

SUPPLEMENTAL MATERIAL

A ABLATION STUDY OF DIFFERENT LOSSES

In Figure 1, we illustrate the adaptive cuboid abstractions for different shapes from the four shape categories. The network with the full loss, *i.e.*, Network I, generates more visual pleasing abstraction results.

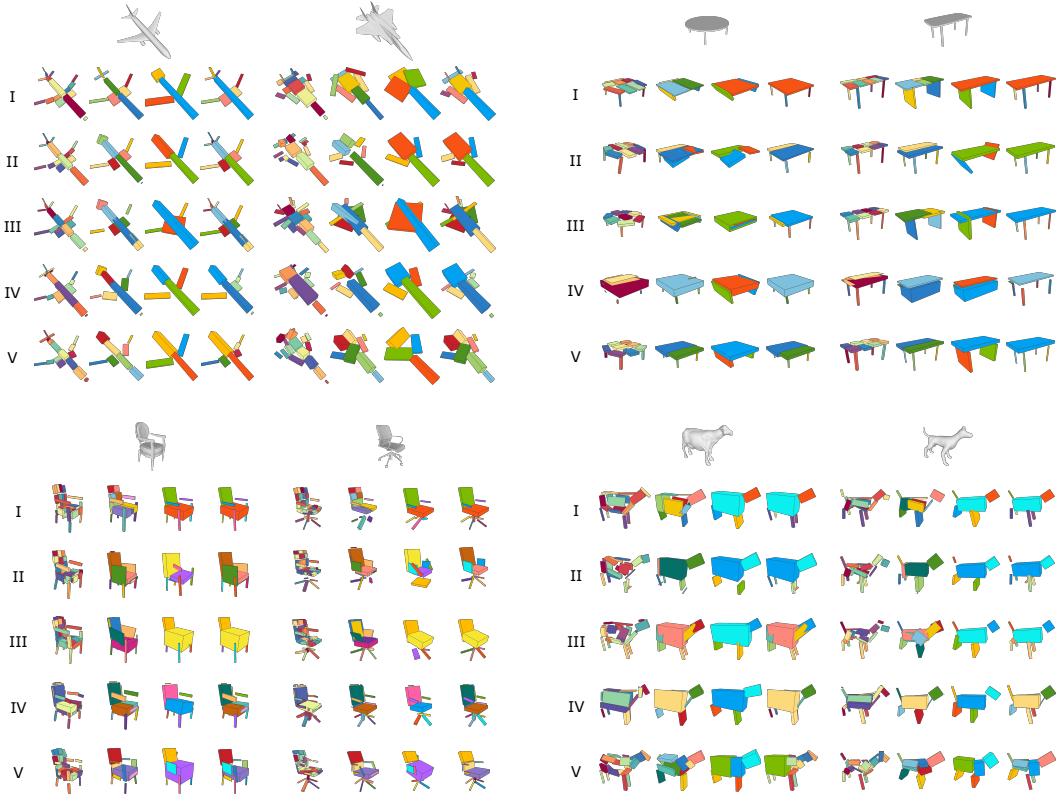


Fig. 1. Visual comparison of the adaptive abstractions predicted by our network using different loss terms. Net II: remove the hierarchical coverage loss; Net III: remove the mutex loss; Net IV: remove the average area loss; Net V: remove the alignment loss and bilateral symmetry loss.

B CHOICES OF CUBOID NUMBER

We trained our network with the cuboid configuration (5, 10, 20). From the visual illustration in Figure 2, we can find that the resulting adaptive abstractions are mostly visually pleasing and achieve similar approximation errors to the results as shown in the paper. However, for airplanes, as there are more candidate cuboids than the default configuration we used, some parts like airframe tend to be approximated by more cuboids. For chairs, as 20 is much smaller than the default number we used – 32, we can see that two legs of the wheelchair is abstracted by a single cuboid due to the insufficient number of cuboids for handling various chair structures.

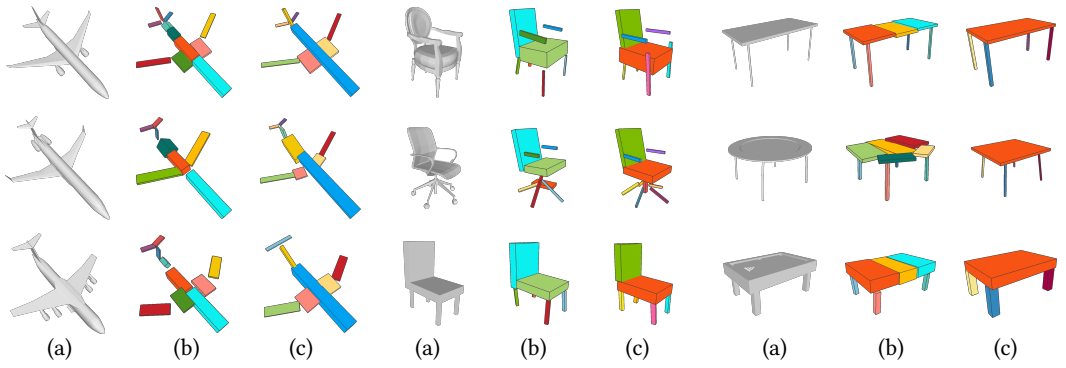


Fig. 2. Visualization of the adaptive abstractions predicted by our network using the default configuration and the cuboid configuration (5,10,20). (a): Input shapes; (b): the results from the network using the cuboid configuration (5,10,20); (c) the results from the network using the default configuration.