

ODIN: A Single Model for 2D and 3D Segmentation

Ayush Jain¹, Pushkal Katara¹, Nikolaos Gkanatsios¹, Adam W. Harley², Gabriel Sarch¹,
 Kriti Aggarwal³, Vishrav Chaudhary³, Katerina Fragkiadaki¹

¹Carnegie Mellon University, ²Stanford University, ³Microsoft

{ayushj2, pkatara ngkanats, gsarch, kfragki2}@andrew.cmu.edu

aharley@cs.stanford.edu, {kragga, vchaudhary}@microsoft.com

Abstract

State-of-the-art models on contemporary 3D segmentation benchmarks like ScanNet consume and label dataset-provided 3D point clouds, obtained through post processing of sensed multiview RGB-D images. They are typically trained in-domain, forego large-scale 2D pre-training and outperform alternatives that featurize the posed RGB-D multiview images instead. The gap in performance between methods that consume posed images versus post-processed 3D point clouds has fueled the belief that 2D and 3D perception require distinct model architectures. In this paper, we challenge this view and propose ODIN (Omni-Dimensional INstance segmentation), a model that can segment and label both 2D RGB images and 3D point clouds, using a transformer architecture that alternates between 2D within-view and 3D cross-view information fusion. Our model differentiates 2D and 3D feature operations through the positional encodings of the tokens involved, which capture pixel coordinates for 2D patch tokens and 3D coordinates for 3D feature tokens. ODIN achieves state-of-the-art performance on ScanNet200, Matterport3D and AI2THOR 3D instance segmentation benchmarks, and competitive performance on ScanNet, S3DIS and COCO. It outperforms all previous works by a wide margin when the sensed 3D point cloud is used in place of the point cloud sampled from 3D mesh. When used as the 3D perception engine in an intractable embodied agent architecture, it sets a new state-of-the-art on the TEACH action-from-dialogue benchmark. Our code and checkpoints can be found at the project website <https://odin-seg.github.io>.

1. Introduction

There has been a surge of interest in porting 2D foundational image features to 3D scene understanding [8, 14, 20, 22, 36, 39, 44–46]. Some methods lift pre-trained 2D image features using sensed depth to 3D feature clouds [8, 36, 39, 45]. Others distill 2D backbones to differentiable paramet-

ric 3D models, e.g., NeRFs, by training them per scene to render 2D feature maps of pre-trained backbones [22, 44]. Despite this effort, and despite the ever-growing power of 2D backbones [4, 51], the state-of-the-art on established 3D segmentation benchmarks such as ScanNet [6] and ScanNet200 [40] *still* consists of models that operate directly in 3D, without any 2D pre-training stage [27, 42]. Given the obvious power of 2D pre-training, why is it so difficult to yield improvements in these 3D tasks?

We observe that part of the issue lies in a key implementation detail underlying these 3D benchmark evaluations. Benchmarks like ScanNet do not actually ask methods to use RGB-D images as input, even though this is the sensor data. Instead, these benchmarks first register all RGB-D frames into a single colored point cloud and reconstruct the scene as cleanly as possible, relying on manually tuned stages for bundle adjustment, outlier rejection and meshing, and ask models to label the *output reconstruction*. While it is certainly viable to scan and reconstruct a room before labelling any of the objects inside, this pipeline is perhaps inconsistent with the goals of embodied vision (and typical 2D vision), which involves dealing with actual sensor data and accounting for missing or partial observations. We therefore hypothesize that method rankings will change, and the impact of 2D pre-training will become evident, if we force the 3D models to take posed RGB-D frames as input rather than pre-computed mesh reconstructions. Our revised evaluation setting also opens the door to new methods, which can train and perform inference in either single-view or multi-view settings, with either RGB or RGB-D sensors.

We propose **Omni-Dimensional INstance** segmentation (ODIN)[†], a model for 2D and 3D object segmentation and labelling that can parse single-view RGB images and/or multiview posed RGB-D images. As shown in Fig. 1, ODIN alternates between 2D and 3D stages in its architecture,

[†]The Norse god Odin sacrificed one of his eyes for wisdom, trading one mode of perception for a more important one. Our approach sacrifices perception on post-processed meshes for perception on posed RGB-D images.

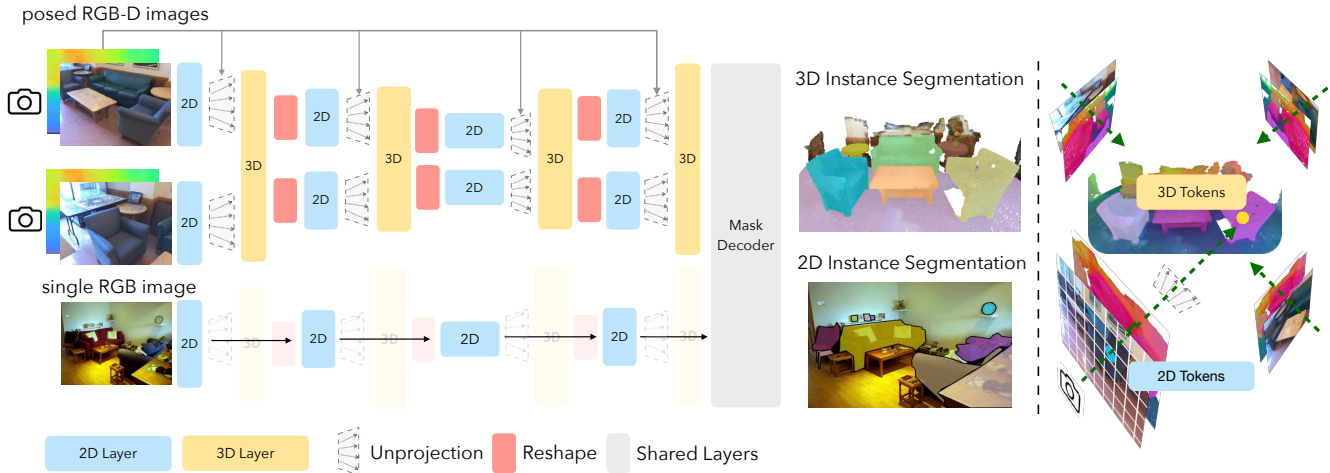


Figure 1. **Omni-Dimensional Instance segmentation (ODIN)** is a model that can parse either a single RGB image or a multiview posed RGB-D sequence into 2D or 3D labelled object segments respectively. **Left:** Given a posed RGB-D sequence as input, ODIN alternates between a within-view 2D fusion and a cross-view 3D fusion. When the input is a single RGB image, the 3D fusion layers are skipped. ODIN shares the majority of its parameters across both RGB and RGB-D inputs, enabling the use of pre-trained 2D backbones. **Right:** At each 2D-to-3D transition, ODIN unprojects 2D feature tokens to their 3D locations using sensed depth and camera intrinsics and extrinsics.

fusing information in 2D within each image view, and in 3D across posed image views. At each 2D-to-3D transition, it unprojects 2D tokens to their 3D locations using the depth maps and camera parameters, and at each 3D-to-2D transition, it projects 3D tokens back to their image locations. Our model differentiates between 2D and 3D features through the positional encodings of the tokens involved, which capture pixel coordinates for 2D patch tokens and 3D coordinates for 3D feature tokens. When dealing with 2D single-view input, our architecture simply skips the 3D layers and makes a forward pass with 2D layers alone.

We test ODIN in 2D and 3D instance segmentation and 3D semantic segmentation on the 2D COCO object segmentation benchmark and the 3D benchmarks of ScanNet [6], ScanNet200 [40], Matterport3D [2], S3DIS [1] and AI2THOR [7, 24]. When compared to methods using pre-computed mesh point cloud as input, our approach performs slightly worse than state-of-the-art on ScanNet and S3DIS, but better on ScanNet200 and Matterport3D. When using real sensor data as input with poses obtained from bundle reconstruction for all methods, our method performs even better, outperforming all prior work by a wide margin, in all datasets. We demonstrate that our model’s ability to jointly train on 3D and 2D datasets results in performance increase on 3D benchmarks, and also yields competitive segmentation accuracy on the 2D COCO benchmark. Our ablations show that interleaving 2D and 3D fusion operations outperforms designs where we first process in 2D and then move to 3D, or simply paint 3D points with 2D features. Stepping toward our broader goal of embodied vision, we also deploy ODIN as the 3D object segmentor of a SOTA embodied

agent model [41] on the simulation benchmark TEACH [35] in the setup with access to RGB-D and pose information from the simulator, and demonstrate that our model sets a new state-of-the-art. We make our code publicly available at <https://odin-seg.github.io>.

2. Related Work

3D Instance Segmentation Early methods in 3D instance segmentation [3, 15, 21, 29, 47, 56] group their semantic segmentation outputs into individual instances. Recently, Mask2Former [4] achieved state-of-the-art in 2D instance segmentation by instantiating *object queries*, each directly predicting an instance segmentation mask by doing dot-product with the feature map of the input image. Inspired by it, Mask3D [42] abandons the grouping strategy of prior 3D models to use the simple decoder head of Mask2Former. MAFT [27] and QueryFormer [33] improve over Mask3D by incorporating better query initialization strategies and/or relative positional embeddings. While this shift to Mask2Former-like architecture brought the 3D instance segmentation architectures closer to their 2D counterparts, the inputs and backbones remain very different: 2D models use pre-trained backbones [16, 32], while 3D methods [42] operate over point clouds and use sparse convolution-based backbones [5], trained from scratch on small-scale 3D datasets. In this work, we propose to directly use RGB-D input and design architectures that can leverage strong 2D backbones to achieve strong performance on 3D benchmarks.

3D Datasets and Benchmarks Most 3D models primarily operate on point clouds, avoiding the use of image-based features partly due to the design of popular benchmarks. These benchmarks generate point clouds by processing raw RGB-D sensor data, involving manual and noisy steps that result in misalignments between the reconstructed point cloud and sensor data. For instance, ScanNet [6] undergoes complex mesh reconstruction steps, including bundle reconstruction, implicit TSDF representation fitting, marching cubes, merging and deleting noisy mesh vertices, and finally manual removal of mesh reconstruction with high misalignments. Misalignments introduced by the mesh reconstruction process can cause methods processing sensor data directly to underperform compared to those trained and tested on provided point clouds. Additionally, some datasets, like HM3D [52] lack access to raw RGB-D data. While mesh reconstruction has its applications, many real-time applications need to directly process sensor data.

2D-based 3D segmentation Unlike instance segmentation literature, several approaches for semantic segmentation like MVPNet [19], BPNNet [17] and DeepViewAgg [39] utilize the sensor point cloud directly instead of the mesh-sampled point cloud. Virtual Multiview Fusion [25] forgoes sensor RGB-D images in favour of rendering RGB-D images from the provided mesh to fight misalignments and low field-of-view in ScanNet images. Similar to our approach, BPNNet and DeepViewAgg integrate 2D-3D information at various feature scales and initialize their 2D streams with pre-trained features. Specifically, they employ separate 2D and 3D U-Nets for processing the respective modalities and fuse features from the two streams through a connection module. Rather than employing distinct streams for featurizing raw data, our architecture instantiates a single unified U-Net which interleaves 2D and 3D layers and can handle both 2D and 3D perception tasks with a single unified architecture. Notably, while these works focus solely on semantic segmentation, our single architecture excels in both semantic and instance segmentation tasks.

Recent advancements in 2D foundation models [23, 38] have spurred efforts to apply them to 3D tasks such as point cloud classification [37, 50, 54], zero-shot 3D semantic segmentation [14, 20, 36] and more recently, zero-shot instance segmentation [45]. Commonly, these methods leverage 2D foundation models to featurize RGB images, project 3D point clouds onto these images, employ occlusion reasoning using depth and integrate features from all views through simple techniques like mean-pooling. Notably, these approaches predominantly focus on semantic segmentation, emphasizing pixel-wise labeling, rather than instance labeling, which necessitates cross-view reasoning to associate the same object instance across multiple views. OpenMask3D [45] is the only method that we are aware of that attempts 3D instance segmentation using 2D foun-

ation models, by training a class-agnostic 3D object segmentor on 3D point clouds and labelling it utilizing CLIP features. Despite their effectiveness in a zero-shot setting, they generally lag behind SOTA 3D supervised methods by 15-20%. Rather than relying on features from foundation models, certain works [10, 12] create 3D pseudo-labels using pre-trained 2D models. Another line of work involves fitting Neural-Radiance Fields (NeRFs), incorporating features from CLIP [22, 46] or per-view instance segmentations from state-of-the-art 2D segmentors [44]. These approaches require expensive per-scene optimization that prohibits testing on all test scenes to compare against SOTA 3D discriminative models. Instead of repurposing 2D foundation models for 3D tasks, Omnivore [13] proposes to build a unified architecture that can handle multiple visual modalities like images, videos and single-view RGB-D image but they only show results for classification tasks. We similarly propose a single unified model capable of performing both single-view 2D and multi-view 3D instance and semantic segmentation tasks while utilizing pre-trained weights for the majority of our architecture.

3. Method

ODIN’s architecture is shown in Fig. 2. It takes either a single RGB image or a set of posed RGB-D images (i.e., RGB images associated with depth maps and camera parameters) and outputs the corresponding 2D or 3D instance segmentation masks and their semantic labels. To achieve this, ODIN alternates between a 2D within-view fusion and a 3D attention-based cross-view fusion, as illustrated in blue blocks and yellow blocks in Fig. 2. A segmentation decoding head predicts instance masks and semantic labels. Notably, ODIN shares the majority of its parameters across both RGB and multiview RGB-D inputs. We detail the components of our architecture below.

Within-view 2D fusion: We start from a 2D backbone, such as ResNet50 [16] or Swin Transformer [32], pre-trained for 2D COCO instance segmentation following Mask2Former [4], a state-of-the-art 2D segmentation model. When only a single RGB image is available, we pass it through the full backbone to obtain 2D features at multiple scales. When a posed RGB-D sequence is available, this 2D processing is interleaved with 3D stages, described next. By interleaving within-view and cross-view contextualization, we are able to utilize the pre-trained features from the 2D backbone while also fusing features across views, making them 3D-consistent.

Cross-view 3D fusion: The goal of cross-view fusion is to make the individual images’ representations consistent across views. As we show in our ablations, cross-view feature consistency is essential for 3D instance segmentation: it enables the segmentation head to realize that a 3D object observed from multiple views is indeed a single instance,

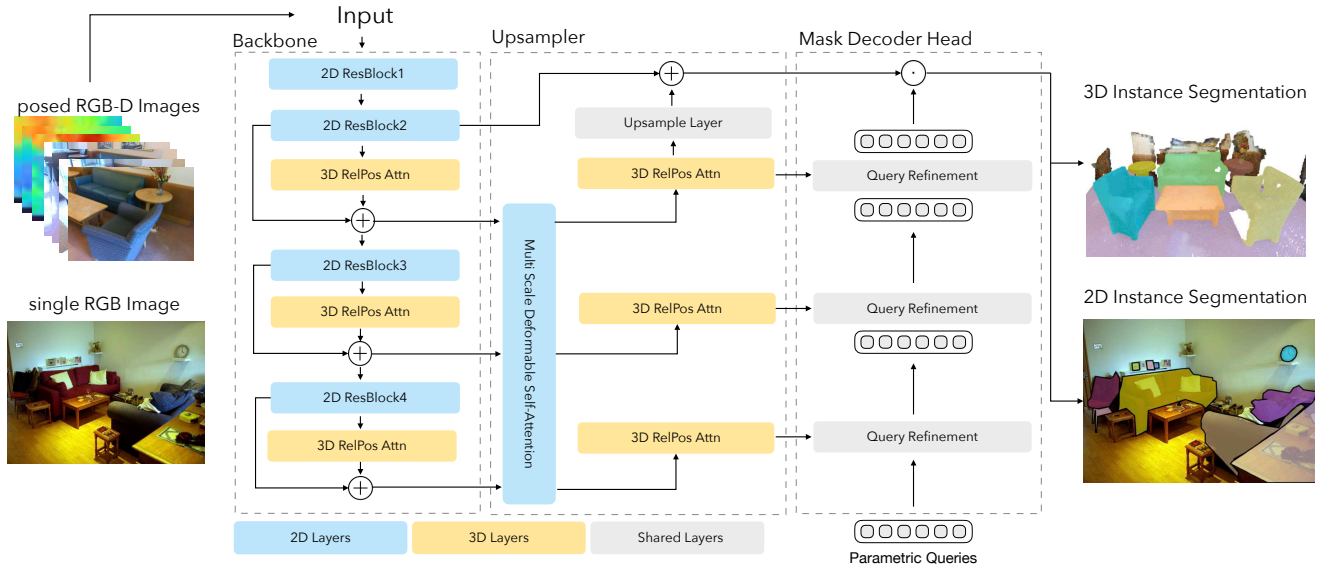


Figure 2. **ODIN Architecture:** The input to our model is either a single RGB image or a multiview RGB-D posed sequence. We feed them to ODIN’s backbone which interleaves 2D within-view fusion layers and 3D cross-view attention layers to extract feature maps of different resolutions (scales). These feature maps exchange information through a multi-scale attention operation. Additional 3D fusion layers are used to improve multiview consistency. Then, a mask decoder head is used to initialize and refine learnable slots that attend to the multi-scale feature maps and predict object segments (masks and semantic classes).

rather than a separate instance in each viewpoint.

1. *2D-to-3D Unprojection:* We unproject each 2D feature map to 3D by lifting each feature vector to a corresponding 3D location, using nearest neighbor depth and known camera intrinsic and extrinsic parameters, using a pinhole camera model. Subsequently, the resulting featurized point cloud undergoes voxelization, where the 3D space is discretized into a volumetric grid. Within each occupied grid cell (voxel), the features and XYZ coordinates are mean-pooled to derive new sets of 3D feature tokens and their respective 3D locations.

2. *3D k -NN Transformer with Relative Positions:* We fuse information across 3D tokens using k -nearest-neighbor attention with relative 3D positional embeddings. This is similar to Point Transformers [49, 55], but we simply use vanilla cross-attention instead of the vector attention proposed in those works. Specifically, in our approach, each 3D token attends to its k nearest neighbors. The positional embeddings in this operation are relative to the query token’s location. We achieve this by encoding the distance vector between a token and its neighbour with an MLP. The positional embedding for the query is simply encoding of the 0 vector. We therefore have

$$\text{query}_{\text{pos}} = \text{MLP}(0); \quad (1)$$

$$\text{key}_{\text{pos}} = \text{MLP}(p_i - p_j), \quad (2)$$

where p_i represents the 3D tokens, shaped $N \times 1 \times 3$, and p_j represents the k nearest neighbors of each p_i , shaped

$N \times k \times 3$. In this way, the attention operation is invariant to the absolute coordinates of the 3D tokens and only depends on their relative spatial arrangements. While each 3D token always attends to the same k neighbors, its effective receptive field grows across layers, as the neighbors’ features get updated when they perform their own attention [11].

3. *3D-to-2D Projection:* After contextualizing the tokens in 3D, we project the features back to their original 2D locations. We first copy the feature of each voxel to all points within that voxel. We then reshape these points back into multiview 2D feature maps, so that they may be processed by the next 2D module. The features vectors are unchanged in this transition; the difference lies in their interpretation and shape. In 2D the features are shaped $V \times H \times W \times F$, representing a feature map for each viewpoint, and in 3D they are shaped $N \times F$, representing a unified feature cloud, where $N = V \cdot H \cdot W$.

Cross-scale fusion and upsampling: After multiple single-view and cross-view stages, we have access to multiple features maps per image, at different resolutions. We merge these with the help of deformable 2D attention, akin to Mask2Former [4], operating on the three lowest-resolution scales (1/32, 1/16, 1/8). When we have 3D input, we apply an additional 3D fusion layer at each scale after the deformable attention, to restore the 3D consistency. Finally, we use a simple upsampling layer on the 1/8 resolution feature map to bring it to 1/4 resolution and add with a skip connection to the 1/4 feature map from the backbone.

Sensor depth to mesh point cloud feature transfer: For 3D benchmarks like ScanNet [6] and ScanNet200 [40], the objective is to label a point cloud derived from a mesh rather than the depth map from the sensor. Hence, on those benchmarks, instead of upsampling the 1/8 resolution feature map to 1/4, we trilinearly interpolate features from the 1/8 resolution feature map to the provided point cloud sampled from the mesh. This means: for each vertex in the mesh, we trilinearly interpolate from our computed 3D features to obtain interpolated features. We additionally similarly interpolate from the unprojected 1/4 resolution feature map in the backbone, for an additive skip connection.

Shared 2D-3D segmentation mask decoder: Our segmentation decoder is a Transformer, similar to Mask2Former’s decoder head, which takes as input upsampled 2D or 3D feature maps and outputs corresponding 2D or 3D segmentation masks and their semantic classes. Specifically, we instantiate a set of N learnable object queries responsible for decoding individual instances. These queries are iteratively refined by a *Query Refinement* block, which consists of cross-attention to the upsampled features, followed by a self-attention between the queries. Except for the positional embeddings, all attention and query weights are shared between 2D and 3D. We use Fourier positional encodings in 2D, while in 3D we encode the XYZ coordinates of the 3D tokens with an MLP. The refined queries are used to predict instance masks and semantic classes. For mask prediction, the queries do a token-wise dot product with the highest-resolution upsampled features. For semantic class prediction, we use an MLP over the queries, mapping them to class logits. We refer readers to Mask2Former [4] for further details.

Open vocabulary class decoder: Drawing inspiration from prior open-vocabulary detection methods [18, 28, 58], we introduce an alternative classification head capable of handling an arbitrary number of semantic classes. This modification is essential for joint training on multiple datasets. Similar to BUTD-DETR [18] and GLIP [28], we supply the model with a *detection prompt* formed by concatenating object categories into a sentence (e.g., “*Chair. Table. Sofa.*”) and encode it using RoBERTa [31]. In the query-refinement block, queries additionally attend to these text tokens before attending to the upsampled feature maps. For semantic class prediction, we first perform a dot-product operation between queries and language tokens, generating one logit per token in the detection prompt. The logits corresponding to prompt tokens for a specific object class are then averaged to derive per-class logits. This can handle multi-word noun phrases such as “*shower curtain*”, where we average the logits corresponding to “*shower*” and “*curtain*”. The segmentation masks are predicted by a pixel-/point-wise dot-product, in the same fashion as described earlier.

Implementation details: We initialize our model with pre-trained weights from Mask2Former [4] trained on COCO [30]. Subsequently, we train all parameters end-to-end, including both pre-trained and new parameters from 3D fusion layers. During training in 3D scenes, our model processes a sequence of N consecutive frames, usually comprising 25 frames. At test time, we input all images in the scene to our model, with an average of 90 images per scene in ScanNet. We use vanilla closed-vocabulary decoding head for all experiments except when training jointly on 2D-3D datasets. There we use our open vocabulary class decoder that lets us handle different label spaces in these datasets. During training, we employ open vocabulary mask decoding for joint 2D and 3D datasets and vanilla closed-vocabulary decoding otherwise. Training continues until convergence on 2 NVIDIA A100s with 40 GB VRAM, with an effective batch size of 6 in 3D and 16 in 2D. For joint training on 2D and 3D datasets, we alternate sampling 2D and 3D batches with batch sizes of 3 and 8 per GPU, respectively. We adopt Mask2Former’s strategy, using Hungarian matching for matching queries to ground truth instances and supervision losses. While our model is only trained for instance segmentation, it can perform semantic segmentation for free at test time like Mask2Former. We refer to Mask2Former [4] for more details.

4. Experiments

4.1. Evaluation on 3D benchmarks

Datasets: First, we test our model on 3D instance and semantic segmentation in the ScanNet [6] and ScanNet200 [40] benchmarks. The goal of these benchmarks is to label the point cloud extracted from the 3D mesh of a scene reconstructed from raw sensor data. ScanNet evaluates on 20 common semantic classes, while ScanNet200 uses 200 classes, which is more representative of the long-tailed object distribution encountered in the real world. We report results on the official validation split of these datasets here and on the official test split in the supplementary.

Evaluation metrics: We follow the standard evaluation metrics, namely mean Average Precision (mAP) for instance segmentation and mean Intersection over Union (mIoU) for semantic segmentation.

Baselines: In *instance segmentation*, our main baseline is the SOTA 3D method Mask3D [42]. For a thorough comparison, we train both Mask3D and our model with sensor RGB-D point cloud input and evaluate them on the benchmark-provided mesh-sampled point clouds. We also compare with the following recent and concurrent works: PBNNet [56], QueryFormer [33] and MAFT [27]. QueryFormer and MAFT explore query initialization and refinement in a Mask3D-like architecture and thus have complementary advantages to ours. Unlike ODIN, these methods

Table 1. **Evaluation on 3D Benchmarks** (§ = trained by us using official codebase).

| (a) ScanNet Instance Segmentation Task. | | | | | (b) ScanNet Semantic Segmentation Task. | | | | |
|--|--------------------------|-------|-------------|-------------|--|------------------------------|-----------------------------|--------------|------|
| | | Model | mAP | mAP50 | mAP25 | | | Model | mIoU |
| Sensor RGBD Point Cloud | Mask3D [§] [42] | | 43.9 | 60.0 | 69.9 | Sensor RGBD Point Cloud | MVPNet [19] | 68.3 | |
| | ODIN-ResNet50 (Ours) | | 47.8 | 69.8 | 83.6 | | BPNNet [17] | 69.7 | |
| | ODIN-Swin-B (Ours) | | 50.0 | 71.0 | 83.6 | | DeepViewAgg [39] | 71.0 | |
| Mesh Sampled Point Cloud | SoftGroup [47] | | 46.0 | 67.6 | 78.9 | Rendered RGBD Point Cloud | ODIN-ResNet50 (Ours) | 73.3 | |
| | PBNNet [56] | | 54.3 | 70.5 | 78.9 | | ODIN-Swin-B (Ours) | 77.8 | |
| | Mask3D [42] | | 55.2 | 73.7 | 83.5 | | VMVF [25] | 76.4 | |
| | QueryFormer [33] | | 56.5 | 74.2 | 83.3 | | Point Transformer v2 [49] | 75.4 | |
| | MAFT [27] | | 58.4 | 75.9 | - | | Stratified Transformer [26] | 74.3 | |
| Zero-Shot | | | | | | Mesh Sampled Point Cloud | OctFormer [48] | 75.7 | |
| | | | | | | | Swin3D-L [53] | 76.7 | |
| | | | | | | | OpenScene [36] | 54.2 | |
| | | | | | | | | | |
| (c) ScanNet200 Instance Segmentation Task. | | | | | (d) ScanNet200 Semantic Segmentation Task. | | | | |
| | | Model | mAP | mAP50 | mAP25 | | | Model | mIoU |
| Sensor RGBD Point Cloud | Mask3D [42] [§] | | 15.5 | 21.4 | 24.3 | Sensor RGBD Point Cloud | ODIN-ResNet50 (Ours) | 35.8 | |
| | ODIN-ResNet50 (Ours) | | 25.6 | 36.9 | 43.8 | | ODIN-Swin-B (Ours) | 40.5 | |
| | ODIN-Swin-B (Ours) | | 31.5 | 45.3 | 53.1 | | Mesh Sampled Point Cloud | LGround [40] | 28.9 |
| Mesh Sampled Point Cloud | Mask3D [42] | | 27.4 | 37.0 | 42.3 | Mesh Sampled Point Cloud | CeCo [57] | 32.0 | |
| | QueryFormer [33] | | 28.1 | 37.1 | 43.4 | | Octformer [48] | 32.6 | |
| | MAFT [27] | | 29.2 | 38.2 | 43.3 | | | | |
| Zero-Shot | OpenMask3D [45] | | 15.4 | 19.9 | 23.1 | | | | |

directly process 3D point clouds and initialize their weights from scratch. As motivated before, utilizing RGB-D input directly has several advantages, including avoiding costly mesh building processes, achieving closer integration of 2D and 3D perception, and leveraging pre-trained features and abundant 2D data.

In *semantic segmentation*, we compare with MVPNet [19], BPNNet [17] and state-of-the-art DeepViewAgg [39] which directly operate on sensor RGB or RGB-D images and point clouds. We also compare with VMVF [25] which operates over rendered RGB-D images from the provided mesh, with heuristics for camera view sampling to avoid occlusions, ensures balanced scene coverage, and employs a wider field-of-view, though we note their code is not publicly available. Similar to ODIN, all of these methods utilize 2D pre-trained backbones. We also compare with Point-Transformer v2 [49], Stratified Transformer [26], OctFormer [48] and Swin3D-L [53] which process the mesh-sampled point cloud directly, without using any 2D pre-training. On the ScanNet200 semantic segmentation benchmark, we compare with SOTA OctFormer [48] and with CeCo [57], a method specially designed to fight class-imbalance in ScanNet200. These methods directly process the point cloud and do not use 2D

image pre-trained weights. We also compare with LGround [40] which uses 2D CLIP pre-training. We also compare with zero-shot 2D foundation model-based 3D models of OpenScene [36] and OpenMask3D [45]. This comparison is unfair since they are not supervised within-domain, but we include them for completeness. The results are presented in Tab. 1. We draw the following conclusions:

Performance drops with sensor point cloud as input (Tab. 1a): Mask3D’s performance drops from 55.2% mAP with mesh point cloud input to 43.9% mAP with sensor point cloud input. This is consistent with prior works [25, 39] in 3D semantic segmentation on ScanNet, which attributes the drop to misalignments caused by noise in camera poses, depth variations and post-processing steps.

ODIN outperforms SOTA 3D methods with sensor point cloud input and underperforms them when baselines use mesh-sampled point clouds (Tab. 1a): Our model significantly outperforms SOTA Mask3D model with sensor point cloud input and achieves comparable performance to methods using mesh-sampled point cloud input on the mAP25 metric while far behind on mAP metric, due to misalignments between the 3D mesh and the sensor point cloud.

ODIN sets a new SOTA in semantic segmentation on ScanNet (Tab. 1b) outperforming all methods on all setups

Table 2. AI2THOR Semantic and Instance Segmentation.

| Model | mAP | mAP50 | mAP25 | mIoU |
|----------------------|-------------|-------------|-------------|-------------|
| Mask3D [42] | 60.6 | 70.8 | 76.6 | - |
| ODIN-ResNet50 (Ours) | 63.8 | 73.8 | 80.2 | 71.5 |
| ODIN-Swin-B (Ours) | 64.3 | 73.7 | 78.6 | 71.4 |

Table 3. Embodied Instruction Following. SR = success rate. GC = goal condition success rate.

| | TEACH | | | | ALFRED | | | |
|----------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | Unseen | | Seen | | Unseen | | Seen | |
| | SR | GC | SR | GC | SR | GC | SR | GC |
| FILM [34] | - | - | - | - | 30.7 | 42.9 | 26.6 | 38.2 |
| HELPER [41] | 15.8 | 14.5 | 11.6 | 19.4 | 37.4 | 55.0 | 26.8 | 41.2 |
| HELPER + ODIN (OURS) | 18.6 | 18.6 | 13.8 | 26.6 | 47.7 | 61.6 | 33.5 | 47.1 |

including the models trained on the sensor, rendered and mesh sampled point clouds.

ODIN sets a new instance segmentation SOTA on the long-tailed ScanNet200 dataset (Tab. 1c) outperforming SOTA 3D models on all setups including the models trained on mesh-sampled point cloud, especially by a large margin in mAP25 metric, while exclusively utilizing sensor RGB-D data. This highlights the contribution of 2D features, particularly in detecting a long tail of class distribution where limited 3D data is available. We show more detailed results with performance on the head, common and tail classes in the appendix.

ODIN sets a new semantic segmentation SOTA on ScanNet200 (Tab. 1d), outperforming SOTA semantic segmentation models that use mesh point clouds.

4.2. Evaluation on multiview RGB-D in simulation

Using the AI2THOR [24] simulation environment with procedural homes from ProcThor [7], we collected RGB-D data for 1500 scenes (1200 training, 300 test) of similar size as ScanNet (more details in appendix). We train and evaluate our model and SOTA Mask3D [42] on the unprojected RGB-D images. As shown in Tab. 2, our model outperforms Mask3D by 3.7% mAP, showing strong performance in a directly comparable RGB-D setup. It suggests that current real-world benchmarks may restrain models that featurizes RGB-D sensor point clouds due to misalignments. We hope this encourages the community to also focus on directly collecting, labeling, and benchmarking RGB-D sensor data.

4.3. Embodied Instruction Following

We apply ODIN in the embodied setups of TEACH [35] and ALFRED [43] where agents have access to RGB, depth and camera poses and need to infer and execute task and

Table 4. Joint Training on Sensor RGB-D point cloud from ScanNet and 2D RGB images from COCO.

| | ScanNet | | | COCO |
|------------------------|-------------|-------------|-------------|-------------|
| | mAP | mAP50 | mAP25 | mAP |
| Mask3D [42] | 43.9 | 60.0 | 69.9 | X |
| Mask2Former [4] | X | X | X | 43.7 |
| ODIN (trained in 2D) | X | X | X | 43.6 |
| ODIN (trained in 3D) | 47.8 | 69.8 | 83.6 | X |
| ODIN (trained jointly) | 49.1 | 70.1 | 83.1 | 41.2 |

action plans from dialogue segments and instructions, respectively. These agents operate in dynamic home environments and cannot afford expensive mesh building steps. Detecting objects well is critical for task success in both cases. Prior SOTA methods [35, 41] run per-view 2D instance segmentation models [4, 9] and link the detected instances using simple temporal reasoning regarding spatial and appearance proximity. Instead, ODIN processes its last N egocentric views and segments objects instances directly in 3D. We equip HELPER [41], a state-of-the-art embodied model, with ODIN as its 3D object detection engine. We evaluate using Task Success Rate (SR) which checks if the entire task is executed successfully, and Goal Conditioned Success Rate (GC) which checks the proportion of satisfied subgoals across all episodes [35, 43]. We perform evaluation on "valid-seen" (houses similar to the training set) and "valid-unseen" (dissimilar houses) splits. In Tab. 3, we observe that HELPER with ODIN as its 3D object detector significantly outperforms HELPER that uses the original 2D detection plus linking perception pipeline.

4.4. Ablations and Variants

We conduct our ablation experiments on the ScanNet dataset in Tab. 4 and Tab. 5. Our conclusions are:

Joint 2D-3D training helps 3D perception We compare joint training of ODIN on sensor RGB-D point clouds from ScanNet and 2D RGB images from COCO to variants trained independently on 2D and 3D data, all initialized from pre-trained COCO weights. Since there are different classes in ScanNet and COCO, we use our open-vocabulary semantic class-decoding head instead of the vanilla closed-vocabulary head. Results in Tab. 4 show that joint training yields a 1.3% absolute improvement in 3D, and causes a similar drop in 2D. This experiment indicates that a single architecture can perform well on both 2D and 3D tasks, thus indicating that we may not need to design vastly different architectures in either domain. However, the drop in 2D performance indicates a potential for further improvements in the architecture design to retain the performance in the 2D domain. Nevertheless, this experiment highlights the benefits of joint training with 2D datasets for 3D segmentation in ODIN. Note that we do not jointly train on

Table 5. Ablations on ScanNet Dataset.

| (a) Cross-View Contextualization. | | | (b) Effect of Pre-Trained Features. | | | (c) Effect of Freezing Backbone. | | | | |
|-----------------------------------|-------------|-------------|-------------------------------------|-------------|-------------|----------------------------------|-------------|-------------|-------------|-------------|
| Model | mAP | mIoU | Model | mAP | mIoU | Model | ResNet50 | | Swin-B | |
| | | | | | | | mAP | mIoU | mAP | mIoU |
| ODIN (Ours) | 47.8 | 73.3 | ODIN (Ours) | 47.8 | 73.3 | ODIN (Ours) | 47.8 | 73.3 | 50.0 | 77.8 |
| No 3D Fusion | 39.3 | 73.2 | Only pre-trained back- | 42.3 | 72.9 | With frozen | 46.7 | 74.3 | 46.2 | 75.9 |
| No interleaving | 41.7 | 73.6 | bone | | | backbone | | | | |
| | | | No pre-trained features | 41.5 | 68.6 | | | | | |

2D and 3D datasets in any of our other experiments due to computational constraints.

Cross-View fusion is crucial for instance segmentation but not for semantic segmentation (Tab. 5a): removing 3D cross-view fusion layers results in an 8.5% mAP drop for instance segmentation, without any significant effect in semantic segmentation. Popular 2D-based 3D open vocabulary works [20, 36] without strong cross-view fusion only focus on semantic segmentation and thus could not uncover this issue. Row-3 shows a 6.1% mAP drop when cross-view 3D fusion happens after all within-view 2D layers instead of interleaving the within-view and cross-view fusion.

2D pre-trained weight initialization helps (Tab. 5b): initializing only the image backbone with pre-trained weights, instead of all layers (except the 3D fusion layers), results in a 5.5% mAP drop (row-2). Starting the entire model from scratch leads to a larger drop of 6.3% mAP (row-3). This underscores the importance of sharing as many parameters as possible with the 2D models to leverage the maximum possible 2D pre-trained weights.

Stronger 2D backbones helps (Tab. 5c): using Swin-B over ResNet50 leads to significant performance gains, suggesting that ODIN can directly benefit from advancements in 2D computer vision.

Finetuning everything including the pre-trained parameters helps (Tab. 5c): ResNet50’s and Swin’s performance increases substantially when we fine-tune all parameters. Intuitively, unfreezing the backbone allows 2D layers to adapt to cross-view fused features better. Thus, we keep our backbone unfrozen in all experiments.

Supplying 2D features directly to 3D models does not help: Concatenating 2D features with XYZ+RGB as input to Mask3D yields 53.8% mAP performance, comparable to 53.3% of the baseline model with only XYZ+RGB as input.

4.5. Limitations

Our experiments reveal the following limitations for ODIN: Firstly, like other top-performing 3D models, it depends on accurate depth and camera poses. Inaccurate depth or camera poses cause a sharp decrease in performance (similar to other 3D models, like Mask3D). In our future work, we

[†]We do not use the expensive DB-SCAN post-processing of Mask3D, and hence it gets 53.3% mAP instead of 55.2% as reported by their paper

aim to explore unifying depth and camera pose estimation with semantic scene parsing, thus making 3D models more resilient to noise. Secondly, in this paper, we limited our scope to exploring the design of a unified architecture without scaling up 3D learning by training on diverse 2D and 3D datasets jointly. We aim to explore this in future to achieve strong generalization to in-the-wild scenarios, akin to the current foundational 2D perception systems. Our results suggest a competition between 2D and 3D segmentation performance when training ODIN jointly on both modalities. Exploring ways to make 2D and 3D training more synergistic is a promising avenue for future work.

5. Conclusion

We presented ODIN, a model for 2D and 3D instance segmentation that can parse 2D images and 3D point clouds alike. ODIN represents both 2D images and 3D feature clouds as a set of tokens that differ in their positional encodings which represent 2D pixel coordinates for 2D tokens and 3D XYZ coordinates for 3D tokens. Our model alternates between within-image featurization and cross-view featurization. It achieves SOTA performance in ScanNet200 and AI2THOR instance segmentation benchmarks, outperforms all methods operating on sensor point clouds and achieves competent performance to methods operating over mesh-sampled pointcloud. Our experiments show that ODIN outperforms alternative models that simply augment 3D point cloud models with 2D image features as well as ablative versions of our model that do not alternate between 2D and 3D information fusion, do not co-train across 2D and 3D and do not pre-train the 2D backbone.

6. Acknowledgements

The authors express gratitude to Wen-Hsuan Chu, Mihir Prabhudesai, and Alexander Swerdlow for their valuable feedback on the early draft of this work. Special thanks to Tsung-Wei Ke for insightful discussions throughout the project. We thank the Microsoft Turing Team for providing us with GPU resources during the initial development phase of this project. This work is supported by Sony AI, DARPA Machine Common Sense, an Amazon faculty award, and an NSF CAREER award.

References

- [1] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. [2](#)
- [2] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. [2](#)
- [3] Shaoyu Chen, Jiemin Fang, Qian Zhang, Wenyu Liu, and Xinggang Wang. Hierarchical aggregation for 3d instance segmentation. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15447–15456, 2021. [2](#)
- [4] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [5] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3075–3084, 2019. [2](#)
- [6] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. [1](#), [2](#), [3](#), [5](#)
- [7] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Winson Han, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Proctor: Large-scale embodied ai using procedural generation. *Advances in Neural Information Processing Systems*, 35:5982–5994, 2022. [2](#), [7](#)
- [8] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. Pla: Language-driven open-vocabulary 3d scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7010–7019, 2023. [1](#)
- [9] Bin Dong, Fangao Zeng, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. Solq: Segmenting objects by learning queries. *Advances in Neural Information Processing Systems*, 34: 21898–21909, 2021. [7](#)
- [10] Shichao Dong, Fayao Liu, and Guosheng Lin. Leveraging large-scale pretrained vision foundation models for label-efficient 3d point cloud segmentation. *arXiv preprint arXiv:2311.01989*, 2023. [3](#)
- [11] Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Graph neural networks with learnable structural and positional representations. *arXiv preprint arXiv:2110.07875*, 2021. [4](#)
- [12] Kyle Genova, Xiaoqi Yin, Abhijit Kundu, Caroline Pantofaru, Forrester Cole, Avneesh Sud, Brian Brewington, Brian Shucker, and Thomas Funkhouser. Learning 3d semantic segmentation with only 2d image supervision. In *2021 International Conference on 3D Vision (3DV)*, pages 361–372. IEEE, 2021. [3](#)
- [13] Rohit Girdhar, Mannat Singh, Nikhila Ravi, Laurens van der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A single model for many visual modalities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16102–16112, 2022. [3](#)
- [14] Huy Ha and Shuran Song. Semantic abstraction: Open-world 3d scene understanding from 2d vision-language models. In *6th Annual Conference on Robot Learning*, 2022. [1](#), [3](#)
- [15] Lei Han, Tian Zheng, Lan Xu, and Lu Fang. Occuseg: Occupancy-aware 3d instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2940–2949, 2020. [2](#)
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#), [3](#)
- [17] Wenbo Hu, Hengshuang Zhao, Li Jiang, Jiaya Jia, and Tien-Tsin Wong. Bidirectional projection network for cross dimension scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14373–14382, 2021. [3](#), [6](#)
- [18] Ayush Jain, Nikolaos Gkanatsios, Ishita Mediratta, and Katerina Fragkiadaki. Bottom up top down detection transformers for language grounding in images and point clouds. In *European Conference on Computer Vision*, pages 417–433. Springer, 2022. [5](#)
- [19] Maximilian Jaritz, Jiayuan Gu, and Hao Su. Multi-view pointnet for 3d scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. [3](#), [6](#)
- [20] Krishna Murthy Jatavallabhula, Alihusein Kuwajerwala, Qiao Gu, Mohd Omama, Tao Chen, Shuang Li, Ganesh Iyer, Soroush Saryazdi, Nikhil Keetha, Ayush Tewari, et al. Conceptfusion: Open-set multimodal 3d mapping. *arXiv preprint arXiv:2302.07241*, 2023. [1](#), [3](#), [8](#)
- [21] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Pointgroup: Dual-set point grouping for 3d instance segmentation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4866–4875, 2020. [2](#)
- [22] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lurf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. [1](#), [3](#)
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. [3](#)
- [24] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017. [2](#), [7](#)
- [25] Abhijit Kundu, Xiaoqi Yin, Alireza Fathi, David Ross, Brian Brewington, Thomas Funkhouser, and Caroline Pantofaru.

- Virtual multi-view fusion for 3d semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, pages 518–535. Springer, 2020. 3, 6
- [26] Xin Lai, Jianhui Liu, Li Jiang, Liwei Wang, Hengshuang Zhao, Shu Liu, Xiaojuan Qi, and Jiaya Jia. Stratified transformer for 3d point cloud segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8500–8509, 2022. 6
- [27] Xin Lai, Yuhui Yuan, Ruihang Chu, Yukang Chen, Han Hu, and Jiaya Jia. Mask-attention-free transformer for 3d instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3693–3703, 2023. 1, 2, 5, 6
- [28] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10965–10975, 2022. 5
- [29] Zhihao Liang, Zhihao Li, Songcen Xu, Mingkui Tan, and Kui Jia. Instance segmentation in 3d scenes using semantic superpoint tree networks. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2763–2772, 2021. 2
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 5
- [31] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 5
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 2, 3
- [33] Jiahao Lu, Jiacheng Deng, Chuxin Wang, Jianfeng He, and Tianzhu Zhang. Query refinement transformer for 3d instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18516–18526, 2023. 2, 5, 6
- [34] So Yeon Min, Devendra Singh Chaplot, Pradeep Kumar Ravikumar, Yonatan Bisk, and Ruslan Salakhutdinov. Film: Following instructions in language with modular methods. In *International Conference on Learning Representations*, 2021. 7
- [35] Aishwarya Padmakumar, Jesse Thomason, Ayush Srivastava, Patrick Lange, Anjali Narayan-Chen, Spandana Gella, Robinson Piramuthu, Gokhan Tur, and Dilek Hakkani-Tur. Teach: Task-driven embodied agents that chat. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2017–2025, 2022. 2, 7
- [36] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 815–824, 2023. 1, 3, 6, 8
- [37] Guocheng Qian, Xingdi Zhang, Abdullah Hamdi, and Bernard Ghanem. Pix4point: Image pretrained transformers for 3d point cloud understanding. 2022. 3
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3
- [39] Damien Robert, Bruno Vallet, and Loic Landrieu. Learning multi-view aggregation in the wild for large-scale 3d semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5575–5584, 2022. 1, 3, 6
- [40] David Rozenberszki, Or Litany, and Angela Dai. Language-grounded indoor 3d semantic segmentation in the wild. In *European Conference on Computer Vision*, pages 125–141. Springer, 2022. 1, 2, 5, 6
- [41] Gabriel Sarch, Yue Wu, Michael J Tarr, and Katerina Fragkiadaki. Open-ended instructable embodied agents with memory-augmented large language models. *arXiv preprint arXiv:2310.15127*, 2023. 2, 7
- [42] Jonas Schult, Francis Engelmann, Alexander Hermans, Or Litany, Siyu Tang, and Bastian Leibe. Mask3d: Mask transformer for 3d semantic instance segmentation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8216–8223. IEEE, 2023. 1, 2, 5, 6, 7
- [43] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10740–10749, 2020. 7
- [44] Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Norman Müller, Matthias Nießner, Angela Dai, and Peter Kotschieder. Panoptic lifting for 3d scene understanding with neural fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9043–9052, 2023. 1, 3
- [45] Ayça Takmaz, Elisabetta Fedele, Robert W Sumner, Marc Pollefeys, Federico Tombari, and Francis Engelmann. Open-mask3d: Open-vocabulary 3d instance segmentation. *arXiv preprint arXiv:2306.13631*, 2023. 1, 3, 6
- [46] Nikolaos Tsagkas, Oisín Mac Aodha, and Chris Xiaoquan Lu. VI-fields: Towards language-grounded neural implicit spatial representations. *arXiv preprint arXiv:2305.12427*, 2023. 1, 3
- [47] Thang Vu, Kookhoi Kim, Tung M Luu, Thanh Nguyen, and Chang D Yoo. Softgroup for 3d instance segmentation on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2708–2717, 2022. 2, 6

- [48] Peng-Shuai Wang. Octformer: Octree-based transformers for 3d point clouds. *arXiv preprint arXiv:2305.03045*, 2023. 6
- [49] Xiaoyang Wu, Yixing Lao, Li Jiang, Xihui Liu, and Hengshuang Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. *Advances in Neural Information Processing Systems*, 35:33330–33342, 2022. 4, 6
- [50] Chenfeng Xu, Shijia Yang, Tomer Galanti, Bichen Wu, Xiangyu Yue, Bohan Zhai, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with 2d image pretrained models. *arXiv preprint arXiv:2106.04180*, 2021. 3
- [51] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. 1
- [52] Karmesh Yadav, Ram Ramrakhya, Santhosh Kumar Ramakrishnan, Theo Gervet, John Turner, Aaron Gokaslan, Noah Maestre, Angel Xuan Chang, Dhruv Batra, Manolis Savva, et al. Habitat-matterport 3d semantics dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4927–4936, 2023. 3
- [53] Yu-Qi Yang, Yu-Xiao Guo, Jian-Yu Xiong, Yang Liu, Hao Pan, Peng-Shuai Wang, Xin Tong, and Baining Guo. Swin3d: A pretrained transformer backbone for 3d indoor scene understanding. *arXiv preprint arXiv:2304.06906*, 2023. 6
- [54] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. 3
- [55] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021. 4
- [56] Weiguang Zhao, Yuyao Yan, Chaolong Yang, Jianan Ye, Xi Yang, and Kaizhu Huang. Divide and conquer: 3d point cloud instance segmentation with point-wise binarization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 562–571, 2023. 2, 5, 6
- [57] Zhisheng Zhong, Jiequan Cui, Yibo Yang, Xiaoyang Wu, Xiojuan Qi, Xiangyu Zhang, and Jiaya Jia. Understanding imbalanced semantic segmentation through neural collapse. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19550–19560, 2023. 6
- [58] Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl, Jianfeng Wang, Lu Yuan, et al. Generalized decoding for pixel, image, and language. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15116–15127, 2023. 5