Few-Shot Generalization for 3D Reconstruction via Priors: Supplementary Material

1. Performance Across Iterations

The iterative variants of our model helped for multi-view reconstruction but were outperformed by the non-iterative version for single-view reconstruction. To understand the behavior of our models (trained for iteration or not) over multiple iterations, we plot the IoUs of the models for 10 iterations on novel categories with various amounts of prior information in Figure 1. In general, the performances from 1-shot priors initially improve across iterations while those with average priors only decrease. Performance from both priors slowly declines after many iterations. This decrease is gradual enough for the 3-iteration model that it maintains a substantial advantage over the baseline throughout in the full-prior case and until the 10th iteration when given a 1-shot prior.

We also observed that the rate and nature of this decrease in performance is correlated to how different the novel category is from the base categories in the training set. When the novel category is similar to those of the training set, the iterations have more consistent performance and initial improvement. When the novel category is very dissimilar to the training set, however, the earlier iterations perform best and the decrease occurs immediately. We hypothesize that this is due to the more unique transfer category shape predictions lying farther outside the training distribution of the model, resulting in instability of predictions. This issue is well known for sequential prediction models where errors can begin to accumulate as the predicted sequence takes the model into unseen parts of the input space, and is exacerbated by our transfer setting.

We observed that the rate and nature of the decrease in performance across iterations varied between categories. Specifically, the novel categories which we identified as simple or similar to the training categories (namely **cabinets** or **benches**) initially increase slightly in peformance during the early iterations. Meanwhile, the elongated categories (i.e. **rifles** and **vessels**) which are very dissimilar to our training categories achieve their best performance with the initial prediction and rapidly decay afterwards. We present per-category performances across iterations for the 3-Iteration model.

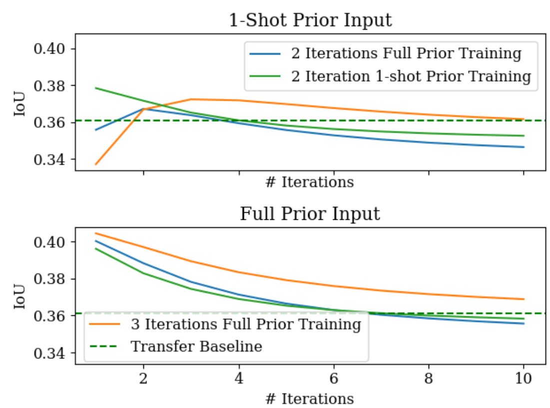

Figure 1. Category-wise IoU across iterations. The green line is the transfer baseline of 0.36. As seen in Table **??**, 2 out of the 3 models initially improve performance at each iteration when given a 1-shot prior. When given a full prior, all models performances decline after the first iteration.

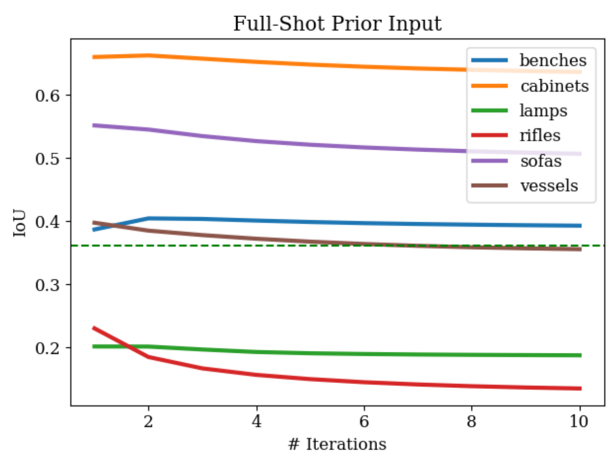

2. Example Predictions:

Here we present 10 random example predictions for each novel category. These are from both the Baseline model and our model that achieved the best performance (1-Shot 1-Iteration). The Baseline model has no shape prior, using only the image as input. The 1-Shot 1-Iteration model is trained using a non-iterative scheme and a random 1-shot shape prior. On runtime, this model is given *one single* random shape example of the new category. Neither model has had any parameter tuning on novel categories.

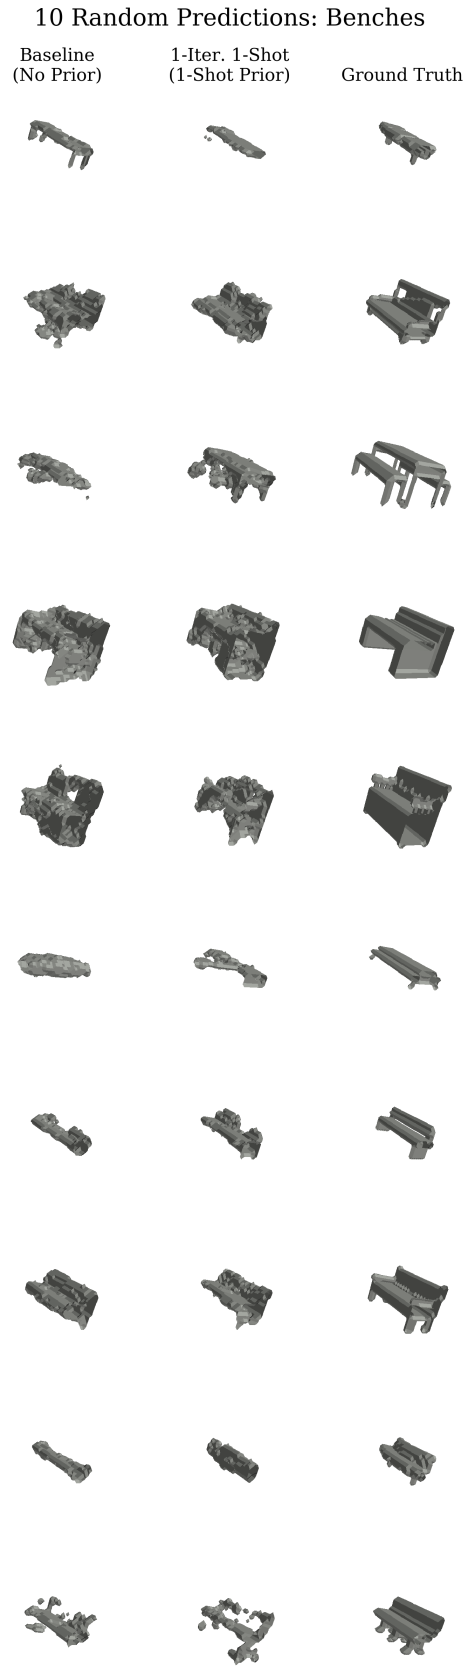

The examples shown are truly random to provide a realistic sample of performance. While qualitatively the reconstructions may seem poor, we note that the task of transfer learning is extremely challenging. Quantitatively, our model peforms almost strictly better than the baseline, sometimes by large margins. This improvement can be seen for example in the third row of **benches**, seventh row of **rifles** and first row of **sofas**. To reiterate, this improvement is from receiving *just a single shape* as categorical information. We also note that on occasion the Baseline model attempts to reconstruct shapes as if they belonged to the training categories. This indicates that the baseline has internalized the shape distributions of the training categories. Examples of this can be seen in the first row of **rifles** and sixth row of **vessels** where the Baseline model appears to reconstruct a **plane** and **display** respectively.

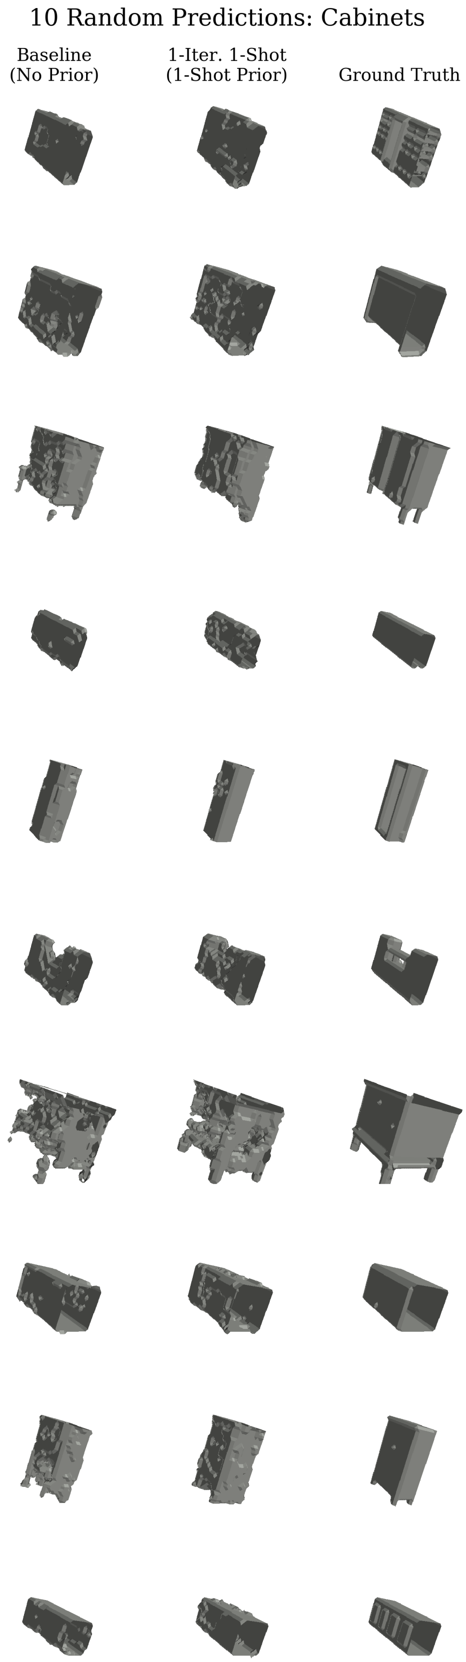
10 Random Predictions: Cabinets
Baseline
(No Prior)
1-Iter. 1-Shot
(1-Shot Prior)
Ground Truth


10 Random Predictions: Lamps
Baseline
(No Prior)
1-Iter. 1-Shot
(1-Shot Prior)
Ground Truth

10 Random Predictions: Sofas

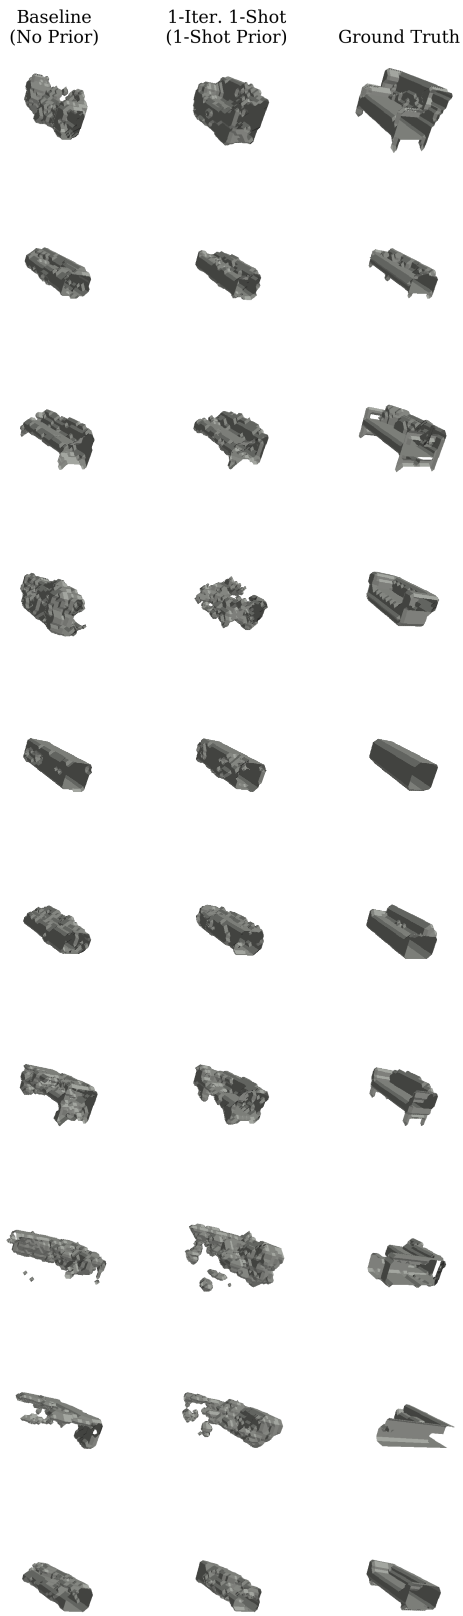

10 Random Predictions: Rifles

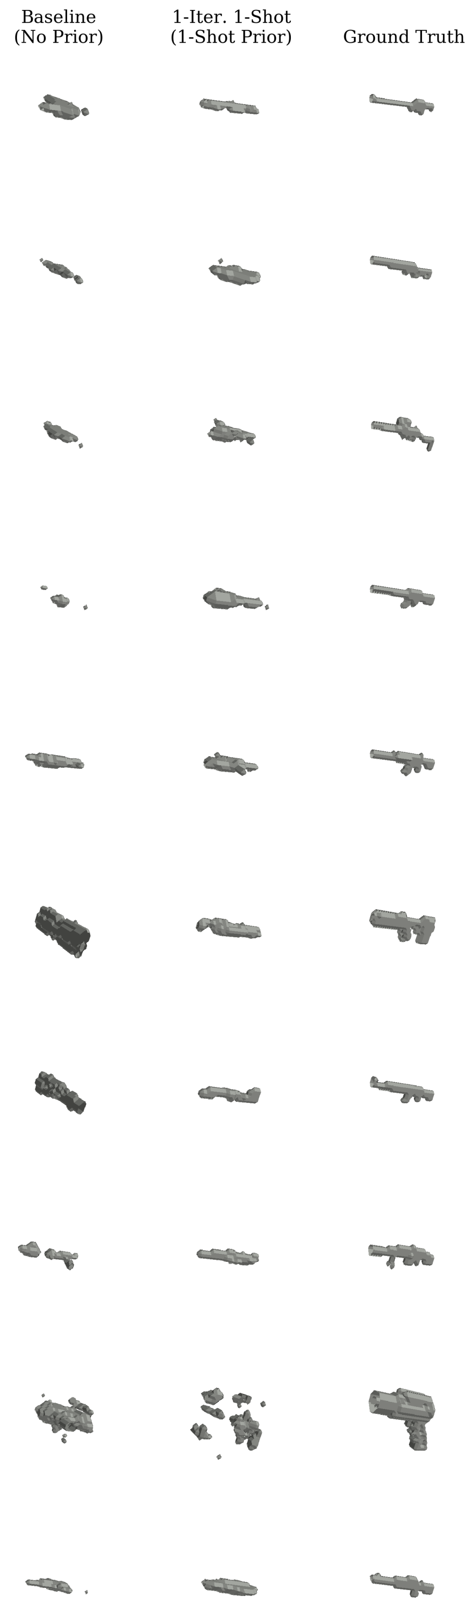

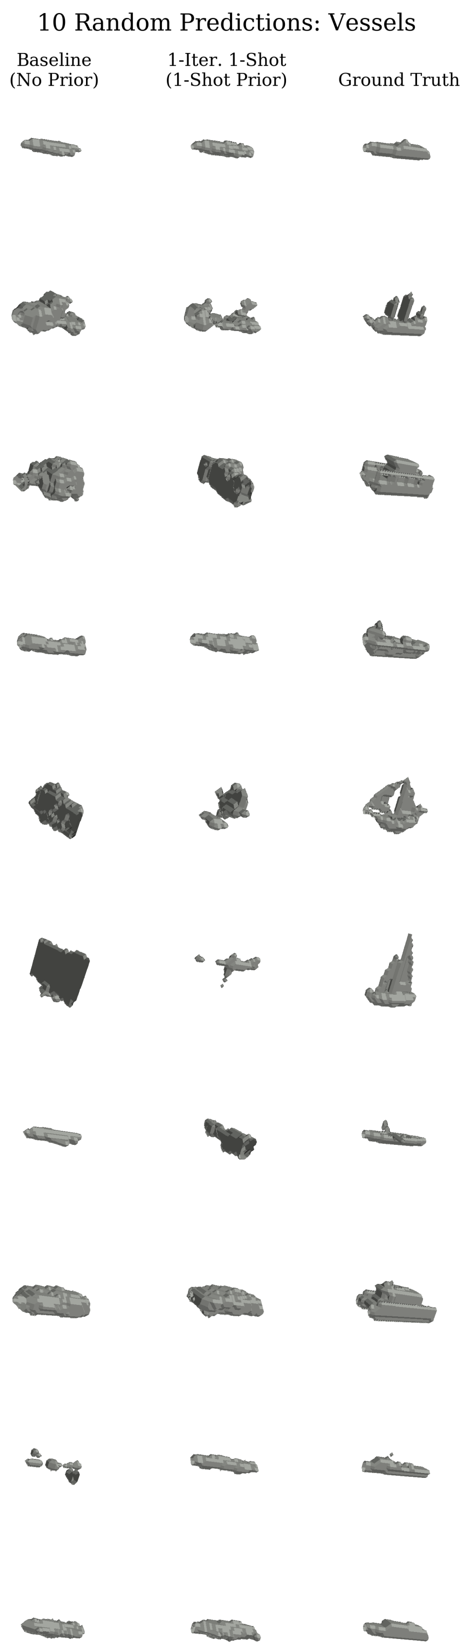
10 Random Predictions: Vessels
Baseline
(No Prior)
1-Iter. 1-Shot
(1-Shot Prior)
Ground Truth