

COLA: COarse Label pre-training for 3D semantic segmentation of sparse LiDAR datasets

Jules Sanchez¹

Jean-Emmanuel Deschaud¹

François Goulette^{2,1}

¹Centre for Robotics, Mines Paris - PSL, PSL University,
75006 Paris, France
firstname.surname@minesparis.psl.eu

²UZIS, ENSTA Paris, Institut Polytechnique de Paris,
1120 Palaiseau, France
firstname.surname@ensta-paris.fr

Abstract—Transfer learning is a proven technique in 2D computer vision to leverage the large amount of data available and achieve high performance with datasets limited in size due to the cost of acquisition or annotation. In 3D, annotation is known to be a costly task; nevertheless, pre-training methods have only recently been investigated. Due to this cost, unsupervised pre-training has been heavily favored.

In this work, we tackle the case of real-time 3D semantic segmentation of sparse autonomous driving LiDAR scans. Such datasets have been increasingly released, but each has a unique label set. We propose here an intermediate-level label set called coarse labels, which can easily be used on any existing and future autonomous driving datasets, thus allowing all the data available to be leveraged at once without any additional manual labeling.

This way, we have access to a larger dataset, alongside a simple task of semantic segmentation. With it, we introduce a new pre-training task: coarse label pre-training, also called COLA.

We thoroughly analyze the impact of COLA on various datasets and architectures and show that it yields a noticeable performance improvement, especially when only a small dataset is available for the finetuning task.

Index Terms—LiDAR, semantic segmentation, pre-training

I. INTRODUCTION

LiDAR-based deep learning has been gaining a lot of traction in recent years, especially for autonomous driving. Thus, an increasing amount of data has been released for scene understanding, such as KITTI [1] and Argoverse 2 [2] for detection or SemanticKITTI [3] and nuScenes [4] for semantic segmentation.

So far, 3D semantic segmentation methods have focused on architecture design and tuning rather than data-driven approaches to improve performance. Some works have started to emerge on *domain adaptation* and *transfer learning* to ensure good performances when the amount of target annotations is limited or nonexistent.

Transfer learning is the ability to reuse previously learned information in a deep learning model. We can decompose transfer learning into two parts: first, a pre-training task for which there is usually a large amount of data available, and then a finetuning task, which is the task we want our model to achieve in the end. The pre-training task helps the network to

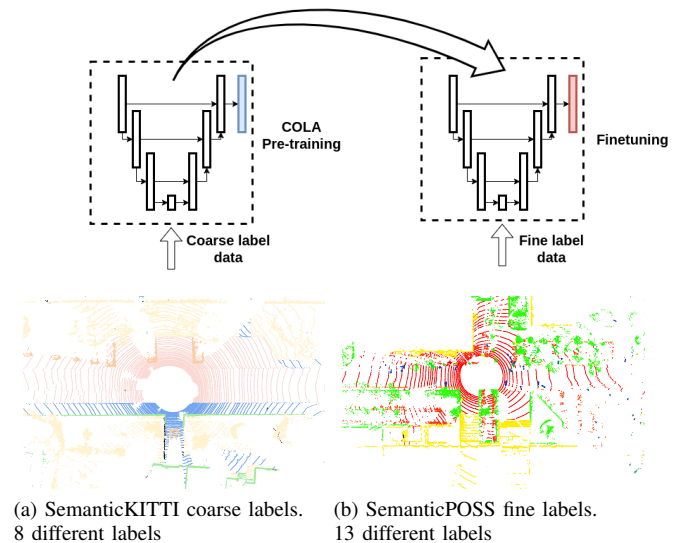


Fig. 1: COLA pipeline.

learn useful representations that should speed up the training and improve the results of the finetuning task.

Domain adaptation is a subset of transfer learning, aiming specifically to transfer knowledge when the pre-training and finetuning tasks are similar. However, it is applied in a specific scope where the pre-training data and the finetuning data are available simultaneously.

Thus far, most of the work done in 3D pre-training has been leveraging contrastive learning (see subsection II-B).

We propose a new pre-training method that leverages a set of heterogeneous source datasets thanks to a simple label unification techniques for LiDAR point cloud semantic segmentation. Contrary to domain adaptation methods, access to the pre-training data is not required when finetuning our models. This way, we can provide off-the-shelf pre-trained models that can be directly finetuned on any autonomous driving semantic segmentation dataset and systematically increase performance

on the finetuning task.

We introduce an intermediate-level label set called coarse labels that can be applied to any autonomous driving semantic segmentation dataset. We demonstrate that using this intermediate label set enables us to use all datasets together to pre-train neural architectures and improve performance after finetuning.

Furthermore, we highlight the limitation of existing contrastive-based methods for the task of autonomous driving semantic segmentation by looking at CSC results. For a fairer comparison, we also train our own contrastive learning-based method by following PointContrast [5] and applying it to sparse autonomous driving data. We call it SPC, for Sparse PointContrast.

The mapping from existing available datasets to the proposed label set alongside links to our pre-trained models using the largest source dataset can be found at <https://github.com/JulesSanchez/CoarseLabel>.

In practice, our contributions are as follows:

- Studying the limitations of current contrastive based methods for autonomous driving 3D semantic segmentation.
- Introducing a label set specific to the task of semantic segmentation of autonomous driving (sparse LiDAR scan) generic enough to leverage every dataset at once.
- Thoroughly experimenting on a new pre-training that leverages this new label set. This method, called COLA (COarse LABEL pre-training) is shown in Figure 1. It is quantitatively and qualitatively studied, showing a mIoU improvement of up to +6.4% on the target dataset.

II. RELATED WORKS

A. 3D autonomous driving semantic segmentation

Semantic segmentation for autonomous driving has two main challenges: achieving real-time inference and dealing with the unordered structure of the LiDAR point clouds.

As such, several types of methods have emerged: using only a permutation invariant operation [6] [7] [8], redefining the convolution [9] [10] [11] [12], projecting the point clouds in 2D with a range-based image [13] [14] and bird’s-eye-view images [15] or leveraging sparse convolution-based architectures as introduced by MinkowskiNet [16], under the name Sparse Residual U-Net (SRU-Net).

SRU-Nets have been used as the backbone for more refined methods [16], [17], [18], [19] that are the best performing methods on the various benchmarks.

B. 3D transfer learning

As mentioned in the introduction, 3D transfer learning has mainly relied on unsupervised contrastive pre-training tasks such as in PointContrast [5], CSC [20], DepthContrast [21] or CrossPoint [22]. PC-FractalDB [23] proposes an alternative to contrastive learning by learning fractal geometry as the pre-training task. Point-BERT [24] looks into pre-training of point cloud transformers specifically.

All those methods look into either shape understanding, indoor dense scene understanding or outdoor dense scene

understanding as they rely on the large availability of such data [25] [26] [27] to perform their pre-training task.

Fewer methods have looked into pre-training for outdoor sparse point clouds. GCC-3D [28] applies a two-step self-supervised method, first learning contrastive features and then leveraging the sequential nature of the data to generate pseudo instances that can be used to learn semantic clustering. ProposalContrast [29] proposes learning contrastive information at the proposal level rather than point or voxel levels. These outdoor methods are specifically designed for object detection and are not evaluated on any semantic segmentation finetuning tasks.

Finally, only SegContrast [30] has looked into the same finetuning task as us, highlighting a gap in the study of the pre-training effect on this case.

C. Leveraging heterogeneous datasets

Using several datasets at once for supervised learning is a nontrivial task. Datasets, even for a similar task, usually have vastly different label sets. This variation can come from the type of elements in the scene or annotation choices.

To leverage as much of these data as possible, several techniques have been implemented. Some approaches have chosen to only leverage common labels across datasets [31] [32] or the union [33], to have one shared encoder and one decoder for each dataset [34] [35], to leverage label relationships [36] [37] [38], finally some have chosen to unify annotations through remapping [39] and re-annotation [40].

Category-shift multi-source domain adaptation, as defined in the proposed taxonomy of [41], aims to tackle a similar issue as it leverages a set of heterogeneously labeled source datasets to improve the performance of a target dataset presenting a different label set, as in Deep Cocktail Network [42] and mDALU [41].

Inspired by the first type of method, COLA focuses on leveraging the annotations of all the pre-training data without any additional human cost. Furthermore, contrary to domain adaptation methods, COLA provides off-the-shelf pre-trained models that can be finetuned without requiring access to the pre-training data.

III. COLA

A. Shortcomings of simple contrastive approaches

Contrastive learning has been the main approach for the pre-training of 3D models. As such, we want to investigate them for our specific finetuning task. As one of the founding and reference methods for dense point clouds, we investigate specifically CSC [20].

In Table I, we show the results after finetuning on three major autonomous driving 3D semantic segmentation datasets when the model is pre-trained with CSC and with a LiDAR-based adaptation Sparse PointContrast (SPC). Whereas CSC is trained on dense point clouds (ScanNetV2 [25]), SPC follows the same training methodology but with 100K LiDAR scans, leveraging mainly KITTI-360 [43]. We see that CSC results in small to moderate improvement (from +0.2% to +1.1%),




Fig. 2: The proposed coarse labels and their mapping to SemanticKITTI fine labels.

whereas SPC pre-training systematically decreases results in the finetuning task.

We have not reported SegContrast [30] results in Table I, even though it is the method working in the closest setup to ours as they are using a different model. One of the differences between SPC and SegContrast is the choice of pre-training data. While SegContrast pre-trains with the same dataset as the finetuning one, here, in SPC we pre-train with datasets different from the finetuning data.

In total, the SPC pre-training took the equivalent of 50 GPU days of training. In comparison, CSC takes approximately 15 GPU days to be trained. This discrepancy stems from the difference in scan sizes regarding number of points and their range in meters. The duration of the SPC pre-training makes it unrealistic to perform it on 1M scans such as the ONCE dataset [44], which is 10 times larger than the dataset used for SPC, as it would amount to 500 GPU days.

	No pre-training	CSC	SPC
nuScenes	67.2	68.3 (+1.1)	66.3 (-0.9)
SemanticKITTI	50.5	50.7 (+0.2)	48.9 (-1.6)
SemanticPOSS	55.0	55.5 (+0.5)	54.4 (-0.6)

TABLE I: Test set mIoU for each target dataset. No pre-training, dense contrastive pre-training (CSC), and sparse contrastive pre-training (SPC) with the SRU-Net architecture.

We see that contrastive pre-training at the moment is not suited for our specific case, and scaling the method to leverage more data would require significant time and computing power. While the shortcomings of CSC, compared to results in their paper, can be explained by the variation of data topology between dense and sparse data. For SPC, we hypothesize that information at a point level is too difficult to extract properly for sparse LiDAR scans and that more meaningful structures should be used to learn insightful geometry, as it is done in SegContrast.

We believe that discarding human annotations to rely solely on unsupervised pre-training is a waste of information, and supervised methods are still relevant.

We propose a novel approach using a supervised pre-train task as opposed to the self-supervised and unsupervised-centered methods that are the typical approaches.

B. Available semantic segmentation data

Since the release of SemanticKITTI in 2019, other autonomous driving semantic segmentation LiDAR datasets have

been available. We focus on five different datasets, which are the most important at the time of the writing due to the variety in size, scene types, and hardware used for acquisition. As seen in Table II, limiting a method to one dataset would be omitting a large amount of data. However, thus far, it is impossible to use them all at once, as there is no unique fine label set.

Dataset	Size	Scene type	#Labels
KITTI-360 [43]	80K	Suburbs	27
nuScenes [4]	35K	Urban	16
SemanticKITTI [3]	23K	Suburbs	19
PandaSet [45]	6K	Urban & Road	37
SemanticPOSS [46]	3K	Campus	13

TABLE II: Summary of the datasets used in the different experiments. The size corresponds to the number of labeled scans. Labels correspond to the labels used at the inference time.

Despite this obvious obstacle, a study of the datasets shows that labels are organized in a tree-like structure with more or less refined labels, similar to the one highlighted by [36]. Moreover, higher levels of trees are more uniform across datasets.

C. Coarse Labels

Based on these observations, we propose a set of **coarse labels** such that we can easily map *fine* labels from any dataset to this one. These coarse labels define categories that retain information relevant to the end task, are consensual across all existing and future such datasets, and display little ambiguity; furthermore, they come at no additional annotation cost.

Besides these technical advantages, we believe that those coarse labels can still represent effectively the major elements of a typical road scene such that a model trained with them would still be relevant for autonomous driving. The coarse labels we propose are as follows:

- Driveable ground covers areas where cars are expected to be driving.
- Other ground corresponds to other large planar surfaces, where pedestrians can be found.
- Structure corresponds to fixed human-made structures, which can be hard to subdivide into smaller groups.
- Vehicles are the typical road users.
- Nature corresponds to the vegetation-like components of the scene.
- Living being groups humans and animals alike, who can interact with all the other elements.

- Dynamic objects are objects that can typically move, such as strollers.
- Static objects are objects that are fixed, such as poles.

An illustration of the labels and the mapping of the fine labels from SemanticKITTI to the proposed label set can be found in Figure 2. All mappings are available in our public GitHub repository¹.

D. COLA (Coarse Label pre-training)

Given this new coarse label set, we can perform semantic segmentation on all the data at once. This supervised learning is the proposed COarse LAbel (COLA) pre-training presented in the introduction. To reuse such a pre-trained model, only the last layer of the model needs to be replaced for finetuning, to account for the new number of labels (see Figure 1). Due to the size of the architecture, no layers need to be frozen at the time of finetuning. In this way, COLA is expected to learn representations tailored for the specific task of autonomous driving semantic segmentation. The validation mIoU of the pre-training reaches between 65% and 78%, depending on the target dataset, which is sufficient to conclude that the architecture has learned relevant features, confirmed in subsection III-E.

In practice, COLA took the equivalent of 10 GPU days of training time, resulting in a much shorter pre-training procedure than SPC.

We expect COLA to perform well for several reasons. First, seeing various sensor and environment settings in its training set will help generalize to new scenes and sensors. Second, aligning the pre-training and finetuning tasks to be very similar ensures the relevance of the learned features, both geometrically and semantically.

E. Pre-training feature analysis

In this section, we proceed to a feature analysis of the last layer, before the classification of a SRU-Net model after the COLA pre-training. We base the analysis on the SemanticPOSS dataset to study COLA when the largest amount of pre-training data is available.

To conduct the analysis, we give point clouds extracted from SemanticPOSS until we reach at least 200 points for each ground truth category and extract their associated last layer features. We then do a t-SNE analysis and plot the 2D projection in Figure 3, where points are colored according to either their coarse ground truth or fine ground truth. We observe the ability of the model to correctly cluster points in both the coarse label set and in the fine label set, even though it was not yet finetuned and had never seen any SemanticPOSS data. This demonstrates the ability of COLA to generalize the segmentation to unseen data and furthermore shows the relevance of learning coarse labels even when fine labels are the end goal.

We focus on a particular case of the COLA t-SNE analysis in Figure 3. In the right image, we highlighted (points

are circled in black) points belonging to either car or bike categories. We observe a highly desirable pattern, which is that COLA learned to distinguish bikes from cars feature-wise, even if they were not annotated as such in the coarse labels. This highlights that the pre-training method leverages both semantics and geometry, and thus, COLA does not rely on a perfect intermediate label set to learn relevant features.

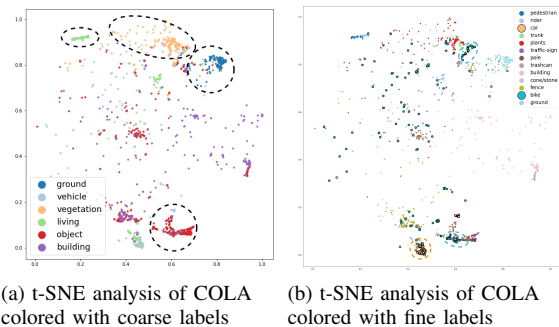


Fig. 3: t-SNE analysis of the last feature layers of SRU-Net trained with COLA. Inference performed on SemanticPOSS.

IV. EXPERIMENTS

In this section, we describe the experimental setup used to evaluate the performance of the proposed coarse label pre-training (COLA).

A. Experimental cases

For all our experiments, we split the data into two categories: the **target** data, which is the data we used for finetuning, on which we compute performance metrics, and the **pre-training** data. KITTI-360 is not used for pre-training in the case of the SemanticKITTI target as there is a significant geographic overlap of their scenes.

The four cases we chose correspond to experimenting on the two reference datasets in the community (nuScenes and SemanticKITTI), as well as on two limited data availability setups (SemanticPOSS and partial nuScenes). Additional details can be found in Table III. Partial nuScenes is a homemade set of datasets, which corresponds to subsets of nuScenes to provide access to a wider range of dataset sizes; otherwise, we would be limited to SemanticPOSS.

As highlighted in Table II, target datasets present different scene types, and furthermore, nuScenes and SemanticKITTI are acquired with sensors of very different resolutions. In this way, we can evaluate the methods in many settings.

We experiment on two architectures, SPVCNN [17] and SRU-Net [16] which are staples of the autonomous driving semantic segmentation field, either as state-of-the-art architecture or typical backbone.

To study COLA, we compute the mean Intersection over Union (mIoU) computed on an extracted test set. All experiments are performed on NVIDIA GeForce RTX 3090 GPUs.

¹<https://github.com/JulesSanchez/CoarseLabel>

Pre-training data	# Scans for pre-training	Target	# Size ratio
SK + K360 + SP + PS	110K	NS	25%
SK + K360 + PS + NS	140K	SP	2%
SP + NS + PS	45K	SK	33%
SK + K360 + SP + PS	110K	PartialNS	2.5%-25%
SK + K360 + SP + PS	110K	PartialSP	0.7%

TABLE III: Datasets used for pre-training depending on the target and the size ratio between the target set and the pre-training set. The various datasets are SemanticKITTI (SK), KITTI-360 (K360), nuScenes (NS), SemanticPOSS (SP) and PandaSet (PS)

B. Partial nuScenes and partial SemanticPOSS

We introduced partial nuScenes previously. Partial nuScenes is a subset of nuScenes, based on the number of scenes. We propose three levels: 10%, 25%, and 50%, which represent 95, 235, and 470 scenes (respectfully 1900, 4700, and 9400 scans), respectively, for training and validation, which are split in the same proportion as the full nuScenes 70/30.

We similarly introduce partial SemanticPOSS, composed of one sequence (00) for training and one for validation.

The objective of introducing such datasets is the ability to study the effect of COLA for a larger variety of scans available during the finetuning stage. We selected SemanticPOSS and nuScenes as they are the setup in which the pre-training data available is the largest.

C. Networks and parameters

We experimented on two architectures (SPVCNN [17] and SRU-Net [16]). All parameters were chosen based on the from-scratch training and simply re-used in the pre-trained setup.

1) *SRU-Net*: For the COLA pre-training, we used the SGD optimizer, with an initial learning rate of 0.4, a momentum of 0.9, and a cosine annealing scheduler. We had a batch size of 48 over 10 epochs.

For finetuning, whether the model has been pre-trained or not, we use the same parameters. We used the SGD optimizer, with an initial learning rate of 0.8, a momentum of 0.9, and a cosine annealing scheduler. We had a batch size of 36 over 30 epochs.

The model used is the same as the one used in [20].

2) *SPVCNN*: For the COLA pre-training and the finetuning, we used the SGD optimizer, with an initial learning rate of 0.24, a momentum of 0.9, and a cosine warmup scheduler. We had a batch size of 16 over 15 epochs.

The model used is the one provided by the original authors in the official GitHub repository [17].

All models are trained with the mixed Lovasz loss function and following the same data augmentation as [16].

V. RESULTS

In this section, we show the results from the thorough experiments described in the previous sections. They are divided into two parts. First, we look at results on full datasets (SemanticKITTI, SemanticPOSS, and nuScenes), and on each network architectures, and then, we investigate more precisely

the case of limited data with partial nuScenes. In each case, we show the mIoU over a test set, which we extract from the training set, representing between 10% and 20% of the available data. We compare all COLA results to the no pre-training case.

All mIoU gains that fall within ± 0.5 are due to the randomness of the training rather than the effect of a method.

A. Results on SemanticKITTI, nuScenes, and SemanticPOSS

	SRU-Net		SPVCNN	
	No pre-training	COLA (Ours)	No pre-training	COLA (Ours)
nuScenes	67.2	69.3 (+2.1)	65.9	66.2 (+0.3)
SemanticKITTI	50.5	52.4 (+1.9)	57.7	58.4 (+0.7)
SemanticPOSS	55	55.8 (+0.8)	49.9	53.8 (+3.9)

TABLE IV: Test set mIoU on each target dataset. No pre-training, and COLA pre-training.

Overall, the models pre-trained with the proposed COLA pre-training show significant improvement compared with the other methods, with a gain of up to +2% for SRU-Net and up to +4% for SPVCNN Table IV, in every case displaying significant improvement when compared with naive contrastive approaches (Table I).

In Figure 4, we study an example where the finetuned model show only a small improvement over the training from scratch. We see that, while the final mIoU are similar, it is reached much faster when pre-trained, and early epochs show some significant improvements.




Fig. 4: Validation mIoU curve for SRU-Net finetuned (green) and trained from scratch (pink) on SemanticPOSS. On the left, mIoU at each iteration, on the right best mIoU so far.

B. Results on partial nuScenes and partial SemanticPOSS

% scenes	# of scans	SRU-Net		SPVCNN	
		No pre-training	COLA (Ours)	No pre-training	COLA (Ours)
10%	1900	46.2	48.4 (+2.2)	44.9	51.3 (+6.4)
25%	4700	58.0	60.0 (+2.0)	52.5	55.6 (+3.1)
50%	9400	62.4	65.0 (+2.6)	56.8	57.2 (+0.4)
100%	16000	67.2	69.3 (+2.1)	65.9	66.2 (+0.3)

TABLE V: Test set mIoU on each partial nuScenes. No pre-training, and COLA pre-training.

We compiled results for partial nuScenes for both architectures in Table V and the ones for partial SemanticPOSS in Table VI. We confirmed the previous conclusions that the proposed COLA pre-training improves the performance significantly, whatever the architecture, when a small dataset is available. This is even more significant for SPVCNN.

# of scans	SRU-Net		SPVCNN	
	No pre-training	COLA (Ours)	No pre-training	COLA (Ours)
500	35.2	37.9 (+2.7)	32.9	36.8 (+3.9)
1500 (full)	55	55.8 (+0.8)	49.9	53.8 (+3.9)

TABLE VI: Test set mIoU on partial SemanticPOSS. No pre-training and COLA pre-training.

C. Qualitative results

We show some qualitative results on SemanticKITTI with SRU-Net in Figure 5. Segmentation errors are red, and we can see a significant decrease in mistakes thanks to COLA pre-training especially for buildings and vegetation. Zoom in for a better visualization.




Fig. 5: Example of results on SemanticKITTI after finetuning with SRU-Net. In blue, points with correct semantic segmentation. In red, errors.

D. Ablation study

In the ablation study presented in Table VII, we analyze the importance of the different steps that are part of our proposed COLA pre-training. We conduct our analysis with SemanticPOSS as the target dataset as this is the one with largest size ratio (see Table III). We are comparing four cases: no pre-training, pre-training with only KITTI-360 labeled with its original fine labels, pre-training with only KITTI-360 labeled with our proposed coarse labels, and finally, applying the full COLA, that is, all the pre-training data labeled with the coarse labels.

First, we arrive at the same conclusions as the one presented in PointContrast [5]. In the case of 3D scene understanding, pre-trainings still provide a substantial gain compared with training from scratch, and the data ceiling is not reached yet, as shown by the increase in performance between COLA and pre-training with only KITTI-360 (+1.4 mIoU) (Table VII).

The second highlight is the impact of the coarse labels compared with the fine labels. We demonstrate an improvement in using the coarse labels rather than the fine labels as a pre-training task (+0.6 mIoU). This means that beside providing a simpler mapping in between label sets, it also serves a purpose model-wise. This stems from the simplicity of the labels provided by the coarse labels. When finetuning the model does not need to unlearn anything, unlike the pre-training with the fine labels, where some labels and objects do not exist anymore.

Finally, we compare our approach (COLA) with a more naive pre-training approach that would have a classification module for each of the pre-training datasets rather than having

a unified classification module. In this way, each dataset is trained with its fine labels rather than using the coarse label set. This method is referred to MH (multi-head) in Table VII. COLA demonstrates a significant improvement (+1.1 mIoU) over MH and highlights the effectiveness of using an intermediate label set rather than naively using fine label sets.

In subsection III-C, we proposed a coarse label set we deemed representative of autonomous driving data. We showed that it managed to learn useful geometry information beyond the annotated coarse label: nevertheless, the choice of the label set could still be discussed. In Table VIII, we study the impact of the coarse label set choice: more precisely, we compare the results when we use coarser category (5 instead of 8 by grouping driveable ground and other ground, and grouping dynamic and static objects and structure) or finer categories (10 categories instead of 8 by subdividing static objects into poles and other objects, and subdividing vehicles into 2-wheeled and 4-wheeled vehicles). We can see that the pre-training is robust to the number and choice of coarse labels, allowing in all cases a gain of about +2% mIoU, with a slight advantage for our 8 categories.

Methodology	mIoU
Only finetuned on SemanticPOSS = (1)	49.9
Pre-trained with KITTI-360 (fine labels) + (1)	51.8 (+1.9)
Pre-trained with KITTI-360 (coarse labels) + (1)	52.4 (+2.5)
Pre-trained with MH (all datasets) + (1)	52.7 (+2.8)
Pre-trained with COLA (all datasets) + (1)	53.8 (+3.9)

TABLE VII: Ablation study of the method focusing on the impact of the coarse labels and the amount of pre-training data with SPVCNN. The target dataset on which the test set mIoU is computed is SemanticPOSS.

Methodology	mIoU
Only finetuned on SemanticPOSS = (1)	49.9
Pre-trained with KITTI-360 with 5 coarse labels + (1)	52.2 (+2.3)
Pre-trained with KITTI-360 with 10 coarse labels + (1)	52.1 (+2.2)
Pre-trained with KITTI-360 with the same labels as COLA + (1)	52.4 (+2.5)

TABLE VIII: Ablation study of the method focusing on the impact of the number of coarse labels with SPVCNN. The target dataset on which the test set mIoU is computed is SemanticPOSS.

VI. CONCLUSION

In this paper, we proposed a new supervised pre-training method for autonomous driving semantic segmentation in order to leverage existing annotations and extract meaningfully semantic and geometric information from the available data.

We demonstrated that the proposed COLA method shows great promises. Specifically, it works very well when the amount of target data is limited, but the proposed method is not restricted to this use case. Furthermore, we have shown the proposed method works on various data setups and manages to improve performance even when the target setup is not seen during the pre-training. COLA needs fewer raw data than typical contrastive approaches followed thus far, in addition to needing significantly less computational power and time.

REFERENCES

- [1] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012.
- [2] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks 2021)*, 2021.
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jürgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9296–9306, 2019.
- [4] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11618–11628, 2020.
- [5] Saining Xie, Jiatao Gu, Demi Guo, Charles R. Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 574–591, Cham, 2020. Springer International Publishing.
- [6] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017.
- [7] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 5105–5114, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [8] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11105–11114, 2020.
- [9] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent convolutions for dense prediction in 3d. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3887–3896, 2018.
- [10] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 90–105, Cham, 2018. Springer International Publishing.
- [11] Zhiyuan Zhang, Binh-Son Hua, and Sai-Kit Yeung. Shellnet: Efficient point cloud convolutional neural networks using concentric shells statistics. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1607–1616, 2019.
- [12] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6410–6419, 2019.
- [13] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet ++: Fast and accurate lidar semantic segmentation. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4213–4220, 2019.
- [14] Chenfeng Xu, Bichen Wu, Zining Wang, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 1–19, Cham, 2020. Springer International Publishing.
- [15] Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9598–9607, 2020.
- [16] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3070–3079, 2019.
- [17] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3d architectures with sparse point-voxel convolution. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 685–702, Cham, 2020. Springer International Publishing.
- [18] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9934–9943, 2021.
- [19] Jianyun Xu, Ruixiang Zhang, Jian Dou, Yushi Zhu, Jie Sun, and Shiliang Pu. Rpnnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16024–16033, October 2021.
- [20] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15582–15592, 2021.
- [21] Ali Thabet, Humam Alwassel, and Bernard Ghanem. Self-supervised learning of local features in 3d point clouds. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4048–4052, 2020.
- [22] Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga Rodrigo. Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9902–9912, June 2022.
- [23] Ryosuke Yamada, Hirokatsu Kataoka, Naoya Chiba, Yukiyasu Domae, and Tetsuya Ogata. Point cloud pre-training with natural 3d structures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21283–21293, June 2022.
- [24] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [25] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2432–2443, 2017.
- [26] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiang Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [27] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiang Xiao. 3d shapenets: A deep representation for volumetric shapes. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1912–1920, 2015.
- [28] Hanxue Liang, Chenhan Jiang, Dapeng Feng, Xin Chen, Hang Xu, Xiaodan Liang, Wei Zhang, Zhenguo Li, and Luc Van Gool. Exploring geometry-aware contrast and clustering harmonization for self-supervised 3d object detection. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3273–3282, 2021.
- [29] Junbo Yin, Dingfu Zhou, Liangjun Zhang, Jin Fang, Cheng-Zhong Xu, Jianbing Shen, and Wenguan Wang. Proposalcontrast: Unsupervised pre-training for lidar-based 3d object detection, 2022.
- [30] Lucas Nunes, Rodrigo Marcuzzi, Xieyuanli Chen, Jens Behley, and Cyrill Stachniss. Segcontrast: 3d point cloud feature representation learning through self-supervised segment discrimination. *IEEE Robotics and Automation Letters*, 7(2):2116–2123, 2022.
- [31] Germán Ros, Simon Stent, Pablo Fernández Alcantarilla, and Tomoki Watanabe. Training constrained deconvolutional networks for road scene semantic segmentation. *ArXiv*, abs/1604.01545, 2016.
- [32] Petra Bevandić, Marin Oršić, Ivan Grubišić, Josip Šarić, and Siniša Šegvić. Multi-domain semantic segmentation with overlapping labels. In

- 2022 *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2422–2431, 2022.
- [33] Xiangyun Zhao, Samuel Schulter, Gaurav Sharma, Yi-Hsuan Tsai, Manmohan Chandraker, and Ying Wu. Object detection with a unified label space from multiple datasets. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 178–193, Cham, 2020. Springer International Publishing.
 - [34] Marco Leonardi, Davide Mazzini, and Raimondo Schettini. Training efficient semantic segmentation cnns on multiple datasets. In Elisa Ricci, Samuel Rota Bulò, Cees Snoek, Oswald Lanz, Stefano Messelodi, and Nicu Sebe, editors, *Image Analysis and Processing – ICIAP 2019*, pages 303–314, Cham, 2019. Springer International Publishing.
 - [35] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Simple multi-dataset detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7571–7580, 2022.
 - [36] Girish Varma, Anbumani Subramanian, Anoop Nambodiri, Manmohan Chandraker, and C.V. Jawahar. Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1743–1751, 2019.
 - [37] Panagiotis Meletis and Gijs Dubbelman. Training of convolutional networks on multiple heterogeneous datasets for street scene semantic segmentation. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1045–1050, 2018.
 - [38] Xiaodan Liang, Eric Xing, and Hongfei Zhou. Dynamic-structured semantic propagation network. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 752–761, 2018.
 - [39] Petra Bevandić, Ivan Krešo, Marin Oršić, and Siniša Šegvić. Simultaneous semantic segmentation and outlier detection in presence of domain shift. In Gernot A. Fink, Simone Frintrop, and Xiaoyi Jiang, editors, *Pattern Recognition*, pages 33–47, Cham, 2019. Springer International Publishing.
 - [40] John Lambert, Zhuang Liu, Ozan Sener, James Hays, and Vladlen Koltun. Mseg: A composite dataset for multi-domain semantic segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2876–2885, 2020.
 - [41] Rui Gong, Dengxin Dai, Yuhua Chen, Wen Li, and Luc Van Gool. mdalu: Multi-source domain adaptation and label unification with partial datasets. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8856–8865, 2021.
 - [42] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018*, 2018.
 - [43] Jun Xie, Martin Kiefel, Ming-Ting Sun, and Andreas Geiger. Semantic instance annotation of street scenes by 3d to 2d label transfer. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3688–3697, 2016.
 - [44] Jiageng Mao, Minzhe Niu, Chenhan Jiang, Xiaodan Liang, Yamin Li, Chaoqiang Ye, Wei Zhang, Zhenguo Li, Jie Yu, Chunjing Xu, et al. One million scenes for autonomous driving: Once dataset. 2021.
 - [45] Scale.AI Hesai. Pandaset open datasets. <https://scale.com/open-datasets/pandaset>, 2020. [Online; accessed 18-February-2022].
 - [46] Yancheng Pan, Biao Gao, Jilin Mei, Sibogeng Geng, Chengkun Li, and Huijing Zhao. Semanticpos: A point cloud dataset with large quantity of dynamic instances. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 687–693, 2020.