# Meta-Sampler: Almost-Universal yet Task-Oriented Sampling for Point Clouds

Ta-Ying Cheng, Qingyong Hu<sup>\*</sup>, Qian Xie, Niki Trigoni, Andrew Markham

Department of Computer Science, University of Oxford

Abstract. Sampling is a key operation in point-cloud task and acts to increase computational efficiency and tractability by discarding redundant points. Universal sampling algorithms (e.g., Farthest Point Sampling) work without modification across different tasks, models, and datasets, but by their very nature are agnostic about the downstream task/model. As such, they have no implicit knowledge about which points would be best to keep and which to reject. Recent work has shown how task-specific point cloud sampling (e.g., SampleNet) can be used to outperform traditional sampling approaches by learning which points are more informative. However, these learnable samplers face two inherent issues: i) overfitting to a model rather than a task, and ii) requiring training of the sampling network from scratch, in addition to the task network, somewhat countering the original objective of down-sampling to increase efficiency. In this work, we propose an *almost-universal* sampler, in our quest for a sampler that can learn to preserve the most useful points for a particular task, yet be inexpensive to adapt to different tasks, models or datasets. We first demonstrate how training over multiple models for the same task (e.g., shape reconstruction) significantly outperforms the vanilla SampleNet in terms of accuracy by not overfitting the sample network to a particular task network. Second, we show how we can train an almost-universal meta-sampler across multiple tasks. This meta-sampler can then be rapidly fine-tuned when applied to different datasets, networks, or even different tasks, thus amortizing the initial cost of training. Code is available at https://github.com/ttchengab/MetaSampler.

Keywords: Point Cloud Sampling, Point Cloud Processing, Meta-Learning

# 1 Introduction

Modern depth sensors such as LiDAR scanners can capture visual scenes with highly dense and accurate points, expanding the real-world applications of point clouds to traditionally challenging 3D vision tasks. However, while existing deep network architectures such as PointNet [34] are capable of consuming these dense point clouds for downstream tasks (*e.g.*, classification, reconstruction), it is standard to downsample initial point cloud to reduce the computational and memory cost, especially for resource-constrained or real-time applications. As such, the

<sup>\*</sup> Corresponding author



Fig. 1. Overview. Top Left: We highlight the points sampled for classification (red) and reconstruction (blue). It is apparent that classification concentrates on more generalised features across the entire point cloud whereas reconstruction focuses on denser aspects for optimisation. Bottom Left: We evaluate the classification performance of two frozen PointNets on 16 sampled points from SampleNet. A large performance gap is observed despite the two models (one adopted during SampleNet training and one unseen) having an identical architecture, implying overfitting onto the model instead of the task itself. Right: An overview of our meta-sampler as a pretrained model that can rapidly adapt with a joint-training mechanism.

objective of extracting a representative subset of points from raw point clouds while maintaining satisfactory performance over various tasks is a key problem.

Early techniques usually adopt Farthest Point Sampling (FPS) [35,27,41]. Inverse Density Importance Sampling (IDIS) [13], or Random Sampling (RS) [19,20] to progressively reduce the resolution of the raw point clouds. Albeit simple and universal, these sampling schemes are inherently heuristic and taskagnostic. Recently, Dovrat et al. [8] and Lang et al. [26] explored a new domain of learning-based, task-specific, and data-driven point cloud sampling strategies. They empirically proved that leveraging the task loss can effectively optimise the sampler to preserve representative and informative features. Although remarkable progress has been achieved in several downstream tasks such as classification and reconstruction, there remain two critical issues to be further explored: 1) The learnt samplers are shown to overfit to a specific task model instead of being generalisable to the task itself — this causes a signifant performance drop when adopting another network for the same task even when the two architectures are identical (as exemplified in Figure 1 Bottom Left); 2) Training a sampler to fit a particular task is both time-consuming and computationally expensive, which counters the original objective of sampling to improve efficiency.

3

To this end, we propose an *almost-universal* sampler (Figure 1 Right) comprising two training alterations to address the aforementioned issues accordingly. First, we suggest jointly training by forwarding the sampled points to multiple models targetting the same task instead of a single model and updating the sampler through a summation of task losses. This kind of ensemble allows us to better simulate the distribution of different task models, encouraging the sampler to truly learn the task rather than a particular instance. Second, we introduce our meta-sampler to learn how to adapt to a specific task, rather than explicitly learning a particular task model. We incorporate a set of tasks, each with multiple task models, for the meta-optimisation. Our meta-sampler can serve as a pretrained module to adhere to any tasks through fine-tuning while being *almost-universal* in the sense that it could be optimised with fewer iterations.

Extensive experimental results justify the performance and versatility of the proposed meta-sampler. In particular, there is a significant improvement in performance for several mainstream tasks with our joint-training technique compared to the best results from the conventional single-task training on SampleNet. Moreover, we thoroughly evaluate the versatility of our meta-sampler by adapting to particular tasks (both included and excluded from the meta-training), model architectures, and datasets. Our meta-sampler adapts rapidly to all challenging scenarios, making it a suitable pretrained candidate for task-specific learning-based samplers.

In summary, the key contributions of this paper are threefold:

- A joint-training scheme for the sampler to truly learn a task rather than simply overfitting to a particular instance (*i.e.*, a specific task model).
- A meta-sampler that can rapidly adapt to downstream point cloud tasks within and beyond the meta-training stage, models of varying architectures, and datasets of different domains.
- Extensive experiments validate the performance and versatility of our metasampler across various tasks, models, and datasets.

# 2 Related Work

### 2.1 Learning with 3D Point Clouds

Earlier pursuit of 3D computer vision tasks is mainly focused on grid-like representations of voxel volumes, as the mature convolutional neural networks (CNNs) can be directly extended to such data representation and easily introduce inductive biases such as translational equivariance [7,38]. However, voxel volume representation has the ingrained drawback of being uniform and low-resolution with densely compacted empty cells consuming vast amount of computational resources. Recently, point-based networks have attracted wide attention with the emergence of PointNet/PointNet++ [34,35]. These architectures pioneered the learning of per-point local features, circumvented the constraint of low resolution and uniform voxel representations, and hence introduced significant flexibilities and inspired a plethora of point-based architectures [13,27,29,40,45]. A number of point cloud based tasks [14] including classification [12,16,28,30,43,48], segmentation [3,19,25,18,17], registration [2,11,22], reconstruction [9,31], and completion [6,46] are extensively investigated. Nevertheless, few arts targeted the fundamental component of point cloud sampling in this deep learning era.

# 2.2 Point Cloud Sampling

Point cloud sampling, a basic component in most point-based neural architectures, is usually used to refine the raw inputs and improve computational efficiency for several downstream tasks. Widely-adopted point cloud sampling methods include RS, FPS [35,33,19], and IDIS [13]. A handful of recent works began to explore advanced and sophisticated sampling schemes [32,5,44]. Nonetheless, despite the remarkable progress in point cloud sampling, these methods are taskagnostic and rather universal, lacking awareness of the important features which a particular task may require.

Recently, Dovrat et al. [8] proposed a learnable, data-driven sampling strategy by imposing a specific task loss to enforce the sampler in learning specificrelated features for a particular task. Later, Lang et al. [26] extended the learning approach by introducing a differentiable relaxation to minimise the training and inference accuracy gap for the sampling operation. Nevertheless, by introducing an additional task loss, sampling ultimately becomes constrained and prone to overfitting on a specific task model instead of the task itself. Additionally, this also requires significant extra training to fit a particular goal.

Our meta-sampler hopes to bring the best of both worlds: being task-oriented yet as universal as possible. Instead of directly overfitting onto a task model, we focus on how to learn a task through incorporating a meta-learning algorithm. By introducing a better approach of learning a particular task through joint-training, our pretrained meta-sampler be rapidly fine-tuned to any task, making it *almost-universal* while easing the computational efficiency to which sampling is targeting in the first place.

## 2.3 Meta-Learning

Meta-learning, the process of learning the learning algorithm itself, has shown to be applicable to several challenging computer vision scenarios such as fewshot and semi-supervised classification [10,36], shape reconstruction [39], and reinforcement learning [15,23] due to its capacity for fast adaptation.

Finn et al. [10] proposed one of the most representative meta-learning methods, termed model-agnostic meta-learning (MAML), that allows the model to quickly adapt to new tasks in different domains such as classification, regression, and reinforcement learning. Later, Antoniou et al. [1] further improved the MAML learning scheme, making the learning more generalisable and stable. Recently, Huang et al. [21] proposed MetaSets, which aims to meta-learn the different geometry of point clouds so that the model can generalize to classification tasks performed on different datasets. In contrast and being analogous to the standard meta-learning problem, our proposed meta-sampler focuses on universal point cloud sampling for different tasks, aiming to achieve fast adaptation to reduce computation efficiency through a training strategy extended from [10,1]. Our fast adaption is not just across tasks within the meta-training, but also across models, datasets and unforeseen tasks.

# 3 Meta-Sampler and Rapid Task Adaptation

Ideally, it is desirable to learn a unified and universal point cloud sampling scheme for different tasks in a data-driven manner — this is most likely unfeasible since different tasks inherently have distinctive preferences of sampling strategies, as shown in Figure 1 top left (for more qualitative comparisons please refer to the Appendices). For example, 3D semantic segmentation pays more attention to the overall geometrical structure, while 3D object detection naturally puts more emphasis on the foreground instance with sparse points [47]. Motivated by this, we take the next-best objective, which is to learn a highly adaptive sampling network that can adapt to a number of tasks with minimal iterations and achieve optimal performances. In particular, this fast adaptation capability allows samplers to be pretrained then quickly fine-tuned, satisfying the ultimate goal of improving computational efficiency.

#### 3.1 Problem Setting

The goal of this paper is to develop a learning-based sampling module  $f_{\theta}$  with trainable parameters  $\theta$ , which takes in a point cloud with m points and outputs a smaller subset of n points (m > n). Apart from the objective of SampleNet to learn task-specific sampling (*i.e.*, particularly suitable for a single task such as shape classification or reconstruction), we take a step further and aim to propose an universal pretrained model, which can be rapidly adapted to a set of different tasks  $S_T = \{T_i\}_{i=1}^{K_T}$ . Formally, we define the ideal adaptation of sampling to a specific task  $T_i$  as capable of achieving satisfactory performance by integrating the sampling module into a set of  $K_{A_i}$  known networks  $S_{A_i} = \{A_{i,j}\}_{j=1}^{K_A}$  (Each  $A_{i,j}$  is one network trained with unsampled point clouds to solve task  $T_i$ ). We split  $S_{A_i}$  into  $S_{A_i}^{train}$  and  $S_{A_i}^{test}$  (*i.e.*, task networks used during training are disjoint to the ones for testing) to make sure that our evaluation on  $f_{\theta}$  is fair and not overfitting to task models instead of the task itself. Note that while  $S_{A_i}^{train}$  is available during training, the weights are frozen when learning our sampler as suggested by [26].

To achieve the dual objectives of high accuracy and rapid convergence, we must first carefully evaluate the best training strategy to better learn each individual task, and then design a training strategy which is adaptive to multiple tasks. We build our sampler  $f_{\theta}$  based on the previous state-of-the-art learnable sampling network — PointNet-based SampleNet architecture [34,26] — and then introduce our training technique in a bottom-up manner.

# 3.2 Single-Task Multi-Model Training

For an individual task  $T_i$ , we hope that the  $f_{\theta}$  learns to sample the best set of points  $\forall A \in S_{A_i}$ .

The conventional way of training the SampleNet uses a single frozen network A' as  $S_A^{train}$  for training by defining a sampling task loss  $\mathcal{L}_{ST_i}$  targeting  $T_i$  as:

$$\mathcal{L}_{ST_i}(f_\theta) = \mathcal{L}_{T_i}(A'(f_\theta)),\tag{1}$$

where  $\mathcal{L}_{T_i}$  is the loss when pretraining A'. We refer to this configuration as singlemodel, single-task training. As mentioned previously, this training method, having accomplished promising results in several tasks, still exhibits a large accuracy discrepancy between the results on A' and  $S_A^{test}$ . In other words, even though A'is frozen during the training of SampleNet, the sampling stage is overfitted onto the task network instead of the task itself.

To alleviate the issue of model-wise overfitting, we extend (1) and create a joint-training approach for a single task. Specifically, we take a set of weight-frozen models  $\{A_{i,j}\}_{i=1}^{k}, 1 < k << K_{A_i}$  as  $S_{A_i}^{train}$  and compute  $\mathcal{L}_{ST_i}$  as:

$$\mathcal{L}_{ST_i}(f_\theta) = \sum_{j=1}^k \mathcal{L}_{T_i}(A_{i,j}(f_\theta)).$$
(2)

It is critical to understand that all the frozen task models are under inference mode (*i.e.*, not significantly sabotaging computation power) and that only a very small number of task models (easily obtainable online or by self-training with different random initial weights) would bring significant improvements to the sampler's performance. We further show in Section 4.2 that a very small k > 1allows the sampling network generalise better across  $S_{A_i}$ , as  $S_{A_i}^{train}$  becomes a vicinity distribution rather than a single specific instance to  $S_{A_i}$ .

In addition to the joint  $\mathcal{L}_{ST_i}$ , we also update the weights with a simplification loss comprising the average and maximum nearest neighbour distance and a projection loss to enforce the probability of projection over the points to be the Kronecker delta function located at the nearest neighbour point (identical to the SampleNet loss [26]).

#### 3.3 Multi-Task Multi-Model Meta-Sampler Training

Instead of restricting ourselves to a single task (e.g., classification), we consider whether training the sampler over multiple tasks could lead to our vision of an almost-universal sampler. Broadly, we aim to extend the sampler beyond multimodel to multi-task, such that given any task  $T_i \in S_T$ , where  $S_T$  is a set of tasks, a good initial starting point could be achieved for the sampler. In this way, adapting or fine-tuning to a particular task (which may even be beyond the known set) will be rapid and cheap.

To tackle this, we draw inspiration from the MAML framework and propose a meta-learning approach for rapidly adaptive sampling [10]. In essence, we aim to utilise the set of  $S_{A_i}^{train}$  to mimic the best gradients in learning a particular



Fig. 2. The pipeline of the proposed meta-sampling. The illustration exemplifies the pretraining with multiple tasks through our meta-training strategy, then fitting onto a single task with our joint-training mechanism.

task  $T_i$  for meta-optimisation, such that given any task  $T_i \in S_T$  or even beyond the known set of tasks, the MAML network can quickly converge within a few iterations and without additional training of the task networks.

The joint-training procedure discussed in the previous section motivates that a particular task is better solved with a set of task networks instead of just one — we transfer this idea to the meta-optimisation such that the sampler is adaptive to a number of tasks instead of just one. Formally, we first optimise the adaptation of  $f_{\theta}$  to  $T_i \in S_T$  by updating the parameters  $\theta$  to  $\theta'_{i,j}$  for every  $A_{i,j}$  through the gradient update:

$$\theta_{i,j}' = \theta - \alpha \nabla \mathcal{L}_{T_i}(A_{i,j}(f_\theta)), \tag{3}$$

where  $\alpha$  is the step size hyperparameter. Similar to MAML, we can directly extend the single gradient update into multi-gradient updates to optimise the effectiveness of  $\theta'_{i,j}$  on  $T_i$ .

With the inner update (3), we then follow the meta-optimisation procedure through a stochastic gradient descent:

$$\theta = \theta - \beta \nabla \sum_{i=1}^{K_T} \sum_{j=1}^k \mathcal{L}_{T_i}(A_{i,j}(f_{\theta'_{i,j}})), \qquad (4)$$

where  $\beta$  is the meta step size hyperparameter that could either be fixed or accompanied with annealings. Note that we apply the single task loss in the inner update (3) but sum all losses from all weights to resemble a task in the metaupdate (4). Section 4.3 shows that our meta-optimisation design is sufficient in learning tasks for rapid adaptation. Simplification and projection losses are also directly optimised at this stage. They are however directly updated rather than included in the meta-update fashion as they are task-agnostic.

# 3.4 Overall Pipeline: Pretrained Meta-Sampler to Task Adaptation

We describe the overall training pipeline of the proposed meta-sampler (Figure 2) as the following:

**Pretrained Meta-Sampler:** Our pipeline begins with training a meta-sampler. First, we take a set of tasks  $S_T$  (*e.g.*, shape classification, reconstruction, retrieval) and their corresponding task networks  $S_{A_i}$  for every  $T_i \in S_T$  (pretrained on the unsampled point clouds). Next, we freeze all their original weights and perform our meta-sampler training as illustrated in Section 3.3 to obtain a pre-trained meta-sampler.

**Rapid Task Adaptation:** The meta-training attempts to optimise  $\theta$  to a position optimal to learn any task  $T_i$ . Therefore, to adapt to a particular task, we can simply take the pretrained weights of the meta-sampler and fine-tune it with the joint-training strategy as illustrated in 3.2 along with the previously proposed simplification and projection loss.

**Disjoint Task Networks for Pretraining and Training:** Realistically, one should be able to directly obtain a pretrained meta-sampler without the task networks and fit to their own networks. To mimic such real-world constraints, we ensure that the meta pretraining and joint-training use disjoint sets of networks — both of which are unseen during testing.

### 4 Experiments

Our empirical studies comprise two major components. First, we evaluate the performance of the proposed joint-training scheme against prior training methodologies on representative individual tasks. Afterward, we justify the versatility and robustness of the meta-sampler by measuring its adaptiveness across different tasks, models, and datasets.

#### 4.1 Experimental Setup

To comprehensively evaluate the performance of our meta-sampler, we extract a set of representative tasks on 3D point clouds, including shape classification, reconstruction, and shape retrieval. Note that all experiments are conducted on the ModelNet40 [42] (except for ShapeNet [4] used in transferring dataset analysis) to ensure fair evaluation (*i.e.*, without introducing additional information during meta-sampler training). The detailed experimental settings (*i.e.*, task network architecture, task loss) are described as follows:

**Shape Classification.** This is a fundamental task in 3D vision to determine the shape categories of a given point cloud. The task network set  $\{A_{i,j}\}$  are pretrained PointNets [34] with random and distinct weight initialisations, and the validation accuracy converges to 89% to 90%.  $\mathcal{L}_{T_i}$  is the vanilla binary cross-entropy (BCE) loss for classification.

**Reconstruction.** This task aims to reconstruct the complete 3D shape from partial point sets. Following [26], the goal of sampling for this task is to preserve

Table 1. Joint v.s. Single Task Network Training. Single and Joint denotes the SampleNet trained through the originally proposed single task network approach [26] and through our proposed multi-model single-task training (k = 3), respectively. Bold texts denote best results.

	Class	sificatio	n (Accura	acy ↑)		Reconstruction (CD $\downarrow)$	
Sampling Ratio	FPS	$\operatorname{SNet}[8]$	Single[26]	Joint	Sampling Ratio	Single <sup>[26]</sup>	Joint
8	70.4%	77.5%	83.7%	88.0%	8	3.29	3.05
16	46.3%	70.4%	82.2%	85.5%	16	3.32	3.15
32	26.3%	60.6%	80.1%	81.5%	32	3.61	3.37
64	13.5%	36.1%	54.1%	61.6%	64	4.43	4.31

	Shape Retrieval (Accuracy $\uparrow$ )							
	4-w	ay	10-way		20-way			
Sampling Ratio	Single[26]	Joint	Single[26]	Joint	Single[26]	Joint		
8	99.6%	99.7%	96.3%	98.3%	95.9%	96.7%		
16	98.7%	99.1%	94.0%	96.7%	89.5%	91.9%		
32	97.2%	97.5%	91.4%	91.5%	82.9%	84.6%		
64	92.5%	94.6%	79.5%	$\mathbf{84.6\%}$	67.0%	71.0%		

*n* key points that could be reconstructed to the original unsampled point clouds. For this task,  $S_{A_i}$  is a set of Point Completion Networks (PCN) [46] trained in an autoencoder fashion to minimise the Chamfer Distance (CD) between the input and output points. We select the PCN architecture owing to its encoder and decoder mechanism that doesn't take in any structural assumptions (e.g.,symmetry), making it suitable for reconstruction even when missing points are randomly distributed upon the entire shape instead of a particular part.  $\mathcal{L}_{T}$ is the two-way CD between the inputs and predicted outputs; all networks are pretrained with the loss to the chamfer distance of around  $3 \times 10^{-4}$ . Unlike other tasks where projection restricts the feature information, reconstruction is benefitted from projections. And thus, our results adopt the pre-projected points for both single and joint training. Using projected points boosts results for both. Shape Retrieval. Given a sampled point cloud, the goal is to match it with the shifted/rotated original point cloud given N options (similar to the N-way evaluation in few-shot settings). Due to the existence of hard negative pairs (point clouds of the same class), this task requires more advanced learning of fine-grained features compared with the pure shape classification. In this case,  $S_{A_i}$  is a set of Siamese PointNets inspired by [24] pretrained on unsampled point clouds matching.  $\mathcal{L}_{T_i}$  is a BCE loss where the ground truth is set to 1 if the point cloud is a shifted/rotated version of the other and vice versa. All networks are pretrained to achieve 100% accuracy on the simple 4-way evaluation.

#### 4.2 Performance Evaluation on Individual Tasks

To justify the effectiveness of the proposed multiple-model training scheme, we present the quantitative comparison of incorporating multiple models training and the traditional SampleNet single-model training strategy in all three individual tasks on the ModelNet40 dataset [42]. We adopt the official train and

	Number of Task Models $(k)$						
	k = 1	k = 2	k = 3	k = 4	k = 5		
Accuracy	80.0%	81.3%	81.7%	82.8%	83.4%		

Table 2. Classification accuracy when Increasing k. Sampling Ratio is 32.

test splits in this dataset, and follow [8,26] to pretrain all task networks with the original point clouds (1024 points by default). All the task models are under inference mode during the training of the sampling network. We evaluate our sampling on different sampling ratios calculated as 1024/n, where *n* is the number of outputted points from the sampler. Note that for reconstruction and shape retrieval tasks, we were adopting a different dataset and task to prior networks. Much work is required for the adaption and thus we only compare with the previously proposed state-of-the-art SampleNet.

Shape Classification. As shown in Table 1, the classification performance achieved with our joint-training scheme consistently outperforms the single SampleNet and previous sampling strategies such as FPS across all sampling ratios. In particular, the classification accuracy achieved with our joint-training mechanism under a sampling ratio of 8 is very close to the upper bound accuracy (88.0% vs. 89.5%) achieved without any sampling, verifying the effectiveness of our joint-training strategy. We also notice that the performance gap between the proposed method and others is widening with a growingly aggressive sampling rate (*e.g.*, at sampling ratio 64 with 16 points left for the point clouds). This further demonstrates the superiority of the proposed training mechanism under aggressive sampling ratios.

**Reconstruction.** The effect of joint-training on reconstruction follows a similar trend to shape classification, outperforming other strategies in terms of the CD across all sampling ratios. The improvement seems to be consistent across all sampling ratios, further exhibiting the effectiveness of joint-training.

**Shape Retrieval.** Our shape retrieval results are presented under the *N*-way few-shot settings (N = 4, 10, 20). It is clear that the joint-training scheme achieves better results compared with the single SampleNet training strategy. Specifically, the advantages of our proposed joint-training scheme is more prominent with the increase of the sampling ratios, suggesting that our sampling schemes can preserve points that have high similarity to the original point cloud.

**Faster Convergence with Multiple Models.** Considering the sampler is exposed to more task networks during training, it is expected for the sampler to converge to stabilised accuracies with a shorter time span of training. Our empirical study generally aligns with this idea. Our joint-training taking usually around 40 epochs to converge as opposed to around 60 for single-task training.

The Impact of the Number of Task Models k. We further dive into the correlation between the number of networks k for joint-training and classification



Fig. 3. The performance comparison in classification (row  $1 \uparrow$ ), reconstruction (row  $2 \downarrow$ ), and shape retrieval (row  $3 \uparrow$ ) with/without the meta-sampler at the initiation of training. All graphs begin after one epoch. Red is ours.

accuracy. Table 2 shows the classification results under the randomly selected sampling ratio of 32 when we slowly increase the number of task networks for ensemble. A clear trend of increments continues as k increases, implying that the wider the set of training networks the better the approximation is to the entire task distribution. Nonetheless, such increments suffer from the trade-off in computational resources time and memory-wise. We stick with 3 networks as the standard for joint-training and in later experiments unless specified.

### 4.3 Versatility of Meta-Sampler

Versatility is a broad term with multiple dimensions requiring evaluation. To fully realize this, we begin with the critical evaluation of our meta-sampler's adaptiveness on tasks included in the meta-optimisation step. We then extend to the more challenging scenarios of changing model architecture, dataset distribution, and ultimately tasks distinct from the ones used for meta-training. All experiments are conducted with hyperparameters  $\alpha$  and  $\beta$  set to 1e-3. All plots begin after one epoch as single SampleNets are not pretrained.



Fig. 4. Accuracy of PointNet++ for classification with joint-training. The results with and without the pretrained meta-sampler on PointNet is presented for sampling ratios at 8 and 16. All graphs begin after one epoch.

**Converging to Meta-Tasks.** To investigate the impact of the meta-sampler for the performance of meta tasks, we conduct several groups of experiments in this section for the three tasks used in our meta-optimisation, including shape classification, reconstruction, and shape retrieval. We compare the task performance achieved with/without the pretrained meta-sampler under different sampling ratios in Figure 3. Specifically, for our meta-sampling, we first deliberately select a bunch of task models unseen during meta-optimisation to fine-tuned the meta-sampler with the joint-training scheme, then evaluate the task performance with the sampled point clouds.

As shown in Figure 3, we separately compare the task performance as the training progresses with/without our pretrained meta-sampler for different meta-tasks. It is clear that as the sampling ratio increases (*i.e.*, the task is more challenging), joint-training without meta-sampler starts at lower performance and requires more iterations to converge to a stable result. By contrast, our pretrained meta-sampler, while acting as a sampler trained without any task loss at 0th epoch (similar to FPS), adapts to the task within one epoch across all sampling ratios and achieves higher accuracies after 10 epochs in most cases.

There are also two intriguing points we would like to address within this empirical study. First, we observe a relatively large fluctuation of the performance for shape classification and shape retrieval — a phenomenon we conjecture to be owing to the distribution shift between training and testing sets. Second, we notice that our meta-sampler not only converges faster, but also pushes the upper bound in some cases. For example, our shape retrieval results at sampling ratio 16 achieved the accuracy of 96.9% (the upper bound of training from scratch is 96.7%) on the 20th epoch (not plotted in the figure). This infers that by learning how to adapt, the sampling model could potentially also be trained to learn better. However, such occurrences do not take place at all times.

**Transferring Model Architectures.** Prior experiments focus on training from the meta-sampler with task networks of identical architecture but different weights. To further explore the versatility of our meta-sampler, we transfer the jointtraining networks from PointNets to PointNet++ [34,35]. All the networks are pretrained until convergence (*i.e.*, around 92% accuracy on unsampled point clouds). Constrained by the original implementation of PointNet++ (*i.e.*, point set abstraction layer) extracting features from the 32 points neighborhoods, we



Fig. 5. Transfer to ShapeNet. We adopt the ModelNet pretrained meta-sampler to fit onto the ShapeNet dataset for classification. All graphs begin after one epoch.

only evaluate our meta-sampler upon the sampling ratio of 8 and 16, where the remaining point cloud size is greater than 32.

The achieved results are plotted in Figure 4. It is apparent that better performance is achieved using our meta-sampler, with a higher starting point and fast convergence speed under both sampling ratios. This further demonstrates the capacity of our meta-sampler in adapting different model architectures. Interestingly, we also notice that the classification accuracy of PointNet++ achieved on sampled points drops significantly compared with that of unsampled point clouds, especially under the sampling ratio of 16. This is likely because FPS is progressively used in each encoding layer. In this case, by adding a SampleNet in front of PointNet++, we are implicitly "double sampling" and leaving very few features for the abstraction layer.

**Transferring Datasets.** To verify that our ModelNet40 [42] pretrained sampler isn't applicable to just the data distribution it was expose to, we measure the effectiveness when using the same pretrained model for the same task but on a different dataset. Specifically, we evaluate the classification performance of our meta-sampler on the ShapeNetCoreV1 Dataset [4], which comprises point cloud objects from 16 different shape categories. Specifically, we still adopt the PointNet [34] architecture and trained three networks following the best practise, while these models can achieve around 98% accuracy on unsampled point clouds. We then show the training progress achieved by using our joint training scheme with/without the meta-sampler. As shown in Figure 5, although the performance is similar when the sampling ratio is small (easier), we can clearly notice that the model without our meta-sampler starts at a much lower performance. By contrast, the model with our pretrained meta-sampler converges much faster (even within one epoch) and is more stable. As such, this empirical study can well prove that our meta-sampler can adapt to a completely disparate dataset distribution and serve as a better and more stabilised starting point.

**Transferring Beyond Meta-Tasks.** Finally, we extend to the most challenging question of whether the proposed meta sampler can generalize to unseen tasks, i.e. tasks that are not included in the meta-optimisation step. This is highly challenging since different tasks inevitably have distinct preferences in sampled points. However, this is also a critical step to validate whether the proposed meta sampler could be the universal pretrained module for all point cloud tasks.



Fig. 6. The performance comparison in point cloud registration with/without the meta-sampler. Left: Rotational error comparison. Right: Standard deviation of rational error per epoch. All graphs begin after one epoch.

We evaluate our meta-sampler on the point cloud registration — the task of finding the spatial transformation between two point clouds. Here, we follow the standard train-test split of PCRNet [37] to obtain pairs of source and template point clouds with templates rotated by three random Euler angles of  $[-45^{\circ}, 45^{\circ}]$ and translated with a value in the range [-1,1].  $\mathcal{L}_{T_i}$  is the CD between the source point cloud and the template point cloud with our predicted transformation. We first train three PCRNets to achieve the rotation error of around 7-9 degrees on unsampled point clouds, then freeze the PCRNet weights and perform the proposed joint-training scheme under the conditions with/without the pretrained meta-sampler under the sampling ratio of 32. We ran each setting three times and show the mean and standard deviation of rotational error during training in Figure 6. Even though the task objective (registration), task network (PCRNet), and even the dataset itself (pairs of point clouds from ModelNet40) are unforeseen during our meta-optimisation, we can still notice two subtle vet solid performance differences adopting our meta-sampler: 1) The pretrained model generally converges faster during the initiation of training and 2) The pretrained model is much more stabilised and improves consistently compared to the model trained from scratch that exhibits a large variance throughout every epoch (other sampling ratios also perform similarly in stabilisation).

# 5 Conclusion

We propose a learnable meta-sampler and a joint-training strategy for taskoriented, almost-universal point cloud sampling. The proposed multi-model joint training scheme on SampleNet achieved promising performance for various point cloud tasks, and the meta-sampler has empirically shown to be effective and stabilising when transferred to any tasks incorporated during meta-optimisation, even extending to unseen model architectures, datasets, and tasks. We hope our pretrained meta-sampler can be used as a plug-and-play module and widely deployed to point cloud downstream tasks to save computational resources.

Acknowledgements. This work was partially supported by Amazon AWS, Oxford Singapore Human-Machine Collaboration Programme, and EPSRC ACE-OPS grant.

# References

- 1. Antoniou, A., Edwards, H., Storkey, A.J.: How to train your MAML. In: ICLR (2019)
- Ao, S., Hu, Q., Yang, B., Markham, A., Guo, Y.: Spinnet: Learning a general surface descriptor for 3D point cloud registration. In: CVPR (2021)
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Gall, J., Stachniss, C.: Towards 3d lidar-based semantic scene understanding of 3d point cloud sequences: The semantickitti dataset. Int. J. Robotics Res. (2021)
- Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: ShapeNet: An information-rich 3D model repository. arXiv preprint arXiv:1512.03012 (2015)
- Chen, S., Tian, D., Feng, C., Vetro, A., Kovacevic, J.: Fast resampling of 3d point clouds via graphs. arXiv preprint arXiv:1702.06397 (2017)
- Chen, X., Chen, B., Mitra, N.J.: Unpaired point cloud completion on real scans using adversarial training. In: ICLR (2020)
- Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV (2016)
- 8. Dovrat, O., Lang, I., Avidan, S.: Learning to sample. In: CVPR (2019)
- 9. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: CVPR (2017)
- Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: Precup, D., Teh, Y.W. (eds.) ICML (2017)
- 11. Gojcic, Z., Zhou, C., Wegner, J.D., Wieser, A.: The perfect match: 3d point cloud matching with smoothed densities. In: CVPR (2019)
- Goyal, A., Law, H., Liu, B., Newell, A., Deng, J.: Revisiting point cloud shape classification with a simple and effective baseline. In: Meila, M., Zhang, T. (eds.) ICML (2021)
- Groh, F., Wieschollek, P., Lensch, H.P.A.: Flex-convolution million-scale pointcloud learning beyond grid-worlds. In: Jawahar, C.V., Li, H., Mori, G., Schindler, K. (eds.) ACCV (2018)
- Guo, Y., Wang, H., Hu, Q., Liu, H., Liu, L., Bennamoun, M.: Deep learning for 3d point clouds: A survey. IEEE transactions on pattern analysis and machine intelligence 43(12), 4338–4364 (2020)
- 15. Gupta, A., Eysenbach, B., Finn, C., Levine, S.: Unsupervised meta-learning for reinforcement learning. arXiv preprint arXiv:1806.04640 (2020)
- 16. Hamdi, A., Giancola, S., Li, B., Thabet, A.K., Ghanem, B.: MVTN: multi-view transformation network for 3d shape recognition. ICCV (2021)
- Hu, Q., Yang, B., Fang, G., Guo, Y., Leonardis, A., Trigoni, N., Markham, A.: Sqn: Weakly-supervised semantic segmentation of large-scale 3d point clouds with 1000x fewer labels. arXiv preprint arXiv:2104.04891 (2021)
- Hu, Q., Yang, B., Khalid, S., Xiao, W., Trigoni, N., Markham, A.: Sensaturban: Learning semantics from urban-scale photogrammetric point clouds. International Journal of Computer Vision pp. 1–28 (2022)
- Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: RandLA-Net: Efficient semantic segmentation of large-scale point clouds. In: CVPR (2020)
- 20. Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: Learning semantic segmentation of large-scale point clouds with random sampling. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)

- 16 T. Cheng et al.
- 21. Huang, C., Cao, Z., Wang, Y., Wang, J., Long, M.: Metasets: Meta-learning on point sets for generalizable representations. In: CVPR (2021)
- 22. Huang, S., Gojcic, Z., Usvyatsov, M., Wieser, A., Schindler, K.: Predator: Registration of 3d point clouds with low overlap. In: CVPR (2021)
- Jabri, A., Hsu, K., Gupta, A., Eysenbach, B., Levine, S., Finn, C.: Unsupervised curricula for visual meta-reinforcement learning. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) NeurIPS (2019)
- Koch, G., Zemel, R., Salakhutdinov, R., et al.: Siamese neural networks for oneshot image recognition. In: ICML Workshop (2015)
- 25. Landrieu, L., Simonovsky, M.: Large-scale point cloud semantic segmentation with superpoint graphs. In: CVPR (2018)
- Lang, I., Manor, A., Avidan, S.: Samplenet: Differentiable point cloud sampling. In: CVPR (2020)
- Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: PointCNN: Convolution on X-transformed points. In: NeurIPS (2018)
- Liu, X., Han, Z., Liu, Y., Zwicker, M.: Point2sequence: Learning the shape representation of 3d point clouds with an attention-based sequence to sequence network. In: AAAI (2019)
- Liu, Z., Tang, H., Lin, Y., Han, S.: Point-voxel CNN for efficient 3d deep learning. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) NeurIPS (2019)
- 30. Ma, X., Qin, C., You, H., Ran, H., Fu, Y.: Rethinking network design and local geometry in point cloud: A simple residual MLP framework. In: ICLR (2022)
- Mandikal, P., L., N.K., Babu, R.V.: 3d-psrnet: Part segmented 3d point cloud reconstruction from a single image. In: Leal-Taixé, L., Roth, S. (eds.) ECCV Workshop (2018)
- 32. Nezhadarya, E., Taghavi, E., Razani, R., Liu, B., Luo, J.: Adaptive hierarchical down-sampling for point cloud classification. In: CVPR (2020)
- Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: ICCV (2019)
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR (2017)
- Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep hierarchical feature learning on point sets in a metric space. In: NeurIPS (2017)
- Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. In: ICLR (2018)
- Sarode, V., Li, X., Goforth, H., Aoki, Y., Arun Srivatsan, R., Lucey, S., Choset, H.: Pcrnet: Point cloud registration network using pointnet encoding (2019)
- Tatarchenko, M., Dosovitskiy, A., Brox, T.: Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In: ICCV (2017)
- Wallace, B., Hariharan, B.: Few-shot generalization for single-image 3d reconstruction via priors. In: ICCV (2019)
- 40. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph CNN for learning on point clouds. ACM TOG (2019)
- Wu, W., Qi, Z., Fuxin, L.: PointConv: Deep convolutional networks on 3D point clouds. In: CVPR (2018)
- 42. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: CVPR (2015)
- 43. Xiang, T., Zhang, C., Song, Y., Yu, J., Cai, W.: Walk in the cloud: Learning curves for point clouds shape analysis. In: ICCV (2021)

- 44. Xu, Q., Sun, X., Wu, C., Wang, P., Neumann, U.: Grid-gcn for fast and scalable point cloud learning. In: CVPR (2020)
- 45. Yu, X., Tang, L., Rao, Y., Huang, T., Zhou, J., Lu, J.: Point-bert: Pretraining 3d point cloud transformers with masked point modeling. arXiv preprint arXiv:2111.14819 (2021)
- 46. Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: PCN: point completion network. In: 3DV (2018)
- 47. Zhang, Y., Hu, Q., Xu, G., Ma, Y., Wan, J., Guo, Y.: Not all points are equal: Learning highly efficient point-based detectors for 3d lidar point clouds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2022)
- 48. Zhong, J.X., Zhou, K., Hu, Q., Wang, B., Trigoni, N., Markham, A.: No pain, big gain: Classify dynamic point cloud sequences with static models by fitting featurelevel space-time surfaces. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2022)