

SPNeRF: Open Vocabulary 3D Neural Scene Segmentation with Superpoints

Weiwen Hu¹, Niccolò Parodi^{1,2}, Marcus Zepp¹, Ingo Feldmann¹, Oliver Schreer¹ and Peter Eisert^{1,3}

¹Fraunhofer Heinrich Hertz Institute, Berlin, Germany

²Technische Universität Berlin, Germany

³Humboldt-Universität zu Berlin, Germany

Keywords: Computer Vision, Neural Radiance Field, Semantic Segmentation, Point Cloud, 3D.

Abstract: Open-vocabulary segmentation, powered by large visual-language models like CLIP, has expanded 2D segmentation capabilities beyond fixed classes predefined by the dataset, enabling zero-shot understanding across diverse scenes. Extending these capabilities to 3D segmentation introduces challenges, as CLIP’s image-based embeddings often lack the geometric detail necessary for 3D scene segmentation. Recent methods tend to address this by introducing additional segmentation models or replacing CLIP with variations trained on segmentation data, which lead to redundancy or loss on CLIP’s general language capabilities. To overcome this limitation, we introduce SPNeRF, a NeRF based zero-shot 3D segmentation approach that leverages geometric priors. We integrate geometric primitives derived from the 3D scene into NeRF training to produce primitive-wise CLIP features, avoiding the ambiguity of point-wise features. Additionally, we propose a primitive-based merging mechanism enhanced with affinity scores. Without relying on additional segmentation models, our method further explores CLIP’s capability for 3D segmentation and achieves notable improvements over original LERF.

1 INTRODUCTION

Traditional segmentation models are often limited by their reliance on closed-set class definitions, which restricts their applicability to dynamic real-world environments, where new and diverse objects frequently appear. Open-vocabulary segmentation, powered by large visual-language models (VLMs), such as CLIP (Radford et al., 2021), overcomes this barrier by enabling zero-shot recognition of arbitrary classes based on natural language queries. This adaptability is crucial in applications like autonomous navigation, augmented reality, and robotic perception, where it is impractical to exhaustively label every possible object. CLIP aligns 2D visual and language features within a shared embedding space, enabling image classification/understanding without the need for extensive retraining.

In 2D segmentation, this flexibility has led to the development of powerful models (Luo et al., 2023; Xu et al., 2022). Some methods, like OpenSeg (Ghiasi et al., 2021) and LSeg (Li et al., 2022), leverage CLIP’s embeddings and additional segmentation annotation to perform dense, pixel-wise 2D segmentation. These methods have demonstrated that open-vocabulary segmentation not only outperforms tradi-

tional closed-set models in adaptability but also provides a scalable solution for handling diverse tasks across various domains. However, transitioning from 2D to 3D segmentation introduces unique challenges, as 3D environments require neural models to interpret complex spatial relationships and geometric structures that 2D models do not address.

To tackle these challenges, recent works such as LERF (Kerr et al., 2023) have embedded CLIP features within 3D representations like Neural Radiance Fields (NeRF) (Mildenhall et al., 2020). These methods aim to bridge 2D VLMs with 3D scene understanding by enabling open-vocabulary querying across 3D spaces. However, due to the image-based nature of CLIP embeddings, which often lack the geometric precision required for fine-grained 3D segmentation, methods either struggle with segmentation in complex scenes (Kerr et al., 2023), or integrate additional segmentation models (Engelmann et al., 2024; Takmaz et al., 2023).

To address these limitations, we propose SPNeRF, a NeRF-based approach specifically designed to incorporate geometric priors directly from the 3D scene. Unlike prior methods that rely solely on CLIP’s image-centric features, SPNeRF leverages geometric primitives to enhance segmentation accuracy.

By partitioning the 3D scene into geometric primitives, SPNeRF creates primitive-wise CLIP embeddings that preserve geometric coherence. This enables the model to better align CLIP’s semantic representations with the underlying spatial structure, mitigating the ambiguities often associated with point-wise features.

Furthermore, SPNeRF introduces a merging mechanism for these geometric primitives, incorporating an affinity scoring system to refine segmentation boundaries. This approach allows SPNeRF to capture semantic relationships between superpoints, resulting in a more accurate and consistent segmentation output. While avoiding additional segmentation models or segmentation-specific training data, our SPNeRF provides a zero-shot architecture for 3D segmentation tasks.

The main contributions of SPNeRF are as follows:

- *Geometric primitives for improved 3D segmentation:* We integrate geometric primitives into NeRF for open-set segmentation, introducing a loss function that maintains consistency within primitive-wise CLIP features, ensuring coherent segmentation across 3D scenes;
- *Primitive-based merging with affinity scoring:* SPNeRF employs a merging mechanism that uses affinity scoring to refine segmentation, capturing semantic relationships among primitives and improving boundary precision;
- *Enhanced segmentation without additional models:* By leveraging primitive-based segmentation and affinity refinement, SPNeRF improves segmentation accuracy over LERF without relying on extra segmentation models, preserving open-vocabulary capabilities with a streamlined architecture;

2 RELATED WORK

2.1 2D Vision-Language Models

CLIP (Radford et al., 2021) has fueled the explosive growth of large vision-language models. It consists of an image encoder and a text encoder, each mapping their respective inputs into a shared embedding space. Through contrastive training on large-scale image-caption pairs, the encoders align encoded image and caption features to the same location in the embedding space if the caption accurately describes the image, otherwise the encoders push them away. 2D segmentation methods building on CLIP have extended its potential. Approaches by (Ghiasi et al., 2021; Li

et al., 2022) achieve open vocabulary segmentation by training or fine tuning on datasets with segmentation info. These datasets tend to have limited vocabulary due to expensive annotation cost of segmentation, which leads to reduced open vocabulary capacity as stated in (Sun et al., 2024; Kerr et al., 2023). The works of (Sun et al., 2024; Lan et al., 2024) explore alternative approaches to maximize CLIP’s potential, achieving competitive semantic segmentation results while preserving its general language capabilities.

2.2 Neural Radiance Fields

Neural Radiance Fields (NeRFs) (Mildenhall et al., 2020) represent 3D geometry and appearance with a continuous implicit radiance field, parameterized by a multilayer perceptron (MLP). They also provide a flexible framework for integrating 2D-based information directly into 3D, supporting complex semantic and spatial tasks. Works, such as (Cen et al., 2024), bring class-agnostic segmentation ability from 2D foundation models to 3D. The method proposed by (Siddiqui et al., 2022) adds multiple branches to NeRF for instance segmentation. Works, like (Engelmann et al., 2024), extend NeRF’s capabilities to 3D scene understanding by leveraging pixel-aligned CLIP features from 2D models like (Li et al., 2022). Our work builds on LERF (Kerr et al., 2023), which utilizes pyramid-based CLIP supervision for open-vocabulary 3D segmentation. However, while LERF’s global CLIP features enable effective language-driven queries, they often lack the precision needed for 3D segmentation - a limitation our method seeks to improve.

2.3 3D Open-Vocabulary Segmentation

Extending open-vocabulary segmentation from 2D to 3D brings challenges, as 2D vision-language models like CLIP struggle with the spatial complexity of 3D scenes. Methods like OpenMask3D (Takmaz et al., 2023) accumulate and average CLIP features obtained from instance-centered image crops. The features are then used to represent the 3D instance for instance segmentation. OpenScene (Peng et al., 2023) projects 2D CLIP features into 3D by aligning point clouds with 2D embeddings using a 3D convolutional network. This enables language-driven queries without labeled 3D data. Other methods, like (Yang et al., 2024), leverage image captioning models (Wang et al., 2022) to generate textual descriptions of images, and align point cloud features with open-text representations. Based on LERF, our method leverages NeRF as a flexible framework for 2D to 3D lift-

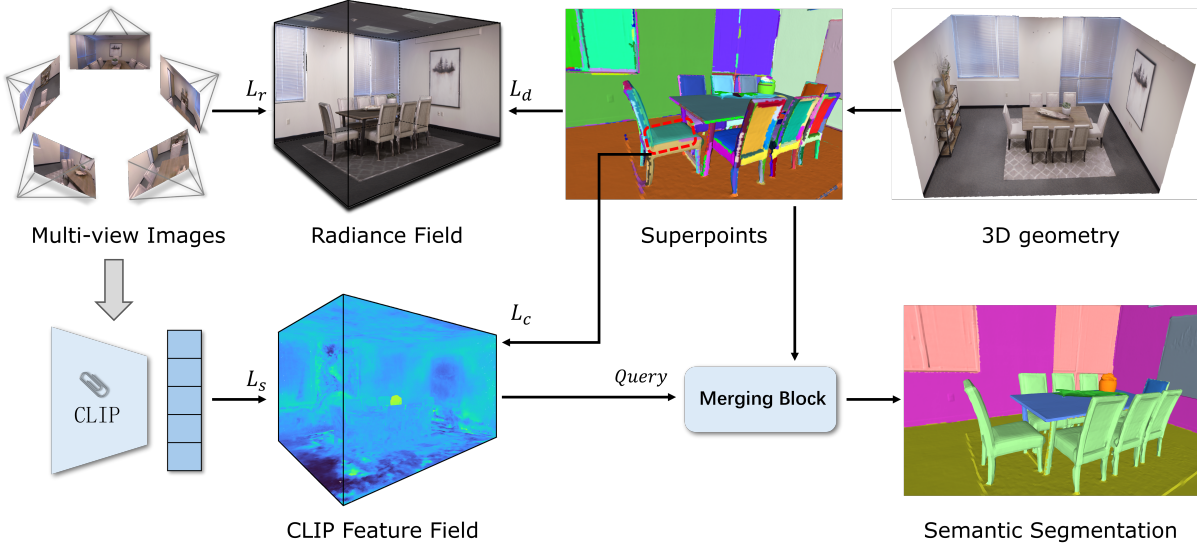


Figure 1: **Overview of SPNeRF Pipeline.** Given 2D posed images as input, SPNeRF optimizes a 3D CLIP feature field by distilling vision-language embeddings from the CLIP image encoder. Simultaneously, the radiance field is trained in parallel. Superpoints, which are extracted from the 3D geometry, are used to enhance both the radiance field and CLIP feature field during optimization, ensuring better alignment of semantic and spatial information. The training process leverages a combination of loss functions \mathcal{L} to refine the consistency and accuracy of the feature representations. The merging block combines query labels with superpoints information to produce semantic segmentation results.

ing, avoiding the geometric consistency limitation of direct 2D projection methods. Furthermore, we take advantage of simple geometric primitives instead of full 3D object masks in (Takmaz et al., 2023) to enhance spatial coherence across the scene.

3 METHOD

In this section, we introduce SPNeRF, our proposed method for zero-shot 3D semantic segmentation. SPNeRF extends NeRF by incorporating CLIP features into an additional feature field, building on principles similar to LERF. We outline the loss functions which are designed to train this feature field, ensuring improved consistency of CLIP features within superpoints. Furthermore, we detail a merging mechanism for robust semantic class score and leverage superpoint affinity scores to refine the segmentation results. A comprehensive overview of the SPNeRF pipeline is presented in Figure 1.

3.1 Preliminary: LERF

We first introduce Language Embedded Radiance Fields (LERF) (Kerr et al., 2023) which SPNeRF is built upon. LERF integrates CLIP embeddings into a 3D NeRF framework, enabling open-vocabulary scene understanding by grounding semantic language

features spatially across the 3D field. Unlike standard NeRF outputs (Mildenhall et al., 2020; Barron et al., 2021), LERF introduces a dedicated language field, which leverages multi-scale CLIP embeddings to capture semantic information across varying levels of detail. This language field is represented by $\mathbf{F}_{\text{lang}}(\mathbf{x}, s)$, where \mathbf{x} is the 3D position and s is the scale.

To supervise this field, LERF uses a precomputed multi-scale feature pyramid of CLIP embeddings as ground truth. The feature pyramid is generated from patches of input multi-view images at different scales. Utilizing volumetric rendering (Max, 1995), the language field can be used to render CLIP embeddings in 2D along each ray $\vec{\mathbf{r}}(t) = \mathbf{o} + t\mathbf{d}$:

$$\phi_{\text{lang}}(\mathbf{r}) = \int T(t)\sigma(t)\mathbf{F}_{\text{lang}}(\mathbf{r}(t), s(t))dt, \quad (1)$$

where $T(t)$ represents accumulated transmittance, $\sigma(t)$ is the volume density, and $s(t)$ adjusts according to the distance from the origin, enabling efficient, scale-aware 3D relevance scoring. The rendered CLIP embedding is then normalized to the unit sphere similar as in (Radford et al., 2021).

The main objective during training is to align the rendered CLIP embeddings with the ground truth CLIP embeddings by minimizing the following loss:

$$\mathcal{L}_{\text{lang}} = -\lambda_{\text{lang}} \sum_i \phi_{\text{lang}} \cdot \phi_{\text{gt}} \quad (2)$$

where ϕ_{lang} denotes the rendered CLIP embedding, ϕ_{gt} is the corresponding target embedding from the

precomputed feature pyramid, and λ_{lang} is the loss weight. This loss encourages CLIP embeddings in the language field to align with its language-driven semantic features, thereby allowing open-vocabulary queries within the 3D scene.

3.2 Geometric Primitive

A core component of SPNeRF is the introduction of geometric primitives. Following recent works (Yin et al., 2024; Yang et al., 2023), we employ a normal-based graph cut algorithm (Felzenszwalb and Huttenlocher, 2004) to over-segment the point cloud $P \in \mathbb{R}^{N \times 3}$ into a collection of superpoints $\{Q_i\}_{i=1}^{N_Q}$, this results in higher-level groupings that better capture the geometric structure of the scene. By aggregating CLIP features at the superpoint level rather than for individual points, we produce more coherent representations, addressing the ambiguities often encountered with point-wise embeddings.

To ensure consistency in the aggregated CLIP features and to align the NeRF representation with the input point cloud, we introduce two complementary loss functions: a *consistency loss* and a *density loss*.

Consistency Loss. To promote consistency across batches of points within superpoints, we employ the consistency loss on sampled pairs of point embeddings following the Huber loss, then average the results across multiple scales. This loss makes the CLIP embeddings more resilient to outliers, allowing the embeddings to align closely with the majority of points in each batch. Given two embeddings, \mathbf{f}_i and \mathbf{f}_j , from a batch of sampled points within a superpoint, the consistency loss for each pair is defined as:

$$\mathcal{L}_c(\mathbf{f}_i, \mathbf{f}_j) = \begin{cases} \frac{1}{2} \|\mathbf{f}_i - \mathbf{f}_j\|^2 & \text{if } \|\mathbf{f}_i - \mathbf{f}_j\| \leq \delta \\ \delta \|\mathbf{f}_i - \mathbf{f}_j\| - \frac{1}{2} \delta^2 & \text{if } \|\mathbf{f}_i - \mathbf{f}_j\| > \delta \end{cases} \quad (3)$$

where δ is a threshold parameter that determines the transition between the quadratic and linear regions of the loss. The overall consistency loss for a batch is then averaged across all scales as follows:

$$\mathcal{L}_{c.\text{batch}} = \frac{1}{N} \sum_{(i,j) \in \text{batch}} \frac{1}{S} \sum_{k=1}^S \mathcal{L}_c(\mathbf{f}_{i_k}, \mathbf{f}_{j_k}) \quad (4)$$

where N is the number of sampled point pairs, and *batch* represents the set of sampled pairs within the superpoints, S is the number of scales. This averaging across scales encourages consistent feature alignment within superpoints.

Density Loss. To ensure that NeRF accurately captures the geometry of the 3D scene, we use a density loss to guide NeRF’s density field based on the point cloud positions. For a given point p_i from the point cloud, we encourage the NeRF density $\sigma(p_i)$ at that location to be close to 1, indicating high occupancy:

$$\mathcal{L}_{\text{density}} = \frac{1}{N} \sum_{i=1}^N (1 - \sigma(p_i))^2 \quad (5)$$

this loss ensures that NeRF correctly represents the occupied regions of the 3D space, aligning the density field with the underlying point cloud geometry.

Progressive Training. To ensure effective optimization of SPNeRF, we employ a progressive training strategy. Initially, we apply the NeRF color rendering loss (Mildenhall et al., 2020) during training to allow the geometry to converge and establish a spatial structure. Then, we introduce the CLIP language embedding loss $\mathcal{L}_{\text{lang}}$, enabling the language field to learn meaningful language features for positions in the 3D field. Finally, we incorporate the consistency loss $\mathcal{L}_{c.\text{batch}}$ and density loss $\mathcal{L}_{\text{density}}$ to enhance the consistency and robustness of the CLIP embeddings within superpoints. This staged training process ensures a balanced and efficient optimization of both the geometric and semantic components of SPNeRF.

3.3 Merging Block

Instead of relying on per-point clustering, SPNeRF assigns segmentation labels based on the relevancy score between superpoint-level CLIP embeddings and the target class label embeddings.

Relevancy Score. After training, we begin by using farthest point sampling to select N_p representative points within each superpoint. Given a sampled point p_i in the superpoint sp_n , $i \in \{1, 2, \dots, N_p\}$, with N_p being the number of sampled points, $n \in \{1, 2, \dots, N\}$ and N the number of superpoints. We retrieve the CLIP embedding \mathbf{f}_{p_i} of point p_i by querying the SPNeRF CLIP feature field at p_i ’s position. These embeddings collectively represent the superpoint sp_n ’s CLIP feature set.

Next, the target class label is encoded using the CLIP text encoder to produce a positive CLIP embedding \mathbf{f}_{pos} . As proposed by LERF (Kerr et al., 2023), we also define a set of negative CLIP embeddings $\{\mathbf{f}_{\text{neg}_k}\}$, $k \in \{1, 2, \dots, K\}$. $\mathbf{f}_{\text{neg}_j}$ represent the encoded features of canonical text like “object” and “things”. For each representative embedding \mathbf{f}_{p_i} within the superpoint, we compute the cosine similarity with both

the positive class embedding and each negative embedding. For each point p_i , its relevancy score R_{p_i} is the minimum score across all negative comparisons after softmax normalization:

$$F(\mathbf{f}_1, \mathbf{f}_2) = \exp(\text{sim}(\mathbf{f}_1, \mathbf{f}_2)) \quad (6)$$

$$R_{p_i} = \min_K \left(\frac{F(\mathbf{f}_{p_i}, \mathbf{f}_{\text{pos}})}{F(\mathbf{f}_{p_i}, \mathbf{f}_{\text{pos}}) + F(\mathbf{f}_{p_i}, \mathbf{f}_{\text{neg}_k})} \right) \quad (7)$$

where \exp is the exponential function, sim is the cosine similarity, and K is the number of negative embeddings.

The relevancy score R_{sp_n} for a given superpoint sp_n is the median relevancy scores of all sampled points, which is robust against outliers. The median point’s CLIP embedding is also used to represent the superpoint’s CLIP embedding \mathbf{f}_{sp_n} .

Affinity Score. To further enhance the effect of relevancy score, we introduce an affinity score. Same as calculating the relevancy score, we need positive and negative embeddings for comparison to define the score. For the given class, we choose N_a superpoints which have the highest relevancy score as positive superpoints sp_{pos_j} , and use their CLIP embeddings as positive embeddings $\{\mathbf{f}_{pos_j}\}$, $j \in \{1, 2, \dots, N_a\}$. We choose another N_a superpoints which have the lowest relevancy score as negative superpoints sp_{neg_k} , and use their CLIP embedding as negative embeddings $\{\mathbf{f}_{neg_k}\}$, $k \in \{1, 2, \dots, N_a\}$. In order to calculate the affinity score $A_{sp_n-sp_{pos_j}}$ between a superpoint sp_n and a positive superpoint sp_{pos_j} , we compare \mathbf{f}_{sp_n} with each positive embedding \mathbf{f}_{pos_j} and the set of N_a negative embeddings $\{\mathbf{f}_{neg_j}\}$, and select the minimum score across all negative comparisons:

$$A_{sp_n-sp_{pos_j}} = \min_{N_a} \left(\frac{F(\mathbf{f}_{sp_n}, \mathbf{f}_{pos_j})}{F(\mathbf{f}_{sp_n}, \mathbf{f}_{pos_j}) + F(\mathbf{f}_{sp_n}, \mathbf{f}_{neg_k})} \right) \quad (8)$$

Then, we use the relevancy score R_{pos_j} of each positive superpoint sp_{pos_j} as weight to average all N_a affinity scores, and acquire the affinity score A_{sp_n} for the superpoint sp_n :

$$A_{sp_n} = \frac{\sum_{j=1}^{N_a} R_{pos_j} \cdot A_{sp_n-sp_{pos_j}}}{N_a} \quad (9)$$

The relevance score R_{sp_n} for superpoint sp_n is then scaled with affinity. The scaled relevancy score $R_{sp_n}^*$ can be calculated as:

$$R_{sp_n}^* = R_{sp_n} \cdot w \cdot (1 + (A_{sp_n} - \min_N(A_{sp_n}))) \quad (10)$$

where w is the affinity weight, and N is the number of superpoints.

For each superpoint sp_n , the class with highest scaled relevancy score is assigned during semantic segmentation.

4 EXPERIMENTS

In this section, we present our experimental evaluation assessing both quantitative and qualitative performance. We compare its performance in zero-shot 3D segmentation with respect to the baseline methods LERF and OpenNeRF. In addition, we conduct an ablation study to analyze the contribution of each of SPNeRF’s components, including the consistency loss and affinity alignment.

4.1 Experiment Setup

We evaluated SPNeRF on the Replica dataset, a standard benchmark for 3D scene understanding. Replica dataset comprises photorealistic indoor scenes with high-quality RGB images and 3D point cloud data, annotated with per-point semantic labels for a variety of object categories. This dataset serves as a robust benchmark for evaluating segmentation in complex, densely populated indoor environments. For each scene in Replica, 200 posed images are used for all experiments.

For image language features extraction, we utilized OpenCLIP ViT-B/16 model. We trained our SPNeRF with posed RGB images and 3D geometry, applied zero-shot semantic segmentation without additional fine-tuned or pre-trained 2D segmentation models. For evaluation, we followed the approach of (Peng et al., 2023). The accuracy of the predicted semantic labels is evaluated using mean intersection over union (mIoU) and mean accuracy (mAcc).

4.2 Method Comparison

Quantitative Evaluation. We compare with OpenNeRF (Engelmann et al., 2024), OpenScene (Peng et al., 2023) and LERF (Kerr et al., 2023). To evaluate LERF, we generated segmentation masks by rendering relevancy maps, projecting them onto Replica point clouds, and assigning each point of the class with the highest score, same as (Engelmann et al., 2024) proposed. OpenNeRF is evaluated with their provided code. In order to provide a comparison of the models’ own effectiveness on segmentation, we did not use NeRF-synthesized novel views to fine-tune the models during comparison. In contrast, SPNeRF and LERF use RGB images as input, while OpenNeRF takes RGB images and corresponding depth maps as input. Results of OpenScene are taken from (Engelmann et al., 2024).

Table 1 summarizes the 3D semantic segmentation performance of SPNeRF relative to baseline methods on the Replica dataset. SPNeRF achieves




Figure 2: **3D segmentation results in comparison to other methods.** Qualitative comparison of 3D semantic segmentation results on the Replica dataset. Rows display results from (top to bottom) ground truth, LERF, OpenNeRF and SPNeRF, across 3 indoor scenes. SPNeRF demonstrates improved boundary coherence and segmentation accuracy in general.

Table 1: Quantitative results on Replica dataset for 3D semantic segmentation.

Method	mIoU	mAcc
LERF	10.5	25.8
OpenScene	15.9	24.6
OpenNeRF	19.73	32.61
SPNeRF (Ours)	17.25	31.07

competitive scores with a mIoU of 17.25% and mAcc of 31.07%. While LERF and SPNeRF both use original CLIP to extract semantic information, SPNeRF improves significantly over the baseline LERF. Although OpenNeRF attains the highest overall performance with the support of a fine-tuned 2D model for segmentation, SPNeRF’s results emphasize its effective integration of superpoint-based feature aggregation and language-driven embeddings without additional 2D segmentation knowledge.

The experimental results demonstrate SPNeRF’s enhanced capability in maintaining feature consistency within superpoints, especially when evaluated

against the LERF baseline. The 6.75% improvement in mIoU over LERF without structural network changes illustrates the impact of our approach in aligning semantic language features spatially across 3D fields. Without any 2D segmentation knowledge, SPNeRF’s results align closely quantitatively with OpenNeRF which is trained with a 2D segmentation model, indicating CLIP’s potential for fine-grained segmentation.

Qualitative Evaluation. Figure 2 illustrates a qualitative comparison of segmentation results between SPNeRF, OpenNeRF, and LERF across various indoor scenes in the Replica dataset. While the other methods’ segmentation tend to splash near boundaries, SPNeRF demonstrates great boundary coherence and spatial consistency, particularly in scenes with complex object arrangements. OpenNeRF, while generally robust in correct class estimation, exhibits minor loss of detail in cluttered environments. SPNeRF’s superpoint-based segmentation mitigates




Figure 3: **Ablation comparison of consistency loss.** Even constrained by the fragmented superpoints, the results w/o loss tend to be consistent due to the image embedding characteristic of CLIP. The consistency loss helps the model to get more precise semantic info, especially for large superpoints like wall surfaces.




Figure 4: **Ablation comparison.** The figure illustrates the improvement by consistency loss and affinity score.

these issues by aggregating features within geometric boundaries, resulting in coherent representations, especially when comparing the wall areas with LERF, SPNeRF learns to concentrate on the correct semantic even using same network structure.

Overall, the qualitative results highlight the ability of SPNeRF to deliver competitive 3D segmentation in complex scenes, complementing its quantitative gains in mIoU and mAcc. The combination of CLIP embeddings and superpoint-based relevancy scoring enables SPNeRF to differentiate structures and maintain semantic consistency across object boundaries, reducing noise and improving clarity in less visible areas.

4.3 Ablation Study

To analyze the impact of SPNeRF’s individual components, we perform an ablation study by systematically removing the primitive consistency loss and affinity-based refinement. Table 2 presents the quantitative results. Removing the primitive consistency loss results in a notable decrease in mIoU (from 17.25 to 15.31) and mAcc (from 31.07 to 26.82), highlighting its importance in preserving coherent embeddings within superpoints. As also shown in Figure 3, consis-

tency loss largely improved the precision of classification, especially for large superpoints like walls, which are more likely to contain different semantic embeddings, consistency loss helps the CLIP feature field to learn the most important and distributed semantics of superpoints. Similarly, excluding the affinity-based refinement slightly reduced the performance numerically. As shown in Figure 4, affinity refinement can improve the segmentation quality by capturing semantic relationships between superpoints, for example chair surfaces, and maintain the possibility to over cover adjacent parts.

Table 2: Ablation study results on the Replica dataset. Both the primitive consistency loss and affinity refinement contribute significantly to SPNeRF’s overall segmentation quality.

Model Variant	mIoU	mAcc
Full SPNeRF	17.25	31.07
w/o Affinity Refinement	17.13	30.77
w/o Consistency Loss	15.31	26.82
w/o both	13.78	24.59

5 CONCLUSION

We introduced SPNeRF, a zero-shot 3D segmentation approach that enhances Neural Radiance Fields (NeRF) through the integration of geometric primitives and visual-language features. Without training on any ground truth labels, our model can semantically segment unseen complex 3D scenes. By embedding superpoint-based geometric structures and applying a primitive consistency loss, SPNeRF overcomes the limitations of CLIP’s image-based embeddings, achieving higher spatial consistency and segmentation quality in 3D environments, while mitigating ambiguities in point-wise embeddings. SPNeRF outperforms LERF and performs competitively with OpenNeRF, while SPNeRF avoids additional 2D segmentation models required by OpenNeRF. While SPNeRF has demonstrated competitive performance, it also inherits limitations from CLIP’s 2D image-based embeddings, leading to occasional ambiguities in details. Future work could explore more efficient alternatives to NeRF, such as Gaussian splatting (Kerbl et al., 2023) or efficiently incorporating 2D foundation models like the Segment Anything Model (SAM) (Kirillov et al., 2023) to enable instance-level segmentation.

ACKNOWLEDGEMENTS

This work has partly been funded by the German Federal Ministry for Digital and Transport (project EConoM under grant number 19OI22009C).

REFERENCES

- Barron, J. T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., and Srinivasan, P. P. (2021). Mipnerf: A multiscale representation for anti-aliasing neural radiance fields. *ICCV*.
- Cen, J., Fang, J., Zhou, Z., Yang, C., Xie, L., Zhang, X., Shen, W., and Tian, Q. (2024). Segment anything in 3d with radiance fields.
- Engelmann, F., Manhardt, F., Niemeyer, M., Tateno, K., Pollefeys, M., and Tombari, F. (2024). OpenNeRF: Open Set 3D Neural Scene Segmentation with Pixel-Wise Features and Rendered Novel Views. In *International Conference on Learning Representations*.
- Felzenszwalb, P. F. and Huttenlocher, D. P. (2004). Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59:167–181.
- Ghiasi, G., Gu, X., Cui, Y., and Lin, T. (2021). Open-vocabulary image segmentation. *CoRR*, abs/2112.12143.
- Kerbl, B., Kopanas, G., Leimkühler, T., and Drettakis, G. (2023). 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. on Graphics*, 42(4).
- Kerr, J., Kim, C. M., Goldberg, K., Kanazawa, A., and Tancik, M. (2023). LERF: Language embedded radiance fields. In *ICCV*.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., and Girshick, R. (2023). Segment anything. *arXiv:2304.02643*.
- Lan, M., Chen, C., Ke, Y., Wang, X., Feng, L., and Zhang, W. (2024). Clearclip: Decomposing clip representations for dense vision-language inference.
- Li, B., Weinberger, K. Q., Belongie, S. J., Koltun, V., and Ranftl, R. (2022). Language-driven semantic segmentation. *CoRR*, abs/2201.03546.
- Luo, H., Bao, J., Wu, Y., He, X., and Li, T. (2023). Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation.
- Max, N. (1995). Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. (2020). Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*.
- Peng, S., Genova, K., Jiang, C. M., Tagliasacchi, A., Pollefeys, M., and Funkhouser, T. (2023). Openscene: 3d scene understanding with open vocabularies.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. (2021). Learning transferable visual models from natural language supervision.
- Siddiqui, Y., Porzi, L., Buló, S. R., Müller, N., Nießner, M., Dai, A., and Kotschieder, P. (2022). Panoptic lifting for 3d scene understanding with neural fields.
- Sun, S., Li, R., Torr, P., Gu, X., and Li, S. (2024). Clip as rnn: Segment countless visual concepts without training endeavor.
- Takmaz, A., Fedele, E., Sumner, R. W., Pollefeys, M., Tombari, F., and Engelmann, F. (2023). OpenMask3D: Open-Vocabulary 3D Instance Segmentation. In *NeurIPS*.
- Wang, P., Yang, A., Men, R., Lin, J., Bai, S., Li, Z., Ma, J., Zhou, C., Zhou, J., and Yang, H. (2022). Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework.
- Xu, J., Mello, S. D., Liu, S., Byeon, W., Breuel, T., Kautz, J., and Wang, X. (2022). Groupvit: Semantic segmentation emerges from text supervision.
- Yang, J., Ding, R., Deng, W., Wang, Z., and Qi, X. (2024). Regionplc: Regional point-language contrastive learning for open-world 3d scene understanding.
- Yang, Y., Wu, X., He, T., Zhao, H., and Liu, X. (2023). Sam3d: Segment anything in 3d scenes.
- Yin, Y., Liu, Y., Xiao, Y., Cohen-Or, D., Huang, J., and Chen, B. (2024). Sai3d: Segment any instance in 3d scenes.