
APPENDIX

1 ADDITIONAL RESULTS

In this section, we provide additional results showing the interpolation behavior of different architectures we studied. Each of the following figures is composed of six blocks, each of which presents a bilinear interpolation of the latent representation of four ground truth images that reside in each corner of the block. Figures 1 and 2 demonstrate bilinear interpolation on a different COIL-100 objects and Figure 3 shows the interpolation results on our synthetic pole dataset.

We demonstrate that our technique produces admissible interpolation while other techniques fail to reconstruct in-between images realistically or to transition smoothly from mode to mode. For example, in Figure 3 we show that the AMR, AAE and the ACAI create a cross-dissolve effect during interpolation while the interpolation of the β -VAE and GAIA method changes abruptly while showing little progression between interpolation frames. Additionally, in Figures 1 and 2 we show that the AAE and AMR exhibit artifacts in interpolated images while the ACAI shows blurry images during reconstruction. The β -VAE and GAIA shows mostly realistic reconstructions; however, the transition is not smooth nor consistent.

2 ABLATION STUDY

We present additional qualitative results of our ablation study in Figures 4 and 5. It is demonstrated that the contribution from all loss components exhibit admissible interpolation. In Figure 6 we present the results of predicting the interpolated alpha value by querying the dataset for the closest image in terms of the L_2 distance to the interpolated image. We present the median for each alpha on all intervals with the corresponding interquartile range. The red line demonstrates the perfect retrieval of the predicted alpha value. It is shown that the absence of the discriminator greatly affects the interpolation faithfulness while the addition of the cycle-consistency and smoothness losses contributes to the consistency of retrieving the correct alpha value.

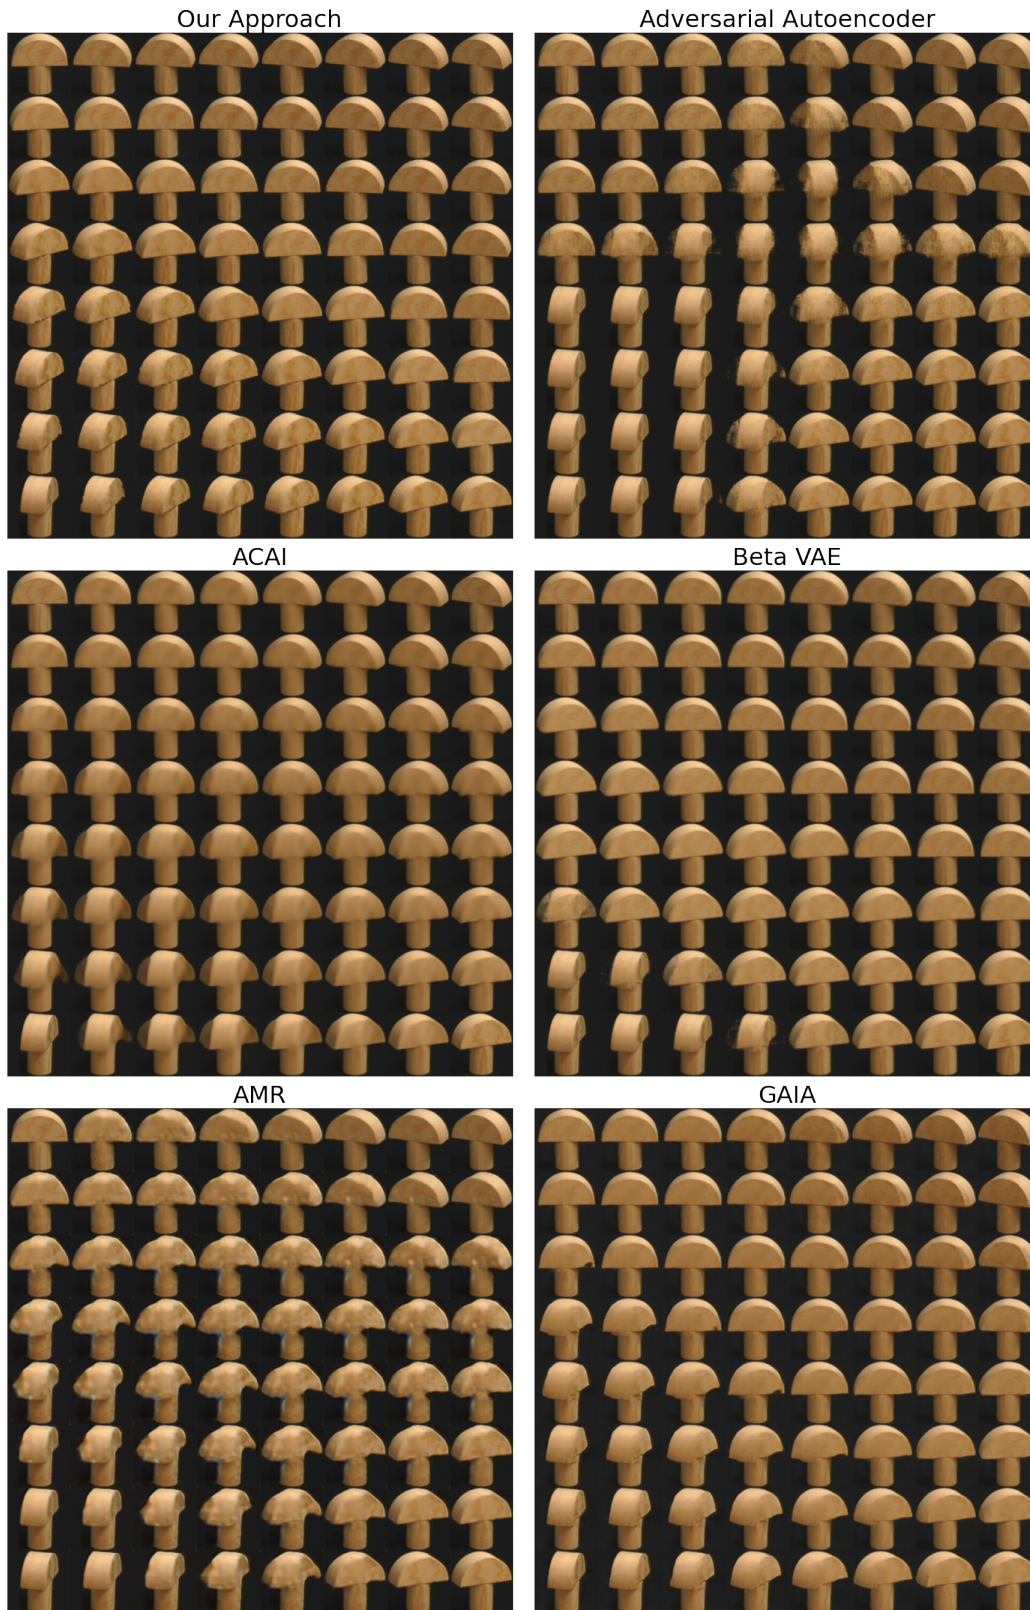


Figure 1: Each of the six blocks shows a bilinear interpolation of four ground truth images that reside in each corner of the block.

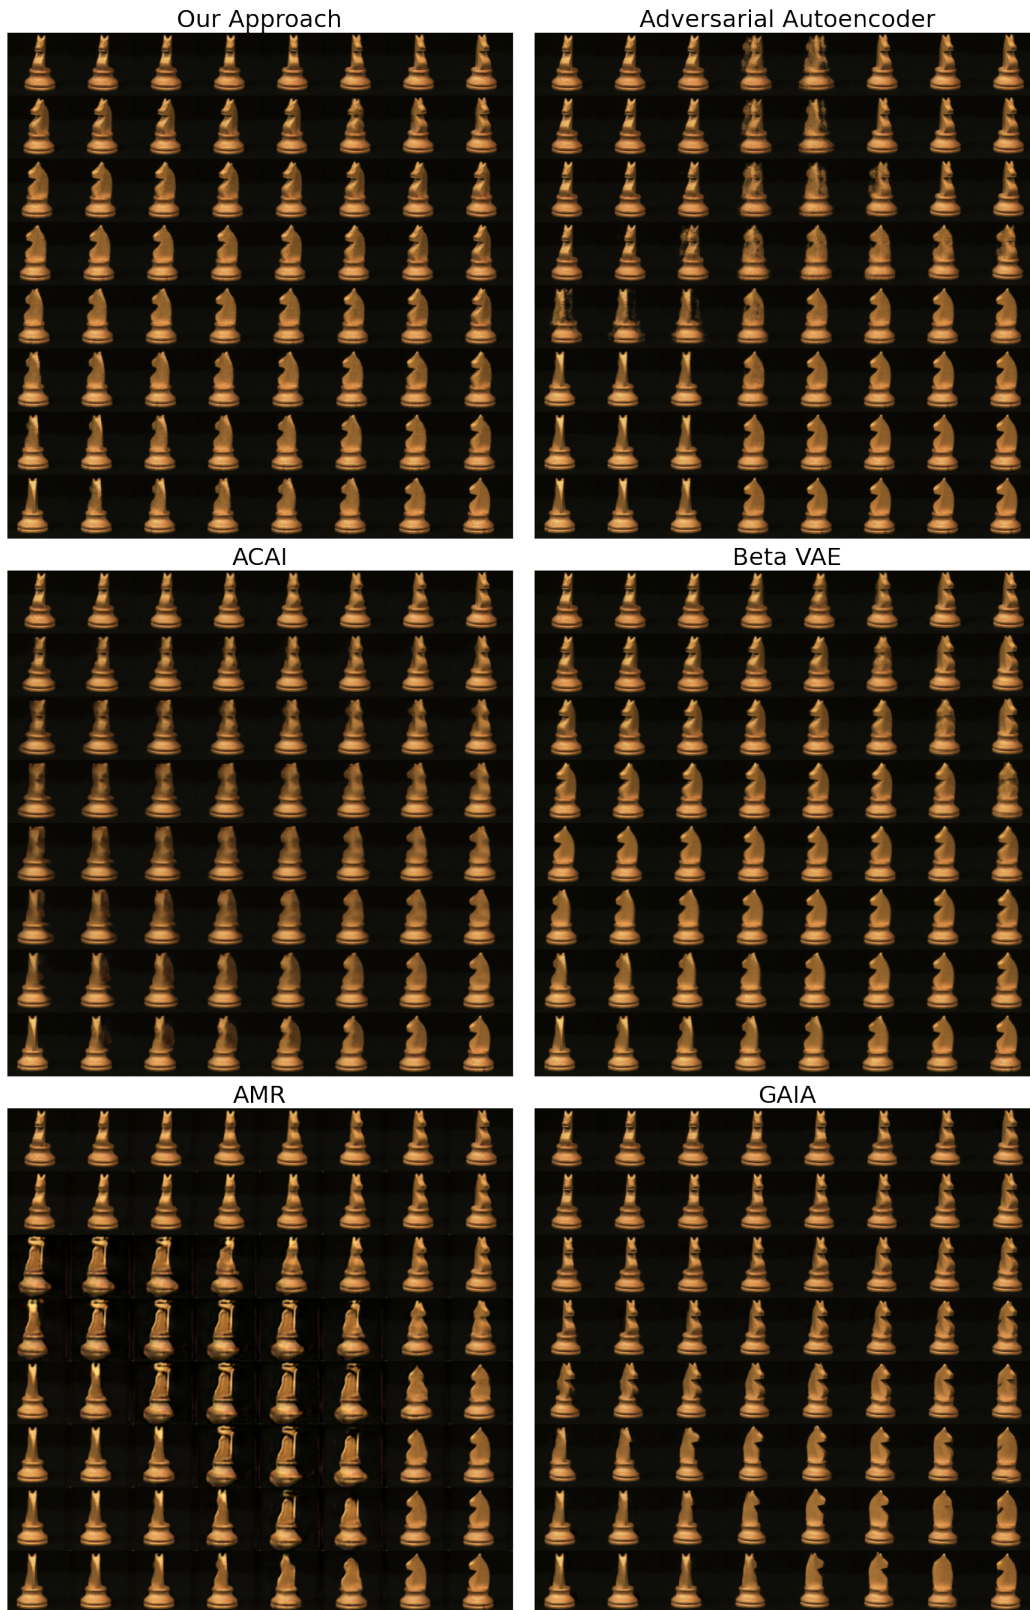


Figure 2: Each of the six blocks shows a bilinear interpolation of four ground truth images that reside in each corner of the block.

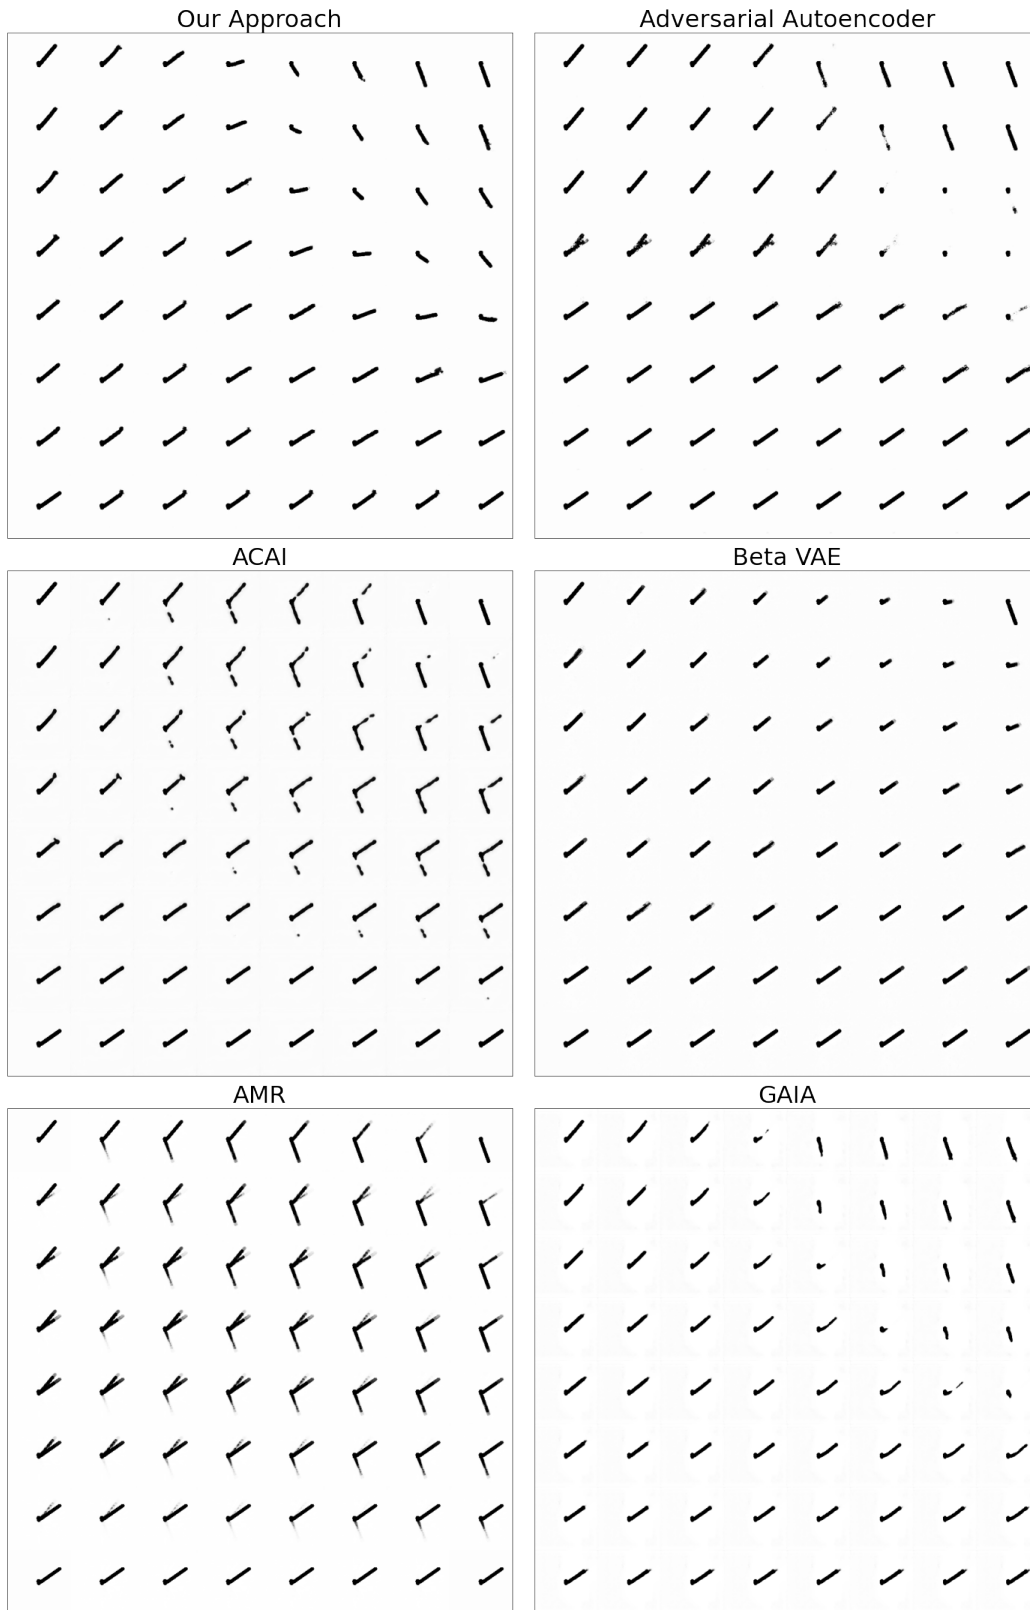


Figure 3: Each of the six blocks shows a bilinear interpolation of four ground truth images that reside in each corner of the block.

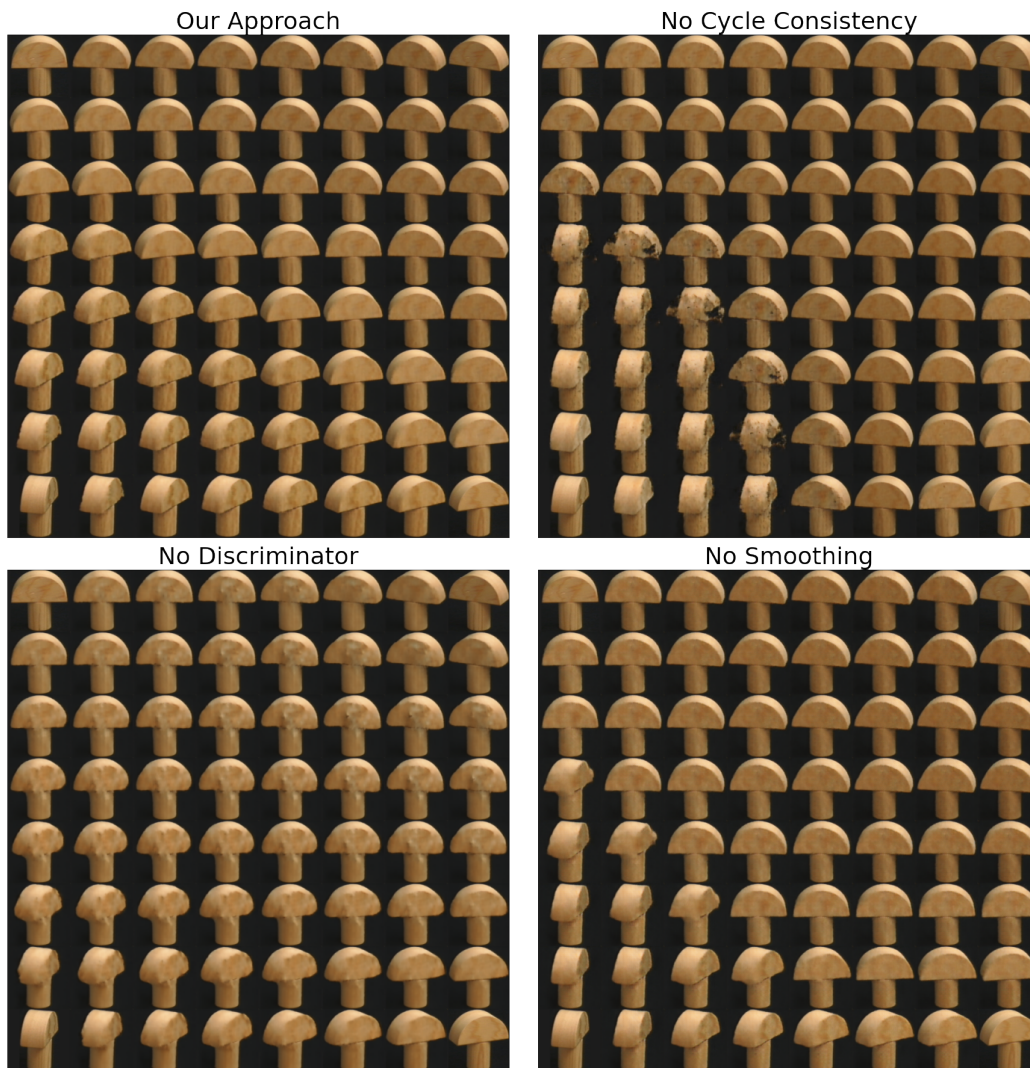


Figure 4: Each of the four blocks presents bilinear interpolation of four ground truth images that reside in each corner of the block. Top left: Bilinear interpolation results of our approach with all loss components. Top right: Removing the cycle-consistency contribution from the loss function. Bottom left: Removing the discriminator contribution. Bottom right: Removing the smoothing contribution.




Figure 5: Each of the four blocks presents bilinear interpolation of four ground truth images that reside in each corner of the block. Top left: Bilinear interpolation results of our approach with all loss components. Top right: Removing the cycle-consistency contribution from the loss function. Bottom left: Removing the discriminator contribution. Bottom right: Removing the smoothing contribution.

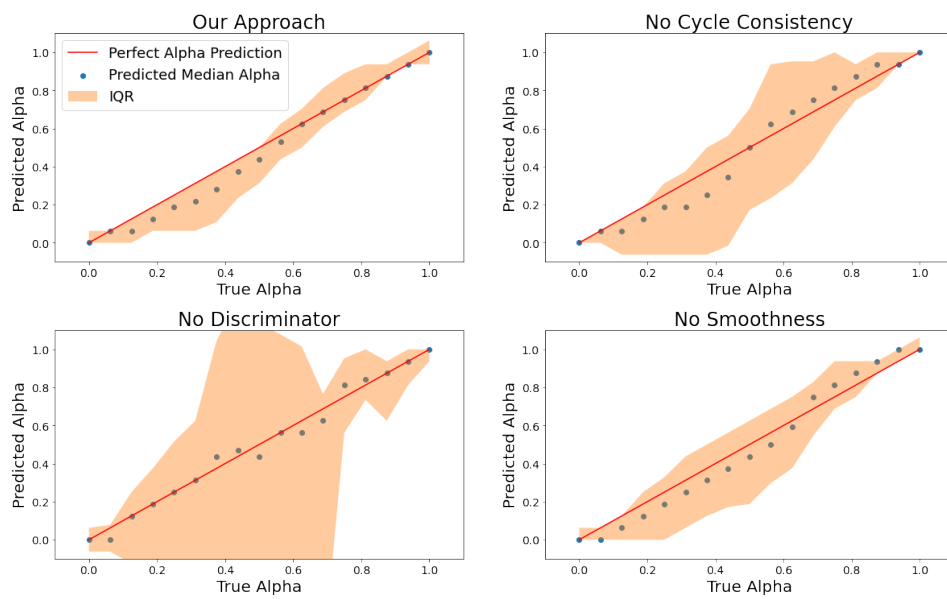


Figure 6: Predicting the interpolated alpha value based on the L_2 distance of the interpolated image to the closest image in the dataset. The dots represent the median and the colored area corresponds to the interquartile range.