

AutoInst: Automatic Instance-Based Segmentation of LiDAR 3D Scans

Cedric Perauer¹, Laurenz Adrian Heidrich¹, Haifan Zhang¹,
Matthias Nießner¹, Anastasiia Kornilova^{2*}, and Alexey Artemov^{1*}

Abstract—Recently, progress in acquisition equipment such as LiDAR sensors has enabled sensing increasingly spacious outdoor 3D environments. Making sense of such 3D acquisitions requires fine-grained scene understanding, such as constructing instance-based 3D scene segmentations. Commonly, a neural network is trained for this task; however, this requires access to a large, densely annotated dataset, which is widely known to be challenging to obtain. To address this issue, in this work we propose to predict instance segmentations for 3D scenes in an unsupervised way, without relying on ground-truth annotations. To this end, we construct a learning framework consisting of two components: (1) a pseudo-annotation scheme for generating initial unsupervised pseudo-labels; and (2) a self-training algorithm for instance segmentation to fit robust, accurate instances from initial noisy proposals. To enable generating 3D instance mask proposals, we construct a weighted proxy-graph by connecting 3D points with edges integrating multi-modal image- and point-based self-supervised features, and perform graph-cuts to isolate individual pseudo-instances. We then build on a state-of-the-art point-based architecture and train a 3D instance segmentation model, resulting in significant refinement of initial proposals. To scale to arbitrary complexity 3D scenes, we design our algorithm to operate on local 3D point chunks and construct a merging step to generate scene-level instance segmentations. Experiments on the challenging SemanticKITTI benchmark demonstrate the potential of our approach, where it attains 13.3% higher Average Precision and 9.1% higher F1 score compared to the best-performing baseline. The code is publicly available at <https://github.com/artonson/autoinst>.

Index Terms—3D mapping, normalized cuts, instance segmentation, unsupervised learning

I. INTRODUCTION

Operation of autonomous systems in real outdoor conditions involves continuously interpreting sensor data representing their 3D environment; commonly, an initial stage involves constructing a holistic 3D environment description by assigning semantic and instance labels to individual 3D points. Typically, neural networks (NNs) [1], [2] are trained for this task, requiring access to a large, densely annotated dataset, which is widely known to be costly and laborious to obtain [3], [4], [5]. Using neural models optimized on existing labeled collections [6], however, is suboptimal due to variations in sensor configurations and acquisition conditions

¹Cedric Perauer, Laurenz Heidrich, Haifan Zhang, Matthias Nießner, and Alexey Artemov are with the Technical University of Munich, Garching, Germany E-mail: artonson@yandex.ru

²Anastasiia Kornilova is with the Skolkovo Institute of Science and Technology, Moscow, Russia.

*Equal senior author contribution.

Anastasiia Kornilova was supported by the Analytical center under the RF Government (subsidy agreement 000000D730321P5Q0002, Grant No. 70-2021-00145 02.11.2021).




Fig. 1: For unsupervised instance segmentation of registered LiDAR 3D scans (a), we integrate multi-modal self-supervised deep features into a weighted proxy-graph, making cuts for generation of instance mask proposals (c) and performing their self-trained refinement (d). Our algorithm is label-free and outperforms unsupervised baselines (b).

across datasets, even when domain adaptation methods [7] are employed.

Approaches based on active or weakly-supervised learning [8] enable using sparse or incomplete labels but commonly need user input to guide initial predictions. Very recently, distilling segmentation foundational models has seen interest for the indoor domain [9], [10], [6] but remains in its infancy for outdoor LiDAR data [11]. In this work, we take an alternative path, proposing an automatic learning-based algorithm for instance-based segmentation of dense aggregated 3D LiDAR point clouds, which requires no semantic or instance annotations during training and produces class-agnostic segmentations that can be mapped to various class nomenclatures and instance granularities for downstream tasks. We make two key design decisions leading to important implications for our method. First, instead of processing individual scans as done by most existing approaches [12], [1], we follow the manual labeling approach [3] and construct dense registered 3D LiDAR acquisitions to improve the quality of our instance segmentations. By improving the sampling density, this allows us to reliably segment smaller objects and exploit dense image-to-point correspondences, enabling effective use of image-based features. As our second design choice, drawing inspiration from the 3D indoor domain [13],

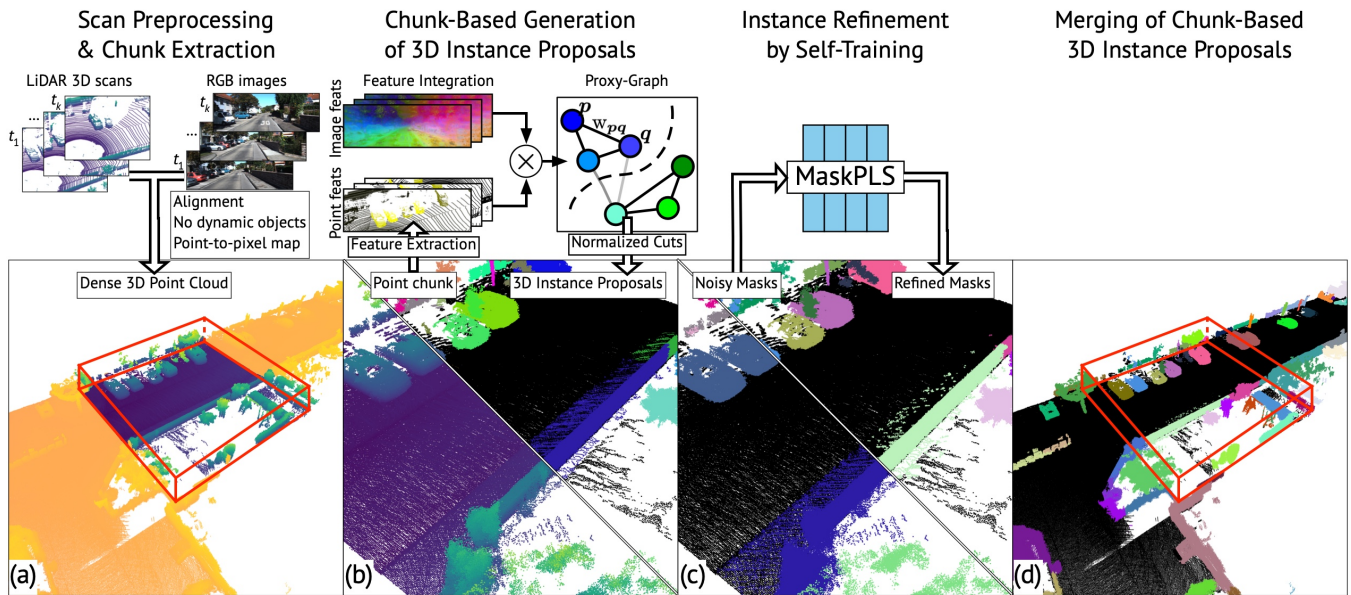


Fig. 2: Overview of our unsupervised 3D instance segmentation framework. We start with a sequence of posed 3D LiDAR scans and RGB images, registering their static segments into a dense 3D map but operate with local overlapping chunks (a). To generate 3D instance mask proposals, we assign multi-modal features to individual 3D points, connecting them into a weighted proxy-graph; we cut the graph to obtain the coarse 3D mask proposals (b). For refinement of 3D instance masks, we start with the coarse proposals and perform multiple rounds of self-training, gradually reintegrating confident instance predictions as ground-truth (c). Map-level segmentation is obtained by merging instances predicted in individual chunks (d).

[14], [15], we design our framework to operate in two stages: *unsupervised generation and self-trained refinement of mask proposals*. For generation of our instance mask proposals, we integrate recent deep self-supervised point- [16] and image-based [17], [18] representations constructed from aligned LiDAR and image captures, providing dense 3D features with high predictive power. We combine these 3D descriptors with spatial priors in a weighted proxy-graph and perform normalized cuts [19] obtaining initial, possibly noisy, instance segmentations. To eliminate spurious or noisy instances and maximize segmentation performance, we refine mask proposals by performing self-training using a procedure similar to [14], [15]. We evaluate our algorithm on SemanticKITTI [3], mapping our segmentations to ground-truth instances. To quantify advantages of our learning-based algorithm, we compare it to non-learnable [20], [10], and learnable unsupervised [12], as well as supervised [1] [10] instance segmentation baselines. We assess contributions of our components and identify an optimal configuration of features for learning-based instance generation and refinement. We contribute:

- A learning-based algorithm for unsupervised instance-based segmentation of outdoor LiDAR 3D point clouds.
- We compare multiple options to build an optimal configuration of self-supervised features for label-free generation and refinement of instance masks, achieving state-of-the-art results against multiple baselines.

II. RELATED WORK

a) Panoptic and Instance 3D Segmentation: Panoptic [21] and instance [22] 3D LiDAR segmentation are

key 3D computer vision tasks with vast applications in autonomous robotics, where they serve to partition an input point cloud into distinct object instances and background regions. Earlier approaches produce instances by spatially clustering [23], [20] segmentations predicted by pretrained scene understanding models but depend on the performance of the backbone [24], [25]. For realtime operation of mobile robots, learning-based methods have been proposed for panoptic [1] and instance-based [22] segmentation of single-scan 3D acquisitions. Alternatively, segmentation of dense 3D point clouds (similar to this work) can aid annotation of 3D data or perform scene analytics. Here, static indoor scenes have seen the most progress with the emergence of panoptic [26], [27] and instance [28] understanding schemes; however instance-based segmentation of dense outdoor LiDAR 3D data remains almost unexplored [29]. All these approaches are supervised and require large-scale annotated 3D LiDAR datasets [3], [30]. We adopt a recent architectural solution [1] in our proposal refinement network.

b) Weakly-, Semi-, and Self-Supervised 3D Segmentation: Multiple works construct and utilize sparse annotations (e.g., a single click or a scribble per instance) [8], or perform self-supervised pre-training of representations coupled with supervised fine-tuning for respective tasks [16], [31]. Still, they predominantly target either only semantic segmentation [31], [8] or only indoor scenes [32]. A recent domain adaptation method [7] transfers a pre-trained panoptic model between source and target domains, minimizing the gap between them. None of these methods can handle new object types emerging during the inference stage. Weakly-

and semi-supervised methods can to some extent reduce the requirements for annotations but are unable to exclude them entirely. We integrate self-supervised point embeddings [16] in our instance proposal stage.

c) Unsupervised 3D Segmentation: To avoid the need for annotated 3D data and enable integrating object types unseen during training, multiple approaches opt to generate and refine 3D instance proposals without associated semantic labels (*i.e.*, *class-agnostic*). Many such approaches originate from clustering algorithms and use various point cloud representations including raw 3D points [20], radial projections, range images [33] and voxel grids [34]. Recently, unsupervised pointwise features [31] were explored as a starting representation for clustering [12]. To ease the segmentation of above ground objects, recent methods leverage ground segmentation [33], [12]; likewise, our approach integrates a robust ground estimator [35] for this purpose. For the indoor domain, recent methods [14], [15] generate 3D instance proposals using deep self-supervised and traditional features, refining proposed masks by a self-trained neural network [28]; our work is conceptually similar but targets outdoor 3D LiDAR data. Multiple recent approaches [10], [9], [6] distill labels from 2D segmentation foundational models [18] into 3D space. Among these, we compare to a simple mask merging baseline [10].

III. METHOD

Our algorithm accepts a set of posed LiDAR scans and RGB images as input, and produces dense class-agnostic 3D instance segmentation as output, requiring no manual annotations during training. To this end, we design a learning-based, unsupervised 3D instance segmentation framework (for overview, see Fig. 2). Our key idea is to adopt pre-trained neural representations to model pairwise point similarity, producing (potentially, noisy) 3D instance masks (Sec. III-B) that are refined by a self-training iteration (Sec. III-C).

A. Scan Preprocessing and Chunk Extraction

We assume that a set of LiDAR 3D scans $\{P_t\}_{t=1}^T$ with their world 3D poses $\{T_t^p\}_{t=1}^T$, as well as a set of images $\{I_t\}_{t=1}^T$ and their world 3D poses $\{T_t^i\}_{t=1}^T$ are produced by a scanning system including one or more LiDAR sensors and RGB cameras. Unless provided, accurate 3D poses can be obtained using GNSS reference poses and further refined with SLAM [36], due to the synchronization of sensors, adjacent LiDAR scans and image captures can be obtained for each 3D pose in the global map. We further obtain the relevant camera features by projecting the pixels onto the aggregated pointcloud and leverage hidden-point removal. Operating with individual 3D LiDAR scans [12] complicates scene understanding tasks as such data rarely has high sampling density [37]; instead, to enable reliable 3D instance segmentation, we compute a dense 3D point cloud. For this, we transform 3D LiDAR and image data to a shared reference frame; dynamic objects can be detected and eliminated by recent effective methods [38]. During this aggregation step, we perform ground segmentation [35],

effectively removing points that do not contain instances and allowing to exploit the geometric separation between objects. To reduce memory constraints, we additionally perform voxel downsampling with a voxel size of 5 cm. Due to the high density of the resulting pointcloud, our method as well the baselines operate on a further down-sampled pointcloud, with predictions being reprojected onto the dense map with 5 cm voxel size for evaluation. To enable using image-based features, we build an inverted index associating each point p to a set of pairs $A_p = \{(I_k, u_k)\}$ where u_k is a pixel in view I_k capturing p . To prevent points occluded in a given view from being assigned its pixels, we identify occluded points w.r.t. each view [39]. For our large-scale scenes, processing entire LiDAR 3D point maps is computationally infeasible as they include billions of sampled points; instead, we extract local *overlapping point chunks* for all subsequent steps. Our chunks are $w \times h \times d$ axis-aligned rectangular parallelepipeds, sampled along the cars trajectory to maximize map coverage. In our experiments, we evaluate the influence of chunk size on performance.

B. Generating 3D Instance Mask Proposals

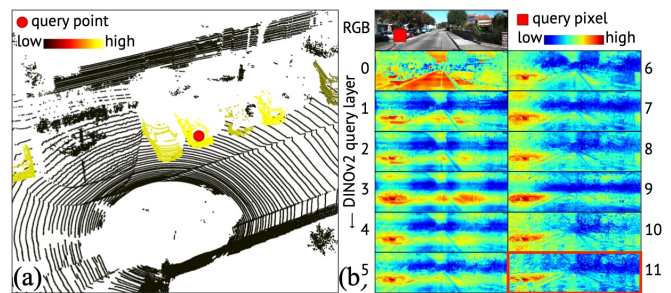


Fig. 3: Pointwise (a) and pixelwise (b) similarity maps (2) for TARL [16] and DINOv2 [17] models, respectively. Following the intuition from prior research [40], we select the output of *query-11* (red box) as our image-based feature map.

To overcome the lack of annotated training data, we adopt an unsupervised grouping approach for generation of 3D instance mask proposals. Specifically, we link neighbouring 3D points into a weighted proxy-graph and cut this graph, isolating groups of points into instances. To effectively model the likelihood of two linked points to belong to the same instance, we weight the graph according to a similarity function integrating multi-modal pointwise representations.

a) Multi-Modal Pointwise Features: To fully leverage 3D LiDAR and image data, we represent each point by a collection of features computed from respective data modalities, specifically, spatial point coordinates (S), point embeddings (P), and image embeddings (I). Building on the availability of pre-trained neural feature extractors, we adopt state-of-the-art self-supervised TARL [16] and DINOv2 [17] representations to serve as our point- and image-based features, respectively, as shown in Fig. 3. TARL [16] learns temporally-consistent dense representations of 3D LiDAR scans by pulling together features computed in spatially close points. DINOv2 [17] is a self-distillation training approach providing dense universal




Fig. 4: Influence of hyperparameters θ^P, θ^I in (2) on distribution of instance mask proposals, as captured by the volume occupied by instances (horizontal axis). $\theta_0^P = 0.5, \theta_0^I = 0.1$.

features at high spatial resolution (e.g., $4\times$ higher compared to a recent contrastive method [41]), enabling fine-grained binding to points. Point-based features \mathbf{x}_p^P are pre-computed by a pre-trained TARL model φ^P on single LiDAR 3D scans directly. To compute an image-based feature \mathbf{x}_p^I at point p , we use a pre-trained DINOv2 model φ^I to extract feature maps $\varphi^I(I)$ at layer *query-11*, that was identified, similarly to [40], by qualitatively analysing the capacity of feature similarity maps to capture instances well (Fig. 3). We compute an average

$$\mathbf{x}_p^I = \frac{1}{|A_p|} \sum_{(I, \mathbf{u}) \in A_p} \varphi^I(I)_{\text{fmap}(\mathbf{u})}, \quad (1)$$

where $\text{fmap}(\mathbf{u})$ is a position in the feature map corresponding to the pixel \mathbf{u} . We complement neural descriptors $\mathbf{x}_p^P, \mathbf{x}_p^I$ by a simple spatial feature $\mathbf{x}_p^S = p$ to express that two *spatially proximal* points may be naturally assigned to the same instance.

b) Linking Points into a Proxy-Graph: To generate instance mask proposals, we start with connecting pairs of points located within a distance $R = 1$ m with weighted edges; we aim to give larger weights to edges within instances and smaller weights to edges crossing instance boundaries. For this, we use features constructed in respective spaces to compute pairwise similarity scores

$$w_{pq}^\mu = e^{-\theta^\mu \|\mathbf{x}_p^\mu - \mathbf{x}_q^\mu\|^2}, \quad (2)$$

where μ denotes a respective feature space and $\theta^\mu > 0$ a scalar parameter. We opt to integrate these quantities into a single score using a multiplicative rule

$$w_{pq} = \prod_{\mu} w_{pq}^\mu, \quad (3)$$

requiring high similarity scores across all modalities for two points to be included in an instance.

c) Instance Isolation by Normalized Cuts: To generate instance mask proposals in a point chunk, we partition the proxy-graph described above using the Normalized Cut (NCut) algorithm [19]. NCut relies on partitioning criteria including both total inter-instance dissimilarity and total intra-instance similarity, naturally matching our task. At each grouping iteration, we check against the cutting eigenvalue

threshold (we use .005 for S+P+I, .03 for S+P, and 0.075 for S) and the minimum share of points in the isolated region (we use 1%) to prevent further cuts. For dense similarity matrices w_{pq} , singular value decomposition (SVD) employed by NCut can be computationally expensive; by avoiding edges between points located outside radius R to each other, we obtain a sparser initial matrix. To additionally accommodate for the computational complexity of the Normalized Cut algorithm, we perform uniform voxel downsampling of chunk point (with voxel size 0.35).

d) Map-Level Instance Segmentation: To predict a map-level instance segmentation, we merge predictions of individual chunks to the global map by exploiting the overlap between adjacent chunks. For each instance in a new chunk, we determine the instance in the global map with the highest bounding box IoU. Consequently, overlapping instances are merged when the IoU is above a threshold of 0.01.

C. Self-Training of Refined 3D Instance Masks

3D masks obtained in Sec. III-B already provide an instance-based scene description but may include noisy, under- or oversegmented detections, see Fig. 2. Instead of treating them as final instances, we take inspiration from recent 3D indoor approaches [14], [15], [42] and seek to refine these initial predictions by a self-trained neural network (NN). Intuitively, such approaches (including ours) build on recent studies [43], [44] revealing that during training, NNs learn cleaner, simpler patterns first, fitting more noisy data samples later. Drawing on these insights, we design a self-training step that learns a clean chunk-based 3D instance segmentation, starting from the initial set of noisy predictions produced by NCut. Our self-trained NN takes xyz point coordinates as input and produces per-point 3D instance masks and confidence values as output. To train it, we use the dataset obtained by applying NCut to the collection of point chunks, treating noisy 3D instance masks as ground-truth labels (in practice, we exclude instances represented by fewer than 100 points). Architecturally, our self-trained NN follows the design of MaskPLS network for panoptic segmentation [1], including a sparse CNN encoder [24] and a transformer decoder [45]. During training, we optimize a weighted combination of dice and cross-entropy losses

$$L = \lambda_{\text{dice}} L_{\text{dice}} + \lambda_{\text{bce}} L_{\text{bce}}. \quad (4)$$

where $\lambda_{\text{dice}} = 5$ and $\lambda_{\text{bce}} = 2$ as in [1]. We train our network for one iteration, *i.e.* until convergence with initial ground-truth labels; we find this to be sufficient for achieving improved instance segmentation performance. Due to the significant density of our aggregated chunk-based point clouds, we preprocess our chunks using uniform downsampling.

IV. EXPERIMENTS

A. Experimental Setup

a) Data Preprocessing and Training Details: We experimentally evaluate our algorithm using SemanticKITTI [3] which provides LiDAR 3D scans with instance annotations and registered RGB images acquired by a moving car




Fig. 5: Comparative instance-based segmentation results on SemanticKITTI [3].

Method	P	R	F1	AP@0.25	AP@0.5	AP	S
HDBSCAN [31]	83.3	68.3	74.8	66.8	58.9	49.5	64.6
3DUIS [22]	79.9	67.0	72.2	63.5	56.1	46.6	62.9
SAM3D [19]	24.7	21.1	22.1	27.5	8.1	1.3	14.3
Ours	91.7	77.8	83.9	79.8	75.8	62.8	72.2

TABLE I: Comparative instance-based segmentation results on SemanticKITTI [3]. All measures are in percentages where higher values correspond to better results. Our method outperforms all baselines across all quantitative measures.

Method	P	R	F1	AP@0.25	AP@0.5	AP	S
MaskPLS (supervised) [6]	94.6	78.2	85.6	79.6	75.8	64.0	71.4
Ours (unsupervised)	86.6	75.8	80.9	76.9	70.3	52.8	68.5

TABLE II: Comparison with the supervised MaskPLS [1].

in outdoor scenes, obtaining scenes on the order of 10^9 points after aggregation. We use sequences 0–10 providing test labels, except sequences 1 and 4 that predominantly consist of dynamic instances. For accurate global alignment of LiDAR 3D scans, we use 3D poses computed using a surfel-based SLAM approach [36]. During mask proposal and refinement stages, we operate on $25 \times 25 \times 25 \text{ m}^3$ point chunks; at inference time, we sample chunk centers along vehicle trajectory spaced 22 m apart (3 m overlap), resulting in 30 chunks per 1K raw LiDAR scans on average. We

extract image-based features from the images obtained by the left camera. For training our refinement network and the supervised baseline, we sample point chunks, our proposals, and ground-truth labels, from KITTI sequences 0–10 with a larger 24 m overlap, obtaining a dataset with 14K chunks. To reduce memory consumption, we uniformly subsample each chunk to roughly 60K points. We train our refinement NN [1] for 7 epochs (12 h) on RTX 4090 GPU using Adam [46] with a learning rate of 10^{-4} .

b) Implementation Details: We set $\theta^I = 0.1$, $\theta^P = 0.5$ for S+P+I. For S+P, the point feature weight is set to $\theta^P = 0.5$, all of our configurations use a spatial weight of $\theta^S = 1$. We use the cutting thresholds as mentioned in Sec. III-B. To generate the final instance masks, we run inference on downsampled input chunks and transfer the predicted labels

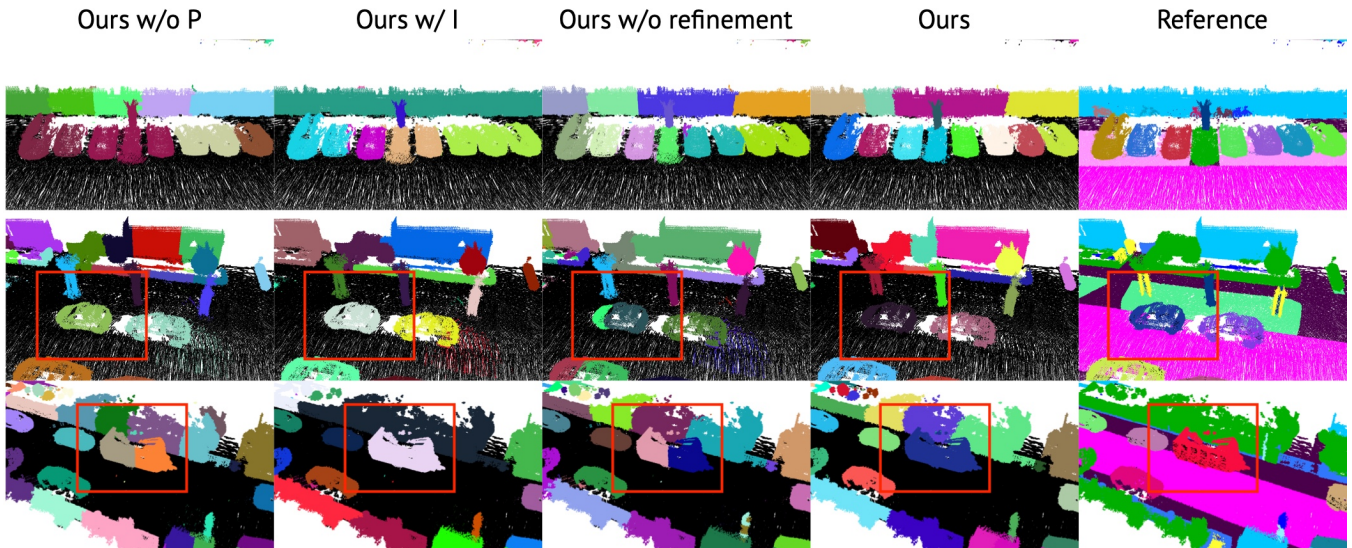


Fig. 6: Ablative instance-based segmentation results on SemanticKITTI [3].

Method	Proposals	Refine?	P	R	F1	AP@0.25	AP@0.5	AP	S
Ours w/o P	NCut (S)	no	84.0	74.6	78.7	76.2	65.3	49.5	70.3
Ours w/o refinement	NCut (S+P)	no	88.0	76.7	81.7	78.1	69.6	52.3	71.1
Ours w/ I	NCut (S+P+I)	no	87.0	62.8	72.6	64.3	54.8	42.1	60.9
Ours w/ dynamic obj.	NCut (S+P)	yes	89.3	77.1	82.6	80.2	75.5	64.2	72.7
Ours w/ HDBSCAN	HDBSCAN [31]	yes	94.0	72.1	80.9	74.1	70.4	59.9	66.5
Ours	NCut (S+P)	yes	91.7	77.8	83.9	79.8	75.8	62.8	72.2

TABLE III: Quantitative comparison of contribution of various components of our method on SemanticKITTI [3]. All measures are in percentages where higher values correspond to better results. “Ours w/o P”, “Ours w/o refinement”, and “Ours w/ I” refer to forming the proxy-graph by weights computed from spatial only, spatial and point, and spatial, point, and image features (Proposals column). “Ours w/ dynamic obj.” performs segmentation without prior removal of dynamic instances. “Ours w/ HDBSCAN” replaces our instance generation step by HDBSCAN [20].

to the denser point cloud by nearest-neighbor assignment. We study the impact of parameter choices in Sec. IV-C.

c) *Metrics*: Our method is class-agnostic; thus, we evaluate it using instance-based annotations and report instance-related metrics, averaging them across available sequences. Specifically, we use Precision (P), Recall (R), and their harmonic mean, F1-score (F1), Average Precision (AP, computed at three IoU thresholds of 0.25, 0.5, and globally), and S_{assoc} [12] (S).

d) *Baselines*: We define two unsupervised and two supervised baselines. HDBSCAN [20] is a density-based clustering algorithm, while 3DUIS [12] is an unsupervised instance segmentation method, refining HDBSCAN instance proposals with deep point-based features [31] and graph cuts. SAM3D [10] is a method merging 2D SAM [9] predictions into 3D instance masks. We pre-process 2D SAM masks by rejecting smaller sub-instances using an IoU threshold of 0.9 and perform hidden point removal. Additionally, we train supervised MaskPLS [1], a state-of-the-art 3D panoptic segmentation method. For experiments in Tab. II, we train on sequences 0–7 and 9, holding out 8 and 10 as the test set. Further experiments (see Tab. I and Tab. III), were tested on

all sequences and do not include supervised training.

B. Comparisons to State-of-the-Art

We start by evaluating our method against baselines; we present quantitative evaluation results against the unsupervised baselines in Tab. I and a quantitative comparison against the supervised MaskPLS method [1] in Tab. II. Our best-performing method (NCut(S + P) + Refine) outperforms point-based methods HDBSCAN and 3DUIS and the supervised method SAM3D without using image-based features. We refer to our ablations IV-C for a comparison of our point-based method (Ncuts(S+P)) with a multi-modal implementation (NCut(S + P + I)). While our approach shows a moderate decrease w.r.t. the fully supervised MaskPLS, it does not require manual labeling and can be used to speed up the annotation for supervised training scenarios. Qualitative comparisons are presented in Fig. 5 for baselines and our algorithm without and with refinement step. One noticeable effect of refinement is that it splits instances that are merged together in initial proposals.

Features	Factor	Runtime	P	R	F1	AP@0.25	AP@0.5	AP	S
S + P	×0.5	11m	81.0	76.4	78.6	85.9	60.8	46.0	75.7
	×1	19m	87.7	86.1	86.9	84.5	75.7	56.6	77.7
	×1.5	72m	80.3	85.5	82.8	79.4	66.9	48.6	76.7
S + P + I	×0.5	45m	81.9	64.2	72.0	70.3	54.2	41.3	65.8
	×1	69m	84.8	72.2	78.0	72.9	55.6	38.3	67.7
	×1.5	130m	78.4	62.7	69.7	66.5	47.1	35.4	61.9

TABLE IV: *Influence of chunk size on performance.* Our chunks provide sufficient context for effective proposal generation at moderate computational cost.

Features	Factor	Runtime	P	R	F1	AP@0.25	AP@0.5	AP	S
S + P	×0.75	55m	82.3	86.1	84.2	85.9	60.8	46.0	78.3
	×1	19m	87.7	86.1	86.9	84.5	75.7	56.6	77.7
	×2	8m	82.5	87.0	84.7	79.4	66.9	48.6	77.2
S + P + I	×0.75	230m	81.5	69.4	75.0	70.3	54.2	41.3	66.2
	×1	69m	84.8	72.2	78.0	72.9	55.6	38.3	67.7
	×2	52m	84.6	71.3	77.4	66.5	47.1	35.4	67.7

TABLE V: *Influence of voxel size on performance.* The base voxel size of 35 cm provides a good balance between computational cost and performance.

Features	Visibility	P	R	F1	AP@0.25	AP@0.5	AP	S
S + P	TRUE	85.7	72.9	78.8	79.1	56.9	43.0	70.0
	FALSE	86.7	72.2	78.8	74.4	60.1	46.6	69.8
S + P + I	TRUE	83.6	52.3	64.4	58.4	45.5	32.1	51.1
	FALSE	86.2	51.9	64.7	55.6	46.4	34.9	50.7

TABLE VI: *Influence of point visibility.* Detecting and eliminating invisible points moderately affects performance.

C. Ablative Studies

a) *Making Effective Instance Proposals:* To identify the optimal ingredients in our instance mask proposal step, we compare applying NCut to point graphs weighted by various combination of features, without any refinement, and present the results in Tab. III (see notation in its caption). We conclude that point embeddings (P) are crucial for obtaining strong performance; in contrast (and to our surprise), image embeddings (I) noticeably decrease performance across measures. To understand whether an alternative proposal generation step may lead to obtaining comparably strong results, we generate instance mask proposals by HDBSCAN and refine them using our self-training step. While this improves upon HDBSCAN, our method achieves better results.

b) *Fine-Grained Tuning of Proposal Hyperparameters:* We assess the influence of chunk and voxel size used by evaluating the performance and runtime of our method, these ablations are performed on a single sequence of KITTI [3]. Results presented in Tabs. IV and V indicate that, for our configuration and dataset, our default chunk size ($25 \times 25 \times 25 \text{ m}^3$) and voxel size (0.35 m) are close to optimal. Larger voxel size leads to coarser instances while finer one yields noisier instances. We additionally made an effort to find an optimal combination of parameters $\theta^S, \theta^P, \theta^I$ in (2). From Fig. 4 we conclude that our default values lead our proposals to match the organic instance distributions.

c) *The Contribution of Refinement:* Convergence of self-supervised refinement for KITTI dataset is depicted in




Fig. 7: *Influence of self-training duration on performance.*

Fig. 7. As demonstrated from the AP, F1 and S_{assoc} results, 7 epochs is enough for the network to be saturated.

d) *Robustness w.r.t. Dynamic Objects:* In Tab. III, we include an additional result to illustrate robustness of our method w.r.t. elimination of dynamic instances (“Ours w/ dynamic obj.”). We aggregate 3D scans without filtering dynamic objects; we then perform instance generation and refinement as usual. We find that training on such noisier data yields almost no drop in performance (e.g., only 1.3% in F1).

e) *Influence of Visibility on Image Embeddings:* We assess whether eliminating points predicted as hidden in all images from the joint point cloud enhances performance. We employ hidden point removal (radius=20) to eliminate invisible points, applying NCut on visible point only. Tab. VI suggests that visibility does not have a significant impact on performance, regardless of the feature configuration.

V. CONCLUSION

We have developed a learning-based, label-free algorithm for instance-based segmentation of outdoor LiDAR 3D point clouds. Our method uses dense, registered 3D and RGB data to perform generation of dense 3D instance mask proposals, constructing and cutting a feature-aware graph. We have discovered that integrating self-supervised point embeddings with simple spatial features results in the best performance across feature combinations. Our algorithm further benefits from a self-training step that does not require ground-truth annotations, instead fitting and refining proposals obtained in an unsupervised way. Our results suggest a better performance compared to clustering-based and learning-based baselines across a variety of instance-sensitive measures.

REFERENCES

- [1] R. Marcuzzi, L. Nunes, L. Wiesmann, J. Behley, and C. Stachniss, “Mask-based panoptic lidar segmentation for autonomous driving,” *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 1141–1148, 2023.
- [2] E. Li, R. Razani, Y. Xu, and B. Liu, “Cpseg: Cluster-free panoptic segmentation of 3d lidar point clouds,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 8239–8245.
- [3] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, “Semantickitti: A dataset for semantic scene understanding of lidar sequences,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9297–9307.
- [4] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, “ScanNet: Richly-annotated 3d reconstructions of indoor scenes,” in *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017.
- [5] C. Yeshwanth, Y.-C. Liu, M. Nießner, and A. Dai, “ScanNet++: A high-fidelity dataset of 3d indoor scenes,” in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2023.

- [6] R. Huang, S. Peng, A. Takmaz, F. Tombari, M. Pollefeys, S. Song, G. Huang, and F. Engelmann, "Segment3d: Learning fine-grained class-agnostic 3d segmentation without manual labels," *arXiv preprint arXiv:2312.17232*, 2023.
- [7] B. Bešić, N. Gosala, D. Cattaneo, and A. Valada, "Unsupervised domain adaptation for lidar panoptic segmentation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3404–3411, 2022.
- [8] H. Shi, J. Wei, R. Li, F. Liu, and G. Lin, "Weakly supervised segmentation on outdoor 4d point clouds with temporal matching and spatial graph propagation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 11 840–11 849.
- [9] H. Guo, H. Zhu, S. Peng, Y. Wang, Y. Shen, R. Hu, and X. Zhou, "Sam-guided graph cut for 3d instance segmentation," *arXiv preprint arXiv:2312.08372*, 2023.
- [10] Y. Yang, X. Wu, T. He, H. Zhao, and X. Liu, "Sam3d: Segment anything in 3d scenes," *arXiv preprint arXiv:2306.03908*, 2023.
- [11] Y. Liu, L. Kong, J. Cen, R. Chen, W. Zhang, L. Pan, K. Chen, and Z. Liu, "Segment any point cloud sequences by distilling vision foundation models," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [12] L. Nunes, X. Chen, R. Marcuzzi, A. Osep, L. Leal-Taixé, C. Stachniss, and J. Behley, "Unsupervised class-agnostic instance segmentation of 3d lidar data for autonomous vehicles," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8713–8720, 2022.
- [13] X. Wang, R. Girdhar, S. X. Yu, and I. Misra, "Cut and learn for unsupervised object detection and instance segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 3124–3134.
- [14] D. Rozenberszki, O. Litany, and A. Dai, "Unscene3d: Unsupervised 3d instance segmentation for indoor scenes," *arXiv preprint arXiv:2303.14541*, 2023.
- [15] Z. Zhang, J. Ding, L. Jiang, D. Dai, and G.-S. Xia, "Freepoint: Unsupervised point cloud instance segmentation," *arXiv preprint arXiv:2305.06973*, 2023.
- [16] L. Nunes, L. Wiesmann, R. Marcuzzi, X. Chen, J. Behley, and C. Stachniss, "Temporal consistent 3d lidar representation learning for semantic perception in autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5217–5228.
- [17] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, *et al.*, "Dinov2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023.
- [18] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," *arXiv:2304.02643*, 2023.
- [19] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 22, no. 8, pp. 888–905, 2000.
- [20] L. McInnes, J. Healy, and S. Astels, "hdbscan: Hierarchical density based clustering," *J. Open Source Softw.*, vol. 2, no. 11, p. 205, 2017.
- [21] J. Behley, A. Milioto, and C. Stachniss, "A benchmark for lidar-based panoptic segmentation based on kitti," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 596–13 603.
- [22] F. Zhang, C. Guan, J. Fang, S. Bai, R. Yang, P. H. Torr, and V. Prisacariu, "Instance segmentation of lidar point clouds," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 9448–9455.
- [23] Y. Zhao, X. Zhang, and X. Huang, "A divide-and-merge point cloud clustering algorithm for lidar panoptic segmentation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 7029–7035.
- [24] C. Choy, J. Gwak, and S. Savarese, "4d spatio-temporal convnets: Minkowski convolutional neural networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 3075–3084.
- [25] X. Zhu, H. Zhou, T. Wang, F. Hong, Y. Ma, W. Li, H. Li, and D. Lin, "Cylindrical and asymmetrical 3d convolution networks for lidar segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 9939–9948.
- [26] S.-C. Wu, J. Wald, K. Tateno, N. Navab, and F. Tombari, "Scenegrph-fusion: Incremental 3d scene graph prediction from rgb-d sequences," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7515–7525.
- [27] M. Kolodiaznyi, A. Vorontsova, A. Konushin, and D. Rukhovich, "Oneformer3d: One transformer for unified point cloud segmentation," *arXiv preprint arXiv:2311.14405*, 2023.
- [28] Z. Schult, F. Engelmann, A. Hermans, O. Litany, S. Tang, and B. Leibe, "Mask3d: Mask transformer for 3d semantic instance segmentation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 8216–8223.
- [29] Z. Yang and C. Liu, "Tupper-map: Temporal and unified panoptic perception for 3d metric-semantic mapping," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1094–1101.
- [30] W. K. Fong, R. Mohan, J. V. Hurtado, L. Zhou, H. Caesar, O. Beijbom, and A. Valada, "Panoptic nusenes: A large-scale benchmark for lidar panoptic segmentation and tracking," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3795–3802, 2022.
- [31] L. Nunes, R. Marcuzzi, X. Chen, J. Behley, and C. Stachniss, "Seg-contrast: 3d point cloud feature representation learning through self-supervised segment discrimination," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2116–2123, 2022.
- [32] J. Chibane, F. Engelmann, T. Anh Tran, and G. Pons-Moll, "Box2mask: Weakly supervised 3d semantic instance segmentation using bounding boxes," in *European Conference on Computer Vision*. Springer, 2022, pp. 681–699.
- [33] D. Zermas, I. Izzat, and N. Papanikolopoulos, "Fast segmentation of 3d point clouds: A paradigm on lidar data for autonomous vehicle applications," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5067–5073.
- [34] J. Behley, V. Steinhage, and A. B. Cremers, "Laser-based segment classification using a mixture of bag-of-words," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4195–4200.
- [35] S. Lee, H. Lim, and H. Myung, "Patchwork++: Fast and robust ground segmentation solving partial under-segmentation using 3d point cloud," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 13 276–13 283.
- [36] J. Behley and C. Stachniss, "Efficient surfel-based slam using 3d laser range data in urban environments," in *Proc. of Robotics: Science and Systems (RSS)*, 2018.
- [37] L. Li, K. N. Ismail, H. P. Shum, and T. P. Breckon, "Durlar: A high-fidelity 128-channel lidar dataset with panoramic ambient and reflectivity imagery for multi-modal autonomous driving applications," in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 1227–1237.
- [38] X. Chen, S. Li, B. Mersch, L. Wiesmann, J. Gall, J. Behley, and C. Stachniss, "Moving object segmentation in 3d lidar data: A learning-based approach exploiting sequential data," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6529–6536, 2021.
- [39] S. Katz, A. Tal, and R. Basri, "Direct visibility of point sets," in *ACM SIGGRAPH 2007 papers*, 2007, pp. 24–es.
- [40] N. Keetha, A. Mishra, J. Karhade, K. M. Jatavallabhula, S. Scherer, M. Krishna, and S. Garg, "Anyloc: Towards universal visual place recognition," *IEEE Robotics and Automation Letters*, 2023.
- [41] X. Wang, R. Zhang, C. Shen, T. Kong, and L. Li, "Dense contrastive learning for self-supervised visual pre-training," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 3024–3033.
- [42] L. Zhang, A. J. Yang, Y. Xiong, S. Casas, B. Yang, M. Ren, and R. Urtasun, "Towards unsupervised object detection from lidar point clouds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9317–9328.
- [43] D. Arpit, S. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, Y. Bengio, *et al.*, "A closer look at memorization in deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 233–242.
- [44] S. Ye, D. Chen, S. Han, and J. Liao, "Learning with noisy labels for robust point cloud segmentation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 6443–6452.
- [45] B. Cheng, A. Schwing, and A. Kirillov, "Per-pixel classification is not all you need for semantic segmentation," *Advances in neural information processing systems*, vol. 34, pp. 17 864–17 875, 2021.
- [46] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.