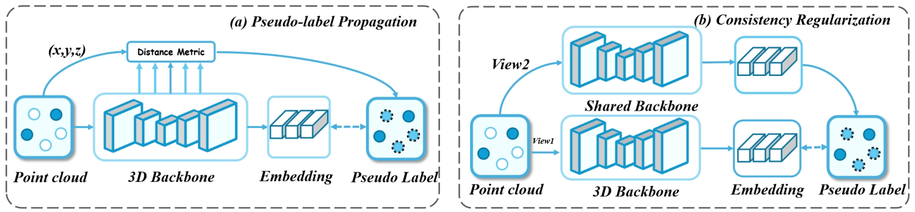
PARI: Policy-Driven Active Residual Intervention for Weakly Supervised Point Cloud Segmentation

Wenhao Gao[1]* Sen Tao[1]* Jiawei Liu[1]† Jiwei Feng[1] Yongchao Xu[1]
Junfeng Wang[1] Tao Jiang[2] Jiangbo Ai[2] Jin Zhang[2] Zheng-Jun Zha[1]

[1] MoE Key Lab of BIPC, University of Science and Technology of China [2] Alibaba Group

Background

We study 3D semantic segmentation under sparse point annotations. Pseudo-label propagation may reinforce incorrect priors, while local consistency regularization does not fully exploit scene context.



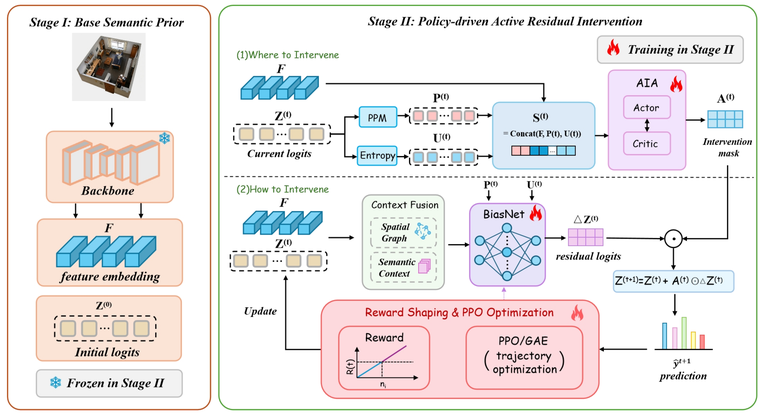
The figure contrasts pseudo-label propagation and cross-view consistency with PARI, which learns an intervention policy and predicts residual corrections for selected points.

Main contributions

1. **Active Intervention Agent (AIA)** formulates point-wise selection as a Markov decision process and learns an actor–critic policy from semantic features, prediction priors and uncertainty.
2. **Bias-aware Refinement Module (BRM)** fuses spatial and feature context through attention and predicts residual logits for the selected points.
3. **Confidence-aware Reward Shaping (CRS)** combines label-driven feedback with exploration rewards on low-confidence points and activation penalties on high-confidence points.

Policy feedback. Label feedback rewards corrected errors and penalizes missed or newly introduced errors. Confidence-aware shaping encourages intervention on uncertain points while suppressing unnecessary activation on reliable predictions.

Methods



Training and optimization

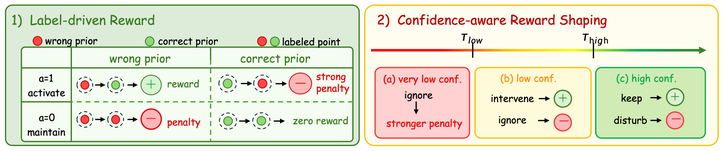
Stage I trains the base segmentation network to obtain semantic features and initial logits.
Stage II freezes this network and jointly optimizes the intervention agent and BiasNet over multi-step trajectories.
The agent is trained with the clipped PPO objective. BiasNet is supervised by sparse labels and high-confidence pseudo-labels and minimizes the following refinement loss:

$$L_{\text{bias}} = L_{\text{CE}} + \lambda_1 L_{\text{KL}} + \lambda_2 L_{\text{ent}} + \lambda_3 L_{\text{neigh}}$$

KL preserves high-confidence priors; entropy reduces uncertainty on activated low-confidence points; neighbor consistency regularizes local predictions.

Reward design

$$r_i^{(t)} = r_{\text{i,lab}}^{(t)} + r_{\text{i,conf}}^{(t)}$$



Experiments

S3DIS Area 5 / ScanNetV2 validation; mIoU (%).
AADNet backbone; matched annotation protocols.

Main comparison

Dataset / labels	Base	PARI	Gain
S3DIS / 1%	68.3	69.4	**+1.1**
S3DIS / 0.1%	67.2	67.9	**+0.7**
S3DIS / 0.01%	60.8	61.4	**+0.6**
ScanNetV2 / 1%	66.8*	67.2	**+0.4**
ScanNetV2 / 20pts	60.2*	60.8	**+0.6**

* Reproduced baselines [Tables 1–2]. Gain: mIoU difference; 20pts: 20 labels per scene.

- **S3DIS sparsity.** PARI improves AADNet by 1.1, 0.7 and 0.6 mIoU at 1%, 0.1% and 0.01% labels. Gains persist as annotation density decreases.
- **ScanNetV2 results.** Gains remain positive at 1% labels (+0.4) and at only 20 labels per scene (+0.6), supporting refinement under extreme sparsity.

Comparison with full supervision

Dataset	Full (100%)	PARI (1%)
S3DIS	69.3	**69.4**
ScanNetV2	66.9	**67.2**

- **Sparse vs. full supervision.** With 1% labels, PARI reaches 69.4/67.2 mIoU on S3DIS/ScanNetV2, matching or slightly exceeding the same-backbone full-supervision references (69.3/66.9). [Tables 1–2]

Module ablation

S3DIS variant	1%	0.1%	0.01%
Base model	68.3	67.2	60.8
+ BRM only	68.9	66.9	60.4
+ BRM + AIA	**69.4**	**67.9**	**61.4**

- **Selection before refinement.** BRM alone gains 0.6 mIoU at 1% labels but loses 0.3/0.4 at 0.1%/0.01%. Adding AIA improves BRM by 0.5/1.0/1.0, indicating that selecting where to intervene is critical under sparse supervision. [Table 5(a)]

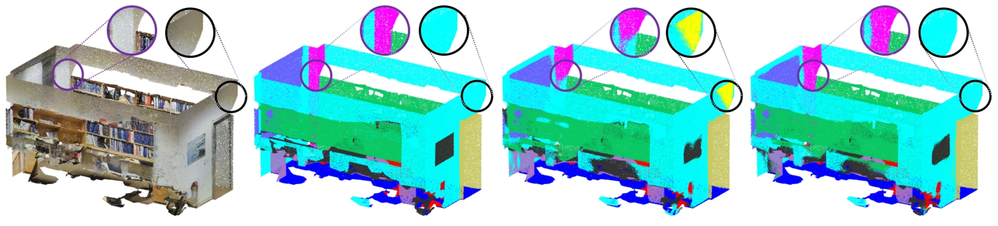
Intervention policy

Selection policy	mIoU	Fix − Break
Random	69.254	+3,384
Entropy top-k	69.323	+4,757
Actor top-k	**69.357**	**+21,792**

- **Learned point selection.** At a 10% local budget, Actor gives 21,792 net corrections (fixed minus newly introduced errors), versus 4,757 for entropy and 3,384 for random selection. The mIoU gain over entropy is 0.034, while net corrections increase 4.6-fold. [Table 6(b)]
- **Budget sensitivity.** A 5%–20% budget gives 69.307–69.395 mIoU, close to the adaptive policy's 69.402. Performance varies little over this range. [Table 6(a)]

Qualitative Comparisons

S3DIS comparison: input, ground truth, baseline and PARI (Fig. 4).



Input point cloud | Ground truth | Baseline | **PARI (ours)**

- **Semantic boundaries.** In the left enlarged region, PARI follows the ground-truth boundary more closely. BRM uses attention to fuse spatial neighbors and semantically similar points before predicting residual logits, supplying geometric and semantic evidence for resolving boundary ambiguity.
- **Wall consistency.** In the right enlarged region, PARI reduces scattered wall errors. BRM combines local geometric support from spatial neighbors with semantic evidence from feature neighbors. The neighbor-consistency loss regularizes local predictions, consistent with the more coherent wall surface.
- **Selective refinement.** AIA selects points from features, prediction priors and uncertainty; its mask gates residual updates. Confidence-aware rewards encourage low-confidence intervention and penalize high-confidence activation, balancing error correction with preservation of reliable priors.