracks those parts which change with different realizations of $c_n$:

$$p(c_n \mid c_{-n}, \mathbf{X}, \mathbf{Y}, b, D, f, \alpha, \eta) \propto \begin{cases} p(c_n \mid D, f, \alpha)\Delta(\mathbf{Y}, \mathbf{X}, b, z(c), \eta) & \text{if } c_n \text{ links } k_1 \text{ and } k_2; \\ p(c_n \mid D, \alpha) & \text{otherwise,} \end{cases}$$

$$\Delta(\mathbf{Y}, \mathbf{X}, b, z(c), \eta) = \frac{\prod_{j=1}^{J} p(Y_{jk_1 \cup k_2} \mid X_{b_j k_1 \cup k_2}, \eta)}{\prod_{j=1}^{J} p(Y_{jk_1} \mid X_{b_j k_1}, \eta) \prod_{j=1}^{J} p(Y_{jk_2} \mid X_{b_j k_2}, \eta)}. \tag{9}$$

Here, $k_1$ and $k_2$ are parts in $z(c_{-n})$. Note that if the mesh segmentation $c$ is the only quantity of interest, the analytically marginalized affine transformations $A_{jk}$ need not be directly estimated. However, for some applications the transformations are of direct interest. Given a sampled segmentation, the part-specific parameters for pose $j$ have the following posterior [10]:

$$p(A_{jk}, \Sigma_{jk} \mid Y_j^k, X^k, \eta) \propto \mathcal{MN}(A_{jk} \mid S_{yx}S_{xx}^{-1}, \Sigma_{jk}, S_{xx})\mathcal{IW}(\Sigma_{jk} \mid N_k + n_0, S_{y|x} + S_0) \tag{10}$$

Marginalizing the noise covariance matrix, the distribution over transformations is then

$$\begin{aligned} p(A_{jk} \mid Y_j^k, X^k, \eta) &= \int \mathcal{MN}(A_{jk} \mid S_{yx}S_{xx}^{-1}, \Sigma_{jk}, S_{xx}) IW(\Sigma_{jk} \mid N_k + n_0, S_{y|x} + S_0)\, d\Sigma_{jk} \\ &= \mathcal{MT}(A_{jk} \mid N_k + n_0, S_{yx}S_{xx}^{-1}, S_{xx}, S_{y|x} + S_0) \end{aligned} \tag{11}$$

where $\mathcal{MT}(\cdot)$ is a matrix-t distribution [11] with mean $S_{yx}S_{xx}^{-1}$, and $N_k + n_0$ degrees of freedom.

# 4 Experimental Results

We now experimentally validate, both qualitatively and quantitatively, our *mesh-ddcrp* model. Because "ground truth" parts are unavailable for the real body pose datasets of primary interest, we propose an alternative evaluation metric based on the prediction of held-out object poses, and show that the mesh-ddcrp performs favorably against competing approaches.

We primarily focus on a collection of 56 training meshes, acquired and aligned [6] from 3D scans of two female subjects in 27 and 29 poses. For quantitative tests, we employ 12 meshes of each of six different female subjects [15] (Figure 4). For each subject, a mesh in a canonical pose is chosen as the reference mesh (Figure 1). These meshes contain about 20,000 faces.

## 4.1 Hyperparameter Specification and MCMC Learning

The hyperparameters that regularize our mesh-ddcrp prior have intuitive interpretations, and can be specified based on properties of the mesh data under consideration. As described in Section 2.1, the ddCRP distances $D$ and $f$ are set to guarantee spatially connected parts. The self-connection parameter is set to a small value, $\alpha = 10^{-8}$, to encourage creation of larger parts.

The matrix normal-inverse-Wishart prior on affine transformations $A_{jk}$, and residual noise covariances $\Sigma_{jk}$, has hyperparameters $\eta = \{n_0, S_0, M, K\}$. The mean affine transformation $M$ is set to