Symmetry Informative and Agnostic Feature Disentanglement for 3D Shapes

Supplementary Material

In the following we provide additional information and details to our paper. In particular, we give additional details about the data generation for BeCoS [17] dataset, discuss implementation details, provide further insights into our boundary loss and show additional qualitative results.

7. BeCoS data generation

χ [53] uses a subsampled version of BeCoS [17] with a train/validation/test split consisting of 990/137/142 shape pairs, we use the full benchmark containing 10185/274/284 shape pairs to train all models and to perform all evaluations.

8. Implementation details

To extract the image features presented in Sec. 3.1, we use a concatenation of DINO-V2 [33] and StableDiffusion [43] foundation models similar to the strategy used in Diff3F [15] and χ [53]. This strategy results in image features of dimension $d = 3968$. The lightweight encoder and decoder described in Sec. 3.1 are implemented as two three-layer MLPs with hidden dimensions 3968. This results in $\mathcal{F}_v^{\mathrm{C}}$ and $\mathcal{F}_v^{\mathrm{S}}$ being of dimension 1 and 3967, respectively. To improve the architecture's autoencoding capabilities, we add skip-connections around the encoder and decoder respectively. Finally, we use $\lambda_1 = 1.0, \lambda_2 = 0.2, \lambda_3 = 10.0, \lambda_4 = 2.0$ as the loss coefficients for the overall loss presented in Sec. 3.2. These coefficients were chosen by hyperparameter search.

9. Further insight into our boundary loss $\mathcal{L}_{\mathrm{bou}}$

In the following, we provide more insights as well as intuitive motivation for our boundary loss $\mathcal{L}_{\mathrm{bou}}$. We assume that the symmetric boundary generally and approximately is a straight line on the surface manifold of a shape. To leverage this knowledge during the training process, we introduce the $\mathcal{L}_{\mathrm{bou}}$ loss. The idea is to regularize the symmetry-informative feature s.t. for each vertex $v \in V$ there should be at least one pair of approximately parallel edges along which the feature does not change. An example can be seen in Fig. 7.

Since we want the symmetric boundary to be a straight line in the surface manifold of the mesh, we first project the neighbouring vertices of each vertex onto the tangent space, as shown in Fig. 8.

10. Additional qualitative results

We include more qualitative results here for intrinsic symmetry detection (Fig. 9), left/right classification (Fig. 10) and shape matching (Fig. 11), to better illustrate the superiority of our proposed framework against other methods.

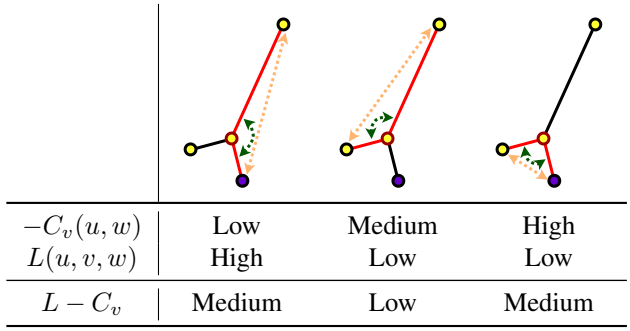

$-C_v(u,w)$	Low	Medium	High
$L(u,v,w)$	High	Low	Low
$L - C_v$	Medium	Low	Medium

Figure 7. For each vertex in the graph induced by the triangle mesh of a shape, we find a pair of edges which minimize the sum of $-C_v$ (green) and L (orange). $-C_v$ is minimal if the edges are close to parallel, while $L(u,v,w)$ is minimal if u, v and w have the same symmetry-informative feature. A low $\mathcal{L}_{\mathrm{bou}}$ loss implies that there exists a straight direction in which the symmetry-informative feature does not change.

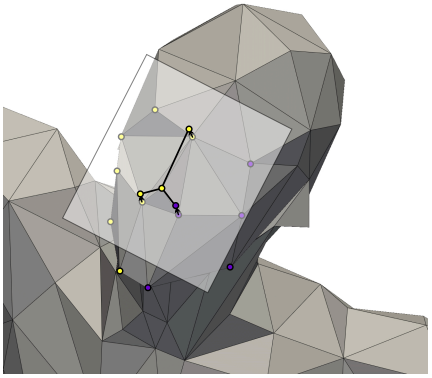

Figure 8. We project all neighbouring vertices onto the tangent space of the center vertex (i.e. the vertex in the picture which is connected by three edges) before loss computation such that our loss is independent of the surface curvature.

In Fig. 9, we provide, to Fig. 4 additional, visualizations of intrinsic symmetries detected with our method and competitors. In addition, we also include qualitative results of another setting (χ [53] (gt) and ours (gt)). We use ground truth left/right annotations to first cluster vertices of a shape into two sets. Then for each vertex, we find its intrinsic symmetry counterpart with different cluster label. We show results in Fig. 9 to further confirm the superiority of our disentangled features compared to features given by χ [53].

In Fig. 10, we include more examples to qualitatively

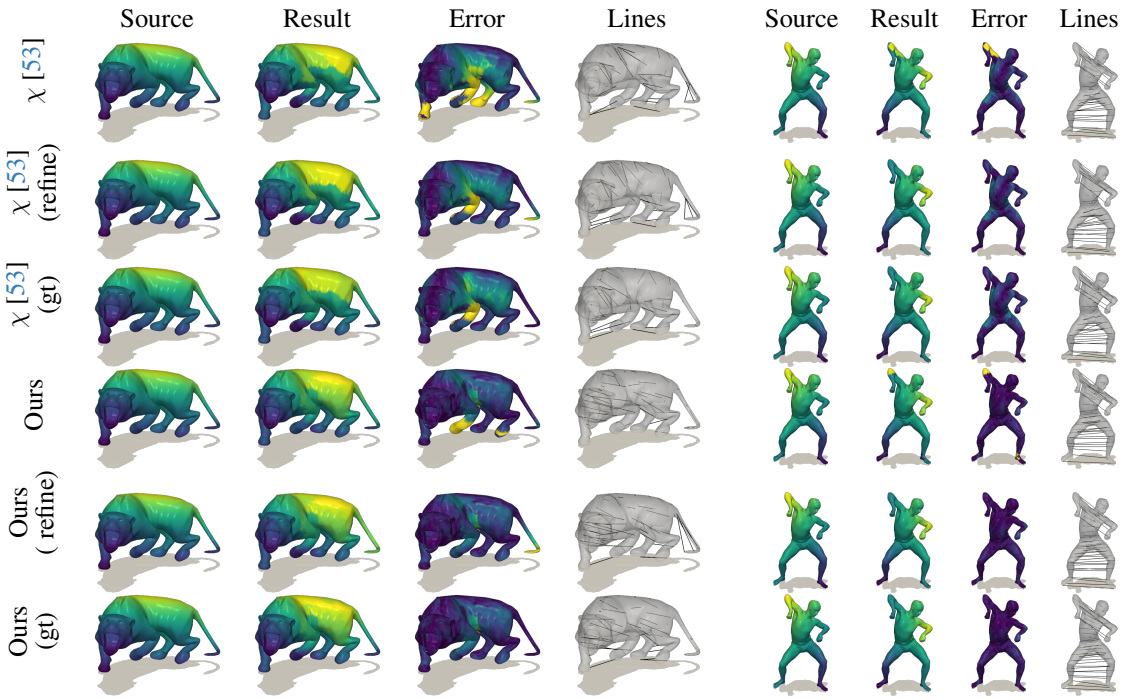

Figure 9. Qualitative results of intrinsic symmetry detection. Examples are chosen to reflect χ's [53] failure modes.

show the left/right classification results. Results on additional shapes in various poses confirm the effectiveness of our symmetry-informative feature.

In Fig. 11, we add qualitative results of more baselines (including χ [53] with refinement and ours without refinement) for a more comprehensive comparison. By comparing methods without refinement in the first four columns and by comparing methods with refinement in the last four columns, we can conclude that our method achieves better matching results compared to χ [53] (which is more prominent in the error visualizations). Further, by comparing χ [53] with and without refinement as well asby comparing ours with and without refinement, we can confirm the effectiveness of our proposed feature refinement technique.

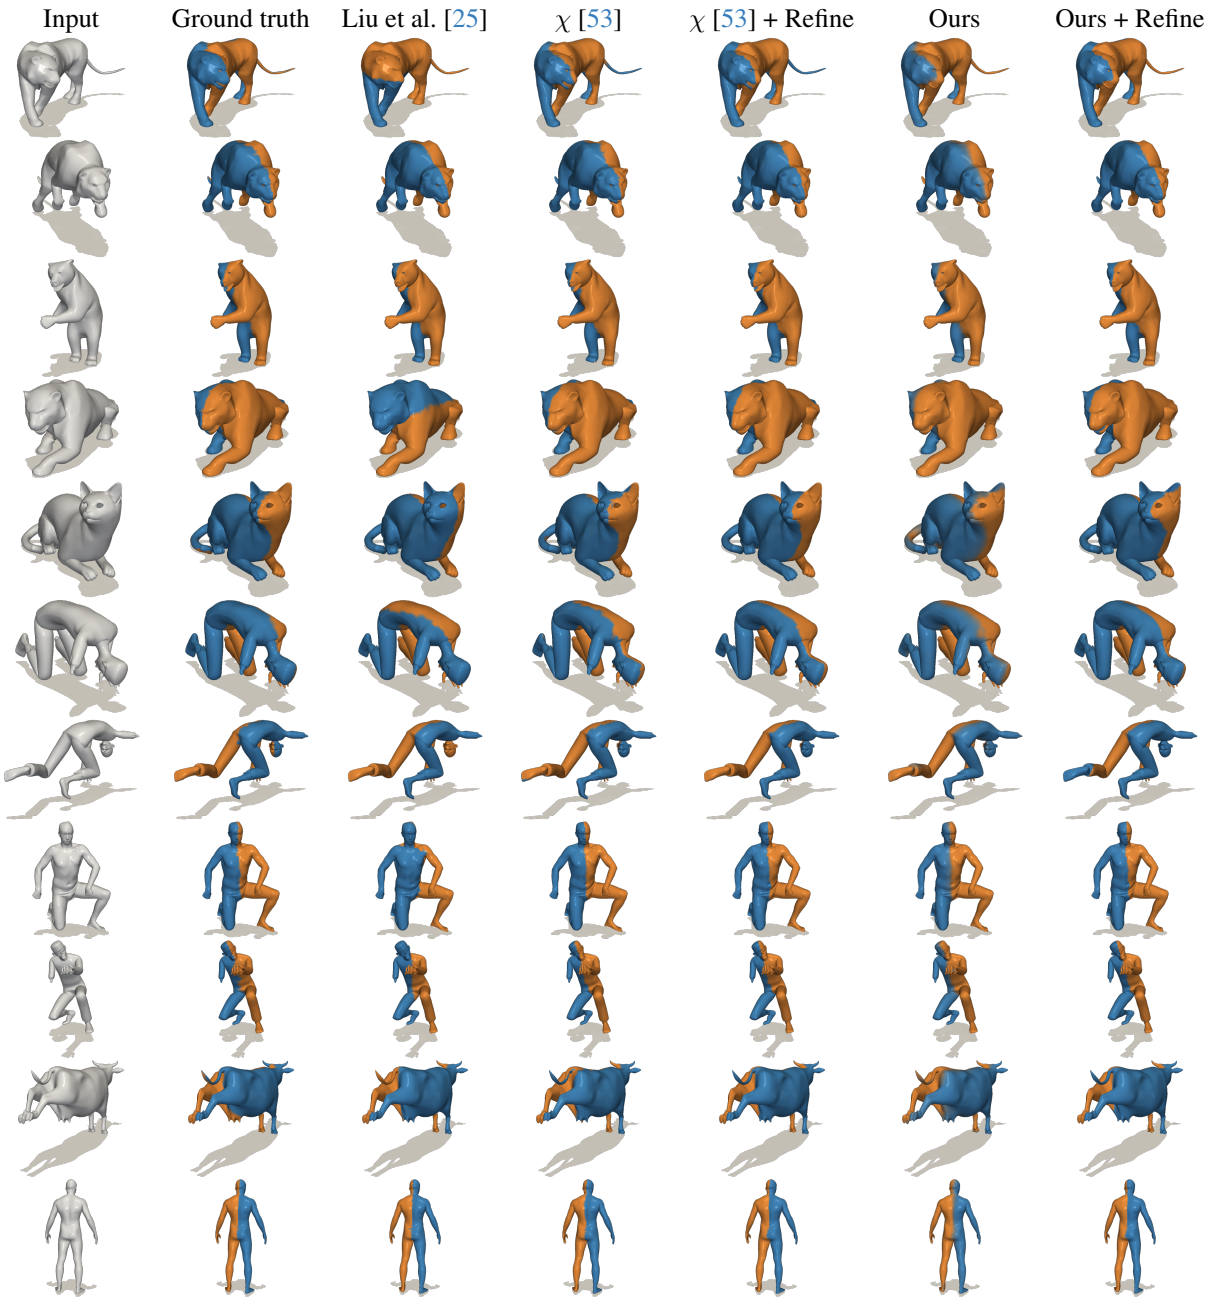

Figure 10. Qualitative results of left/right classification.

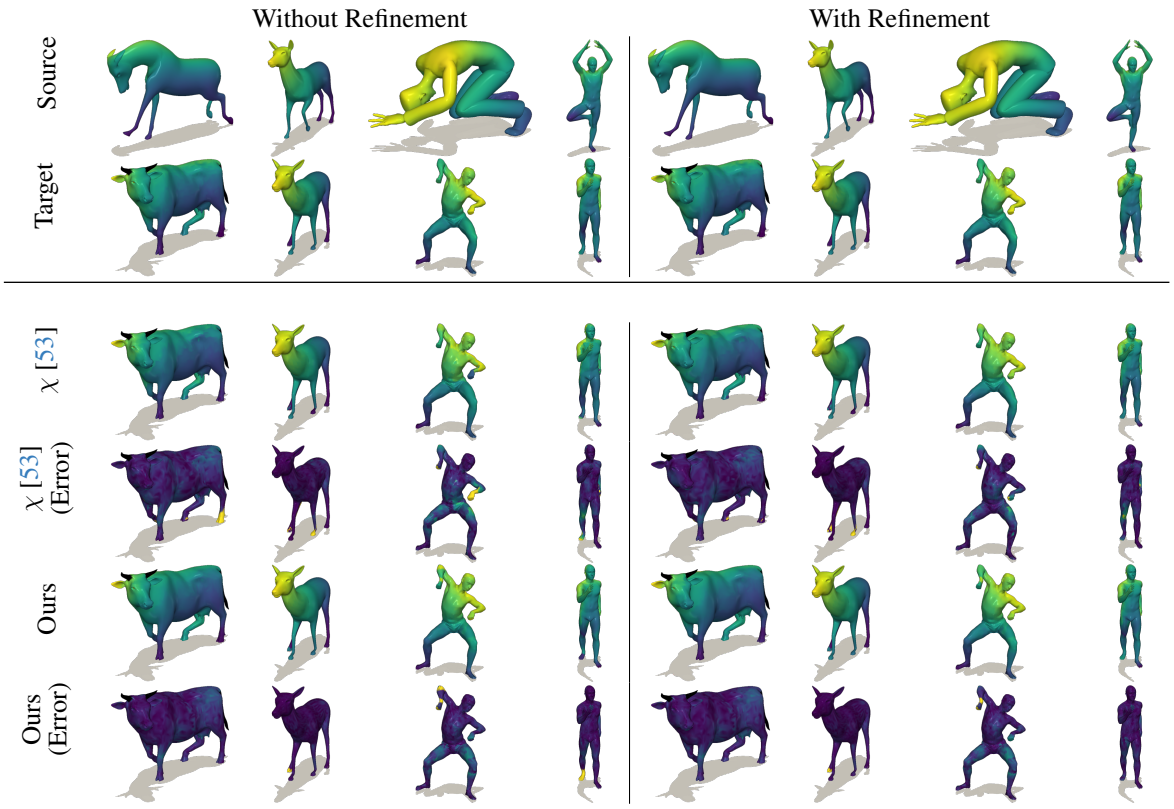

Figure 11. Qualitative results of shape matching.