

PanopticRecon: Leverage Open-vocabulary Instance Segmentation for Zero-shot Panoptic Reconstruction

Xuan Yu¹, Yili Liu¹, Chenrui Han¹, Sitong Mao², Shunbo Zhou², Rong Xiong¹, Yiyi Liao¹, Yue Wang¹

Abstract—Panoptic reconstruction is a challenging task in 3D scene understanding. However, most existing methods heavily rely on pre-trained semantic segmentation models and known 3D object bounding boxes for 3D panoptic segmentation, which is not available for in-the-wild scenes. In this paper, we propose a novel zero-shot panoptic reconstruction method from RGB-D images of scenes. For zero-shot segmentation, we leverage open-vocabulary instance segmentation, but it has to face *partial labeling* and *instance association* challenges. We tackle both challenges by propagating partial labels with the aid of dense generalized features and building a 3D instance graph for associating 2D instance IDs. Specifically, we exploit partial labels to learn a classifier for generalized semantic features to provide complete labels for scenes with dense distilled features. Moreover, we formulate instance association as a 3D instance graph segmentation problem, allowing us to fully utilize the scene geometry prior and all 2D instance masks to infer global unique pseudo 3D instance ID. Our method outperforms state-of-the-art methods on the indoor dataset ScanNet V2 and the outdoor dataset KITTI-360, demonstrating the effectiveness of our graph segmentation method and reconstruction network.

I. INTRODUCTION

Panoptic reconstruction that builds the 3D geometry of the scene labeled with semantics and instance, is an important task in 3D scene understanding. The representation is ideal for localization and planning of autonomous robots, as well as interaction with human users, thus drawing the attention of the community.

Pioneering methods aim at semantic reconstruction [1] design rules for integration of scans labeled by semantic segmentation [2]. One weakness of these methods is the separate optimization of geometry and semantics. Recent years witnesses great success in implicit neural representation [3], [4], which allows for end-to-end semantic reconstruction [5], [6], [7]. Empowered by foundation visual language model (VLM) [8], [9], the open-vocabulary semantic cues are further employed for zero-shot semantic reconstruction to eliminate human annotation and segmentation network training for unseen class [10], [11].

PNF [12] and panoptic NeRF [13] pioneeringly extend the semantic reconstruction method to panoptic reconstruction by replacing the semantic labels with 2D panoptic labels and 3D instance bounding box, and further assuming the known

This work was supported by the National Nature Science Foundation of China under Grant 62373322, Zhejiang Provincial Natural Science Foundation of China under Grant No. LD24F030001.

¹Xuan Yu, Yili Liu, Chenrui Han, Rong Xiong, Yiyi Liao and Yue Wang are with Zhejiang University, Hangzhou, Zhejiang, China. Yue Wang is the corresponding author wangyue@iipc.zju.edu.cn

²Sitong Mao and Shunbo Zhou are with Huawei Cloud Computing Technologies Co., Ltd., Shenzhen, China.

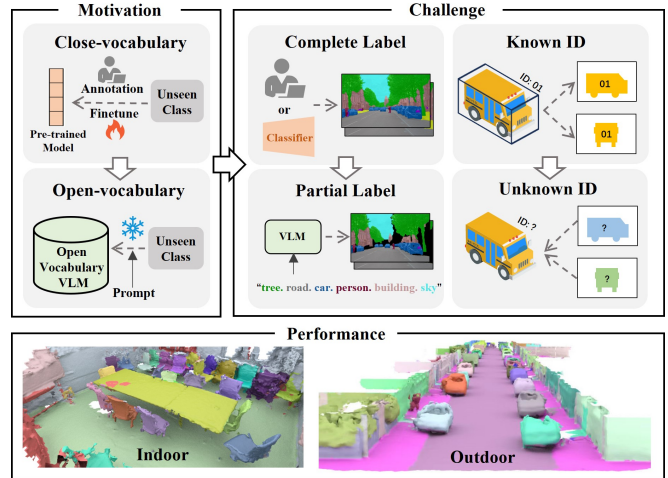


Fig. 1. Zero-shot panoptic reconstruction by leveraging open-vocabulary instance segmentation faces two challenges: 1) 2D semantic labels provided by text prompt based VLMs are not complete. 2) No object-level instance 3D pseudo ID makes 2D instance ID inconsistent. We supplement blank pixel labels with distilled DINOv2 features and establish a graph to infer 3D instance pseudo IDs.

data association between unique 3D instance ID and per-frame 2D instance ID. For zero-shot panoptic reconstruction, since there is no open-vocabulary panoptic segmentation VLM, PVLFF [14] proposes to distill features of two VLMs for open-vocabulary semantic segmentation (LSeg [8]) and instance segmentation (SAM [15]). Since SAM is not able to achieve object-level segmentation, PVLFF struggles to segment semantically aligned object-level instances.

In contrast to SAM which may provide over-segmented masks, Grounded-SAM [16] is a promising alternative to provide open-vocabulary semantic and instance labels of the desired granularity given a few text queries. But there are still two remaining challenges as shown in Fig 1: First, language-conditioned instance segmentation VLM yields *partial labeling* [17] that only pixels corresponding to the class prompt have a valid prediction. In this way, the labeling result can be regarded as semantic segmentation with additional classes of unknown, which may cause lower accuracy for known classes. Moreover, only pixels with labels belonging to known classes can be supervised, thus insufficient for labeling the full 3D space. Second, *instance association* across multiple frames is challenging. In [18], the association is derived by tracking [19], which fails for large frame in-between gap and long-term re-identification. In [14], [20], contrastive learning is employed to build the embedding space of the instance, avoiding the association. When clustering the embedding space for instance segmentation, it is

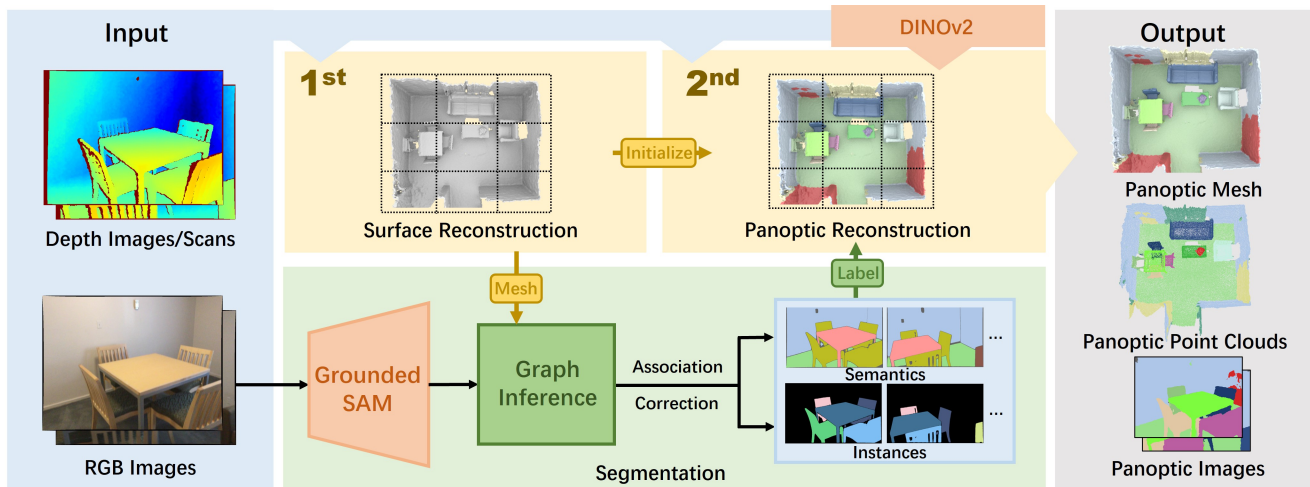


Fig. 2. PanopticRecon consists of a reconstruction task and a segmentation task. The first step of the reconstruction task realizes the implicit surface reconstruction through RGB-D observations to provide the scene geometry for the segmentation task. Secondly, the segmentation task builds a graph from the normal of mesh, and infers 3D pseudo IDs to associate the 2D instance IDs by instance mask of Grounded SAM. In addition, 3D instance ID corrects some of the erroneous semantic labels. Then, the second reconstruction step realizes 2D-3D labeling supervised by consistent semantic and instance labels, and finally obtains the panoptic mesh, point cloud, and novel view images of the scene.

hard to tune the coefficients e.g. instance count. Panoptic Lifting [21] is proposed to infer the instance association using linear assignment, which may lose the global uniqueness due to the initial randomization sensitivity.

In this work, we propose a novel zero-shot panoptic reconstruction method, named PanopticRecon, by combining open vocabulary instance segmentation, Grounded-SAM [16], and neural implicit surface reconstruction [7]. We tackle both challenges mentioned above, by 1) propagating partial labels to the full scene via a generalizable network and 2) building a 3D instance graph from 2D instance masks to establish globally consistent association. Specifically, to address the challenge of partial labeling, we propose to distill the *dense* semantic VLM features to the 3D space, and learn a generalizable point-level classifier supervised by the *partial* labels, effectively enabling label propagation. For the second challenge of instance association, we first reconstruct the scene geometry without labeling, and then leverage the reconstruction and 2D instance-level segmentation to build pseudo instance GTs in the 3D space to guarantee the ID uniqueness globally. The pseudo GT is acquired by graph inference, which is stated as the edge cutting to maximize the 3D-2D instance mask consistency. Assembling the improved prediction and association into the neural implicit representation, PanopticRecon arrives at zero-shot panoptic surface reconstruction without human annotation, which achieves superior accuracy in geometry and semantics as demonstrated in experiments. In summary, the contributions involve:

- We propose a novel neural implicit representation-based method to leverage Grounded-SAM for zero-shot 3D dense panoptic reconstruction.
- We propagate partial labels with the aid of DINOv2, generate instance association by instance graph inference and correct semantic labels based on unique instance IDs, to address the limitations of Grounded-SAM results.

- Both indoor and outdoor sequence experiments to evaluate the effectiveness of outcomes including panoptic segmentation and mesh geometric quality.

II. RELATED WORKS

A. Semantic Neural Fields

Initially, NeRF [3] modeled only the radiance and geometry of 3D scenes, and some subsequent studies have used neural implicit representations to extend semantic information to the tasks of 3D scene modeling. Semantic NeRF [5] first explored encoding semantics into NeRF by adding a separate branch to predict semantic labels, fusing noisy 2D semantic segmentation into a consistent volumetric model. NeSF [6] then predicts semantic fields by feeding density fields to a 3D semantic segmentation model. To achieve open-vocabulary semantic segmentation, NIVLFF [11] combines NeRF and features from a foundation model, distilling visual features into an implicit representation model and achieving semantic segmentation of open-vocabulary scenes by mapping class texts embedding to distilled features.

B. Panoptic Neural Fields

Kirillov et al [22] defined the 2D panoptic segmentation task and its metrics for the first time. In recent years, some works [12], [13], [20], [21] have adopted NeRF to solve the 3D scene panoptic segmentation problem. PNF [12] first proposed a method to obtain panoptic radiation fields of dynamic 3D scenes from images. The method takes the 3D bounding boxes of dynamic instances with known motion trajectories as the sampling space for each instance, with each instance having its own MLP. Panoptic NeRF [13] also uses the 3D bounding box to assist panoptic segmentation which focuses more on the refinement of the 3D polygon labels, requiring more manual annotation in the early stage. To automatically obtain the 3D bounding boxes of objects to assist in supervising instances, Instance NeRF [23] proposes to determine the 3D bounding box of the object through the

density reconstructed in advance, but it is only suitable for relatively simple scenes. Since 3D supervision is not easy to obtain, Panoptic Lifting [21] proposes to add linear assignments in the optimization process to align the instance IDs of multiple viewpoints, while losing the global uniqueness. Contrastive Lift [20] designs a contrastive formulation to enhance 2D segmentation to 3D by feature clustering.

All the above methods apply a 2D pre-trained 2D segmentation model, which can only handle close-vocabulary scenes. For open-vocabulary scenes, PVLFF [14] distills 2D features from two different VLMs for the semantic task and the instance task respectively, and realizes the mask label alignment in 3D by making the features in the SAM [15] mask similar. However, two different foundation models lead to inconsistent labeling of semantics and instances, and the characteristics of SAM itself cause PVLFF to be unable to segment object-level instances accurately.

III. METHOD

A. Overview

Given a set of posed multi-view images and depth of a scene, the proposed zero-shot panoptic reconstruction can construct an implicit representation that encodes the SDF, color, semantics, and instance, with the aid of Grounded-SAM. The pipeline is fully automatic as shown in Fig. 2, of which the only information provided by the human user is the class text. We divide the pipeline into two main components, reconstruction and segmentation. In segmentation, we employ Grounded-SAM to label the instance with given classes and associate 2D instances by 3D mesh based graph inference, resulting in zero-shot multi-view images with semantic and consistent instance labels for panoptic reconstruction. In reconstruction, we have two stages of reconstruction. In the 1st stage, we build an initial 3D pure geometric map to help instance association. In the 2nd stage, we build a panoptic map by jointly learning SDF, color, semantics, and instances, using labeled images and DINOv2 as supervision, yielding the final panoptic segmented mesh and novel view images.

B. Zero-shot Panoptic Label Preparation

We first identify a set of interested classes as text prompts for Grounded-SAM to label images with instances belonging to these classes. When pixels are labeled with more than one instance, we assign them the label with the highest score. We group the instances with the same class into semantic masks and divide all semantic classes into *thing* and *stuff* categories. Note that the per-frame instances are not associated, the semantic masks may be incomplete and Grounded-SAM may miss or misclassify instances in a single frame, which are the main issues to address for panoptic reconstruction.

C. Instance Association and Label Correction

Graph Construction: We begin by building an initial 3D pure geometric map for instance association. The technical detail for the implicit surface reconstruction is shown in III-E. Given the 3D map and the 2D instance labeled images, we




Fig. 3. The points in the upper graph in (a) are the nodes (superpoints) of the graph. The color region corresponding to each node is the surface. We determine the nodes in an instance mask of a frame based on the overlap between the instance mask provided by Grounded SAM and the mask projected by the surface, and vote for the edges between the selected nodes. Similarly, we reduce the votes for the edges between nodes corresponding to masks of different instances in that frame. The edges with non-positive votes are finally cut and the nodes connected form an instance as shown in (b). Once the 3D instance pseudo IDs are obtained, we associate 2D instance IDs while correcting incorrect semantic labels.

state the association problem as a graph inference problem. First, we apply [24] to compute the similarity between mesh vertices based on their normal directions, upon which the map is clustered into surfaces to reduce complexity. The center point of each surface is called a superpoint, which are modeled the nodes \mathcal{V} of the scene graph $G = (\mathcal{V}, \mathcal{E})$. Initialize that edges \mathcal{E} exist between all nodes and that each edge has zero vote count. The final vote counts determine whether or not to cut off the edges.

For an image I , we define the set of nodes with valid projection masks as $\mathcal{V}_I \subset \mathcal{V}$. Then we define the overlap U_{ij} as the percentage of the intersection area between the projection mask of the surface corresponding to the node $V_i \in \mathcal{V}_I$, and the 2D instance mask M_j in the image I . For the set of nodes $\mathcal{V}_I^j \subset \mathcal{V}_I$ with non-zero U_{ij} , we define the center node of M_j :

$$V_c^j = \arg \max_{V_i \in \mathcal{V}_I^j} U_{ij} \quad (1)$$

For remaining nodes $V_k \in \mathcal{V}_I^j$, if the node has U_{kj} greater than a threshold, we vote the edge between V_k and V_c^j . For nodes in $\mathcal{V}_I - \mathcal{V}_I^j$, we deduct a vote of their edges to \mathcal{V}_I^j . By iterating all images, we derive a resultant graph, in which we reserve only edges with positive vote counts.

Association and Correction: As shown in Fig 3, on the resultant graph, we apply the graph clustering and regard the clusters as 3D instances. By assigning scene level globally unique IDs to 3D instances, we project them onto all 2D images and align them with the 2D instance by evaluating the IoUs with the projection masks, so that the 2D instance in each image can be associated with a unique 3D ID. Since 3D instances satisfy multi-view consistency, we determine their instance class by voting from labels of all associated 2D instance classes and furthermore correct the 2D labels to refine the semantic segmentation.

D. Label Propagation

Given the image with partial semantic labeling corrected after association, we further propagate the labels to the whole




Fig. 4. The flow of the segmentation of the previous work is shown as (a) and (b), and ours is shown as (c).

3D scene. By integrating the semantic images in the 3D volume, a 3D point not labeled in one semantic image, may be compensated by other images, thus the 3D scene labeling is expected to be improved. The remaining problem is to propagate the labels to points that are unknown in all images.

DINOv2 Aided Semantics: The propagation is achieved by utilizing the image cues. Specifically, given an RGB image I , its corresponding DINOv2 feature is denoted as V . We reduce the feature dimension of V by PCA to 64-dimensional compressed features V_{64} . Using these features as the supervision, we distill them to the intermediate layers of the semantic branch. As shown in Fig. 5, concatenating the distilled feature with the 3D point position encoding, the semantic labeling of this point is actually stated as a per-point classification problem. The classifier is MLP, and the input is the feature co-supervised by DINOv2 and semantic labels. Since DINOv2 feature is believed to be semantic meaningful, the generalization to unknown points is expected.

Discussion on Classifier: In contrast, existing open-vocabulary semantic segmentation methods [10], [14], [11] predict the labels by correlating the class language feature with distilled VLM features. As shown in Fig. 4, in the perspective of per-point classification, these methods employ a linear classifier for labeling. Considering the nonlinear 2D semantic decoder is usually required after the original 2D image VLM feature [8], [9], such linear classification of the 3D point may be too simple for an accurate result even with distilled VLM features.

E. Neural Implicit Representation

With the corrected labels, instance association, and the DINOv2 aided semantic label propagation, we derive the panoptic reconstruction of the scene by learning an implicit neural representation. As shown in Fig. 5, the network of PanopticRecon consists of hierarchical hash encoding [25] and multitask branches. We design independent feature volumes for geometry, color, semantics, and instances and optimize the features jointly, where the task branches interact with each other to complement the accuracy.

Feature Volumes: As shown in Fig. 5, our representation consists of 4 feature volumes including geometry, color, semantics and instance, denoted as Ψ_c , $c \in \{G, C, S, I\}$. Given a 3D point, the feature is yielded by querying the corresponding volume. All these features are learnable by the supervision of multiple tasks.

Volume Rendering: The representation is supervised at the ray level, thus we introduce the volume rendering to generate the prediction of the pixel corresponding to the ray. Based on the camera pose, we define the origin and direction of a ray (o, d) passing through a pixel x in image, depth, semantic or instance segmentation image. Along the ray, we

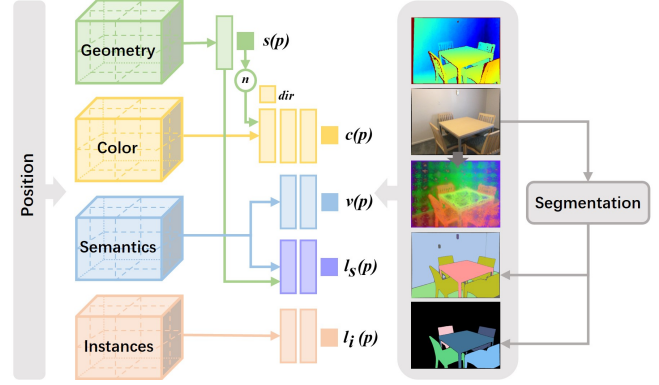


Fig. 5. The network architecture which is described in detail in Sec.E.

sample N successive 3D spatial points $\{p_i\}$ as $p_i = o + \rho_i d$, where ρ_i is the distance. Based on the feature volumes and the downstream network processing as shown in Fig. 5, each point is assigned with a prediction $f(p_i)$, which can be SDF s , color c , semantic probability l_s , distilled DINOv2 feature v , instance probability l_i , and depth d . We then have the volume rendering as:

$$u_f(x) = \sum_{i=1}^N T_i \alpha_i f(p_i) \quad (2)$$

where $T_i = \prod_{m=1}^{i-1} (1 - \alpha_m)$, α_m is the discrete opacity value defined under the S-density function assumption [4] as:

$$\alpha_i = \max \left(\frac{\Phi(s_i) - \Phi(s_{i+1})}{\Phi(s_i)}, 0 \right) \quad (3)$$

where s is SDF of the point, $\Phi(s)$ is Sigmoid function $\Phi(s) = (1 + e^{-\xi s})^{-1}$ with a temperature coefficient ξ .

Network Architecture: Given a 3D point p_i , the geometry feature volume maps p_i to a feature vector $\Psi_G(p_i)$. A small MLP h is applied to yield the SDF value:

$$s(p_i) = h_s(\Psi_G(p_i)) \quad (4)$$

Its normal is derived as:

$$n(p_i) = \frac{\partial s}{\partial p_i} \quad (5)$$

Then the color is derived by querying the color volume feature $\Psi_C(p_i)$, the direction encoding $\lambda(p_i)$, and the normal $n(p_i)$, which are passed through a small MLP h_c as:

$$c(p_i) = h_c(\Psi_C(p_i), \lambda(p_i), n(p_i)) \quad (6)$$

Together with the semantic label probability $l_s(p_i)$, distilled DINOv2 $v(p_i)$ is also generated by querying the semantics volume feature $\Psi_S(p_i)$ as well as geometric volume $\Psi_G(p_i)$, which are passed through a small MLP h_s as:

$$l_s(p_i), v(p_i) = h_s(\Psi_S(p_i), \Psi_G(p_i)) \quad (7)$$

Finally, the instance label probability $l_i(p_i)$ is generated by querying the instance volume feature $\Psi_I(p_i)$, which is passed through a small MLP h_i as:

$$l_i(p_i) = h_i(\Psi_I(p_i)) \quad (8)$$

TABLE I
PANOPTIC SEGMENTATION QUALITY USING DIFFERENT METHODS ON INDOOR AND OUTDOOR DATASETS

Method	ScanNet(V2)							KITTI-360						
	PQ ^s ↑	SQ↑	RQ↑	mIoU↑	mAcc↑	mCov↑	mW-Cov↑	PQ ^s ↑	SQ↑	RQ↑	mIoU↑	mAcc↑	mCov↑	mW-Cov↑
<i>Close-vocabulary:</i>														
GT + Panoptic NeRF	52.12	59.50	71.35	52.65	61.28	61.79	72.04	42.61	51.38	53.01	46.90	59.73	57.91	83.50
<i>Open-vocabulary:</i>														
GS + Panoptic Lifting	47.96	52.11	64.79	61.79	71.76	46.61	64.77	20.86	20.86	27.27	41.10	49.29	6.63	14.64
PVLFF	25.14	36.08	34.55	51.84	58.97	44.94	48.15	27.46	32.61	40.25	40.94	46.20	12.22	21.34
Ours	62.39	64.25	74.69	70.19	84.79	65.67	77.87	35.95	46.43	45.54	43.08	62.65	36.80	57.08

Losses: With all point-level predictions, we can generate pixel-level predictions by (2). We set up geometric loss, photometric loss, semantic loss, and instance loss to supervise the panoptic reconstruction.

The geometric loss consists of three parts:

$$L_{SDF}(p_i) = \begin{cases} |s(p_i) - b(p_i)| & |b| \leq \tau \\ \max(0, e^{-\beta s(p_i)} - 1, s(p_i) - b(p_i)) & o.w. \end{cases} \quad (9)$$

where τ is a threshold to truncate SDF, $b(p_i) = D(x) - \rho_i$ is the distance between the p_i and the observed depth of pixel x along the ray as an approximated SDF. In addition, eikonal loss is used as regularization:

$$L_{eik}(p_i) = \|1 - |n(p_i)|\|^2 \quad (10)$$

Depth is also supervised at the pixel level for depth generated by volume rendering:

$$L_{Depth}(x) = |D(x) - u_d(x)| \quad (11)$$

The photometric loss is built between the rendered color and the ground-truth color:

$$L_{Color}(x) = \|I(x) - u_c(x)\|^2 \quad (12)$$

The semantic loss and instance loss share the same form. Each ray yields a semantic and instance probability distribution which are supervised by cross-entropy of u_{l_s, l_i} with respect to instance and semantic image labels I_s and I_i as pseudo ground truth Γ_r :

$$L_{S/I}(x) = - \sum_{k \in K_{s,i}} I_{s,i}(x) \log u_{l_{s,i}}(x) \quad (13)$$

where K_s is the total number of semantic classes and K_i instances. In addition, the DINOv2 loss for label propagation is built between rendering of distilled DINOv2 and the compressed features V_{64} as:

$$L_{DN}(x) = \|V_{64}(x) - u_v(x)\|^2 \quad (14)$$

In a nutshell, our total loss L is:

$$L = L_{Depth} + L_{SDF} + L_{eik} + L_{Color} + L_{S/I} + L_{DN} \quad (15)$$

Two-stage Reconstruction: The neural representation learning is conducted for two rounds. At the first time, no semantic and instance supervision is applied since there is no instance association and the corrected labels i.e. no L_S , L_{dino} and L_I . But this first-round reconstruction helps the instance association and label correction as shown above. After that, given all supervision available, we have the second-round reconstruction, which finally yields the panoptic reconstruction of the scene.

IV. EXPERIMENTAL RESULTS

We validate the effectiveness of our method and compare with others in both indoor scenes and outdoor scenes to demonstrate the panoptic reconstruction performance. We also conduct an ablation study on the network components.

A. Setup

Dataset: Our study involves experiments on an indoor datasets: Scannet V2 [26] and an outdoor dataset: KITTI-360 [27]. ScanNet is an RGBD dataset that contains a large number of real indoor scenes with perspective-rich images captured by iPad Air2 and reconstructed geometry[28] as ground truth, which makes it suitable for evaluation tests. We use 3 scenes for our evaluation. KITTI-360 is a large-scale driving dataset with 2D/3D manually annotated ground truth and multimodal observations captured using a Velodyne HDL-64 LiDAR and 4 cameras. We select a part of seq00 with a length of 200m for evaluation. We use depth observations and one perspective camera in our experiments.

Metrics: We evaluate the panoptic reconstruction quality using scene-level panoptic quality (PQ^s) proposed by[21], which merges the segments belonging to the same instance identifier in all frames of a scene and compute the panoptic quality of the merged segments for evaluation. We also evaluated the $mIoU$, $mAcc$, $mCov$ and $mWCov$ for 2D semantic segmentation and instance segmentation, respectively, to show the effectiveness of different branches.

Implementation: We set the text prompt of Grounded-SAM by one or more related words per class. Since some classes are hard to discriminate even by human, we follow [21] to merge some classes into one class, such as *sofa* and *armchair*, *ground* and *sidewalk*, *vegetation* and *terrain*. For network architecture, the decoders for the geometric, semantic, and instance are all one-layer MLP of 128 hidden units, and the decoders for the color and DINOv2 branches are all 3-layer MLPs of 128 hidden units. All experiments run on a single A100 GPU.

Baselines: We adopt the Panoptic NeRF [13] as close-vocabulary model aided method, Grounded-SAM (GS) + Panoptic Lifting [16], [13], and PVLFF [14] methods as open-vocabulary model aided methods. The three state-of-the-art methods are capable of panoptic reconstruction using RGBD inputs. Regarding Panoptic NeRF, we provide the ground-truth 3D instance bounding box (GT), making it a reference method. The original Panoptic Lifting employs the close-vocabulary panoptic segmentation method, we replace it with results of Grounded-SAM, making it a comparative open-vocabulary model aided method.



Fig. 6. Comparison of the quality of semantic segmentation and instance segmentation of different methods on ScanNet.

B. Comparative Study

Semantic Segmentation: We first compare the semantic segmentation. Panoptic Nerf uses close-vocabulary model PSPNet [29] to provide 2D semantic images. For KITTI-360, we use the segmentation model pre-trained on CityScape [30] as that in [13]. For ScanNet, as there is no reference model, we use the segmentation model pre-trained on ADE20k [31] that is also collected in indoor environment. As shown in Tab. I, our semantic segmentation metrics are comparable to those of Panoptic Nerf on KITTI-360, and superior to those on ScanNet, reflecting the limitation of close-vocabulary semantic segmentation performance on novel dataset that even collected from similar scene. Compared with other open-vocabulary segmentation based methods, our method shows better performance, validating the effectiveness of DINOv2 to propagate the partial labels to the whole image.

Instance Segmentation: We then compare the instance segmentation in Tab. I. As Panoptic Nerf employs GT 3D bounding boxes, its performance is unsurprisingly far better than the others on KITTI-360. In ScanNet, the bounding boxes of different instances overlap with each other more

frequently than in KITTI-360, degenerating the instance segmentation performance of Panoptic Nerf. The error in semantic segmentation further affects the extraction of instance masks. In contrast, thanks to the explicit graph inference to associate the instances and correct the class labels, our method achieves better performance. Compared with other open-vocabulary methods, our method also outperforms. Note that the non-uniqueness causes inferior performance of Panoptic Lifting, especially on KITTI-360, which has a large number of instances. For the PVLFF, it is hard to correctly cluster instances with hand-crafted coefficients and rules.

Panoptic Segmentation: We finally compare the panoptic segmentation in Tab. I. On KITTI-360, Panoptic Nerf is naturally the best method with the aid of both ground truth bbox and semantic segmentation model trained on a similar dataset. Our method is slightly inferior due to the instance segmentation. For ScanNet, with the aid of open-vocabulary model and proposed techniques, our method performs best, which shows its zero-shot capability when close-vocabulary model are constrained by the generalization gap.

Compared with other open-vocabulary model aided meth-

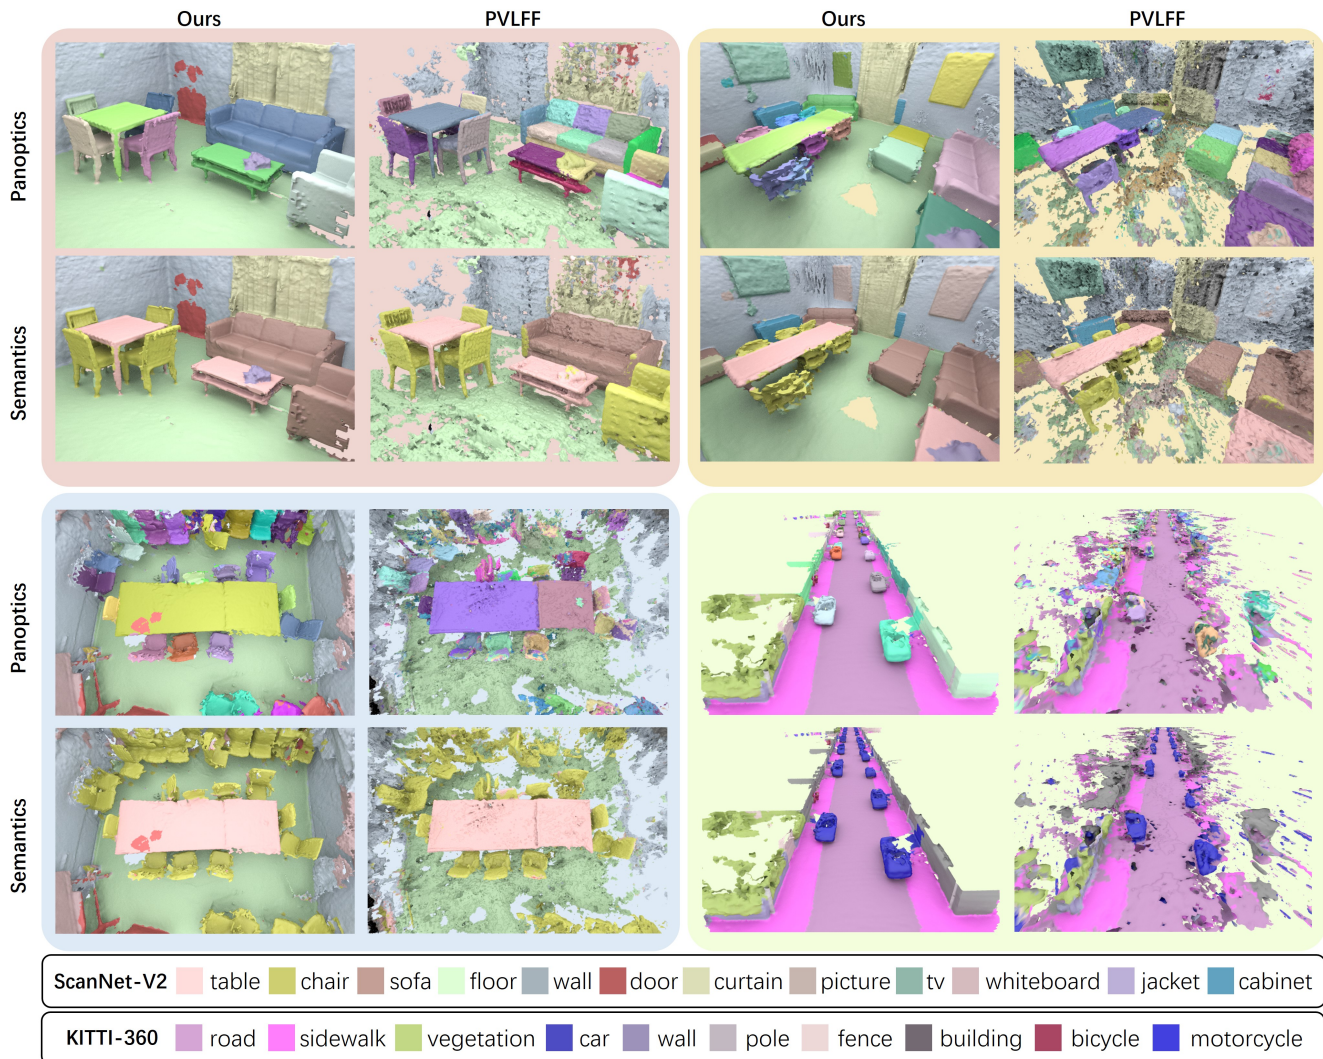


Fig. 7. Comparison of the quality of panoptic reconstruction and semantic reconstruction on ScanNet and KITTI-360 by different methods.

TABLE II
MAPPING QUALITY USING DIFFERENT METHODS ON SCANNET

Method	Comp (cm)↓	Acc (cm)↓	C-L1 (cm)↓
PVLFF	4.55	14.04	9.30
Ours	1.84	1.75	1.80

ods, the benefits of our method are discussed in the semantic and instance segmentation above. Addressing instance association and partial labeling overcomes the limitations of an open-vocabulary segmentation model, leading to improved performance in zero-shot panoptic surface reconstruction.

Mesh Reconstruction: To evaluate the reconstruction accuracy, we compare the surface accuracy with PVLFF which also employs depth supervision. As shown in Tab. II, our method demonstrates the best performance, which is explained by the SDF formulation, depth observations, as well as better labeling. As several cases are shown in Fig. 7, our method can recover better surface details.

C. Ablation Study

We then present the ablation study for components to show its contribution to the whole pipeline, including label

TABLE III
ABLATION TO THE SEMANTIC SEGMENTATION QUALITY ON SCANNET

DINOv2	Correction	mIoU↑	mAcc↑
✗	✗	66.10	79.25
✗	✓	67.24	80.53
✓	✗	68.43	83.08
✓	✓	70.19	84.79

propagation, instance association, and label correction.

Effect of Label Propagation: To evaluate the label propagation, we compare our network with and with the aid of DINOv2. As shown in Tab. III, with the label propagation, the partial labeling commonly occurring in open-vocabulary segmentation is relieved.

Effect of Instance Association: To evaluate the instance association, we replace the graph inference with two existing options: label tracking [18] and linear assignment [21]. To track the label across frames, we employ DAVA [19]. The results are shown in Tab. IV, which tells the advantage of an explicit graph inference on the global scene: uniqueness and multi-view consistency, compared with linear assignment and local tracking. Note that the linear assignment performs worse than Panoptic Lifting due to the lower randomness of

TABLE IV

ABLATION OF INSTANCE ASSOCIATION METHODS TO THE QUALITY OF
INSTANCE SEGMENTATION ON SCANNET

Method	mCov \uparrow	mW-Cov \uparrow
Tracking (DAVA)	55.22	59.93
Linear Assignment	31.81	50.93
Graph Segmentation (Ours)	65.67	77.87

the initial ID distribution when using a feature volume than a large MLP.

Effect of Label Correction: To evaluate the contribution of label correction, we compare with the method using labels without correction. Since prediction error is unavoidable in open-vocabulary methods, an instance may be assigned with different classes across frames, which divides the 3D instance into different labels, resulting in weaker performance.

V. CONCLUSION

In this paper, we propose to endow the neural surface reconstruction with the open-vocabulary segmentation method Grounded-SAM, enabling the zero-shot panoptic reconstruction. As a cost, we have to address the partial labeling and instance association. Accordingly, we propose graph inference, label correction and label propagation, and build a neural implicit representation based panoptic reconstruction pipeline. The results demonstrate the effectiveness of PanopticRecon over existing zero-shot reconstruction methods and close-vocabulary prediction aided reconstruction methods.

REFERENCES

[1] X. Chen, A. Milioto, E. Palazzolo, P. Giguère, J. Behley, and C. Stachniss, “SuMa++: Efficient LiDAR-based Semantic SLAM,” in *Proceedings of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2019, pp. 4530–4537.

[2] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girshick, “Masked-attention mask transformer for universal image segmentation,” in *CVPR*, 2022.

[3] B. Mildenhall, P. P. Srinivasan, M. Tanck, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, 2021.

[4] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang, “Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction,” *arXiv preprint arXiv:2106.10689*, 2021.

[5] S. Zhi, T. Laidlow, S. Leutenegger, and A. J. Davison, “In-place scene labelling and understanding with implicit scene representation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

[6] S. Vora*, N. Radwan*, K. Greff, H. Meyer, K. Genova, M. S. M. Sajjadi, E. Pot, A. Tagliasacchi, and D. Duckworth, “Nesf: Neural semantic fields for generalizable semantic segmentation of 3d scenes,” *Transactions on Machine Learning Research*, 2022.

[7] X. Yu, Y. Liu, S. Mao, S. Zhou, R. Xiong, Y. Liao, and Y. Wang, “Nf-atlas: Multi-volume neural feature fields for large scale lidar mapping,” *IEEE Robotics and Automation Letters*, vol. 8, no. 9, 2023.

[8] B. Li, K. Q. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, “Language-driven semantic segmentation,” in *International Conference on Learning Representations*, 2022.

[9] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, R. Howes, P.-Y. Huang, H. Xu, V. Sharma, S.-W. Li, W. Galuba, M. Rabbat, M. Assran, N. Ballas, G. Synnaeve, I. Misra, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, “Dinov2: Learning robust visual features without supervision,” 2023.

[10] S. Peng, K. Genova, C. M. Jiang, A. Tagliasacchi, M. Pollefeys, and T. Funkhouser, “Openscene: 3d scene understanding with open vocabularies,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

[11] K. Blomqvist, F. Milano, J. J. Chung, L. Ott, and R. Siegwart, “Neural implicit vision-language feature fields,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 1313–1318.

[12] A. Kundu, K. Genova, X. Yin, A. Fathi, C. Pantofaru, L. J. Guibas, A. Tagliasacchi, F. Dellaert, and T. Funkhouser, “Panoptic neural fields: A semantic object-aware neural scene representation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 12 871–12 881.

[13] X. Fu, S. Zhang, T. Chen, Y. Lu, L. Zhu, X. Zhou, A. Geiger, and Y. Liao, “Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation,” in *International Conference on 3D Vision (3DV)*, 2022.

[14] H. Chen, K. Blomqvist, F. Milano, and R. Siegwart, “Panoptic vision-language feature fields,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 9, no. 3, pp. 2144–2151, 2024.

[15] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment anything,” *arXiv:2304.02643*, 2023.

[16] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, Z. Zeng, H. Zhang, F. Li, J. Yang, H. Li, Q. Jiang, and L. Zhang, “Grounded sam: Assembling open-world models for diverse visual tasks,” 2024.

[17] C. Liu, K. Wang, J. Shi, Z. Qiao, and S. Shen, “Fm-fusion: Instance-aware semantic mapping boosted by vision-language foundation models,” *IEEE Robotics and Automation Letters*, vol. 9, no. 3, 2024.

[18] M. Ye, M. Danelljan, F. Yu, and L. Ke, “Gaussian grouping: Segment and edit anything in 3d scenes,” *arXiv preprint arXiv:2312.00732*, 2023.

[19] H. K. Cheng, S. W. Oh, B. Price, A. Schwing, and J.-Y. Lee, “Tracking anything with decoupled video segmentation,” in *ICCV*, 2023.

[20] Y. Bhalgat, I. Laina, J. F. Henriques, A. Zisserman, and A. Vedaldi, “Contrastive lift: 3d object instance segmentation by slow-fast contrastive fusion,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[21] Y. Siddiqui, L. Porzi, S. R. Bulò, N. Müller, M. Nießner, A. Dai, and P. Kotschieder, “Panoptic lifting for 3d scene understanding with neural fields,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023.

[22] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár, “Panoptic segmentation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 9404–9413.

[23] Y. Liu, B. Hu, J. Huang, Y.-W. Tai, and C.-K. Tang, “Instance neural radiance field,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.

[24] P. F. Felzenszwalb and D. P. Huttenlocher, “Efficient graph-based image segmentation,” *International journal of computer vision*, vol. 59, pp. 167–181, 2004.

[25] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant neural graphics primitives with a multiresolution hash encoding,” *ACM Transactions on Graphics (ToG)*, vol. 41, no. 4, pp. 1–15, 2022.

[26] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, “ScanNet: Richly-annotated 3d reconstructions of indoor scenes,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5828–5839.

[27] Y. Liao, J. Xie, and A. Geiger, “Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 3292–3310, 2022.

[28] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt, “Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration,” *ACM Transactions on Graphics (ToG)*, vol. 36, no. 4, p. 1, 2017.

[29] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, jul 2017, pp. 6230–6239.

[30] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.

[31] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, “Scene parsing through ade20k dataset,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.