Don't Starve the Boundaries: Boundary-Constrained Label Propagation for Weakly Supervised 3D Segmentation

Shuwei Wu, Shuo Jin, Zhijin He, Qiufeng Wang, Siyue Yu*, Eng Gee Lim, and Jimin Xiao

Xi'an Jiaotong-Liverpool University 西交利物浦大学

ECCV 2026 | MALMÖ | SEPT 8–12

Abstract— Although fully supervised methods have substantially advanced segmentation in boundary areas of point cloud, effective weakly supervised approaches remain scarce. This is primarily because limited supervision rarely reaches boundary regions, leaving them lacking reliable supervision. We propose a novel pseudo-label propagation paradigm in native 3D that does not rely on the model's own predictions or any external foundation models, yet is able to generate high purity pseudo-labels. Compared with SAM based 2D-3D projection, our pseudo-labels are purer and more uniformly distributed. Even under 1 pt/obj setting on S3DIS, our initial offline propagation achieves >94.4% accuracy (≈93k pts per scene). Experiments show that our method outperforms existing state-of-the-art methods.

Introduction

3D scene understanding relies on costly annotations, motivating the rise of Weakly Supervised 3D Semantic Segmentation (WS3DSS). However, due to the scarcity of annotations at class transitions, precise boundary segmentation remains a challenge. While consistency-based training is standard, it struggles with boundary delineation. Experiments reveal that enforcing strict consistency at boundaries degrades performance by blending semantic distinctions—we term *boundary starvation*. As illustrated in the Figure 1, the red boxes highlight the model's poor discrimination at class boundaries. Furthermore, existing pseudo-labeling based methods exacerbate this issue by filtering out valid but low-confidence points. Thereby further contributing to *boundary starvation.*

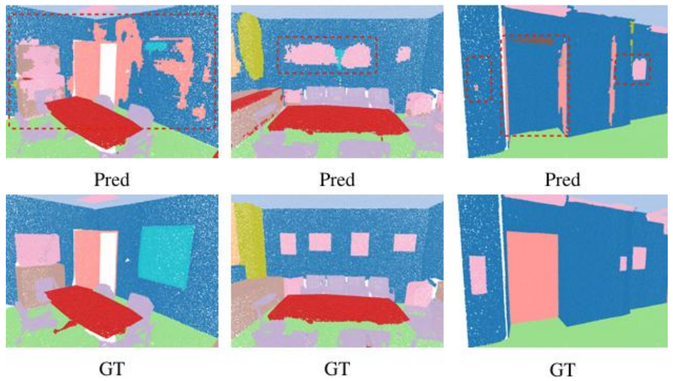

Figure 1. Systematic boundary distortion exhibited by previous method[1] under 0.01% annotation.

To address this, we propose Bound3D, a framework that prioritizes high-quality boundary pseudo-labels. We introduce a Boundary Illumination Module to extract boundary cues, which guide a Boundary-constrained Pseudo Label Generator to propagate reliable labels.

Method

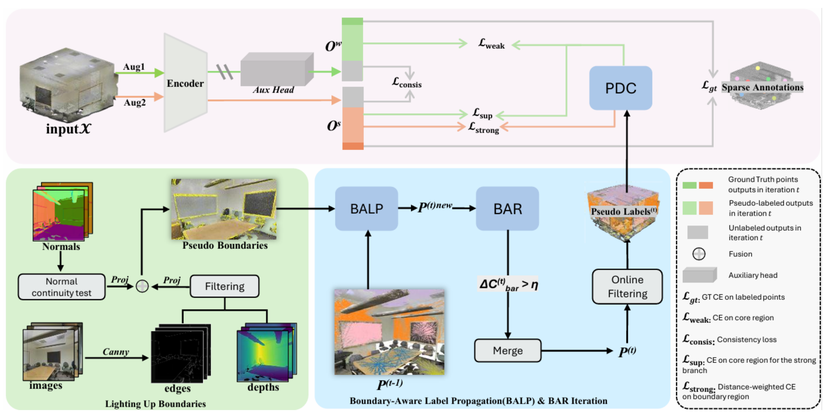

Figure 2. The pipeline of Bound3D. The input point cloud is augmented twice with different perturbation strengths, where Aug2 denotes the stronger augmentation.

1) We first feed the given point cloud together with its multi-view RGB images, depth maps, and normal maps into BIM for boundary detection, where 2D boundary is extracted by combining depth-filtered Canny edges with normal-discontinuity masks and then projected into the corresponding 3D scene, yielding high-quality 3D boundary points for our BPLG.

2) After that, the detected boundary points guide our BPLG to generate comprehensive pseudo-labels; within BPLG, the BCLP mechanism produces the labels via boundary-constrained propagation, while boundary arrival rate is computed to controls online iterations and supplies new seed points for each BCLP iteration.

3) Finally, PADC dispatches the pseudo-labels to a dual-branch consistency network in a divide-and-conquer manner. After training, the inference process requires only the raw point cloud for final segmentation results. Note here that no 2D information is needed.

Results and Discussion

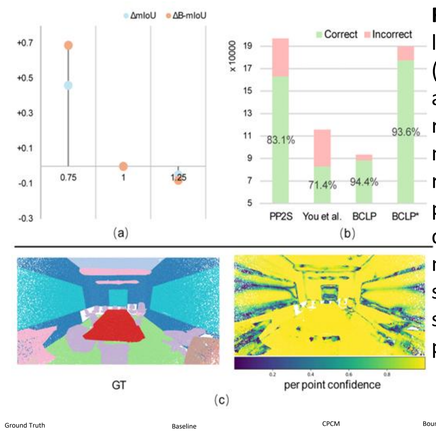

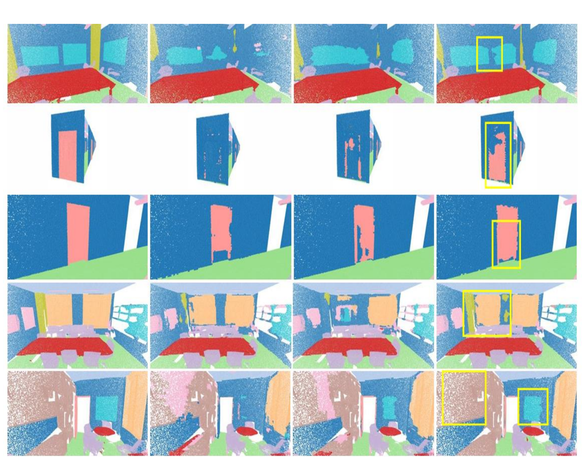

Figure 3. (a): Vary the consistency loss weight w on boundary regions (with w=1.0 as the baseline). The y-axis reports ΔmIoU and ΔB-mIoU relative tow=1.0. B-mIoU denotes mIoU computed within boundary regions. (b): BCLP vs. SAM-based pseudo-label generation[2,3]: quantity and accuracy. ∗denotes a more aggressive hyper parameter setting. (c): GT points in the same scene and the corresponding predicted confidence.

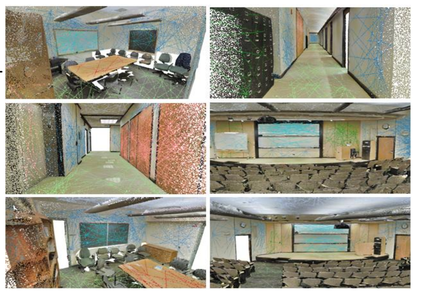

0.01%	SQN [8]	ECCV 2022	45.3
	CPCM [17]	ICCV 2023	59.3
	CPCM† [17]	ICCV 2023	59.8
	PointMatch [37]	Comput. Graph.2023	59.9
	AADNet [20]	AAAI 2025	60.8
	Bound3D (Ours)	Submission	**61.2**
OTOC (0.02%)	OTOC [19]	CVPR 2021	50.1
	DAT [38]	ECCV 2022	56.5
	RAC-Net [39]	IJCV 2024	58.4
	You et al. (PTv3) [44]	arXiv 2025	64.4
1 pt/obj	**Bound3D (Ours)**	Submission	**61.9**
	AO (PTv2) [12]	CVPR 24	62.7

Table 1. SOTA comparison on S3DIS Area5.

Figure 4. Visualizes the pseudo-labels generated by BPLG during training on S3DIS under the 0.01% supervision setting.

Figure 5 presents a comparison between our method, the Ground Truth (GT), the baseline, and the SOTA method. The yellow boxes highlight Bound3D's superior segmentation capability at class boundaries, demonstrating its ability to identify object boundaries even under extreme annotation budgets.

Conclusion and Future Work

In this paper, we identify and analyze *boundary starvation*, a long-standing and cross-paradigm issue in WS3DSS, and introduce BCLP, a native 3D propagation method. BCLP departs from traditional reliance on model predictions or costly 2D strong priors (e.g., SAM), producing purer and more uniformly distributed pseudo-labels. Furthermore, we design a Position-aware Divide-and-Conquer strategy to break the uniform regularization that suppresses boundary learning, thereby further mitigating *boundary starvation*. We hope this work will motivate the community to prioritize discriminative feature learning at point cloud class boundaries, rather than solely pursuing overall mIoU metrics that are dominated by interior and dominant classes.