
Supplementary Material: Segment Any Point Cloud Sequences by Distilling Vision Foundation Models

Anonymous Author(s)

Affiliation

Address

email

1 In this file, we supplement the following materials to support the findings and observations in the
2 main body of this paper:

- 3 • Section A elaborates on additional implementation details to facilitate reproduction.
- 4 • Section B provides the complete quantitative results of our experiments.
- 5 • Section C includes more qualitative results to allow better visual comparisons.
- 6 • Section D acknowledges the public resources used during the course of this work.

7 A Additional Implementation Detail

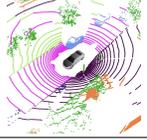
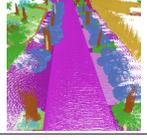
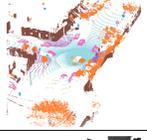
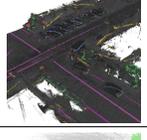
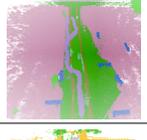
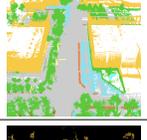
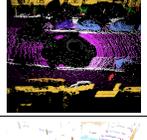
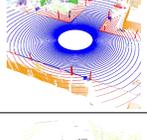
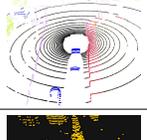
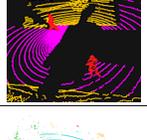
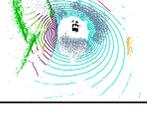
8 A.1 Datasets

9 In this work, we conduct extensive experiments from a wide range of point cloud datasets to verify
10 the effectiveness of the proposed *Seal* framework. A summary of the detailed configurations and
11 emphases of these datasets is shown in Table A.

- 12 • ¹**nuScenes** [7]: The nuScenes¹ dataset offers a substantial number of samples collected
13 by the LiDAR, RADAR, camera, and IMU sensors from Boston and Singapore, allowing
14 machine learning models to learn useful multi-modal features effectively. For the point
15 cloud segmentation task, it consists of 1000 scenes of a total number of 1.1B annotated
16 points, collected by a Velodyne HDL32E LiDAR sensor. It also includes image data from six
17 cameras, which are synchronized with the LiDAR sensor. In this work, we use the LiDAR
18 point clouds and image data from nuScenes for model pretraining. We also conduct detailed
19 fine-tuning experiments to validate the effectiveness of representation learning. More details
20 of this dataset can be found at <https://www.nuscenes.org/nuscenes>.
- 21 • ²**SemanticKITTI** [2]: The SemanticKITTI dataset is a comprehensive dataset designed
22 for semantic scene understanding of LiDAR sequences. This dataset aims to advance the
23 development of algorithms and models for autonomous driving and scene understanding
24 using LiDAR data. It provides 22 densely labeled point cloud sequences that cover urban
25 street scenes, which are captured by a Velodyne HDL-64E LiDAR sensor. In this work,
26 we use the LiDAR point clouds from SemanticKITTI as a downstream task to validate
27 the generalizability of pertaining methods. More details of this dataset can be found at
28 <http://semantic-kitti.org/>.
- 29 • ³**Waymo Open** [19]: The Waymo Open dataset is a large-scale collection of real-world
30 autonomous driving data. The 3D semantic segmentation subset contains 1150 scenes,

¹Here we refer to the *lidarseg* subset of the nuScenes dataset. Know more details about this dataset at the official webpage: <https://www.nuscenes.org/nuscenes>.

Table A: The sensor configuration and data statistics for the *eleven* datasets used in our experiments.

Dataset	Illustration	Sensor Setup	Statistics	Type
nuScenes [7]		1× LiDAR (32-beam) 6× RGB Camera 5× RADAR 1× IMU & GPS	16 semantic classes 29130 training samples 6019 validation samples 6008 testing samples	Real-world Low-resolution point cloud Dense annotation Multi-modality
SemanticKITTI [2]		1× LiDAR (64-beam) 1× Stereo Camera 1× IMU & GPS	19 semantic classes 19130 training samples 4071 validation samples 20351 testing samples	Real-world High-resolution point cloud Dense annotation Multi-modality
Waymo Open [19]		1× LiDAR (64-beam) 5× RGB Camera 1× IMU & GPS	23 semantic classes 23691 training samples 5976 validation samples 2982 testing samples	Real-world High-resolution point cloud Dense annotation Multi-modality
ScribbleKITTI [20]		1× LiDAR (64-beam) 1× Stereo Camera 1× IMU & GPS	19 semantic classes 19130 training samples 4071 validation samples 20351 testing samples	Real-world High-resolution point cloud Sparse annotation Weakly-supervised learning
RELLIS-3D [9]		1× LiDAR (64-beam) 1× LiDAR (32-beam) 1× Stereo Camera 1× IMU & GPS	20 semantic classes 7800 training samples 2413 validation samples 3343 testing samples	Real-world High-resolution point cloud Dense annotation Multi-modality
SemanticPOSS [16]		1× LiDAR (40-beam) 1× RGB Camera 1× IMU & GPS	14 semantic classes 2488 training samples 500 validation samples	Real-world High-resolution point cloud Dense annotation Dynamic instance
SemanticSTF [22]		1× LiDAR (64-beam) 1× LiDAR (32-beam) 1× Stereo Camera 1× RADAR	21 semantic classes 1326 training samples 250 validation samples 500 testing samples	Real-world High-resolution point cloud Dense annotation Adverse weather
SynLiDAR [21]		1× LiDAR (64-beam) 1× simulation suite	32 semantic classes 198396 total samples	Synthetic High-resolution point cloud Dense annotation Transfer learning
Synth4D [17]		1× LiDAR (64-beam) 1× LiDAR (32-beam) 1× simulation suite	22 semantic classes 10000 training samples 10000 validation samples	Synthetic Low-resolution point cloud Dense annotation Transfer learning
DAPS-3D [11]		3× LiDAR (64-beam) 1× simulation suite	4 semantic classes 19061 training samples 3995 validation samples	Semi-synthetic High-resolution point cloud Dense annotation Transfer learning
nuScenes-C [12]		1× LiDAR (32-beam) 6× RGB Camera 5× RADAR 1× IMU & GPS	16 semantic classes 144456 validation samples	Synthetic Low-resolution point cloud Dense annotation Robustness

- 31 split into 798 training, 202 validation, and 150 testing scenes. This subset contains 23691
 32 training scans, 5976 validation scans, and 2982 testing scans, respectively, with semantic
 33 segmentation labels from 23 classes. The data are captured by five LiDAR sensors: one
 34 mid-range LiDAR sensor truncated to a maximum of 75 meters, and four short-range LiDAR
 35 sensors truncated to a maximum of 20 meters. In this work, we use the LiDAR point clouds
 36 from Waymo Open as a downstream task to validate the generalizability of pertaining
 37 methods. More details of this dataset can be found at <https://waymo.com/open>.
- 38 • ⁴**ScribbleKITTI** [20]: The ScribbleKITTI dataset is a recent variant of the SemanticKITTI
 39 dataset, with weak supervisions annotated by line scribbles. It shares the exact same amount
 40 of training samples with SemanticKITTI, *i.e.*, 19130 scans collected by a Velodyne HDL-
 41 64E LiDAR sensor, where the total number of valid semantic labels is 8.06% compared to
 42 the fully-supervised version. Annotating the LiDAR point cloud in such a way corresponds
 43 to roughly a 90% time-saving. In this work, we use the LiDAR point clouds from Scrib-
 44 bleKITTI as a downstream task to validate the generalizability of pertaining methods. More
 45 details of this dataset can be found at <https://github.com/ouenal/scribblekitti>.
 - 46 • ⁵**RELLIS-3D** [9]: The RELLIS-3D dataset is a multimodal dataset collected in an off-road
 47 environment from the Rellis Campus of Texas A&M University. It consists of 13556 LiDAR
 48 scans from 5 traversal sequences. The point-wise annotations are initialized by using the
 49 camera-LiDAR calibration to project the more than 6000 image annotations onto the point
 50 clouds. In this work, we use the LiDAR point clouds from RELLIS-3D as a downstream
 51 task to validate the generalizability of pertaining methods.
 - 52 • ⁶**SemanticPOSS** [16]: The SemanticPOSS dataset is a relatively small-scale point cloud
 53 dataset with an emphasis on dynamic instances. It includes 2988 scans collected by a Hesai
 54 Pandora LiDAR sensor, which is a 40-channel LiDAR with 0.33 degree vertical resolution,
 55 a forward-facing color camera, 4 wide-angle mono cameras covering 360 degrees around
 56 the ego-car. The data in this dataset was collected from the campus of Peking University. In
 57 this work, we use the LiDAR point clouds from SemanticPOSS as a downstream task to
 58 validate the generalizability of pertaining methods. More details of this dataset can be found
 59 at <http://www.unmannedlab.org/research/RELLIS-3D>.
 - 60 • ⁷**SemanticSTF** [22]: The SemanticSTF dataset is a small-scale collection of 2076 scans,
 61 where the data are borrowed from the STF dataset [4]. The scans are collected by a Velodyne
 62 HDL64 S3D LiDAR sensor and covered various adverse weather conditions, including 694
 63 snowy, 637 dense-foggy, 631 light-foggy, and 114 rainy scans. The whole dataset is split
 64 into three sets: 1326 scans for training, 250 scans for validation, and 500 scans for testing.
 65 All three splits have similar proportions of scans from different weather conditions. In this
 66 work, we use the LiDAR point clouds from SemanticSTF as a downstream task to validate
 67 the generalizability of pertaining methods. More details of this dataset can be found at
 68 <https://github.com/xiaoaoran/SemanticSTF>.
 - 69 • ⁸**SynLiDAR** [21]: The SynLiDAR dataset contains synthetic point clouds captured from
 70 constructed virtual scenes using the Unreal Engine 4 simulator. In total, this dataset contains
 71 13 LiDAR point cloud sequences with 198396 scans. As stated, the virtual scenes in
 72 SynLiDAR are constituted by physically accurate object models that are produced by expert
 73 modelers with the 3D-Max software. In this work, we use the LiDAR point clouds from
 74 SynLiDAR as a downstream task to validate the generalizability of pertaining methods. More
 75 details of this dataset can be found at <https://github.com/xiaoaoran/SynLiDAR>.
 - 76 • ⁹**Synth4D** [17]: The Synth4D dataset includes two subsets with point clouds captured by
 77 simulated Velodyne LiDAR sensors using the CARLA simulator. We use the Synth4D-
 78 nuScenes subset in our experiments. It is composed of around 20000 labeled point clouds
 79 captured by a virtual vehicle navigating in town, highway, rural area, and city. The label
 80 mappings are mapped to that of the nuScenes dataset. In this work, we use the LiDAR point
 81 clouds from Synth4D-nuScenes as a downstream task to validate the generalizability of
 82 pertaining methods. More details of this dataset can be found at <https://github.com/saltoricristiano/gipso-sfouda>.
 - 84 • ¹⁰**DAPS-3D** [11]: The DAPS-3D dataset consists of two subsets: DAPS-1 and DAPS-2;
 85 while the former is a semi-synthetic one with a larger scale, the latter is recorded during
 86 a real field trip of the cleaning robot to the territory of the VDNH Park in Moscow. Both
 87 subsets are with scans collected by a 64-line Ouster OS0 LiDAR sensor. We use the DAPS-1

88 subset in our experiments, which contains 11 LiDAR sequences with more than 23000
89 labeled point clouds. In this work, we use the LiDAR point clouds from DAPS-1 as a
90 downstream task to validate the generalizability of pertaining methods. More details of this
91 dataset can be found at <https://github.com/subake/DAPS3D>.

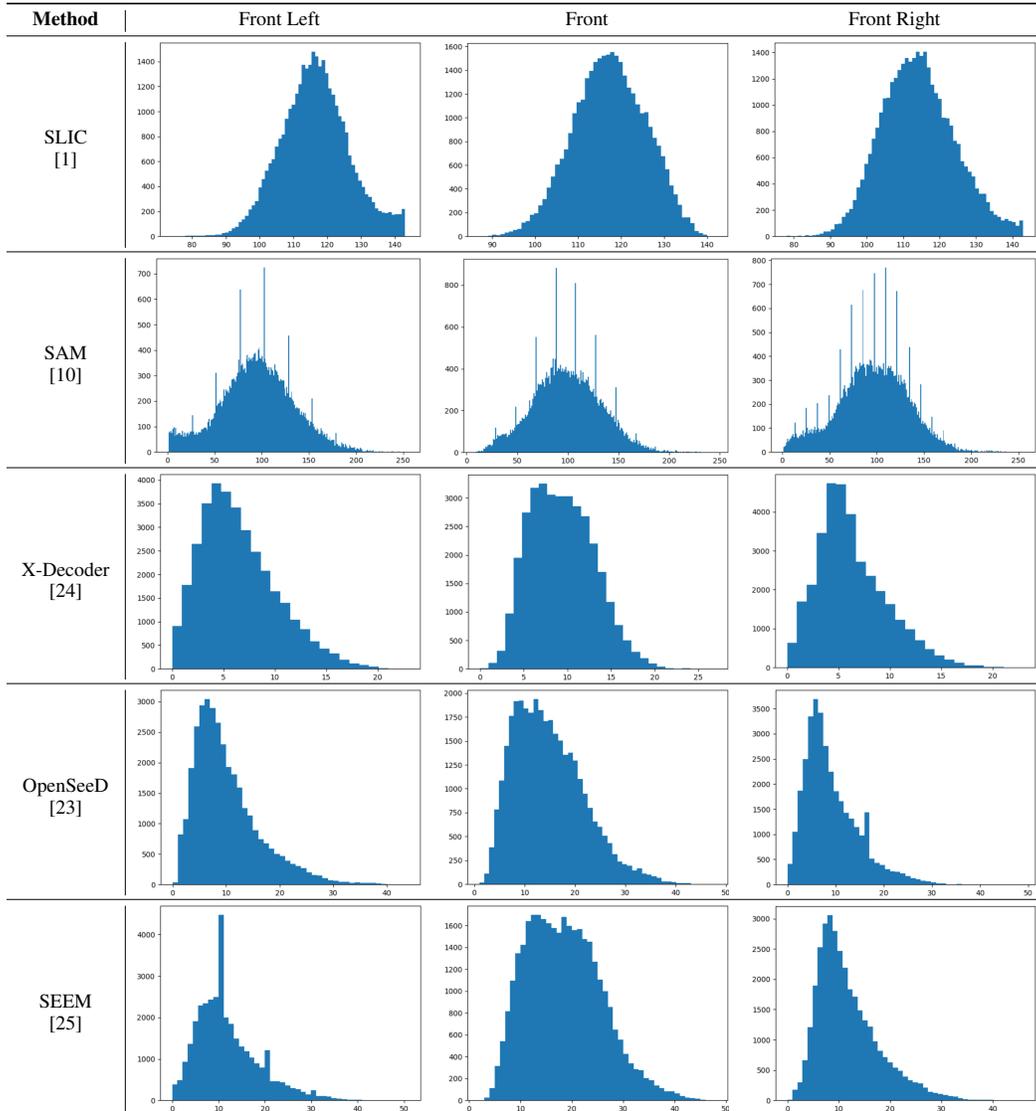
- 92 • ¹¹**nuScenes-C** [12]: The nuScenes-C dataset is one of the corruption sets in the Robo3D
93 benchmark, which is a comprehensive benchmark heading toward probing the robustness of
94 3D detectors and segmentors under out-of-distribution scenarios against natural corruptions
95 that occur in real-world environments. A total number of eight corruption types, stemming
96 from severe weather conditions, external disturbances, and internal sensor failure, are
97 considered, including ‘fog’, ‘wet ground’, ‘snow’, ‘motion blur’, ‘beam missing’, ‘crosstalk’,
98 ‘incomplete echo’, and ‘cross-sensor’ scenarios. These corruptions are simulated with rules
99 constrained by physical principles or engineering experiences. In this work, we use the
100 LiDAR point clouds from nuScenes-C as a downstream task to validate the robustness of
101 pertaining methods under out-of-distribution scenarios. More details of this dataset can be
102 found at <https://github.com/ldkong1205/Robo3D>.

103 A.2 Vision Foundation Models

104 In this work, we conduct comprehensive experiments on analyzing the effects brought by different
105 vision foundation models (VFMs), compared to the traditional SLIC [1] method. Some statistical
106 analyses of these different visual partition methods are shown in Table B and Table C.

- 107 • **SLIC** [1] (traditional method): The SLIC model, which stands for ‘simple linear iterative
108 clustering’, is a popular choice for visual partitions of RGB images. It adapts a k-means
109 clustering approach to generate superpixels, in an efficient manner, and offers good un-
110 supervised partition abilities for many downstream tasks. The pursuit of adherence to
111 boundaries and computational and memory efficiency allows SLIC to perform well on
112 different image collections. In this work, we follow SLidR [18] and use SLIC to generate
113 superpixels, with a fixed quota of 150 superpixels per image, and compare our framework
114 with previous ones using SLIC superpixels. More details of this model can be found at
115 <https://github.com/valeoai/SLidR>.
- 116 • **SAM** [10]: The Segment Anything Model (SAM) is a recent breakthrough towards zero-shot
117 transferable visual understanding across a wide range of tasks. This model is trained on SA-
118 1B, a large-scale dataset with over 1 billion masks on 11M licensed and privacy-respecting
119 images. As a result, SAM is able to segment images, with either point, box, or mask prompts,
120 across different domains and data distributions. In this work, we use a fixed SAM model
121 with the ViT-H backbone (termed as ViT-H SAM model) to generate superpixels. We use
122 this pretrained model directly without any further fine-tuning. More details of this model
123 can be found at <https://github.com/facebookresearch/segment-anything>.
- 124 • **X-Decoder** [24]: The X-Decoder model is a generalized decoding framework that can
125 predict pixel-level segmentation and language tokens seamlessly. This model is pretrained
126 on three types of data, including panoptic segmentation, image-text pairs, and referring
127 segmentation. For the panoptic segmentation task, the model is trained on COCO2017,
128 which includes around 104k images for model training. In this work, we use a fixed X-
129 Decoder model termed as BestSeg-Tiny to generate superpixels. We use this pretrained
130 model directly without any further fine-tuning. More details of this model can be found at
131 <https://github.com/microsoft/X-Decoder>.
- 132 • **OpenSeeD** [23]: The OpenSeeD model is designed for open-vocabulary segmentation
133 and detection, which jointly learns from different segmentation and detection datasets.
134 This model consists of an image encoder, a text encoder, and a decoder with foreground,
135 background, and conditioned mask decoding capability. The model is trained on COCO2017
136 and Objects365, under the tasks of panoptic segmentation and object detection, respectively.
137 In this work, we use a fixed OpenSeeD model termed as openseed-swint-lang to generate
138 superpixels. We use this pretrained model directly without any further fine-tuning. More
139 details of this model can be found at <https://github.com/IDEA-Research/OpenSeeD>.
- 140 • **SEEM** [25]: The SEEM model contributes a new universal interactive interface for image
141 segmentation, where ‘SEEM’ stands for ‘segment everything everywhere with multi-modal

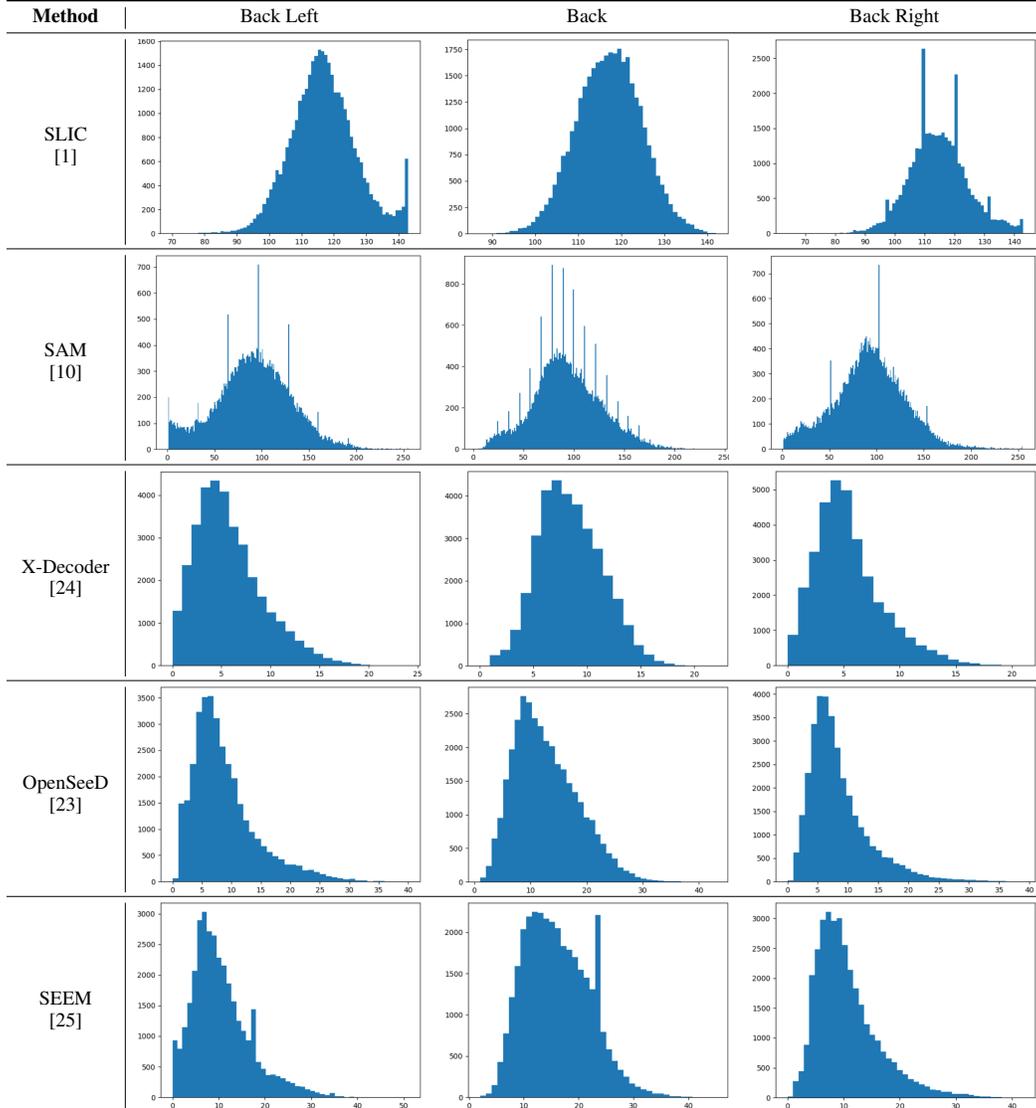
Table B: The **statistics of superpixels** (for front-view cameras) generated by SLIC [1] and different vision foundation modes [10, 24, 23, 25]. The **horizontal axis** denotes the number of superpixels per image. The **vertical axis** denotes the frequency of occurrence.



142 prompts all at once’. The newly designed prompting scheme can encode various user intents
 143 into prompts in a joint visual-semantic space, which possesses properties of versatility,
 144 compositionality, interactivity, and semantic awareness. As such, this model is able to
 145 generalize well to different image datasets, under a zero-shot transfer manner. Similar to
 146 X-Decoder, SEEM is trained on COCO2017, with a total number of 107k images used
 147 during model training. In this work, we use a fixed SEEM model termed as SEEM-oq101
 148 to generate superpixels. We use this pretrained model directly without any further fine-
 149 tuning. More details of this model can be found at [https://github.com/UX-Decoder/](https://github.com/UX-Decoder/Segment-Everything-Everywhere-All-At-Once)
 150 [Segment-Everything-Everywhere-All-At-Once](https://github.com/UX-Decoder/Segment-Everything-Everywhere-All-At-Once).

151 The quality of superpixels directly affects the performance of self-supervised representation learning.
 152 Different from the previous paradigm [18, 15], our *Seal* framework resorts to the recent VFMs
 153 for generating superpixels. Compared to the traditional SLIC method, these VFMs are able to
 154 generate semantically-aware superpixels that represent coherent semantic meanings of objects and
 155 backgrounds around the ego-vehicle.

Table C: The **statistics of superpixels** (for back-view cameras) generated by SLIC [1] and different vision foundation modes [10, 24, 23, 25]. The **horizontal axis** denotes the number of superpixels per image. The **vertical axis** denotes the frequency of occurrence.



156 As been verified in our experiments, these semantic superpixels have the ability to ease the “over-
 157 segment” problem in current self-supervised learning frameworks [18, 15] and further improve the
 158 performance for both linear probing and downstream fine-tuning.

159 The histograms shown in Table B and Table C verify that the number of superpixels per image of
 160 VFMs is much smaller than that of SLIC. This brings two notable advantages: *i*) Since semantically
 161 similar objects and backgrounds are grouped together in semantic superpixels, the “self-conflict”
 162 problem in existing approaches is largely mitigated, which directly boosts the quality of representation
 163 learning. *ii*) Since the embedding length D of the superpixel embedding features $\mathbf{Q} \in \mathbb{R}^{M \times D}$ and
 164 superpoint embedding features $\mathbf{K} \in \mathbb{R}^{M \times D}$ directly relates to computation overhead, a reduction
 165 (*e.g.*, around 150 superpixels per image in SLIC [1] and around 25 superpixels per image in X-
 166 Decoder [24], OpenSeeD [23], and SEEM [25]) on D would allow us to train the segmentation
 167 models in a more efficient manner.

168 Some typical examples of our generated superpixels and their corresponding superpoints are shown
 169 in Fig. A, Fig. B, Fig. C, and Fig. D. As will be shown in Section B, our use of semantic superpixels

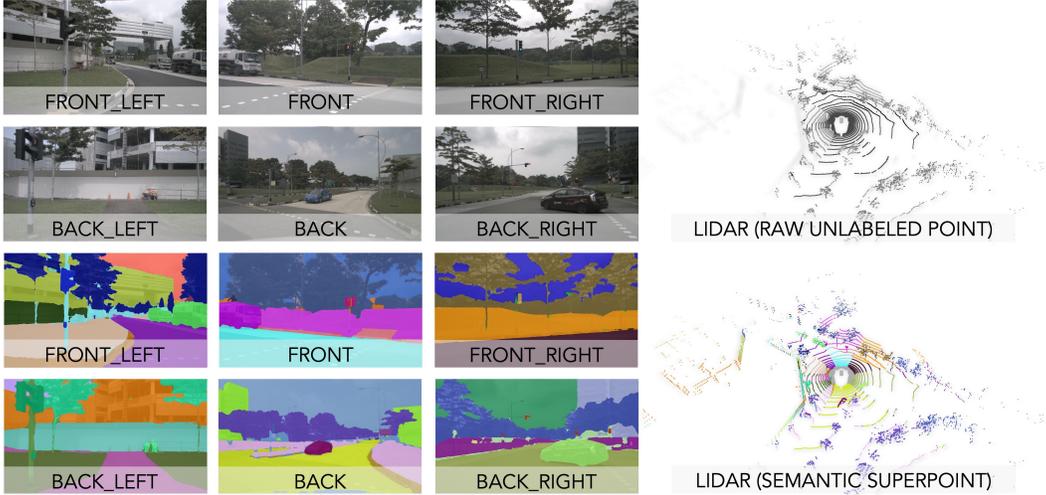


Figure A: Illustration of the **semantic superpixel-to-superpoint transformation** in the proposed *Seal* framework. [Row 1 & 2] The raw data captured by the multi-view camera and LiDAR sensors. [Row 3 & 4] The semantic superpixel on the camera images and superpoint formed by projecting superpixel into the point cloud via camera-LiDAR correspondence. The superpixels are generated using SEEM [25]. Each color represents one distinct segment. Best viewed in color.

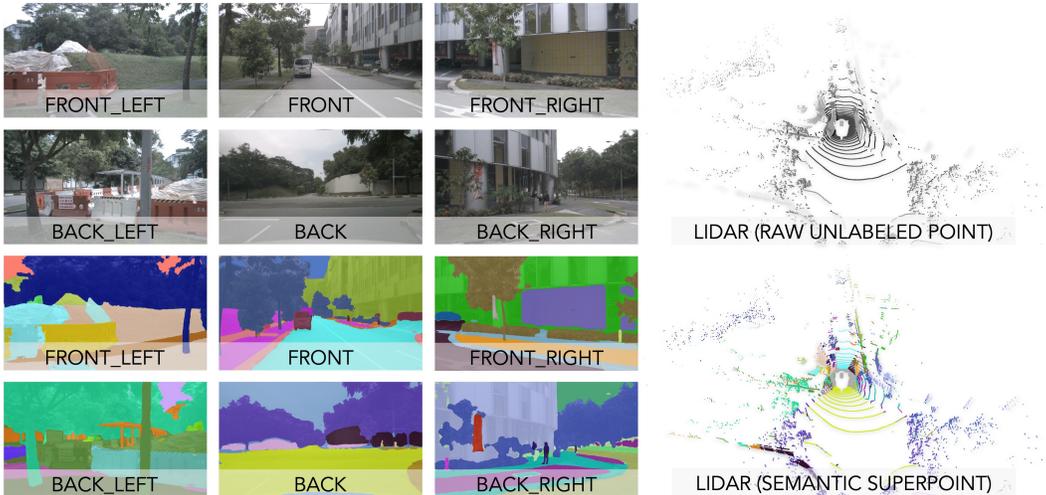


Figure B: Illustration of the **semantic superpixel-to-superpoint transformation** in the proposed *Seal* framework. [Row 1 & 2] The raw data captured by the multi-view camera and LiDAR sensors. [Row 3 & 4] The semantic superpixel on the camera images and superpoint formed by projecting superpixel into the point cloud via camera-LiDAR correspondence. The superpixels are generated using SEEM [25]. Each color represents one distinct segment. Best viewed in color.

170 generated by VFMs brings not only performance gains but also a much faster convergence rate during
 171 the model pretraining stage.

172 A.3 Implementation Details

173 A.3.1 Data Split

174 For **model pertaining**, we follow the SLiDR protocol [18] in data splitting. Specifically, the nuScenes
 175 [7] dataset consists of 700 training scenes in total, 100 of which are kept aside, which constitute the
 176 SLiDR mini-val split. All models are pretrained using all the scans from the 600 remaining training
 177 scenes. The 100 scans in the mini-val split are used to find the best-possible hyperparameters. The

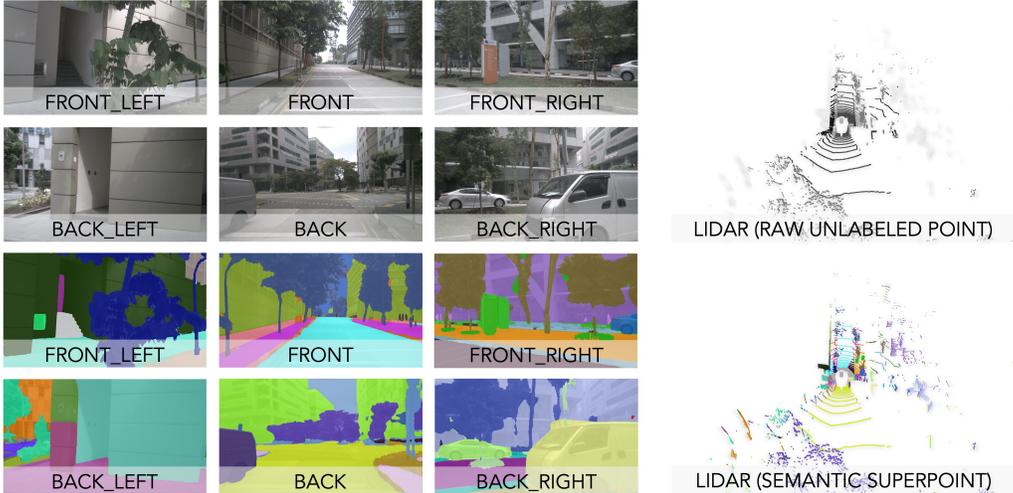


Figure C: Illustration of the **semantic superpixel-to-superpoint transformation** in the proposed *Seal* framework. [Row 1 & 2] The raw data captured by the multi-view camera and LiDAR sensors. [Row 3 & 4] The semantic superpixel on the camera images and superpoint formed by projecting superpixel into the point cloud via camera-LiDAR correspondence. The superpixels are generated using SEEM [25]. Each color represents one distinct segment. Best viewed in color.

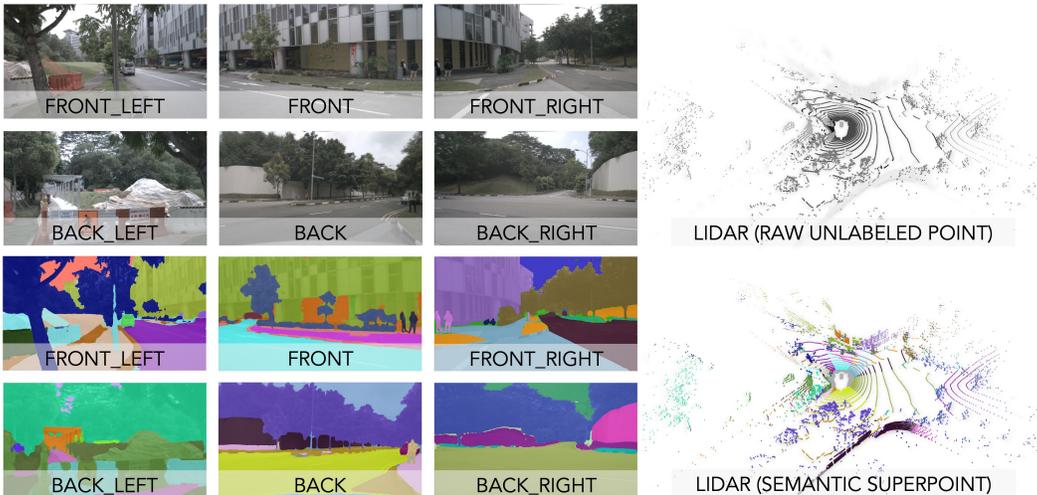


Figure D: Illustration of the **semantic superpixel-to-superpoint transformation** in the proposed *Seal* framework. [Row 1 & 2] The raw data captured by the multi-view camera and LiDAR sensors. [Row 3 & 4] The semantic superpixel on the camera images and superpoint formed by projecting superpixel into the point cloud via camera-LiDAR correspondence. The superpixels are generated using SEEM [25]. Each color represents one distinct segment. Best viewed in color.

178 trained models are then validated on the official nuScenes validation set, without any kind of test-time
 179 augmentation or model ensemble. This is to ensure a fair comparison with previous works and also
 180 in line with the practical requirements.

181 For **linear probing**, the pretrained 3D network F_{θ_p} is frozen with a trainable point-wise linear
 182 classification head which is trained for 50 epochs on a A100 GPU with a learning rate of 0.05, and
 183 batch size is 16 on the nuScenes train set for all methods.

184 For **downstream fine-tuning** tasks, we stick with the common practice in SLiDR [18] whenever
 185 possible. The detailed data split strategies are summarized as follows.

- 186 • For fine-tuning on **nuScenes** [7], we follow the SLiDR protocol to split the train set of
187 nuScenes to generate 1%, 5%, 10%, 25%, and 100% annotated scans for the training subset.
- 188 • For fine-tuning on **DAPS-3D** [11], we take sequences 38-18_7_72_90 as the training set
189 and 38-18_7_72_90, 42-48_10_78_90, and 44-18_11_15_32 as the validation set.
- 190 • For fine-tuning on **SynLiDAR** [21], we use the sub-set which is a uniformly downsampled
191 collection from the whole dataset.
- 192 • For fine-tuning on **SemanticPOSS** [16], we use sequences 00 and 01 as *half* of the annotated
193 training scans and use sequences 00 to 05, except 02 for validation to create *full* of the
194 annotated training samples.
- 195 • For fine-tuning on **SemanticKITTI** [2], **Waymo Open** [19], **ScribbleKITTI** [20], **RELLIS-**
196 **3D** [9], **SemanticSTF** [22], and **Synth4D** [17], we follow the SLiDR protocol to create
197 1%, 10%, *half*, or *full* split of the annotated training scans, *e.g.*, one scan is taken every 100
198 frame from the training set to get 1% of the labeled training samples. Notably, the point
199 cloud segmentation performance in terms of IoU is reported on the official validation sets
200 for all the above-mentioned datasets.

201 A.3.2 Experimental Setup

202 In our experiments, we fine-tune the entire 3D network on the semantic segmentation task using a
203 linear combination of the cross-entropy loss and the Lovász-Softmax loss [3] as training objectives
204 on a single A100 GPU. For the few-shot semantic segmentation tasks, the 3D networks are fine-tuned
205 for 100 epochs with a batch size of 10 for the SemanticKITTI [2], Waymo Open [19], ScribbleKITTI
206 [20], RELLIS-3D [9], SemanticSTF [22], SemanticPOSS [16], DAPS-3D [11], SynLiDAR [21], and
207 Synth4D [17] datasets.

208 For the nuScenes [7] dataset, we fine-tune the 3D network for 100 epochs with a batch size of 16
209 while training on the 1% annotated scans. The 3D network train on the other portions of nuScenes is
210 fine-tuned for 50 epochs with a batch size of 16. We adopt different learning rates on the 3D backbone
211 F_{θ_p} and the classification head, except for the case that F_{θ_p} is randomly initialized. The learning rate
212 of F_{θ_p} is set as 0.05 and the learning rate of the classification head is set as 2.0, respectively, for all
213 the above-mentioned datasets except nuScenes. On the nuScenes dataset, the learning rate of F_{θ_p} is
214 set as 0.02.

215 We train our framework using the SGD optimizer with a momentum of 0.9, a weight decay of 0.0001,
216 and a dampening ratio of 0.1. The cosine annealing learning rate strategy is adopted which decreases
217 the learning rate from its initial value to zero at the end of the training.

218 A.3.3 Data Augmentation

219 For the **model pretraining**, we apply two sets of data augmentations on the point cloud and the
220 multi-view image, respectively, and update the point-pixel correspondences after each augmentation
221 by following SLiDR [18].

- 222 • Regarding the point cloud, we adopt a random rotation around the z -axis and flip the x -axis
223 or y -axis with a 50% probability. Besides, we randomly drop the cuboid where the length of
224 each side covers no more than 10% range of point coordinates on the corresponding axis,
225 and the cuboid center is located on a randomly chosen point in the point cloud. We also
226 ensure that the dropped cuboid retains at least 1024 pairs of points and pixels; otherwise,
227 we select another new cuboid instead. For the temporal frame, we apply the same data
228 augmentation to it.
- 229 • Regarding the multi-view image, we apply a horizontal flip with a 50% probability and a
230 cropped resize which reshapes the image to 416×224 . Before resizing, the random crop
231 fills at least 30% of the available image space with a random aspect ratio between 14 : 9 and
232 17 : 9. If this random cropping does not preserve at least 1024 or 75% of the pixel-point
233 pairs, a different crop is chosen.

234 For the **downstream fine-tuning** tasks, we apply a random rotation around the z -axis and flip the x -
235 axis or y -axis with a 50% probability for all the points in the point cloud. If the results are mentioned
236 with LaserMix [13] augmentation, we augment the point clouds via the LaserMix before the above

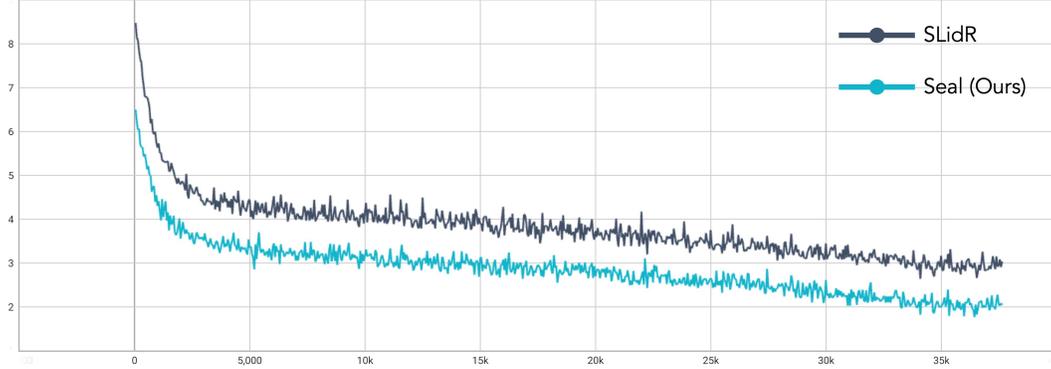


Figure E: The convergence rate comparison between SLidR [18] and the proposed *Seal* framework.

237 data augmentations. We report both results in the main paper so that we can fairly compare the point
 238 cloud segmentation performance with previous works.

239 A.3.4 Model Configuration

240 For the **model pretraining**, the 3D backbone F_{θ_p} is a Minkowski U-Net [6] with $3 \times 3 \times 3$ kernels;
 241 while the 2D image encoder G_{θ_i} is a ResNet-50 [8] initialized with 2D self-supervised pretrained
 242 model of MoCoV2 [5]. These configurations are kept the same as SLidR [18]. The channel dimension
 243 of G_{θ_i} head and F_{θ_p} head is set to 64. For the **linear probing** task, we adopt a linear classification
 244 head, as mentioned in previous sections. For the **downstream fine-tuning** tasks, the same 3D
 245 backbone F_{θ_p} , *i.e.*, the Minkowski U-Net [6] with $3 \times 3 \times 3$ kernels, is used.

246 B Additional Quantitative Result

247 B.1 Self-Supervised Learning

248 We report the complete results (*i.e.*, the class-wise IoU scores) for the **linear probing** and **downstream**
 249 **fine-tuning** tasks shown in the main paper. We report the official results from previous works
 250 whenever possible. We also report our reproduced results for random initialization, PPKT [14], and
 251 SLidR [18]. Specifically, the complete results on the nuScenes [7], SemanticKITTI [2], Waymo Open
 252 [19], and Synth4D [17] datasets are shown in Table D, Table E, Table F, and Table G, respectively.
 253 We observe that *Seal* constantly outperforms prior methods for most semantic classes on all datasets.

254 We also compare the convergence rate of *Seal* with SLidR [18] in Fig. E. As can be seen, our
 255 framework is able to converge faster with the use of semantic superpixels, where these higher
 256 qualitative contrastive samples gradually form a much more coherent optimization landscape.

257 B.2 Robustness Probing

258 We report the complete results of the robustness evaluation experiments in the main paper, including
 259 the per-corruption CE scores and the per-corruption RR scores. Specifically, the complete results on
 260 the nuScenes-C [12] dataset are shown in Table H and Table I. We observe that *Seal* is superior to
 261 previous methods in terms of both CE and RR metrics.

262 C Additional Qualitative Result

263 C.1 Cosine Similarity

264 We provide more examples for the cosine similarity study in Fig. F, Fig. G, and Fig. H.

265 C.2 Downstream Tasks

266 We provide more qualitative results for downstream fine-tuning tasks in Fig. I, Fig. J, and Fig. K.

Table D: The **per-class IoU scores** of different pretraining methods pretrained on *nuScenes* [7] and linear probed or fine-tuned on different proportions (1%, 5%, 10%, 25%, and Full) of the *nuScenes* [7] data. Symbol \clubsuit denotes our reproduced results and the remaining are reported scores. All IoU scores are given in percentage (%). The **best mIoU score** is highlighted in **bold**.

Method	mIoU	barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	driv. surf.	other flat	sidewalk	terrain	manmade	vegetation
Linear Probing																	
Random \clubsuit	8.4	0.5	0.0	0.0	3.9	0.0	0.0	0.0	6.4	0.0	3.9	60.4	0.0	0.1	16.2	30.6	12.2
PointContrast	21.9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
DepthContrast	22.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
PPKT	35.9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
SLiDR \clubsuit	39.2	44.2	0.0	30.8	60.2	15.1	22.4	47.2	27.7	16.3	34.3	80.6	21.8	35.2	48.1	71.0	71.9
ST-SLiDR	40.5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seal (Ours)	45.0	54.7	5.9	3.6	61.7	18.9	28.8	48.1	31.0	22.1	39.5	83.8	35.4	46.7	56.9	74.7	74.7
Fine-tuning (1%)																	
Random	30.3	0.0	0.0	8.1	65.0	0.1	6.6	21.0	9.0	9.3	25.8	89.5	14.8	41.7	48.7	72.4	73.3
PointContrast	32.5	0.0	1.0	5.6	67.4	0.0	3.3	31.6	5.6	12.1	30.8	91.7	21.9	48.4	50.8	75.0	74.6
DepthContrast	31.7	0.0	0.6	6.5	64.7	0.2	5.1	29.0	9.5	12.1	29.9	90.3	17.8	44.4	49.5	73.5	74.0
PPKT	37.8	0.0	2.2	20.7	75.4	1.2	13.2	45.6	8.5	17.5	38.4	92.5	19.2	52.3	56.8	80.1	80.9
SLiDR	38.8	0.0	1.8	15.4	73.1	1.9	19.9	47.2	17.1	14.5	34.5	92.0	27.1	53.6	61.0	79.8	82.3
ST-SLiDR	40.8	0.0	2.7	16.0	74.5	3.2	25.4	50.9	20.0	17.7	40.2	92.0	30.7	54.2	61.1	80.5	82.9
Seal (Ours)	45.8	0.0	9.4	32.6	77.5	10.4	28.0	53.0	25.0	30.9	49.7	94.0	33.7	60.1	59.6	83.9	83.4
Fine-tuning (5%)																	
Random	47.8	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Random \clubsuit	44.5	50.1	2.9	57.3	70.3	1.1	6.1	39.1	18.3	17.3	44.8	92.3	38.6	54.9	61.1	80.3	77.9
PPKT \clubsuit	52.7	56.1	8.2	65.3	79.0	9.1	15.5	54.3	34.5	26.7	58.6	93.2	44.1	63.1	64.8	85.1	83.6
SLiDR	52.5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ST-SLiDR	54.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seal (Ours)	55.6	61.0	7.4	70.4	82.4	11.9	30.4	59.2	34.0	33.6	61.1	94.7	46.2	63.4	63.9	85.7	84.9
Fine-tuning (10%)																	
Random	56.2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Random \clubsuit	53.2	56.8	5.2	66.3	74.5	5.7	36.1	49.5	38.2	29.2	54.4	94.0	47.7	61.4	67.6	82.8	81.1
PPKT \clubsuit	60.3	64.0	12.0	67.8	77.6	16.0	56.8	63.3	49.8	28.3	56.3	94.1	62.7	66.4	68.7	85.9	85.4
SLiDR	59.8	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ST-SLiDR	60.8	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seal (Ours)	63.0	64.9	15.1	78.9	83.5	22.5	61.0	63.0	51.5	37.4	65.2	95.2	59.4	67.7	69.6	86.2	86.4
Fine-tuning (25%)																	
Random	65.5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Random \clubsuit	63.0	64.9	15.1	78.9	83.5	22.5	61.0	63.0	51.5	37.4	65.2	95.2	59.4	67.7	69.6	86.2	86.4
PPKT \clubsuit	67.1	63.7	12.1	87.4	85.2	42.0	61.2	69.6	54.7	50.1	74.9	95.8	64.6	69.9	70.4	87.4	85.1
SLiDR	66.9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ST-SLiDR	67.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seal (Ours)	68.4	67.4	15.5	90.5	85.1	40.0	62.3	68.1	58.3	54.5	76.0	95.8	65.6	69.8	71.0	87.9	86.7
Fine-tuning (Full)																	
Random	74.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Random \clubsuit	74.9	77.4	34.4	90.4	86.5	51.1	78.4	77.0	64.9	67.8	77.7	96.6	72.2	74.3	73.2	89.6	86.7
PPKT \clubsuit	74.5	75.4	37.4	91.7	86.1	45.3	77.3	75.6	65.1	65.0	79.9	96.6	72.1	74.6	73.4	89.2	86.6
SLiDR	74.8	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ST-SLiDR	75.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seal (Ours)	75.6	76.7	26.0	93.4	86.1	55.8	82.7	77.3	66.0	67.9	83.9	96.5	73.6	74.3	73.2	89.3	87.1

Table E: The **per-class IoU scores** of different pretraining methods pretrained on *nuScenes* [7] and fine-tuned on 1% of the *SemanticKITTI* [2] data. Symbol \clubsuit denotes our reproduced results and the remaining are reported scores. All IoU scores are given in percentage (%). The **best mIoU score** is highlighted in **bold**.

Method	mIoU	car	bicycle	motorcycle	truck	other-vehicle	person	bicyclist	motorcyclist	road	parking	sidewalk	other-ground	building	fence	vegetation	trunk	trunk	trunk	trunk
Random	39.5	91.2	0.0	9.4	8.0	10.7	21.2	0.0	0.0	89.4	21.4	73.0	1.1	85.3	41.1	84.9	50.1	71.4	55.4	37.6
PPKT	43.9	91.3	1.9	11.2	23.1	12.1	27.4	37.3	0.0	91.3	27.0	74.6	0.3	86.5	38.2	85.3	58.2	71.6	57.7	40.1
SLiDR	44.6	92.2	3.0	17.0	22.4	14.3	36.0	22.1	0.0	91.3	30.0	74.7	0.2	87.7	41.2	85.0	58.5	70.4	58.3	42.4
ST-SLiDR	44.7	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seal	46.6	92.3	14.9	18.7	16.1	23.7	43.0	34.4	0.0	91.3	27.2	75.3	0.7	85.7	38.8	85.1	61.9	71.3	57.7	47.7

Table F: The **per-class IoU scores** of different pretraining methods pretrained on *nuScenes* [7] and fine-tuned on 1% of the *Waymo Open* [19] data. Symbol \ddagger denotes our reproduced results and the remaining are reported scores. All IoU scores are given in percentage (%). The **best** mIoU score is highlighted in **bold**.

Method	mIoU	car	truck	bus	other vehicle	motorcyclist	bicyclist	pedestrian	sign	traffic light	pole	construction	bicycle	motorcycle	building	vegetation	tree trunk	curb	road	lane marker	other ground	walkable	sidewalk	
Random	39.4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
PPKT \ddagger	47.6	92.3	48.8	34.8	6.9	0.0	22.0	73.0	58.3	17.6	55.8	20.4	25.5	7.8	91.5	87.1	53.8	55.1	89.1	38.6	34.7	71.4	62.6	
SLiDR	47.1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ST-SLiDR	44.9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seal	49.3	92.5	52.6	33.5	3.7	0.0	31.9	73.1	61.0	23.5	57.0	31.4	20.1	12.4	91.4	87.1	53.2	57.2	89.5	38.8	34.9	72.1	68.7	

Table G: The **per-class IoU scores** of different pretraining methods pretrained on *nuScenes* [7] and fine-tuned on 1% of the *Synth4D* [17] data. Symbol \ddagger denotes our reproduced results and the remaining are reported scores. All IoU scores are given in percentage (%). The **best** mIoU score is highlighted in **bold**.

Method	mIoU	building	fences	other	pedestrian	pole	roadlines	road	sidewalk	vegetation	vehicle	wall	traffic sign	sky	ground	bridge	rail track	guardrail	traffic light	static	dynamic	water	terrain
Random	20.2	33.2	13.8	0.0	16.3	13.9	0.0	88.6	51.0	48.9	95.7	27.6	0.0	0.0	0.0	0.0	0.1	19.4	0.0	0.2	0.0	0.0	36.2
PPKT \ddagger	61.1	84.4	67.1	62.9	77.4	75.0	2.2	92.0	76.3	92.5	99.2	73.4	67.3	0.0	62.2	39.6	83.5	66.8	0.0	63.6	55.3	37.4	66.0
SLiDR \ddagger	63.1	83.7	66.3	64.9	77.4	76.4	6.5	92.8	78.7	92.5	99.0	72.5	64.2	0.0	74.3	48.9	85.3	67.1	0.0	67.1	60.0	47.1	63.6
Seal	64.5	84.8	70.5	64.8	80.3	76.3	9.3	92.9	79.8	92.7	98.9	73.0	60.7	0.0	75.2	55.3	84.6	67.0	0.0	68.2	60.7	53.2	70.9

Table H: The **Corruption Error (CE) scores** of different pretraining methods pretrained on *nuScenes* [7] and probed under the eight out-of-distribution corruptions in the *nuScenes-C* dataset from the Robo3D benchmark [12]. All CE scores are given in percentage (%). The **best** CE score for each corruption type is highlighted in **bold**.

	Initial	Backbone	mCE \downarrow	Fog	Wet	Snow	Move	Beam	Cross	Echo	Sensor
LP	PPKT	MinkUNet ₁₈	183.44	149.59	247.53	120.50	266.03	213.54	109.62	199.03	161.65
	SLiDR	MinkUNet ₁₈	179.38	140.47	237.29	112.93	276.44	210.65	107.86	189.27	160.16
	Seal (Ours)	MinkUNet ₁₈	166.18	135.18	219.36	117.47	234.01	189.70	108.54	172.16	153.03
Full	Random	PolarNet	115.09	90.10	115.33	58.98	208.19	121.07	80.67	128.17	118.23
	Random	FIDNet	122.42	75.93	122.58	68.78	192.03	164.84	57.95	141.66	155.56
	Random	CENet	112.79	71.16	115.48	64.31	156.67	159.03	53.27	129.08	153.35
	Random	WaffleIron	106.73	94.76	99.92	84.51	152.35	110.65	91.09	106.41	114.15
	Random	Cylinder3D	105.56	83.22	111.08	69.74	165.28	113.95	74.42	110.67	116.15
	Random	SPVCNN ₁₈	106.65	88.42	105.56	98.78	156.48	110.11	86.04	104.26	103.55
	Random	SPVCNN ₃₄	97.45	95.21	99.50	97.32	95.34	98.73	97.92	96.88	98.74
	Random	MinkUNet ₁₈	112.20	79.90	112.50	74.64	181.47	120.76	93.22	111.58	123.53
	PPKT	MinkUNet ₁₈	105.64	77.63	104.22	68.60	160.95	114.81	86.71	108.96	123.20
	SLiDR	MinkUNet ₁₈	106.08	74.61	106.13	73.75	165.09	118.02	79.08	107.38	124.57
	Seal (Ours)	MinkUNet ₁₈	92.63	58.97	98.47	56.63	127.25	108.20	57.97	110.95	122.63

Table I: The **Resilience Rate (RR) scores** of different pretraining methods pretrained on *nuScenes* [7] and probed under the eight out-of-distribution corruptions in the *nuScenes-C* dataset from the Robo3D benchmark [12]. All RR scores are given in percentage (%). The **best** RR score for each corruption type is highlighted in **bold**.

	Initial	Backbone	mRR \uparrow	Fog	Wet	Snow	Move	Beam	Cross	Echo	Sensor
LP	PPKT	MinkUNet ₁₈	78.15	85.38	98.66	78.33	81.36	91.42	54.37	78.02	57.69
	SLiDR	MinkUNet ₁₈	77.18	89.90	98.17	84.12	68.14	86.93	53.63	81.29	55.26
	Seal (Ours)	MinkUNet ₁₈	75.38	83.05	95.15	66.59	83.94	89.70	45.18	83.94	55.48
Full	Random	PolarNet	76.34	81.59	97.95	90.82	62.49	86.75	57.12	75.16	58.86
	Random	FIDNet	73.33	90.78	95.29	82.61	68.51	67.44	80.48	68.31	33.20
	Random	CENet	76.04	91.44	95.35	84.12	79.57	68.19	83.09	72.75	33.82
	Random	WaffleIron	72.78	73.71	97.19	65.19	78.16	85.70	43.54	80.86	57.85
	Random	Cylinder3D	78.08	83.52	96.57	79.41	76.18	87.23	61.68	81.55	58.51
	Random	SPVCNN ₁₈	74.70	79.31	97.39	55.22	78.44	87.85	49.50	83.72	66.14
	Random	SPVCNN ₃₄	75.10	72.95	96.70	54.79	97.47	90.04	36.71	84.84	67.35
	Random	MinkUNet ₁₈	72.57	84.33	94.63	74.31	69.26	83.06	42.27	79.88	52.79
	PPKT	MinkUNet ₁₈	76.06	85.90	97.71	79.28	76.72	85.72	48.77	81.31	53.10
	SLiDR	MinkUNet ₁₈	75.99	87.46	96.68	74.89	74.97	84.06	56.08	81.78	52.01
	Seal (Ours)	MinkUNet ₁₈	83.08	96.11	98.29	87.59	87.49	87.25	75.98	79.19	52.71

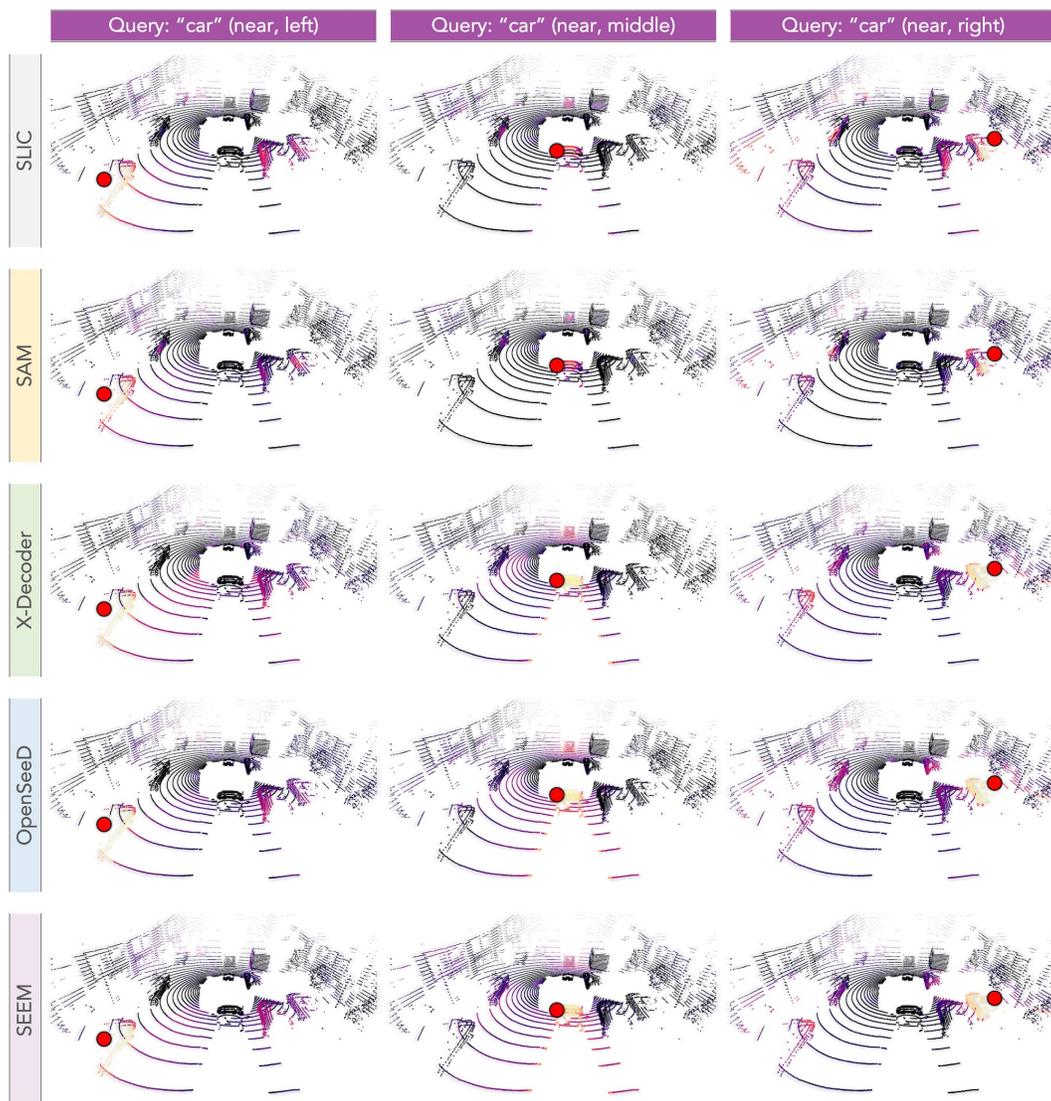


Figure F: The **cosine similarity** between the query point (denoted as the **red dot**) and the feature learned with SLIC [1] and different VFMs [10, 24, 23, 25]. The color goes from **violet** to **yellow** denoting **low** and **high** similarity scores, respectively. Best viewed in color.

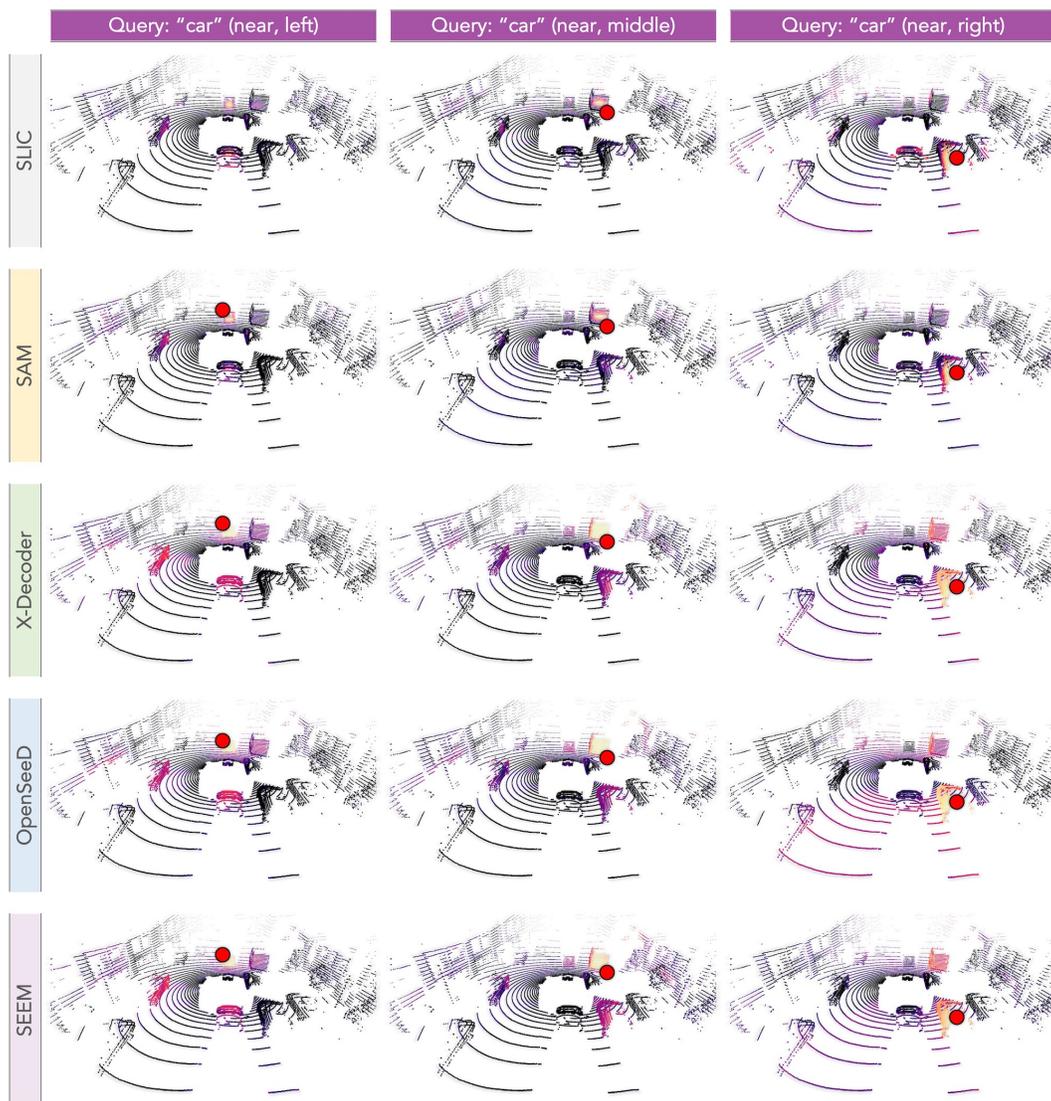


Figure G: The **cosine similarity** between the query point (denoted as the **red dot**) and the feature learned with SLIC [1] and different VFMs [10, 24, 23, 25]. The color goes from **violet** to **yellow** denoting **low** and **high** similarity scores, respectively. Best viewed in color.

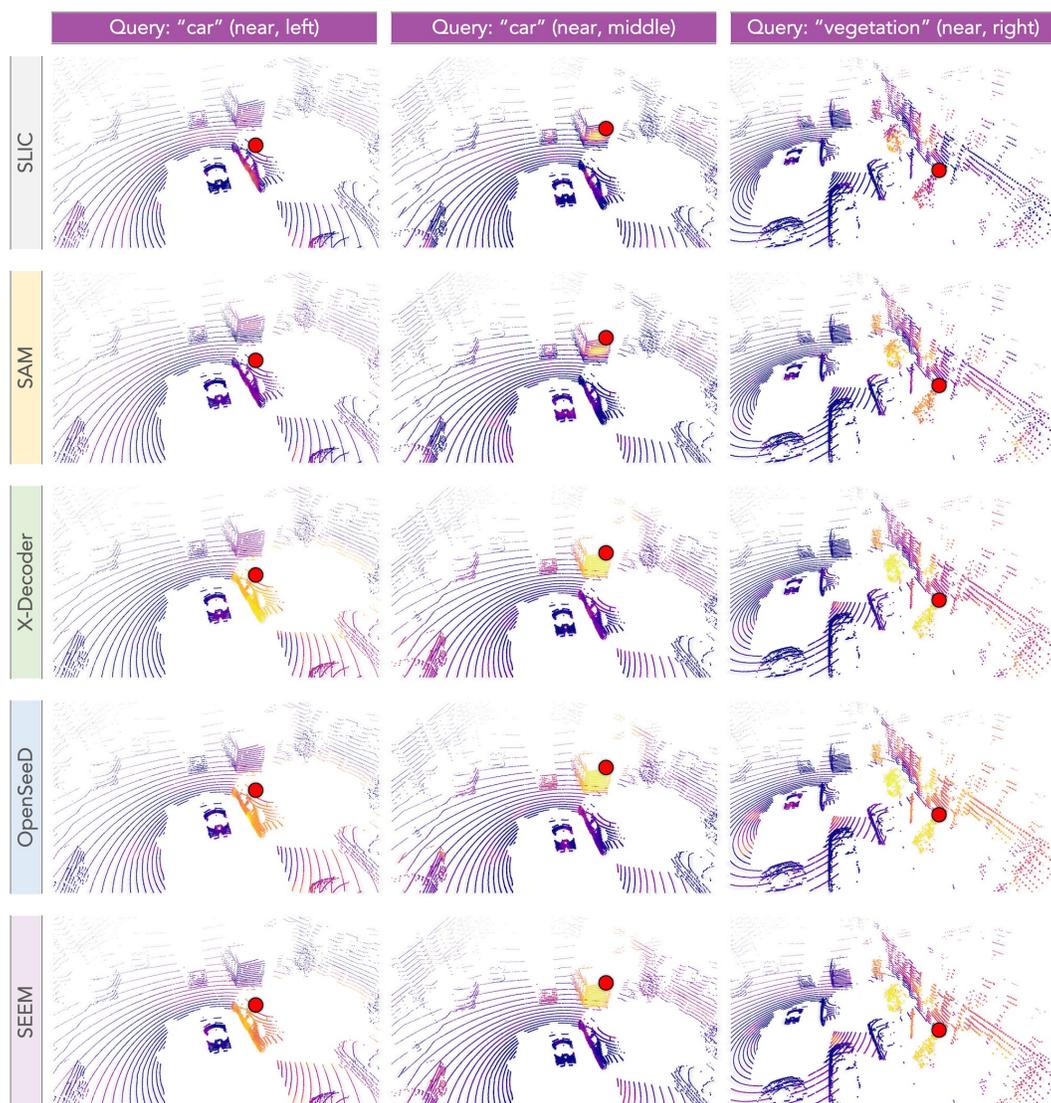


Figure H: The **cosine similarity** between the query point (denoted as the **red dot**) and the feature learned with SLIC [1] and different VFMs [10, 24, 23, 25]. The color goes from **violet** to **yellow** denoting **low** and **high** similarity scores, respectively. Best viewed in color.

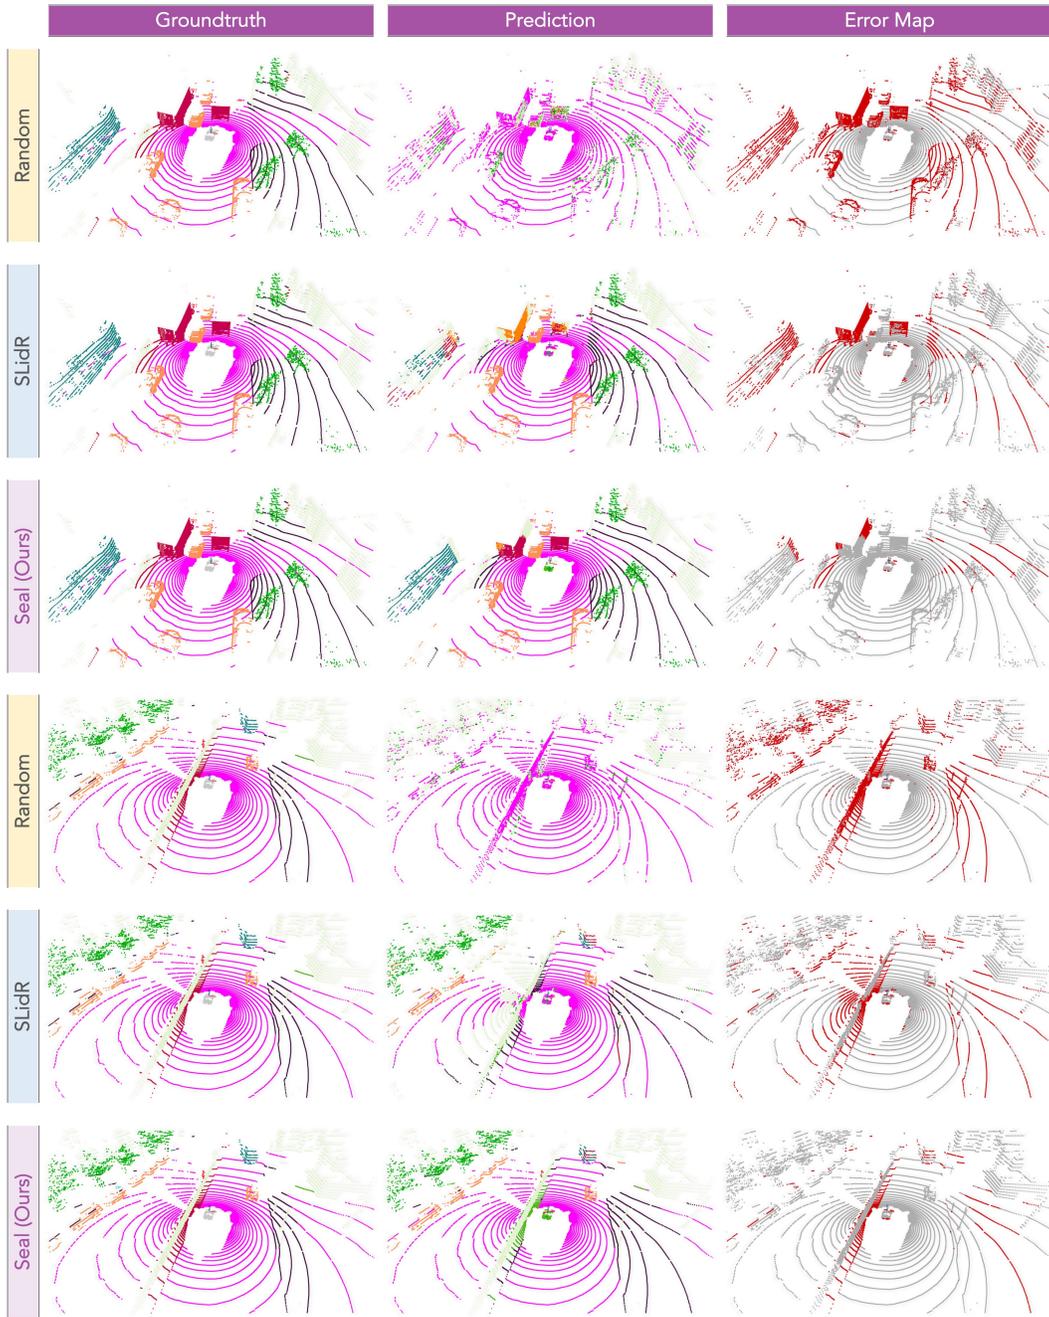


Figure I: The **qualitative results** of different point cloud pretraining approaches pretrained on the raw data of *nuScenes* [7] and fine-tuned with 1% labeled data. To highlight the differences, the **correct / incorrect** predictions are painted in **gray / red**, respectively. Best viewed in color.

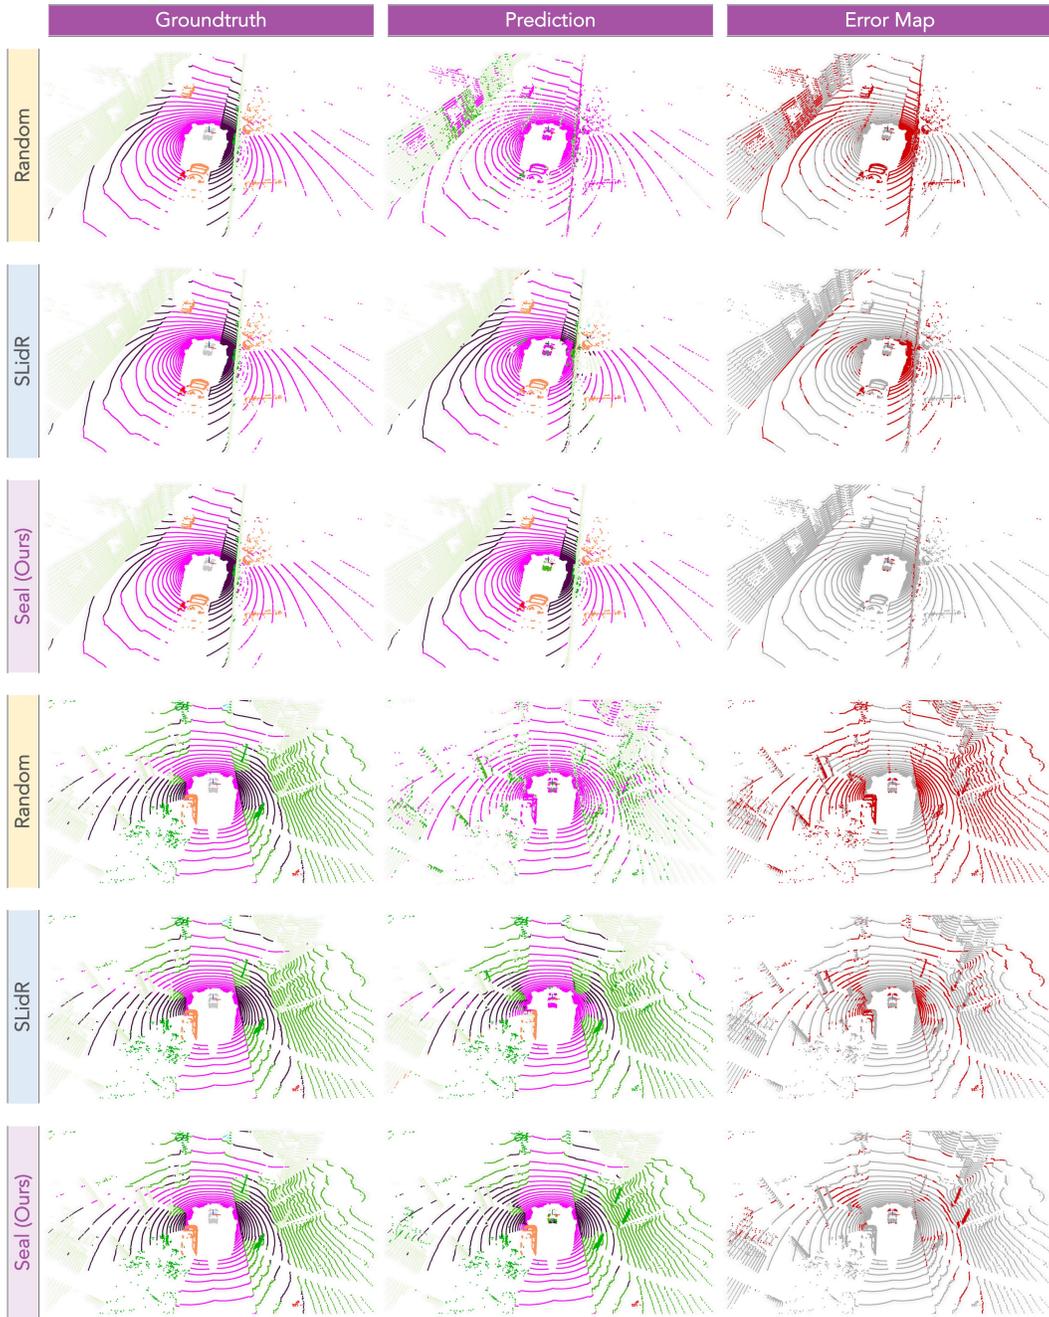


Figure J: The **qualitative results** of different point cloud pretraining approaches pretrained on the raw data of *nuScenes* [7] and fine-tuned with 1% labeled data. To highlight the differences, the **correct / incorrect** predictions are painted in **gray / red**, respectively. Best viewed in color.

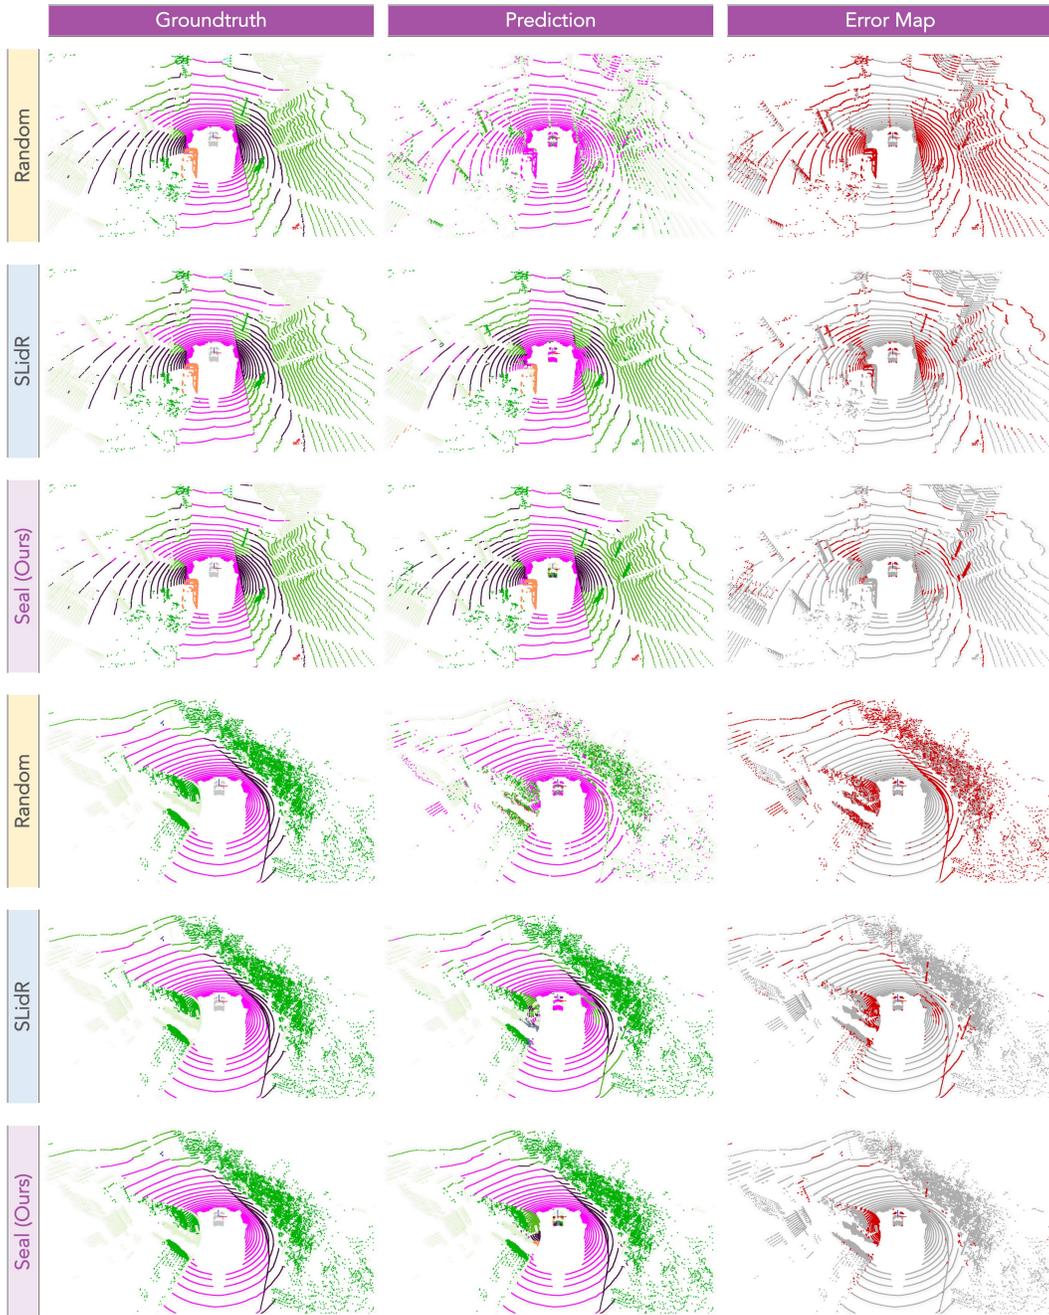


Figure K: The **qualitative results** of different point cloud pretraining approaches pretrained on the raw data of *nuScenes* [7] and fine-tuned with 1% labeled data. To highlight the differences, the **correct** / **incorrect** predictions are painted in **gray** / **red**, respectively. Best viewed in color.

267 **D Public Resources Used**

268 We acknowledge the use of the following public resources, during the course of this work:

269	• nuScenes ²	CC BY-NC-SA 4.0
270	• nuScenes-devkit ³	Apache License 2.0
271	• SemanticKITTI ⁴	CC BY-NC-SA 4.0
272	• SemanticKITTI-API ⁵	MIT License
273	• Waymo Open Dataset ⁶	Waymo Dataset License
274	• ScribbleKITTI ⁷	Unknown
275	• RELLIS-3D ⁸	CC BY-NC-SA 3.0
276	• SemanticPOSS ⁹	Unknown
277	• SemanticSTF ¹⁰	CC BY-NC-SA 4.0
278	• SynLiDAR ¹¹	MIT License
279	• Synth4D ¹²	GNU General Public License 3.0
280	• DAPS-3D ¹³	MIT License
281	• nuScenes-C ¹⁴	CC BY-NC-SA 4.0
282	• MinkowskiEngine ¹⁵	MIT License
283	• SLiDR ¹⁶	Apache License 2.0
284	• spvnas ¹⁷	MIT License
285	• Cylinder3D ¹⁸	Apache License 2.0
286	• LaserMix ¹⁹	CC BY-NC-SA 4.0
287	• mean-teacher ²⁰	Attribution-NonCommercial 4.0 International
288	• PyTorch-Lightning ²¹	Apache License 2.0
289	• mmdetection3d ²²	Apache License 2.0

²<https://www.nuscenes.org/nuscenes>.

³<https://github.com/nutonomy/nuscenes-devkit>.

⁴<http://semantic-kitti.org>.

⁵<https://github.com/PRBonn/semantic-kitti-api>.

⁶<https://waymo.com/open>.

⁷<https://github.com/ouenal/scribblekitti>.

⁸<http://www.unmannedlab.org/research/RELLIS-3D>.

⁹<http://www.poss.pku.edu.cn/semanticposs.html>.

¹⁰<https://github.com/xiaoaoaran/SemanticSTF>.

¹¹<https://github.com/xiaoaoaran/SynLiDAR>.

¹²<https://github.com/saltoricristiano/gipso-sfouda>.

¹³<https://github.com/subake/DAPS3D>.

¹⁴<https://github.com/ldkong1205/Robo3D>.

¹⁵<https://github.com/NVIDIA/MinkowskiEngine>.

¹⁶<https://github.com/valeoai/SLiDR>.

¹⁷<https://github.com/mit-han-lab/spvnas>.

¹⁸<https://github.com/xinge008/Cylinder3D>.

¹⁹<https://github.com/ldkong1205/LaserMix>.

²⁰<https://github.com/CuriousAI/mean-teacher>.

²¹<https://github.com/Lightning-AI/lightning>.

²²<https://github.com/open-mmlab/mmdetection3d>.

290 References

- 291 [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic
292 superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and*
293 *Machine Intelligence*, 34(11):2274–2282, 2012.
- 294 [2] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Juergen
295 Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *IEEE/CVF*
296 *International Conference on Computer Vision*, pages 9297–9307, 2019.
- 297 [3] Maxim Berman, Amal Rannen Triki, and Matthew B Blaschko. The lovász-softmax loss: a tractable
298 surrogate for the optimization of the intersection-over-union measure in neural networks. In *IEEE/CVF*
299 *Conference on Computer Vision and Pattern Recognition*, pages 4413–4421, 2018.
- 300 [4] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix
301 Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather.
302 In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11682–11692, 2020.
- 303 [5] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive
304 learning. *arXiv preprint arXiv:2003.04297*, 2020.
- 305 [6] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski
306 convolutional neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
307 pages 3075–3084, 2019.
- 308 [7] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and
309 Abhinav Valada. Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking.
310 *IEEE Robotics and Automation Letters*, pages 3795–3802, 2022.
- 311 [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition.
312 In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- 313 [9] Peng Jiang, Philip Osteen, Maggie Wigness, and Srikanth Saripallig. Rellis-3d dataset: Data, benchmarks
314 and analysis. In *IEEE International Conference on Robotics and Automation*, 2021.
- 315 [10] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao,
316 Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything.
317 *arXiv preprint arXiv:2304.02643*, 2023.
- 318 [11] Alexey Klokov, Di Un Pak, Aleksandr Khorin, Dmitry Yudin, Leon Kochiev, Vladimir Luchinskiy, and
319 Vitaly Bezuglyj. Daps3d: Domain adaptive projective segmentation of 3d lidar point clouds. *Preprint*,
320 2023.
- 321 [12] Lingdong Kong, Youquan Liu, Xin Li, Runnan Chen, Wenwei Zhang, Jiawei Ren, Liang Pan, Kai Chen,
322 and Ziwei Liu. Robo3d: Towards robust and reliable 3d perception against corruptions. *arXiv preprint*
323 *arXiv:2303.17597*, 2023.
- 324 [13] Lingdong Kong, Jiawei Ren, Liang Pan, and Ziwei Liu. Lasermix for semi-supervised lidar semantic
325 segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21705–21715,
326 2023.
- 327 [14] Yueh-Cheng Liu, Yu-Kai Huang, Hung-Yueh Chiang, Hung-Ting Su, Zhe-Yu Liu, Chin-Tang Chen, Ching-
328 Yu Tseng, and Winston H. Hsu. Learning from 2d: Contrastive pixel-to-point knowledge transfer for 3d
329 pretraining. *arXiv preprint arXiv:2104.0468*, 2021.
- 330 [15] Anas Mahmoud, Jordan SK Hu, Tianshu Kuai, Ali Harakeh, Liam Paull, and Steven L. Waslander. Self-
331 supervised image-to-point distillation via semantically tolerant contrastive loss. In *IEEE/CVF Conference*
332 *on Computer Vision and Pattern Recognition*, 2023.
- 333 [16] Yancheng Pan, Biao Gao, Jilin Mei, Sibogeng, Chengkun Li, and Huijing Zhao. Semanticpos: A point
334 cloud dataset with large quantity of dynamic instances. In *IEEE Intelligent Vehicles Symposium*, pages
335 687–693, 2020.
- 336 [17] Cristiano Saltori, Evgeny Krivosheev, Stéphane Lathuilière, Nicu Sebe, Fabio Galasso, Giuseppe Fiameni,
337 Elisa Ricci, and Fabio Poiesi. Gippo: Geometrically informed propagation for online adaptation in 3d lidar
338 segmentation. In *European Conference on Computer Vision*, pages 567–585, 2022.
- 339 [18] Corentin Sautier, Gilles Puy, Spyros Gidaris, Alexandre Boulch, Andrei Bursuc, and Renaud Marlet.
340 Image-to-lidar self-supervised distillation for autonomous driving data. In *IEEE/CVF Conference on*
341 *Computer Vision and Pattern Recognition*, pages 9891–9901, 2022.

- 342 [19] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James
343 Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao,
344 Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens,
345 Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open
346 dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020.
- 347 [20] Ozan Unal, Dengxin Dai, and Luc Van Gool. Scribble-supervised lidar semantic segmentation. In
348 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2697–2707, 2022.
- 349 [21] Aoran Xiao, Jiaying Huang, Dayan Guan, Fangneng Zhan, and Shijian Lu. Transfer learning from synthetic
350 to real lidar point cloud for semantic segmentation. In *AAAI Conference on Artificial Intelligence*, pages
351 2795–2803, 2022.
- 352 [22] Aoran Xiao, Jiaying Huang, Weihao Xuan, Ruijie Ren, Kangcheng Liu, Dayan Guan, Abdulmoteleb El
353 Saddik, Shijian Lu, and Eric Xing. 3d semantic segmentation in the wild: Learning generalized models for
354 adverse-condition point clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
355 2023.
- 356 [23] Hao Zhang, Feng Li, Xueyan Zou, Shilong Liu, Chunyuan Li, Jianfeng Gao, Jianwei Yang, and Lei Zhang.
357 A simple framework for open-vocabulary segmentation and detection. *arXiv preprint arXiv:2303.08131*,
358 2023.
- 359 [24] Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl,
360 Jianfeng Wang, Lu Yuan, Nanyun Peng, Lijuan Wang, Yong Jae Lee, and Jianfeng Gao. Generalized
361 decoding for pixel, image, and language. In *IEEE/CVF Conference on Computer Vision and Pattern
362 Recognition*, 2023.
- 363 [25] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment
364 everything everywhere all at once. *arXiv preprint arXiv:2304.06718*, 2023.