Appendices for Meta-Sampler

Ta-Ying Cheng, Qingyong Hu*, Qian Xie, Niki Trigoni, Andrew Markham

Department of Computer Science, University of Oxford

1 Limitations

While this work shows the effectiveness of incorporating multiple tasks in the meta-training step, the question of finding the best combination and the number of tasks for the best model is indeed intriguing and important. Intuitively, more tasks will supposedly lead to more exposure to different task varieties. Simultaneously, this could also lead to tasks potentially contradicting each other and thus diminishing the effectiveness of the pretrained meta-sampler.

Another interesting direction is the sampling strategy for tasks such as large-scale segmentation, where the conventional sampling approaches are progressively used in the network, instead of simply using it once at the beginning. How to adapt our meta-sampler to such tasks efficiently and effectively is worth exploring.

2 Qualitative Results

2.1 Sampling for Different Tasks

We provide additional qualitative analyses on the different points sampled from different tasks for classification, reconstruction, and shape retrieval (Figure 1). The results further validate our assumption that different tasks have different preferences in terms of feature sampling (*e.g.,* classification often require the sampled points to be scattered to give an overview of the object; reconstruction focuses on denser points to minimise the Chamfer Distance loss during reconstruction; shape retrieval often preserve fine-grained features like corners to distinguish hard negative of the same class), and hence an almost-universal sampler is indeed the best way to tackle the task of sampling.

2.2 Sampling from Single and Multi-model training

We visualise the points sampled from single-task single-model (blue) and our proposed multi-task-multi-model training (red), as shown in Figure 2. It is apparent that while there exists some overlapping points (which is intuitive as they are targeting the same task), the majority of points are different. Based on the accuracy improvements from our main paper, we can thus infer that our joint-training scheme allows us to capture features that are more universal towards a particular task.

* Corresponding author

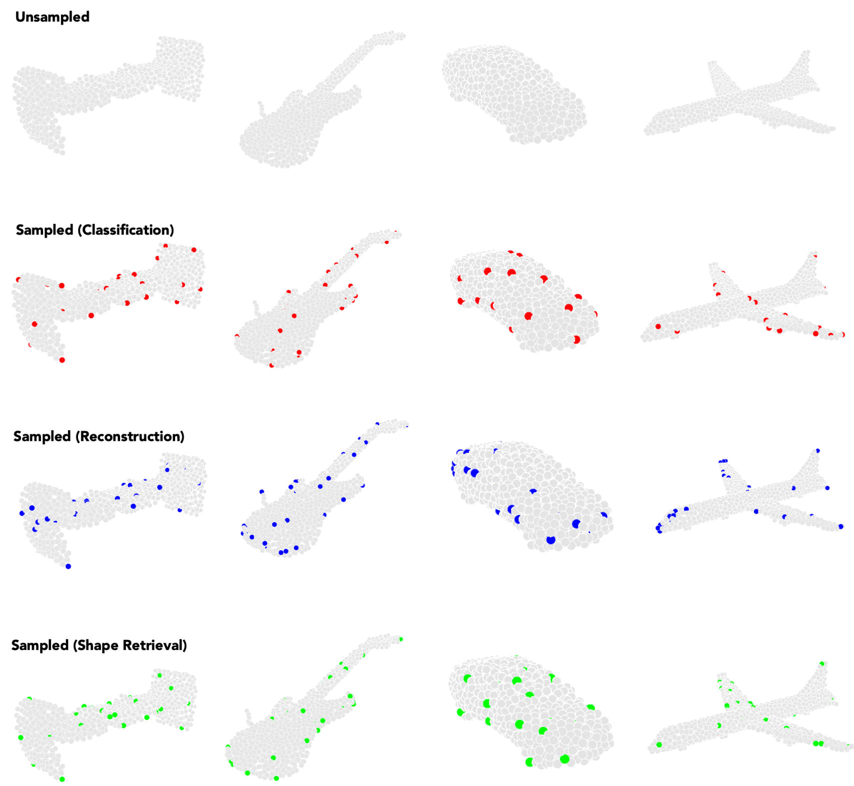

Fig. 1. Visualisations of sampled points for classification (red), reconstruction (blue), and shape retrieval (green) at sampling ratio of 32.

3 Reproducibility

All our training (both meta-sampler and task adaptation) has a batch size of 24. Meta-training uses 5 gradient steps in the inner update. The rapid task adaptation uses the Adam optimiser with a learning rate of 1e-3. The ratio of task, simplification, and projection losses for task adaption is 1:1:1.

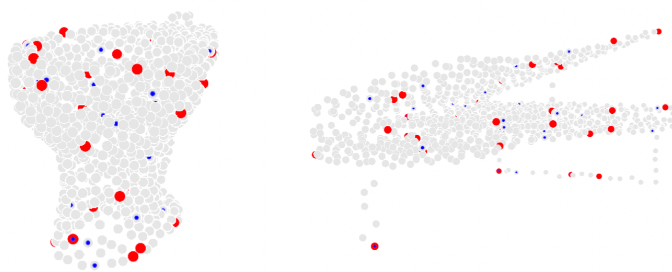

Fig. 2. Visualisations for SampleNet trained with single-task single-model (blue) and with single task multi-model (red) at sampling ratio of 32.