

Point-SRA: Self-Representation Alignment for 3D Representation Learning

Lintong Wei¹, Jian Lu^{1*}, Haozhe Cheng², Jihua Zhu², Kaibing Zhang³

¹School of Electronics and Information, Xi'an Polytechnic University

²School of Software, Xi'an Jiaotong University

³School of Computer Science, Xi'an Polytechnic University
lujian@xpu.edu.cn

Abstract

Masked autoencoders (MAE) have become a dominant paradigm in 3D representation learning, setting new performance benchmarks across various downstream tasks. Existing methods with fixed mask ratio neglect multi-level representational correlations and intrinsic geometric structures, while relying on point-wise reconstruction assumptions that conflict with the diversity of point cloud. To address these issues, we propose a 3D representation learning method, termed **Point-SRA**, which aligns representations through self-distillation and probabilistic modeling. Specifically, we assign different masking ratios to the MAE to capture complementary geometric and semantic information, while the **MeanFlow Transformer (MFT)** leverages cross-modal conditional embeddings to enable diverse probabilistic reconstruction. Our analysis further reveals that representations at different time steps in MFT also exhibit complementarity. Therefore, a **Dual Self-Representation Alignment** mechanism is proposed at both the MAE and MFT levels. Finally, we design a **Flow-Conditioned Fine-Tuning Architecture** to fully exploit the point cloud distribution learned via MeanFlow. Point-SRA outperforms Point-MAE by 5.37% on ScanObjectNN. On intracranial aneurysm segmentation, it reaches 96.07% mean IoU for arteries and 86.87% for aneurysms. For 3D object detection, Point-SRA achieves 47.3% AP@50, surpassing MaskPoint by 5.12%.

Introduction

MAE has emerged as a leading framework in Self-Supervised Representation Learning (SSRL). In the 3D community, methods such as Point-MAE (Pang et al. 2022), Point-M2AE (Zhang et al. 2022), and MaskPoint (Liu, Cai, and Lee 2022) have successfully adapted this paradigm and achieved strong performance across various tasks. By reconstructing masked regions from sparse visible points, 3D MAE learns robust geometric representations with strong generalization ability. Despite its effectiveness, most existing methods adopt fixed masking ratio based on empirical settings, lacking theoretical insight into how the ratio influence the learned representations. This leads to an essential but underexplored question: *Do different masking ratios inherently yield representations of varying richness, and can*

*Corresponding author.

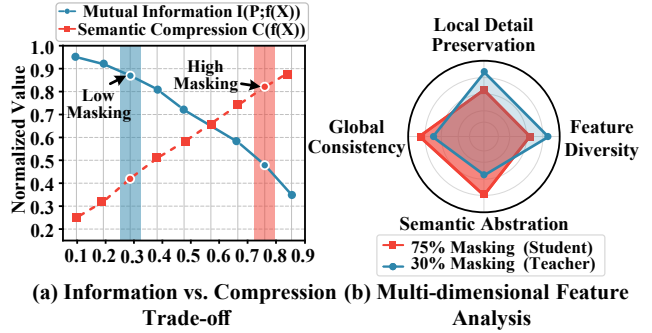


Figure 1: Complementarity under varying mask ratios. (a) Mutual information drops and semantic compression rises with increasing mask ratio. (b) 30% mask (teacher) preserves fine-grained details; 75% mask (student) yields better semantic abstraction.

these differences be leveraged to improve overall representation quality?

Our investigation reveals a fundamental principle: **masking ratio complementarity**. As shown in Figure 1, we discover that masking ratio choice creates a systematic trade-off between geometric detail preservation and semantic abstraction. Specifically, as masking ratio increases, mutual information with the input decreases while semantic compression improves, leading to representations that capture different aspects of 3D structure. Low masking ratios ($\leq 30\%$) excel at preserving fine-grained geometric features, while high masking ratios ($\geq 75\%$) force the model to learn abstract semantic patterns. This complementarity is not incidental but fundamental to the information bottleneck induced by masking. However, current 3D MAE methods suffer from two fundamental limitations that hinder their ability to utilize masking ratio complementarity effectively:

1. Fixed masking strategies limit the model’s ability to benefit from the complementary representations available under different masking ratios.

2. Under high masking ratios, point-wise deterministic reconstruction in 3D MAE often leads to mismatches, as it fails to account for the inherent structural diversity. As illustrated in Figure 2, the same visible region may correspond to multiple plausible results.

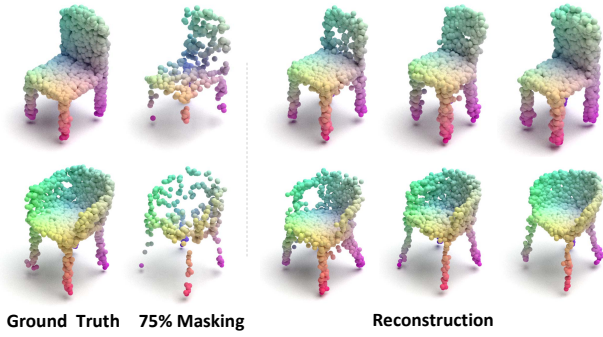


Figure 2: Under high masking ratios, diverse chair reconstructions from the same masked input (identical visible regions) reflect inherent uncertainty in geometric generation and plausible variations in leg shape, backrest angle, and seat thickness.

To fully leverage the discovered masking ratio complementarity, we incorporate **Self-Representation Alignment (SRA)** as an internal knowledge transfer mechanism. Meanwhile, we investigate the probabilistic reconstruction capabilities of MeanFlow and analyze representation information across different time steps. To this end, the **Mean-Flow Transformer (MFT)** is designed to effectively leverage these properties. Moreover, MeanFlow probabilistic formulation naturally accommodates the multi-solution nature of 3D reconstruction. The trajectory-based learning process enables alignment of representations across temporal states, while its support for conditional generation allows flexible incorporation of multi-modal information. Therefore, we design a **Dual Self-Representation Alignment (Dual SRA)** mechanism: **MAE-Level Self-Representation Alignment (MAE-SRA)** integrates representations obtained under different masking ratios, facilitating the fusion of geometric and semantic knowledge. **MFT-Level Self-Representation Alignment (MFT-SRA)** aligns representations across time steps using trajectories learned via MeanFlow, capturing the evolution of point cloud distributions. Finally, we propose **Point-SRA**, a framework that systematically exploits masking ratio complementarity for 3D representation learning. In addition, the standard Transformer fine-tuning network is enhanced by integrating MeanFlow-based updates. Our main contributions are summarized as follows:

- We conduct a systematic theoretical analysis of masking ratio complementarity, reconstruction uncertainty, and numerical stability of MeanFlow, and leverage these insights to design the Point-SRA framework.
- We present a unified Dual SRA mechanism for aligning representations across both masking ratios and temporal states, enabling fully self-contained knowledge transfer.
- We propose the MFT to address the limitations of point-wise reconstruction in 3D MAE, enabling representation guidance and probabilistic reconstruction.
- We develop a Flow-Conditioned Fine-Tuning Architecture that leverages distributional knowledge learned during pre-training for downstream tasks.

Observations

The point cloud $\mathcal{P} \in \mathcal{R}^{N \times 3}$ is partitioned into G local regions $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_G$ via Farthest Point Sampling (FPS). A subset of $\lfloor r \cdot G \rfloor$ regions are randomly selected for masking, where the masking ratio $r \in [0, 1]$. Let \mathcal{V} denote the index set of visible regions, so the visible regions can be expressed as $\mathcal{X}_r = \{\mathcal{N}_i : i \in \mathcal{V}\}$. The encoder $f_{\theta_r} : \mathcal{X}_r \rightarrow \mathcal{R}^d$ maps the visible regions to a d -dimensional representation space.

Theorem A: Masking Ratio Complementarity. For masking ratios $r_l < r_h$, under the information bottleneck framework, the corresponding optimal encoders $f_{\theta_l^*}$ and $f_{\theta_h^*}$ satisfy:

$$\mathcal{I}(\mathcal{P}; f_{\theta_l^*}(\mathcal{X}_{r_l})) > \mathcal{I}(\mathcal{P}; f_{\theta_h^*}(\mathcal{X}_{r_h})), \quad (1)$$

$$\mathcal{C}(f_{\theta_h^*}(\mathcal{X}_{r_h})) > \mathcal{C}(f_{\theta_l^*}(\mathcal{X}_{r_l})), \quad (2)$$

where $\mathcal{I}(\cdot; \cdot)$ denotes mutual information and $\mathcal{C}(\cdot)$ denotes the semantic compression degree defined as:

$$\mathcal{C}(Z) = \frac{\mathcal{H}(Z)}{\mathcal{I}(\mathcal{P}_{\text{semantic}}; Z)}, \quad (3)$$

where $\mathcal{P}_{\text{semantic}}$ denotes semantic information of the point cloud. $\mathcal{H}(Z)$ is the entropy of representation Z . A detailed proof is provided by the Supplementary Material.

Corollary: Representation Complementarity. There exist projection functions: $\pi_{geo} : \mathcal{R}^d \rightarrow \mathcal{R}^{d_{geo}}$ and $\pi_{sem} : \mathcal{R}^d \rightarrow \mathcal{R}^{d_{sem}}$, such that:

$$\begin{aligned} \|\pi_{geo}(f_{\theta_l^*}(\mathcal{X}_{r_l})) - \pi_{geo}(\mathcal{P})\|_F \\ < \|\pi_{geo}(f_{\theta_h^*}(\mathcal{X}_{r_h})) - \pi_{geo}(\mathcal{P})\|_F, \end{aligned} \quad (4)$$

$$\begin{aligned} \|\pi_{sem}(f_{\theta_h^*}(\mathcal{X}_{r_h})) - \pi_{sem}(\mathcal{P})\|_F \\ < \|\pi_{sem}(f_{\theta_l^*}(\mathcal{X}_{r_l})) - \pi_{sem}(\mathcal{P})\|_F, \end{aligned} \quad (5)$$

This provides the theoretical foundation for our Dual SRA mechanism, indicating that a high masking rate is conducive to preserving semantic information, while a low masking rate tends to retain geometric information.

Theorem B: Reconstruction Uncertainty. Traditional MAE methods are based on a point-wise reconstruction assumption, i.e., there exists a unique optimal reconstruction target:

$$\mathcal{L}_{det} = E_{\mathcal{P} \sim p_{data}} [\|\mathcal{P}_M - G_{en}(\mathcal{P}_V)\|^2], \quad (6)$$

where G_{en} is the generative function, and \mathcal{P}_M and \mathcal{P}_V denote the masked and visible regions, respectively. However, point cloud geometric reconstruction inherently exhibits ambiguity. Given visible regions \mathcal{P}_V , the plausible configurations of the masked region form a conditional probability distribution:

$$p(\mathcal{P}_M | \mathcal{P}_V) = \int_{\Theta} p(\mathcal{P}_M | \mathcal{P}_V, \omega) p(\omega | \mathcal{P}_V) d\omega, \quad (7)$$

where Θ denotes geometric parameters such as curvature and density. When the masking ratio lies within a reasonable range, the entropy of the reconstruction distribution satisfies:

$$H(p(\mathcal{P}_M | \mathcal{P}_V)) > 0. \quad (8)$$

Naturally, we adopt MeanFlow to learn a continuous transformation from a noise distribution to the true data distribution, which can handle geometric uncertainty and generate diverse, geometrically consistent reconstructions.

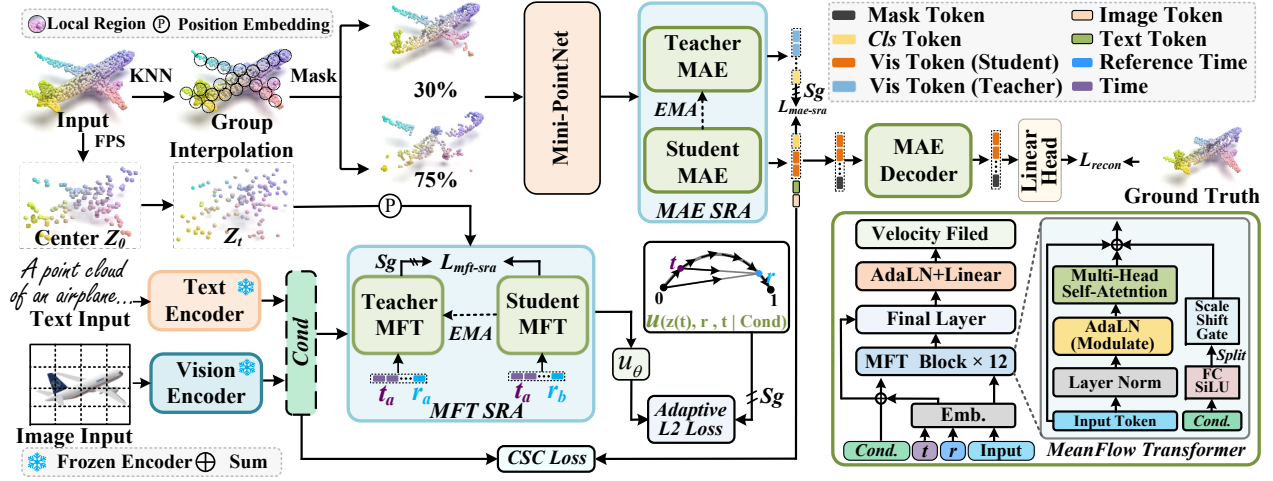


Figure 3: Overview of the Point-SRA. The input point is partitioned via FPS and KNN. Image and text inputs are encoded by pre-trained Vision/Text Transformers. A dual self-representation alignment mechanism adopts a teacher-student MAE. The teacher uses a 30% masking ratio to retain geometry, and the student uses 75% to learn semantics. The teacher is updated via EMA. The MFT reconstructs probabilistic distributions and aligns representations across time. Sg denotes gradient stop.

Methodology

Network Architecture

As illustrated in Figure 3, the proposed Point-SRA framework consists of four tightly integrated components that work collaboratively to enhance 3D representation. The MAE serves as the foundational module, which learns geometric features by reconstructing the masked regions of the input. To address the inherent ambiguity in point cloud geometry, the MFT is designed to model the data distribution through continuous probabilistic trajectories, capturing geometric uncertainty more effectively. Building upon these two modules, the MAE-SRA aligns feature representations obtained under different masking ratios, enabling the fusion of geometric detail and semantic abstraction. Furthermore, the MFT-SRA aligns the probabilistic flow representations across temporal states, allowing the model to capture the dynamic evolution of the distribution.

MAE. A masking operation is first applied to randomly retain a subset of visible points $\mathcal{P}_{vis} = \{p_i | \mathcal{M}(p_i) = 0\}$, while the rest are masked as $\mathcal{P}_{mask} = \{p_i | \mathcal{M}(p_i) = 1\}$. The encoder E_{MAE} processes the visible points to extract their feature representations:

$$h_{vis} = E_{MAE}(\mathcal{P}_{vis}). \quad (9)$$

At the same time, learnable mask tokens are generated for the masked points and concatenated with the visible features to form a complete feature set h_{full} . The decoder D_{MAE} then takes this full representation as input to reconstruct the coordinates of all points $\hat{\mathcal{P}}$:

$$\hat{\mathcal{P}} = D_{MAE}(h_{full}). \quad (10)$$

To improve reconstruction quality, the Chamfer Distance (Fan, Su, and Guibas 2017) is adopted as the reconstruction

loss:

$$\begin{aligned} \mathcal{L}_{recon} = & \frac{1}{|\mathcal{P}|} \sum_{p_i \in \mathcal{P}} \min_{q_j \in \hat{\mathcal{P}}} \|p_i - q_j\|_2^2 \\ & + \frac{1}{|\hat{\mathcal{P}}|} \sum_{q_j \in \hat{\mathcal{P}}} \min_{p_i \in \mathcal{P}} \|q_j - p_i\|_2^2 \end{aligned} \quad (11)$$

where p_i denotes a point from the ground-truth and q_j denotes a point from the predicted point cloud.

Conditional Distribution Modeling with MeanFlow.

Given a point cloud and its associated multi-modal conditional information, including image features f_{img} and text features f_{text} , the conditional representations are first extracted using dedicated image and text encoders. A continuous trajectory $\{z_t\}_{t \in [0,1]}$ is then defined. Where z_0 is the target point cloud and z_1 is sampled from a standard normal distribution. The trajectory is constructed via linear interpolation:

$$z_t = (1 - t) \cdot z_0 + t \cdot z_1. \quad (12)$$

Subsequently, the MeanFlow method is employed to predict the average velocity field along this trajectory. Given a current time step t and a reference time step r with $r < t$, MeanFlow estimates the average velocity from t to r as:

$$u_\theta(z_t, r, t|c) \approx \frac{z_r - z_t}{r - t}, \quad (13)$$

where c denotes the conditional feature vector incorporating multi-modal information, constructed as follows:

$$c = e_t(t) + e_r(r) + W_{img}f_{img} + W_{text}f_{text}, \quad (14)$$

where $e_t(\cdot)$ and $e_r(\cdot)$ are time embedding functions, and W_{img} , W_{text} are projection matrices. The MFT is adopted as the backbone network to predict the mean velocity field, as illustrated in Figure 3.

$$\mathcal{L}_{MFM} = E_{t,r,z_t,z_r,c} [\|u_\theta(z_t, r, t|c) - u_{target}\|^2]. \quad (15)$$

The target velocity u_{target} is derived from theoretical principles, not only considering the basic finite difference $\frac{z_r - z_t}{r - t}$, but also incorporating the time derivative of the instantaneous velocity field:

$$u_{target} = v_t - (t - r) \cdot \frac{d}{dt} v_t(z_t), \quad (16)$$

where $v_t = z_1 - z_0$, and the time derivative $\frac{d}{dt} v_t(z_t)$ is computed using the Jacobian-vector product (JVP):

$$\frac{d}{dt} v_t(z_t) = \nabla_{z_t} v_t(z_t) \cdot v_t(z_t) + \frac{\partial v_t(z_t)}{\partial t}. \quad (17)$$

To stabilize training, we apply an Adaptive L2 loss, which dynamically adjusts the loss weight based on prediction error:

$$\mathcal{L}_{MFM} = E [sg(w) \cdot \|u_\theta - u_{target}\|^2], \quad (18)$$

where $w = \frac{1}{(\|u_\theta - u_{target}\|^2 + \epsilon)^p}$, $sg(\cdot)$ denotes the stop-gradient operation, $p = 1 - \gamma$ is the power exponent, ϵ is a small constant for numerical stability, u_θ is the predicted average velocity field by MFT, and u_{target} is the theoretical target computed via the MeanFlow identity.

Cross-modal Joint Conditioning. The geometric features of the point cloud are extracted via the MAE encoder, while image and text features are obtained as *ImageTransformer* and *TextTransformer*, respectively. These features are propagated across modules as follows:

$$f_{img-token} = W_{img}(f_{cls}), \quad (19)$$

$$f_{text-token} = W_{text}(f_{cls}), \quad (20)$$

where f_{cls} denotes the $[CLS]$ token feature from the MAE encoder, and W_{img} , W_{text} are linear projection matrices. A Cross-modal Semantic Consistency (CSC) Loss is defined to align the point cloud encoder's representations with those from the image and text modalities:

$$\mathcal{L}_{CSC} = \mathcal{L}_{SmoothL1}(f_{img-token}, f_{img}) + \mathcal{L}_{SmoothL1}(f_{text-token}, f_{text}). \quad (21)$$

Dual Self-Representation Alignment. As illustrated in Figure 3, Self-Representation Alignment (SRA) is applied at both the MAE level and the MFT level to facilitate effective knowledge transfer. The teacher model parameters are updated from the student model via Exponential Moving Average (EMA):

$$\theta_{teacher} \leftarrow m \cdot \theta_{teacher} + (1 - m) \cdot \theta_{student}, \quad (22)$$

where $m \in [0, 1]$ is a momentum coefficient controlling the update rate.

MAE-SRA. At the MAE level, a teacher-student architecture is built, where the student model is trained with a high masking ratio, and the teacher model with a low masking ratio. Formally, we apply two masking patterns \mathcal{M}_h and \mathcal{M}_l corresponding to high and low masking ratios, respectively, yielding masked inputs \mathcal{P}_h and \mathcal{P}_l . The student model processes the high-mask input:

$$h_{student} = E_{MAE}(\mathcal{P}_h). \quad (23)$$




Figure 4: Flow-conditioned Fine-Tuning Architecture of Point-SRA. The pretrained MFT (with frozen parameters) takes the center coordinates and outputs flow vectors sampled at different time steps.

The teacher model processes the low-mask input:

$$h_{teacher} = E_{MAE}^{EMA}(\mathcal{P}_l), \quad (24)$$

where E_{MAE}^{EMA} denotes the teacher encoder. The MAE-SRA loss is defined as the cosine similarity loss between the feature representations of the student and teacher models:

$$\mathcal{L}_{mae-sra} = 1 - \frac{h_{student} \cdot h_{teacher}}{|h_{student}| \cdot |h_{teacher}|}. \quad (25)$$

MFT-SRA. At the MFT level, a temporal alignment strategy is designed to align representations at different time points and capture the evolution of the probabilistic flow. Specifically, we select two distinct time steps $t_a > t_b$, where the student model processes the representation at t_a , and the teacher model processes that at t_b . Given the condition c , and the corresponding noisy point cloud states z_{t_a} and z_{t_b} , the feature representations are extracted as:

$$h_{t_a} = F_{MF}(z_{t_a}, t_a, c), h_{t_b} = F_{MF}^{EMA}(z_{t_b}, t_b, c), \quad (26)$$

where F_{MF} and F_{MF}^{EMA} denote the feature extractors of the student and teacher MFT networks, respectively. The MFT-SRA loss is then defined as the cosine similarity loss between the two representations:

$$\mathcal{L}_{mft-sra} = \|h_{t_a} - sg(h_{t_b} + u_\theta(z_{t_b}, t_a, t_b|c) \cdot (t_a - t_b))\|^2. \quad (27)$$

Joint Loss Optimization. The final training objective is formulated as a weighted sum of multiple loss components:

$$\mathcal{L}_{total} = \mathcal{L}_{recon} + \lambda_{flow} \mathcal{L}_{MFM} + \mathcal{L}_{CSC} + \lambda_{mae-sra} \mathcal{L}_{mae-sra} + \lambda_{mft-sra} \mathcal{L}_{mft-sra}, \quad (28)$$

where λ_{flow} is set to 0.5, and both $\lambda_{mae-sra}$ and $\lambda_{mft-sra}$ are set to 0.2.

Flow-Conditioned Fine-Tuning Architecture. As illustrated in Figure 4, during fine-tuning, The pre-trained MFT is employed to compute the flow vector for each point cloud group. Unlike the pre-training stage, fine-tuning utilizes only point cloud geometry and does not rely on image or text modalities:

$$F_{u_\theta} = MFT_{frozen}(Center, t, r), \quad (29)$$

where $Center \in \mathcal{R}^{G \times 3}$ denotes the center coordinates of each point cloud group, and MFT_{frozen} indicates the pre-trained MFT module with frozen parameters. The time parameters t and r are sampled as in pre-training. To project

Method	ScanObjectNN			Params (M)
	BG	ONLY	RS	
<i>Supervised Learning Only</i>				
PointNet (2017a)	73.30	79.20	68.00	3.5
PointNet++ (2017b)	82.30	84.30	77.90	1.5
PointNeXt (2022)	–	–	87.7±0.4	1.4
P2P (2022)	–	–	89.30	195.8
<i>Single-Modal Self-Supervised Representation Learning</i>				
Point-BERT (2022)	87.43	88.12	83.07	22.1
Point-MAE (2022)	90.02	88.29	85.18	22.1
Point-M2AE (2022)	91.22	88.81	86.43	15.3
PointMamba (2024)	94.32	92.60	89.31	12.3
PointDif (2024)	93.29	91.91	87.61	–
Point-FEMAE (2024)	95.18	93.29	90.22	41.5
Point-JEPA (2025)	93.20	91.90	87.60	–
RI-MAE (2025)	91.90	–	–	–
<i>Cross-Modal Self-Supervised Representation Learning</i>				
ACT (2023)	93.29	91.91	88.21	22.1
I2P-MAE(2023)	94.15	91.57	90.11	15.3
ReCon (2023)	95.18	93.29	90.63	44.3
Point-SRA (Ours)	95.53	93.31	90.77	40.1

Table 1: Classification accuracy (%) on three subsets of ScanObjectNN. “–” indicates that the result is not available.

the 3D flow vectors into the feature space, we design a dedicated projection layer:

$$F_{cond} = MLP(F_{u_\theta} + \|F_{u_\theta}\|_2), \quad (30)$$

where $\|F_{u_\theta}\|_2$ is the L2 norm of the flow vector applied element-wise or broadcast to match dimensions. In addition, an adaptive gating mechanism is introduced to control the influence of flow conditioning on the original group features:

$$g = \sigma(MLP_{gate}(F_{cond})), \quad (31)$$

$$H_e = H_g \odot (1 + \alpha g) + \beta F_{cond}, \quad (32)$$

where H_g is the original group feature, $g \in [0, 1]^{G \times d}$ is the learned gating value, and α, β are trainable modulation parameters. The operator \odot denotes element-wise modulation.

Experiments

Object Classification. The pre-trained model is evaluated on the real-world object classification benchmark. ScanObjectNN (Uy et al. 2019) comprises 15K point cloud objects across 15 categories and is divided into three subsets: OBJ_BG (objects with background), OBJ_ONLY (objects only), and PB_T50_RS (objects with background and human-induced perturbations). As shown in Table 1, our method achieves overall accuracies of 95.53%, 93.31%, and 90.77% on OBJ_BG, OBJ_ONLY, and PB_T50_RS respectively, surpassing Point-MAE (Pang et al. 2022) by +5.51%, +5.02%, and +5.59%.

Intracranial Aneurysm Classification and Segmentation. The IntraA (Yang et al. 2020) contains 1,909 vessel segments, including 1,694 healthy segments and 215 aneurysmal segments. The F1-Score (F1) is used as the primary

Method	Classification			Segmentation			
	V%	A%	F1%	IoU%		DSC%	
				V	A	V	A
PointNet(2017a)	93.7	69.5	69.2	74.2	37.8	84.2	49.6
PointNet++(2017b)	98.8	87.3	90.2	93.2	76.2	96.4	84.6
SO-Net(2018)	98.9	83.9	88.5	94.5	81.4	97.1	88.8
PointCNN(2018)	99.0	85.8	90.4	93.6	73.6	96.6	81.4
Point-MAE(2022)	97.9	89.4	93.7	89.1	67.7	93.9	75.6
PointDif(2024)	97.0	63.6	83.1	75.2	33.7	85.1	43.5
ACT(2023)	97.3	84.9	91.4	85.4	57.7	91.7	67.9
ReCon(2023)	99.1	93.9	96.8	95.7	84.7	97.8	91.2
Point-SRA	100	95.1	97.7	96.1	86.9	98.0	92.7

Table 2: Classification and segmentation results for vascular and aneurysm segments on the IntraA.

Method	5-way		10-way	
	10-shot	20-shot	10-shot	20-shot
Point-BERT(2022)	94.6±3.1	96.3±2.7	91.0±5.4	92.7±5.1
MaskPoint(2022)	95.0±3.7	97.2±1.7	91.4±4.0	93.4±3.5
Point-MAE(2022)	96.3±2.5	97.8±1.8	92.6±4.1	95.0±3.0
Point-M2AE(2022)	96.8±1.8	98.3±1.4	92.3±4.5	95.0±3.0
RMRL(2024)	97.2±2.3	98.7±1.2	93.2±3.4	95.6±2.6
PointDif(2024)	96.4±1.8	98.0±1.4	92.3±4.5	95.2±3.0
ACT(2023)	96.8±2.3	98.0±1.4	93.3±4.0	95.6±2.8
ReCon(2023)	97.3±1.9	98.9±1.2	93.3±3.9	95.8±3.0
Point-SRA(Ours)	97.6±2.2	99.0±1.2	93.3±4.5	95.9±2.7

Table 3: Few-shot classification results on ModelNet40 are presented. The average accuracy (%) is reported under both 5-way and 10-way settings with 10-shot and 20-shot configurations.

evaluation metric. V and A represent the classification accuracy for healthy intracranial vessels and aneurysms, respectively. IoU and Dice Similarity Coefficient (DSC) are employed to assess segmentation performance for vessel and aneurysm regions. As shown in Table 2, Point-SRA achieves an F1-Score of 97.7, outperforming the unimodal PointMAE by 4, and the multi-modal ACT (Dong et al. 2023) method by 6.3. In the segmentation task, Point-SRA attains an IoU of 86.9% and a DSC of 92.7% for aneurysms, both significantly surpassing mainstream segmentation baselines.

Few-Shot Classification. Systematic experiments are conducted on ModelNet40 (Wu et al. 2015) under few-shot learning scenarios. As shown in Table 3, under the 5-way configuration, Point-SRA achieves an average accuracy of 97.6% and 99.0% in the 10-shot and 20-shot settings, respectively, outperforming the Point-MAE (Pang et al. 2022) baseline by 1.3% and 1.2%. In the 10-way configuration, Point-SRA reaches 93.3% and 95.9% accuracy in the 10-shot and 20-shot settings, respectively.

3D Object Detection. In the indoor scene detection task on the ScanNetV2 (Dai et al. 2017), a 3D sparse convolu-

Method	Pre Dataset	ScanNet(mIoU%)
VoteNet(2019)	-	35.5
STRL(2021)	ScanNet	38.4
PointContrast(2020)	ScanNet	38.0
DepthContrast(2021)	ScanNet-vid	42.9
Point-BERT(2022)	ScanNet-Medium	38.3
MaskPoint(2022)	ScanNet-Medium	42.1
Point-MAE(2022)	ShapeNet	42.8
PointDif(2024)	ShapeNet	43.7
Point-SRA(Ours)	ShapeNet	47.4

Table 4: Object detection results (AP@50) on the ScanNetV2.

Method	Pre-training	mIoU(%)	mAcc(%)
PointNet(2017a)	✗	41.1	49.0
PointCNN(2018)	✗	57.3	63.9
SegGCN(2020)	✗	63.6	70.4
Pix4Point(2024)	✗	69.6	75.2
MKConv(2023)	✗	67.7	75.1
PointNeXt(2022)	✗	68.5	75.1
Point-BERT(2022)	✓	68.9	76.1
MaskPoint(2022)	✓	68.6	74.2
Point-MAE(2022)	✓	68.4	76.2
PointDif(2024)	✓	70.0	77.1
Point-SRA(Ours)	✓	71.8	79.3

Table 5: Semantic segmentation results on S3DIS Area-5. mIoU and mAcc represent mean Intersection over Union and mean accuracy, respectively.

tional network based on MinkowskiEngine (Choy, Gwak, and Savarese 2019) is adopted as the backbone, and the Point-SRA pre-trained encoder is integrated into the feature extraction pipeline. The ScanNetV2 consists of 1,513 fully scanned indoor scenes, covering 18 common object categories. We follow the standard Average Precision (AP) metric and evaluate under an IoU threshold of 0.5 (AP@50). Table 4 show the Point-SRA pre-trained encoder achieves an AP@50 of 47.4%. This improvement mainly comes from more accurate localization and recognition of objects with complex geometric structures, demonstrating the effectiveness of Point-SRA’s geometric feature representations in modeling spatial complexity.

Indoor Scene Segmentation. The S3DIS (Armeni et al. 2016) consists of 6 large indoor areas with a total of 271 rooms, covering 13 semantic categories and furniture objects. We follow the standard Area-5 evaluation protocol, using Area-5 as the test set and the remaining areas for training. In our experiments, a sparse convolutional network backbone based on MinkowskiEngine (Choy, Gwak, and Savarese 2019) is adopted, keeping the same decoder archi-

	OBJ_BG	OBJ_ONLY	PB_T50_RS
Baseline	90.02	88.29	85.18
MeanFlow	95.18	92.77	90.63
MAE-SRA	95.01	92.77	89.69
MFT-SRA	95.35	92.91	90.01
Point-SRA	95.53	93.31	90.77

Table 6: Ablation study results on core components.

	BG	MN40(1k)	Segpart(C.mIoU)
Transformer	95.01	93.92	84.27
+Linear Proj.	95.11	93.97	84.32
+Gate Fusion	95.35	94.01	84.64
+Full	95.53	94.16	84.88

Table 7: Flow-Conditioned Fine-Tuning Architecture.

ture and segmentation head, while the encoder is replaced with the pre-trained Point-SRA model. During training, the Point-SRA encoder is frozen, and only the decoder and segmentation head are fine-tuned. The AdamW optimizer is used with an initial learning rate of 0.006 and weight decay of 0.05, and training is conducted for 100 epochs. As shown in Table 5, the results demonstrate that Point-SRA achieves 71.8% mIoU and 79.3% mAcc on the S3DIS test set.

Ablation Study

Effectiveness of Core Components. Key components of Point-SRA are progressively removed to quantify each module’s contribution to the final performance. Point-MAE serves as the baseline, with all components utilizing the Flow-Conditioned Fine-Tuning Architecture. Table 6 presents the performance under different component combinations. MeanFlow brings a significant improvement, boosting accuracy on PB_T50_RS by 5.45%, validating the effectiveness of probabilistic reconstruction over point-wise reconstruction. The complete Point-SRA framework achieves substantial gains across all tasks compared to the baseline MAE, demonstrating the effectiveness of component integration and the soundness of the overall design.

Flow-Conditioned Fine-Tuning Architecture. Table 7 presents the specific impact of the flow-conditioned enhancement on downstream performance. We first use a flow-conditioned projection layer to transform the pre-trained MFT’s geometric distribution into a feature-aligned space. Then, a learnable gating fusion mechanism adaptively controls the fusion strength, allowing the network to dynamically leverage pre-trained knowledge. While using projection or gating alone brings limited gains, their combination (FULL) boosts segmentation performance to 84.88%, improving by 0.61% over the baseline.

Comparison of Probabilistic Modeling Methods. To validate the superiority of MeanFlow over other probabilistic generative methods, multiple probabilistic modeling

	OBJ_BG	OBJ_ONLY	PB_T50_RS
Deterministic MAE	90.02	88.29	85.18
+DDPM	93.29	91.91	87.61
+Rectified Flow	94.84	92.60	89.60
+MeanFlow(Ours)	95.18	92.77	90.63

Table 8: Performance comparison of different probabilistic modeling methods.

Loss	Sampling strategy	OBJ_BG
MSE	Uniform Sampling	95.01
	Lognormal Sampling	94.84
Adaptive L2	Uniform Sampling	95.53
	Lognormal Sampling	95.35

Table 9: Comparison of time sampling strategies and loss functions.

techniques are compared under the same experimental settings. Table 8 presents the performance results of different methods. In terms of classification accuracy, MeanFlow achieves the highest accuracy, outperforming diffusion models and other flow matching variants. Regarding Numerical Stability, the time-interval MeanFlow effectively reduces gradient variance during pre-training, resulting in a more stable and reliable training process. Specifically, it satisfies the inequality:

$$\text{Var}[\mathbf{u}_{t,s}(\mathbf{z}_t)] \leq \frac{1}{(s-t)^2} \int_t^s \text{Var}[\mathbf{v}_\tau(\mathbf{z}_\tau)] d\tau, \quad (33)$$

where $\mathbf{v}_\tau(\mathbf{z}_\tau)$ denotes the instantaneous velocity field at time τ . A detailed proof is provided by **Theorem C** in the Supplementary Material.

MFT Hyperparameter. We systematically analyze key design choices of the MFT, including network block, time sampling strategy, and loss function. As shown in Figure 5(a), 12-layer MFT blocks achieve the best trade-off between performance and computational cost. Table 9 compares different time scheduling methods, where uniform sampling over the interval $[0, 1]$ provides the most stable training signals for MeanFlow and outperforms log-normal sampling. Regarding the loss function, the adaptive L2 loss leverages a weighting mechanism to better handle sample-wise loss variations, offering more stable and efficient optimization compared to MSE loss.

Mask Ratio Configuration. Figure 5(b) shows the impact of various mask ratio settings on performance. The optimal configuration uses a 30% mask ratio for the teacher and 75% for the student, achieving the best results. This setup allows the teacher model to retain rich geometric details while the student focuses on learning higher-level semantic abstractions. An approximate mask ratio difference of 0.45 achieves the best balance; smaller differences fail to provide sufficient complementarity, while larger differences increase alignment difficulty and degrade knowledge transfer effectiveness.




Figure 5: MFT and teacher-student hyperparameter analysis.

Related Work

Point-MAE (Pang et al. 2022) and Point-M2AE (Zhang et al. 2022) apply MAE to point cloud with fixed or multi-scale masking, but struggle with geometric diversity and point-wise reconstruction. MaskPoint (Liu, Cai, and Lee 2022) introduces geometry-aware masking, while PTM (Cheng et al. 2024) and Point-FEMAE (Zha et al. 2024) propose density or a parallel masking strategy, though all involve complex tuning and limited semantic understanding. ReCon (Qi et al. 2023) propose tri-modal contrastive learning frameworks to enhance semantic representation capabilities of 3D MAE. Flow Matching (Lipman et al. 2023) enables direct learning of probability flows, and MeanFlow (Geng et al. 2025) improves training stability via average velocity prediction. PointFM (Cheng et al. 2025) applies rectified flow matching to point cloud representation learning but requires careful design of flow predictors and suffers from training instability. Methods like REPA (Yu et al. 2024) demonstrates the benefits of self-distillation and cross-stage alignment. I-DAE (Chen et al. 2025) demonstrates that hidden states of diffusion models can learn discriminative representations. Our work builds on these insights to propose a unified Dual Self-Representation Alignment mechanism across masking ratios and temporal states.

Conclusion

This paper proposes Point-SRA, a self-supervised representation learning framework. Firstly, it reveals for the complementary features of different mask ratios through systematic study. Secondly, we innovatively explore MeanFlow for 3D SSRL by probabilistic modeling and representation guidance. In addition, a dual self-representation alignment mechanism is designed by combining mask ratio alignment and MeanFlow temporal alignment. Finally, we build a dedicated fine-tuning network that leverages geometric distribution knowledge learned during pre-training via flow vector condition fusion. Extensive experiments validate the effectiveness of Point-SRA, achieving strong results on many standard benchmarks. Future work will focus on exploring efficient latent states knowledge distillation methods to further advance 3D SSRL.

Acknowledgements

This paper was funded in part by the Applied Technology Research and Development Project of Beilin District, Xi'an City (Grant No. GX2305), the Natural Science Basic Research Plan of Shaanxi Province, China (Grant No. 2025JC-JCQN-091), and the Technology Innovation Leading Program of Shaanxi Province (Program No. 2024QY-SZX-23).

References

- Armeni, I.; Sener, O.; Zamir, A. R.; Jiang, H.; Brilakis, I.; Fischer, M.; and Savarese, S. 2016. 3d semantic parsing of large-scale indoor spaces. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 1534–1543.
- Chen, X.; Liu, Z.; Xie, S.; and He, K. 2025. Deconstructing denoising diffusion models for self-supervised learning. In *Proc. Int. Conf. Learn. Represent. (ICLR)*.
- Cheng, H.; Wei, L.; Wang, W.; Lu, J.; Xiao, J.; and Zhu, J. 2025. PointFM: Point Cloud Understanding by Flow Matching. *IEEE Rob. Autom. Lett.*
- Cheng, H.; Zhu, J.; Hu, N.; Chen, J.; and Yan, W. 2024. Ptm: Torus masking for 3d representation learning guided by robust and trusted teachers. *IEEE Trans. Circuits Syst. Video Technol.*
- Choy, C.; Gwak, J.; and Savarese, S. 2019. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 3075–3084.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 5828–5839.
- Ding, Z.; Han, X.; and Niethammer, M. 2019. Votenet: A deep learning label fusion method for multi-atlas segmentation. In *Proc. Med. Image Comput. Comput. Assist. Interv. (MICCAI)*, 202–210.
- Dong, R.; Qi, Z.; Zhang, L.; Zhang, J.; Sun, J.; Ge, Z.; Yi, L.; and Ma, K. 2023. Autoencoders as Cross-Modal Teachers: Can Pretrained 2D Image Transformers Help 3D Representation Learning? In *Proc. Int. Conf. Learn. Represent. (ICLR)*.
- Fan, H.; Su, H.; and Guibas, L. J. 2017. A point set generation network for 3d object reconstruction from a single image. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 605–613.
- Geng, Z.; Deng, M.; Bai, X.; Kolter, J. Z.; and He, K. 2025. Mean flows for one-step generative modeling. *arXiv preprint arXiv:2505.13447*.
- Huang, S.; Xie, Y.; Zhu, S.-C.; and Zhu, Y. 2021. Spatio-temporal self-supervised representation learning for 3d point clouds. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 6535–6545.
- Lei, H.; Akhtar, N.; and Mian, A. 2020. Seggcn: Efficient 3d point cloud segmentation with fuzzy spherical kernel. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 11611–11620.
- Li, J.; Chen, B. M.; and Lee, G. H. 2018. So-net: Self-organizing network for point cloud analysis. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 9397–9406.
- Li, Y.; Bu, R.; Sun, M.; Wu, W.; Di, X.; and Chen, B. 2018. Pointcnn: Convolution on x-transformed points. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*.
- Liang, D.; Zhou, X.; Xu, W.; Zhu, X.; Zou, Z.; Ye, X.; Tan, X.; and Bai, X. 2024. Pointmamba: A simple state space model for point cloud analysis. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow matching for generative modeling. In *Proc. Int. Conf. Learn. Represent. (ICLR)*.
- Liu, H.; Cai, M.; and Lee, Y. J. 2022. Masked discrimination for self-supervised learning on point clouds. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 657–675. Springer.
- Pang, Y.; Wang, W.; Tay, F. E.; Liu, W.; Tian, Y.; and Yuan, L. 2022. Masked autoencoders for point cloud self-supervised learning. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 604–621. Springer.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*.
- Qi, Z.; Dong, R.; Fan, G.; Ge, Z.; Zhang, X.; Ma, K.; and Yi, L. 2023. Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. In *Proc. Int. Conf. Mach. Learn. (ICML)*, 28223–28243. PMLR.
- Qian, G.; Hamdi, A.; Zhang, X.; and Ghanem, B. 2024. Pix4Point: Image Pretrained Standard Transformers for 3D Point Cloud Understanding. In *2024 International Conference on 3D Vision (3DV)*, 1280–1290.
- Qian, G.; Li, Y.; Peng, H.; Mai, J.; Hammoud, H.; Elhoseiny, M.; and Ghanem, B. 2022. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 23192–23204.
- Saito, A.; Kudeshia, P.; and Poovvancheri, J. 2025. Pointjepa: A joint embedding predictive architecture for self-supervised learning on point cloud. In *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, 7348–7357. IEEE.
- Su, K.; Wu, Q.; Cai, P.; Zhu, X.; Lu, X.; Wang, Z.; and Hu, K. 2025. RI-MAE: Rotation-Invariant Masked AutoEncoders for Self-Supervised Point Cloud Representation Learning. In *Proc. Conf. Artif. Intell. (AAAI)*, volume 39, 7015–7023.
- Uy, M. A.; Pham, Q.-H.; Hua, B.-S.; Nguyen, T.; and Yeung, S.-K. 2019. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 1588–1597.
- Wang, C.; Zha, Y.; He, J.; Yang, W.; and Zhang, T. 2024. Rethinking Masked Representation Learning for 3D Point Cloud Understanding. *IEEE Trans. Image Process.*

Wang, Z.; Yu, X.; Rao, Y.; Zhou, J.; and Lu, J. 2022. P2p: Tuning pre-trained image models for point cloud analysis with point-to-pixel prompting. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 14388–14402.

Woo, S.; Lee, D.; Hwang, S.; Kim, W. J.; and Lee, S. 2023. MKConv: Multidimensional feature representation for point cloud analysis. *Pattern Recognit.*, 143: 109800.

Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 1912–1920.

Xie, S.; Gu, J.; Guo, D.; Qi, C. R.; Guibas, L.; and Litany, O. 2020. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 574–591. Springer.

Yang, X.; Xia, D.; Kin, T.; and Igarashi, T. 2020. Intra: 3d intracranial aneurysm dataset for deep learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2656–2666.

Yu, S.; Kwak, S.; Jang, H.; Jeong, J.; Huang, J.; Shin, J.; and Xie, S. 2024. Representation alignment for generation: Training diffusion transformers is easier than you think. In *Proc. Int. Conf. Learn. Represent. (ICLR)*.

Yu, X.; Tang, L.; Rao, Y.; Huang, T.; Zhou, J.; and Lu, J. 2022. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 19313–19322.

Zha, Y.; Ji, H.; Li, J.; Li, R.; Dai, T.; Chen, B.; Wang, Z.; and Xia, S.-T. 2024. Towards compact 3d representations via point feature enhancement masked autoencoders. In *Proc. Conf. Artif. Intell. (AAAI)*, volume 38, 6962–6970.

Zhang, R.; Guo, Z.; Gao, P.; Fang, R.; Zhao, B.; Wang, D.; Qiao, Y.; and Li, H. 2022. Point-m2ae: multi-scale masked autoencoders for hierarchical point cloud pre-training. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 27061–27074.

Zhang, R.; Wang, L.; Qiao, Y.; Gao, P.; and Li, H. 2023. Learning 3d representations from 2d pre-trained models via image-to-point masked autoencoders. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 21769–21780.

Zhang, Z.; Girdhar, R.; Joulin, A.; and Misra, I. 2021. Self-supervised pretraining of 3d features on any point-cloud. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 10252–10263.

Zheng, X.; Huang, X.; Mei, G.; Hou, Y.; Lyu, Z.; Dai, B.; Ouyang, W.; and Gong, Y. 2024. Point cloud pre-training with diffusion models. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 22935–22945.