

Figure 1: Rendering of the robotic manipulation environment. We show four frames from a single trajectory.

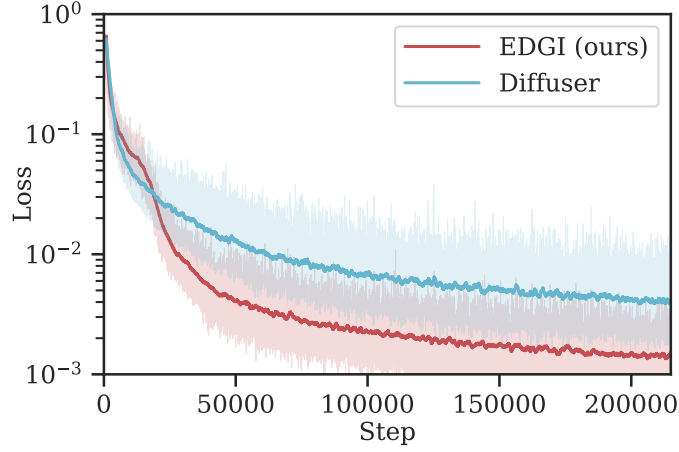


Figure 2: Loss curves on the robotic manipulation task. In terms of optimizer steps, EDGI converges much faster than the Diffuser baseline.

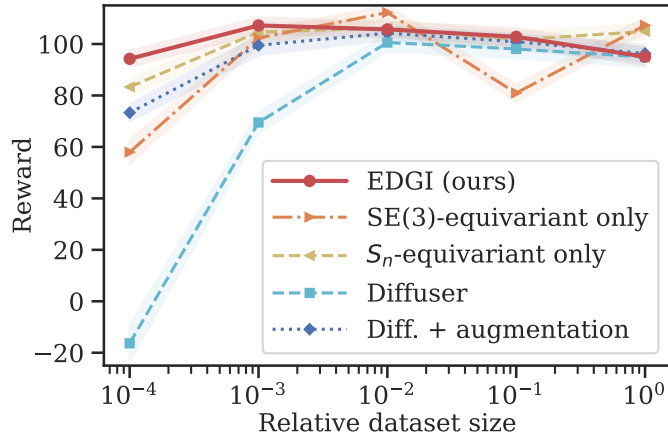


Figure 3: Symmetry group ablations and a data augmentation baseline. We show normalized reward on the navigation task as a function of training samples. First, training a Diffuser model with $SO(3)$ data augmentation (blue dashed line) leads to better data efficiency than Diffuser, but EDGI remains the most sample-efficient model. Second, ablation models (yellow dashed, orange dot-dashed) that are equivariant to $SE(3)$ but not S_n , and vice versa, also improve over the Diffuser baseline but do not quite reach the performance of EDGI for very small training sets.