

StruMamba3D: Exploring Structural Mamba for Self-supervised Point Cloud Representation Learning

Chuxin Wang^{1,2}, Yixin Zha¹, Wenfei Yang^{1,2,*}, Tianzhu Zhang^{1,2}

¹University of Science and Technology of China

²National Key Laboratory of Deep Space Exploration, Deep Space Exploration Laboratory
{wxc0602, zyxcn}@mail.ustc.edu.cn, {yangwf, tzhang}@ustc.edu.cn

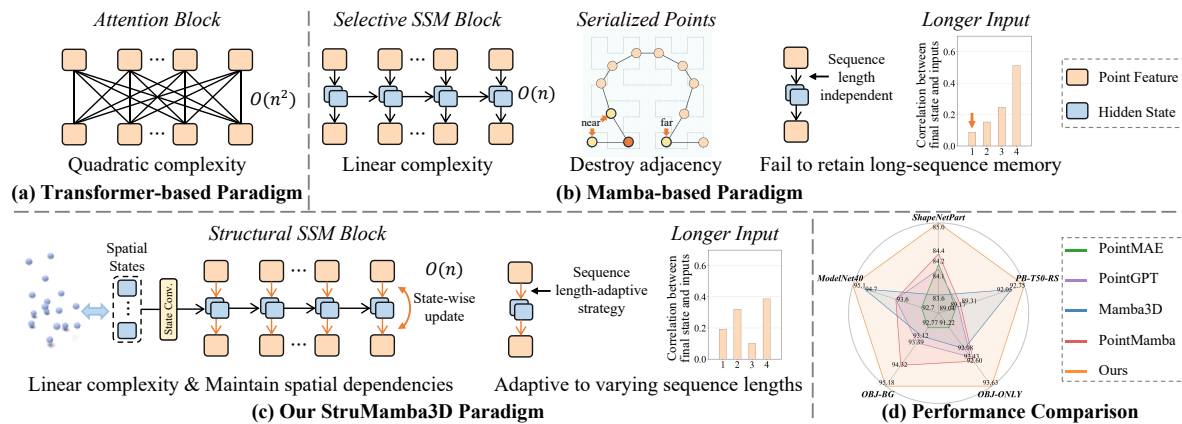


Figure 1. (a): Transformer-based paradigm uses the attention block with quadratic complexity to model the dependencies between points. (b): Mamba-based paradigm uses the selective SSM with linear complexity. However, serialized points destroy the adjacency of 3D points, and the pre-trained selection mechanism fails to retain long-sequence memory. (c): Our StruMamba3D paradigm uses the structural SSM to maintain the spatial dependencies among points and the sequence length-adaptive strategy to retain long-sequence memory. (d): Our StruMamba3D significantly outperforms existing Transformer-based and Mamba-based methods.

Abstract

Recently, Mamba-based methods have demonstrated impressive performance in point cloud representation learning by leveraging State Space Model (SSM) with the efficient context modeling ability and linear complexity. However, these methods still face two key issues that limit the potential of SSM: Destroying the adjacency of 3D points during SSM processing and failing to retain long-sequence memory as the input length increases in downstream tasks. To address these issues, we propose StruMamba3D, a novel paradigm for self-supervised point cloud representation learning. It enjoys several merits. First, we design spatial states and use them as proxies to preserve spatial dependencies among points. Second, we enhance the SSM with a state-wise update strategy and incorporate a lightweight convolution to facilitate interactions between spatial states for efficient structure modeling. Third, our method reduces the sensitivity of pre-trained Mamba-based models to varying input lengths by introducing a sequence length-adaptive

strategy. Experimental results across four downstream tasks showcase the superior performance of our method. In addition, our method attains the SOTA 95.1% accuracy on ModelNet40 and 92.75% accuracy on the most challenging split of ScanObjectNN without voting strategy.

1. Introduction

Point cloud representation learning aims to extract geometric and semantic features from point clouds to support downstream tasks such as classification, segmentation, and detection. As a foundational task in Computer Vision (CV), it is crucial for various real-world applications, including autonomous driving [4, 20], AR [13], and robotics [26, 32]. However, unlike images or text with structured arrangements, point clouds are inherently unordered and sparse, presenting challenges for designing effective models.

To solve this problem, numerous point cloud representation learning methods have been proposed, which can be broadly divided into three categories according to model architectures: Point-based methods [19, 27, 28, 30, 31,

*Corresponding author

35], Transformer-based methods [12, 22, 25, 42–44], and Mamba-based methods [15, 21, 47]. Point-based methods directly process point clouds by modeling the relationships between individual points, such as through set abstraction [27, 28] and graph convolution [19, 35]. These methods offer a simple and computationally efficient design. However, their limited receptive field hinders capturing global information, constraining complex structure modeling. Transformer-based methods leverage attention mechanisms to capture long-range dependencies, expanding the receptive field. Moreover, inspired by self-supervised learning paradigms in NLP [6] and CV [16], many transformer-based methods [22, 25, 42, 44] utilize large unlabeled point cloud datasets for pre-training, achieving impressive results in downstream tasks. However, these methods have limited scalability due to the quadratic complexity of the attention mechanism, as shown in Fig. 1(a). To overcome this limitation, Mamba-based methods [15, 21, 47] first convert the 3D points into a 1D sequence and then utilize selective State Space Model (SSM) [9] with linear complexity to extract point features. Leveraging the robust information aggregation ability of selective SSM, Mamba-based methods have attained promising performance.

Although existing Mamba-based methods [15, 21, 47] have achieved comparable performance to transformer-based methods, two key issues limit the potential of SSM: **(1) Distort the adjacency of points.** As shown in Fig. 1(b), although methods [21, 47] use serialization strategies to align input points, they cannot ensure that spatially adjacent points remain neighbors in the 1D sequence. Unlike text and images, points inherently lack context information and rely on local structures as base feature units. This spatial distortion hampers Mamba to model fine-grained structure information. **(2) Fail to retain long-sequence memory in downstream tasks.** Mamba-based methods [15, 21] use Masked Point Modeling (MPM) [25, 42] as a pre-training task, which creates a significant discrepancy in input sequence lengths between pre-training and downstream tasks. Mamba leverages a selection mechanism to inductively learn new information while selectively forgetting outdated knowledge. However, the selection mechanism does not account for variations in input sequence length. Consequently, when Mamba is pre-trained on short sequences, it tends to update states more frequently. As the input sequence length increases in downstream tasks, frequent state updates lead to Mamba struggling to retain long-sequence memory, as shown in Fig. 1(b). The deficiency in long-sequence memory adversely affects the long-range semantics modeling, resulting in suboptimal performance.

To unleash the potential of Mamba for point cloud representation learning, we explore solutions to address the above two issues. **(1) How to maintain spatial dependencies among points?** Maintaining spatial relationships

among 3D points within a 1D input sequence is challenging. Mamba efficiently aggregates contextual features from the 1D input sequence via hidden states. Is it possible to design a structural Mamba that models spatial relationships among points through hidden states? **(2) How to retain long-sequence memory for longer inputs?** When handling longer input sequences, states need to retain memory for an extended period to effectively capture the contextual information across the entire sequence. If the model adjusts the frequency of state updates when the sequence length changes, it can preserve memory for longer sequences.

To achieve the above goals, we propose a novel StruMamba3D paradigm for self-supervised point cloud representation learning, which consists of two core components: structural SSM block and sequence length-adaptive strategy. **In structural SSM block,** we leverage states in SSM to model the spatial relationships among points. We begin by assigning positional attributes to the states, using them as proxies to capture the local structure of point clouds. Next, we propose a state-wise update strategy that incorporates the positional relationship between input points and spatial states during the state update and propagation process. Based on spatial states and state-wise update strategy, we can maintain spatial dependencies among points during the SSM process. To address the limitation of standard Mamba ignoring inter-state interactions, we introduce a lightweight convolution module for spatial states. **In sequence length-adaptive strategy,** the adaptive state update mechanism is first employed to adjust the state update frequency based on sequence length, ensuring that the states retain long-sequence memory as the sequence length grows. Additionally, we introduce a spatial state consistency loss in the pre-training task, which enforces consistency in the updated states across inputs of varying sequence lengths. The spatial state consistency loss enhances the robustness of the Mamba-based pre-trained model to varying input lengths.

To sum up, the main contributions of this work are: (1) We propose the novel StruMamba3D for self-supervised point cloud representation learning, which is the first to model the structure information by latent states in SSM. (2) We design the structural SSM to maintain spatial dependencies among points during state information selection and propagation, while the sequence length-adaptive strategy reduces the sensitivity of pre-trained Mamba-based models to varying input lengths. (3) Experimental results on four downstream tasks showcase the superior performance of our method. Additionally, comprehensive ablation studies validate the effectiveness of our designs.

2. Related Work

We briefly review methods for point cloud representation learning and self-supervised representation learning.

Point Cloud Representation Learning. To effectively

leverage the sparse, unordered and irregular point cloud, various point cloud representation learning methods have been proposed, which can be broadly categorized into three groups: Point-based methods [19, 27, 28, 30, 31, 35], Transformer-based methods [12, 22, 25, 33, 42, 44], and Mamba-based methods [15, 21, 47]. Point-based methods directly process raw point clouds to learn geometric features. For example, PointNet [27] employs MLPs to each point independently, followed by global max pooling to aggregate features. PointNet++ [28] extends PointNet by introducing a hierarchical architecture for multi-scale feature extraction. Subsequent methods [19, 24, 30, 31, 35] explore convolutional techniques for point clouds. DGCNN [35] leverages graph convolution to model structural information among neighboring points, while KPConv [31] introduces a convolution method based on dynamic kernel points to improve feature learning. To capture long-range dependencies, methods [12, 37, 38, 49] have explored transformer architectures for point clouds. PCT [12] applies global attention to points, while Point Transformer [49] introduces local attention to capture local geometry features. Recent variants [37, 38] further refine the architectures to improve performance and efficiency. These supervised methods [12, 30, 31, 47, 49] have demonstrated outstanding performance in specific tasks. However, these methods are inherently limited by the specific data domains on which they are trained, leading to poor generalization to other domains.

Self-supervised Representation Learning. To solve the above issues and fully utilize a large number of unlabeled data, self-supervised representation learning methods [3, 22, 25, 40, 42, 44, 48] have emerged. Some contrastive learning-based methods [40, 48] train models to learn distinguishing features by differentiating between positive and negative samples. For instance, Point-Contrast [40] leverages feature consistency across different views of the same point cloud. Inspired by the success of BERT [6] and MAE [16], many transformer-based methods [3, 22, 25, 42, 44] adopt MPM task for pre-training, encouraging models to infer masked patches from visible ones. PointBERT [42] predicts discrete tokens for masked points, while PointMAE [25] reconstructs masked point clouds. Additionally, recent studies [29, 45] integrate cross-modal information, such as images and text, to enhance pre-training. However, transformer-based methods suffer from quadratic complexity, posing challenges for handling long sequences and making them unsuitable for resource-constrained devices. Recently, SSMs [5, 9–11, 17, 34] have emerged as a prominent research focus due to their linear complexity and powerful information aggregation ability. In the field of point cloud representation learning, several Mamba-based methods [15, 21, 47] have been proposed. PointMamba [21] employs space-filling curves to serialize point clouds, establishing a simple yet effective base-

line. PCM [47] introduces multiple serialization strategies and a bidirectional scanning mechanism to capture comprehensive point cloud structures, but also introduces larger computational overhead. Mamba3D [15] also leverages a bidirectional scanning mechanism and proposes a local norm pooling operation to enhance local geometric features. However, these methods introduce structural distortions when converting 3D points into 1D sequences, limiting model performance, especially in tasks requiring fine-grained geometry like part segmentation. Moreover, they neglect the impact of varying input lengths between pre-training and downstream tasks. Unlike these methods, we propose StruMamba3D, which leverages states as proxies to maintain spatial dependencies among points. Additionally, we introduce a sequence-length adaptive strategy to ensure that the states retain long-sequence memory as the sequence length grows in downstream tasks.

3. Method

Below, we review the previous SSMs (Sec. 3.1), followed by an overview of our StruMamba3D (Sec. 3.2). Then, we introduce the details of Structural SSM Block (Sec. 3.3) and Sequence Length-adaptive Strategy (Sec. 3.4).

3.1. Preliminaries

State Space Model (SSM) is a continuous system that maps the input $x(t)$ to $y(t)$ through implicit states $h(t) \in \mathbb{R}^M$:

$$h'(t) = \mathbf{A}h(t) + \mathbf{B}x(t), y(t) = \mathbf{C}h'(t), \quad (1)$$

where the matrix \mathbf{A} governs the evolution of states, while the matrices \mathbf{B} and \mathbf{C} are used for modeling the relationship between inputs and outputs through the latent states. To handle discrete-time sequence data, the zero-order hold discretization rule is used:

$$h_t = \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, y_t = \mathbf{C}h_t, \quad (2)$$

$$\bar{\mathbf{A}} = \exp(\Delta\mathbf{A}), \bar{\mathbf{B}} = (\Delta\mathbf{A})^{-1}(\exp(\Delta\mathbf{A}) - \mathbf{I})\Delta\mathbf{B},$$

where Δ represents the sampling interval between consecutive inputs x_{t-1} and x_t . However, due to the linear time-invariant nature of SSM, the parameters $(\Delta, \mathbf{A}, \mathbf{B}, \mathbf{C})$ remain fixed across all time steps, which limits the expressive capacity of SSM. To overcome this limitation, Mamba [9] treats the parameters $(\Delta, \mathbf{B}, \mathbf{C})$ as functions of input x_t , transforming the SSM into a time-varying model:

$$h_t = \phi_{\bar{\mathbf{A}}}(x_t)h_{t-1} + \phi_{\bar{\mathbf{B}}}(x_t)x_t, y_t = \phi_{\mathbf{C}}(x_t)h_t, \quad (3)$$

where $\phi_*(x_t)$ denotes the linear projection of input x_t . Additionally, Mamba introduces a hardware-aware scan algorithm to achieve near-linear complexity. Our StruMamba3D adopts the same hardware-aware scan algorithm, maintaining linear complexity while exhibiting strong structural modeling capabilities for point clouds.

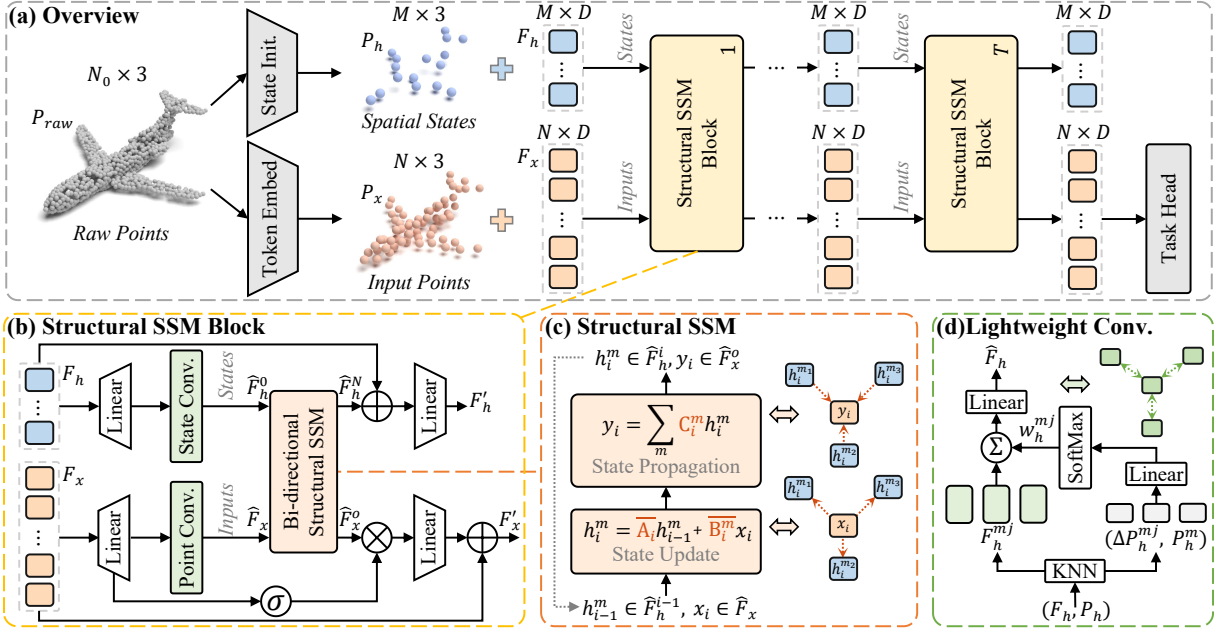


Figure 2. **Overview of StruMamba3D.** (a): Architecture of StruMamba3D. (b): Details of Structural SSM Block. We employ structural SSM to capture dependencies between spatial states and input points. Additionally, we integrate a lightweight convolution module to introduce local interactions among spatial states or input points. (c): Details of Structural SSM. (d): Details of Lightweight Convolution.

3.2. Overview

Despite Mamba achieving remarkable results on sequential data, maintaining the spatial relationships among 3D points remains challenging. We aim to model the spatial relationships by latent states and design StruMamba3D for point cloud representation learning, as illustrated in Fig. 2(a). Given raw point cloud $P_{raw} \in \mathbb{R}^{N_0 \times 3}$, we obtain the input points $P_x \in \mathbb{R}^{N \times 3}$ using farthest point sampling (FPS) and K -nearest neighbor clustering (KNN). Then, we use a lightweight PointNet [27] to extract the input token embeddings $F_x \in \mathbb{R}^{N \times D}$. Meanwhile, we initialize the spatial states $P_h \in \mathbb{R}^{M \times 3}$ as the proxies for local structures from the raw point cloud. Finally, the spatial states and input token embeddings are fed into a feature encoder composed of multiple structural SSM blocks, where the spatial states act as a bridge to connect the inputs and outputs. Notably, we use the output spatial states from the previous structural SSM block as the input spatial states for the current block.

3.3. Structural SSM Block

The structural SSM block is designed to capture complex structural features by preserving spatial dependencies among points throughout SSM processing. As shown in Fig. 2(b), the block takes spatial states F_h and token embeddings F_x as input and produces updated spatial states F'_h and token embeddings F'_x as output. To clarify the details of structural SSM block, we describe each component in sequence: Spatial State Initialization, State-wise Update Mechanism, Structural SSM, and Lightweight Convolution. **Spatial State Initialization.** In original Mamba, latent

states do not contain geometric information, making it difficult to model the local structure of point clouds. To address this, we introduce positional attributes $P_h \in \mathbb{R}^{M \times 3}$ to the latent states, enabling each state to focus on distinct regions of the point cloud. Due to the sparsity of point clouds, we initialize state positions P_h based on raw points P_{raw} to ensure comprehensive coverage of point clouds. Specifically, we employ FPS and KNN clustering to segment the point cloud into groups $\{\mathcal{G}_m\}_{m=1}^M$, calculating the centroid:

$$P_h^m = \frac{1}{|\mathcal{G}_m|} \sum_{P_i \in \mathcal{G}_m} P_i, \quad (4)$$

where P_h^m represents the positional attribute of the m -th state. We refer to states with positional attributes as spatial states and then use linear mapping ϕ_h to embed positional attributes as the initial state features.

$$F_h = \phi_h(P_h) \in \mathbb{R}^{M \times D}. \quad (5)$$

The state features F_h are then used as the initial latent states in SSM and updated along with input token embeddings.

State-wise Update. As illustrated in Eq. (3), the state update and propagation equations of the original Mamba rely solely on input tokens, disregarding the spatial relationships between input points and spatial states. This limitation makes it challenging for Mamba to effectively model the structural information of point clouds. To address this, we modify the state update and propagation equations to explicitly model the spatial dependencies between states and inputs. Specifically, we first calculate the relative offsets

between the input points P_x and the state positions P_h :

$$\Delta P_i^m = P_x^i - P_h^m, \quad (6)$$

where ΔP_i^m denotes the relative offset between the i -th input point and the m -th state. Next, we map these offsets to the SSM parameters (\mathbf{B}, \mathbf{C}) , which govern state update and propagation. It is important to note that considering only spatial relationships is insufficient, as the significance of information varies across different points. Thus, we design the following state update and propagation mechanism:

$$(\mathbf{B}_i^m, \mathbf{C}_i^m) = \phi(x_i) + \text{MLP}(\Delta P_i^m), \quad (7)$$

where MLP represents a multi-layer perceptron, ϕ is a linear projection and x_i denotes the feature of the i -th input point. Based on the modified parameters, the spatial states can selectively update features using points within the same regions, while the input points can acquire local structural information from the states, as shown in Fig. 2(c).

Structural SSM. By leveraging spatial states and state-wise updates, we effectively model the structural information of point clouds. However, due to unidirectional scanning mechanism, bidirectional information exchange between input points cannot be achieved within a single forward pass. To address this, we adopt a bidirectional scanning mechanism inspired by VisionMamba [25]. Specifically, we apply both forward and backward structural SSMs, with the backward SSM processing input points in reverse order. Both use spatial states \hat{F}_h as initial states \hat{F}_h^0 , generating updated features for input points and spatial states. Finally, we fuse these outputs using a linear layer ϕ_o to obtain the final features F'_x and F'_h for input points and spatial states.

$$F'_x, F'_h = \phi_o(\text{SSM}_f(\hat{F}_x, \hat{F}_h) + \text{SSM}_b(\hat{F}_x, \hat{F}_h)), \quad (8)$$

where f and b denote the forward and backward processes.

Lightweight Convolution. The proposed structural SSM effectively models the dependencies between spatial states P_h and input points P_x . However, we observe that the spatial states P_h in SSMs remain isolated, lacking direct interaction with each other. To address this, we introduce a lightweight convolution module, as shown in Fig. 2(d). Inspired by graph convolution [35], this module aggregates features from neighboring states based on relative position information. Specifically, for each spatial state P_h^m , we first identify its k -nearest neighbors $\mathcal{N}(m)$ and compute the relative positions $\Delta P_h^{mj}, j \in \mathcal{N}(m)$. We then generate attention weights $\{w_{mj}\}$ for each neighbor using a linear layer ϕ_w , followed by a softmax operation:

$$w_h^{mj} = \text{softmax}(\phi_w(\Delta P_h^{mj}, P_h^m)). \quad (9)$$

The output state feature \hat{F}_h^m for each state P_h^m is obtained through weighted aggregation:

$$\hat{F}_h^m = \phi_c\left(\sum_{j \in \mathcal{N}(m)} w_h^{mj} F_h^{mj}\right), \quad (10)$$

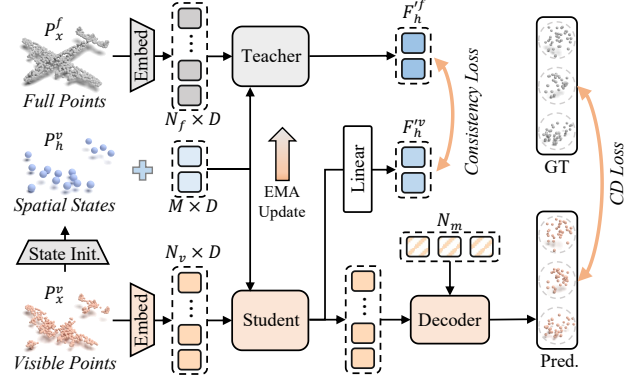


Figure 3. **Overview of the pre-training framework.** We pre-train the StruMamba3D (student model) using the MPM task and the spatial state consistency loss. The teacher model has the same structure as the student model, but its parameters are updated using the exponential moving average strategy.

where F_h^{mj} denotes the neighboring state features and ϕ_c is a linear layer. By extending the receptive field of spatial states, this lightweight convolution enhances their ability to capture global semantic information. Furthermore, we replace the causal 1D convolution in the original Mamba block with this module, making it better suited for processing input points with complex geometric structures.

3.4. Sequence Length-adaptive Strategy

To address the issue of insufficient long-sequence memory in downstream tasks, we propose a sequence length-adaptive strategy that integrates an adaptive state update mechanism and spatial state consistency loss.

Adaptive State Update Mechanism. According to Eq. (2), the parameter sampling interval Δ controls the frequency of state updates by influencing the parameters $(\bar{\mathbf{A}}, \bar{\mathbf{B}})$ in SSMs. A larger Δ causes the model to update state features more frequently, whereas a smaller Δ allows the model to maintain its current state with minimal changes. To preserve memory for longer sequences in downstream tasks, we propose an adaptive state update mechanism. This mechanism adjusts the sampling interval Δ according to the input sequence length, allowing the model to maintain a consistent total sampling time $\Delta_{all} = \sum_{i=1}^N \Delta_i$ across different sequence lengths. Specifically, we design a learnable parameter τ , which regulates the total sampling time Δ_{all} for sequences of varying lengths. The sampling interval Δ_i for each token in the inputs is calculated as follows:

$$\Delta_i = \frac{\tau \times \Delta_i}{\sum_{i=1}^N \Delta_i}. \quad (11)$$

Based on the adaptive state update mechanism, the model can maintain a consistent total sampling time $\Delta_{all} = \tau$ across different sequence lengths.

Spatial State Consistency Loss. To further ensure consistency in updated state features across different input se-

Table 1. **Shape Classification on ScanObjectNN and ModelNet40 Datasets.** For ScanObjectNN, we report the classification accuracy(%) over the three subsets: OBJ-BG, OBJ-ONLY, and PB-T50-RS. And † indicates using simple rotational augmentation [7] for fine-tuning.

| Method | Backbone | Param. (M) | FLOPs (G) | ScanObjectNN | | | MODELNET40 | |
|--|----------------|------------|-----------|--------------|--------------|--------------|-------------|-------------|
| | | | | OBJ-BG | OBJ-ONLY | PB-T50-RS | w/o Voting | w/ Voting |
| <i>Supervised learning only</i> | | | | | | | | |
| PointNet [27] | - | 3.5 | 0.5 | 73.3 | 79.2 | 68.0 | 89.2 | - |
| PointNet++ [28] | - | 1.5 | 1.7 | 82.3 | 84.3 | 77.9 | 90.7 | - |
| PointCNN [19] | - | 0.6 | - | 86.1 | 85.5 | 78.5 | 92.2 | - |
| DGCNN [35] | - | 1.8 | 2.4 | 82.8 | 86.2 | 78.1 | 92.9 | - |
| PRANet [8] | - | 2.3 | - | - | - | 81.0 | 93.7 | - |
| PointNeXt [30] | - | 1.4 | 3.6 | - | - | 87.7 | 94.0 | - |
| PointMLP [24] | - | 12.6 | 31.4 | - | - | 85.4 | 94.5 | - |
| DeLA [2] | - | 5.3 | 1.5 | - | - | 88.6 | 94.0 | - |
| PCM [47] | - | 34.2 | 45.0 | - | - | 88.1 | 93.4 | - |
| <i>Pre-training using single-modal information</i> | | | | | | | | |
| PointBERT [42] | Transformer | 22.1 | 4.8 | 87.43 | 88.12 | 83.07 | 92.7 | 93.2 |
| MaskPoint [22] | Transformer | 22.1 | 4.8 | 89.30 | 88.10 | 84.30 | - | 93.8 |
| PointM2AE [44] | Transformer | 12.7 | 7.9 | 91.22 | 88.81 | 86.43 | 92.9 | 93.4 |
| PointMAE† [25] | Transformer | 22.1 | 4.8 | 92.77 | 91.22 | 89.04 | 92.7 | 93.8 |
| PointGPT-S† [3] | Transformer | 29.2 | 5.7 | 93.39 | 92.43 | 89.17 | 93.3 | 94.0 |
| PointMamba† [21] | Mamba | 12.3 | 3.1 | 94.32 | 92.60 | 89.31 | 93.6 | 94.1 |
| Mamba3D† [15] | Mamba | 16.9 | 3.9 | 93.12 | 92.08 | 92.05 | 94.7 | 95.1 |
| Ours† | Structural SSM | 15.8 | 4.0 | 95.18 | 93.63 | 92.75 | 95.1 | 95.4 |
| <i>Pre-training using cross-modal information</i> | | | | | | | | |
| ACT† [7] | Transformer | 22.1 | 4.8 | 93.29 | 91.91 | 88.21 | 93.7 | 94.0 |
| Joint-MAE [14] | Transformer | 22.1 | - | 90.94 | 88.86 | 86.07 | - | 94.0 |
| I2P-MAE† [46] | Transformer | 15.3 | - | 94.15 | 91.57 | 90.11 | 93.7 | 94.1 |
| ReCon† [29] | Transformer | 43.6 | 5.3 | 95.18 | 93.29 | 90.63 | 94.5 | 94.7 |

quence lengths, we introduce a spatial state consistency loss to the pre-training task. As shown in Fig. 3, we present an overview of the pre-training framework. Following the previous works [15, 21, 25], we reconstruct the points in the masked regions and use the reconstruction loss \mathcal{L}_{cd} to guide the model to learn discriminative features from the visible points. For the spatial state consistency loss, we first use a teacher model to update the spatial states with the full input tokens $P_x^f \in \mathbb{R}^{N_f \times 3}$. The state features $F_h'^f$ output by the teacher model serve as pseudo-labels. We then enforce a consistency loss \mathcal{L}_{ssc} between the updated state features of the student model $F_h'^v$ and the pseudo-labels $F_h'^f$:

$$\mathcal{L}_{ssc} = \text{Smooth L1}(F_h'^v, F_h'^f). \quad (12)$$

The spatial state consistency loss \mathcal{L}_{ssc} not only enhances the robustness of our StruMamba3D to varying input sequence lengths but also improves the ability to infer the structural information of the complete point cloud from incomplete point clouds. The final pre-training loss is:

$$\mathcal{L}_{total} = \mathcal{L}_{cd} + \lambda \times \mathcal{L}_{ssc}, \quad (13)$$

where λ is the weight of the spatial state consistency loss.

4. Experimental Results

4.1. Implementation Details

Following the previous works [15, 21, 25], we employ $T=12$ structural SSM blocks, each with a feature dimension of $D=384$. The number of spatial states M is set to 16. For the lightweight convolution in the structural SSM block, we set the number of neighbors to 4 for spatial states and 8 for input points. Following previous works [25, 42, 44], we pre-train the proposed model on ShapeNet dataset [1], which contains 52472 unique 3D models across 55 common object categories. Each input raw point cloud, containing $N_0=1024$ points, is divided into 64 patches with each consisting of 32 points. A random masking ratio of 0.6 is applied, with the decoder composed of 4 standard Mamba blocks, and the λ in the pre-training loss is set to 2. More details can be found in the supplementary material.

4.2. Comparison on Downstream Tasks

We present the experimental results of our StruMamba3D across various downstream tasks and compare them with previous state-of-the-art (SOTA) methods.

Shape classification. ScanObjectNN [41] is a real-world dataset containing roughly 15000 objects across 15 categories, some of which have cluttered backgrounds or partial scans. For this dataset, we apply rotation as a data aug-

Table 2. **Part Segmentation on ShapeNetPart.** We report the mean IoU for all classes (mIoU_c) and all instances (mIoU_i).

| Method | Architecture | mIoU _c | mIoU _i |
|--|--------------|-------------------|-------------------|
| <i>Supervised learning only</i> | | | |
| PointNet [27] | Single-scale | 80.4 | 83.7 |
| PointNet++ [28] | Multi-scale | 81.9 | 85.1 |
| APES [36] | Multi-scale | 83.7 | 85.8 |
| DeLA [2] | Multi-scale | 85.8 | 87.0 |
| PCM [47] | Multi-scale | 85.3 | 87.0 |
| <i>Pre-training using single-modal information</i> | | | |
| MaskPoint [22] | Single-scale | 84.6 | 86.0 |
| PointBERT [42] | Single-scale | 84.1 | 85.6 |
| PointMAE [25] | Single-scale | 84.2 | 86.1 |
| PointM2AE [44] | Multi-scale | 84.9 | 86.5 |
| PointGPT-S [3] | Single-scale | 84.1 | 86.2 |
| PointMamba [21] | Single-scale | 84.4 | 86.2 |
| Mamba3D [15] | Single-scale | 83.6 | 85.6 |
| Ours | Single-scale | 85.0 | 86.7 |

mentation strategy [7] and use 2048 points as input. ModelNet40 [39] is a synthetic dataset comprising of 12311 clean 3D CAD models from 40 categories. For this dataset, we use 1024 points as input and apply scale and translation augmentations [27]. As shown in Tab. 1, we conduct classification experiments on both datasets. Compared to transformer-based methods, our model achieves significant performance improvements across multiple metrics on both datasets. Even when compared to the best PointGPT-S [3], our model outperforms it by 3.58% on ScanObjectNN and 1.8% on ModelNet40. Additionally, on the ScanObjectNN dataset, StruMamba3D surpasses PointMamba [21] by 0.86%, 1.03%, and 3.44% across the three splits, while also outperforming Mamba3D [15] by 2.06%, 1.55%, and 0.70%. Furthermore, on the ModelNet40 dataset, StruMamba3D exceeds PointMamba by 1.5% and outperforms Mamba3D by 0.4%. Overall, these results demonstrate that StruMamba3D outperforms all existing unimodal architectures, achieving multiple SOTA results and highlighting the effectiveness of structural modeling within Mamba.

Part Segmentation on ShapeNetPart. ShapeNetPart [41] dataset comprises 16880 models spanning 16 shape categories and annotated with 50 part labels. Compared to classification tasks, part segmentation poses a greater challenge, as it requires assigning a label to each point, making it highly dependent on the intrinsic structural information of point clouds. Tab. 8 presents our part segmentation results for all classes and instances on the ShapeNetPart dataset. For single-scale models, our method achieves substantial performance improvements, surpassing Mamba3D [15] by 1.4% in mIoU_c and 1.2% in mIoU_i, and outperforming PointMamba [21] by 0.6% in mIoU_c and 0.5% in mIoU_i. Notably, Mamba3D without a serialization strategy performs poorly, falling behind the baseline PointMAE [25], whereas PointMamba, which incorporates serialization, exhibits improved performance. This observation suggests the crucial role of serialization in enhancing structural mod-

Table 3. **Few-Shot Classification on ModelNet40.** We report overall accuracy and standard deviation without a voting strategy.

| Method | 5-way | | 10-way | |
|-----------------|-----------------|-----------------|-----------------|-----------------|
| | 10-shot | 20-shot | 10-shot | 20-shot |
| PointBERT [42] | 94.6±3.6 | 93.9±3.1 | 86.4±5.4 | 91.3±4.6 |
| MaskPoint [22] | 95.0±3.7 | 97.2±1.7 | 91.4±4.0 | 92.7±5.1 |
| PointMAE [25] | 96.3±2.5 | 97.8±1.8 | 92.6±4.1 | 93.4±3.5 |
| PointM2AE [44] | 96.8±1.8 | 98.3±1.4 | 92.3±4.5 | 95.0±3.0 |
| PointGPT-S [3] | 96.8±2.0 | 98.6±1.1 | 92.6±4.6 | 95.2±3.4 |
| PointMamba [21] | 96.9±2.0 | 99.0±1.1 | 93.0±4.4 | 95.6±3.2 |
| Mamba3D [15] | 96.4±2.2 | 98.2±1.2 | 92.4±4.1 | 95.2±2.9 |
| Ours | 97.5±2.3 | 99.1±1.4 | 93.5±3.7 | 96.1±3.5 |

eling within Mamba-based architectures. In contrast, our StruMamba3D leverages spatial states to effectively capture structural information, thereby eliminating the need for serialization and achieving 85.0% in mIoU_c and 86.7% in mIoU_i. Besides, StruMamba3D even outperforms the multi-scale self-supervised method PointM2AE [44]. Although multi-scale approaches typically yield better performance in part segmentation tasks, they often come with significantly higher computational costs (e.g., PCM [47] requires 45.0G FLOPs). In this work, we focus on designing an efficient and effective Mamba-based framework for structural modeling in point clouds. While incorporating multi-scale designs could further improve part segmentation performance, such strategies may introduce heuristic dependencies. Overall, our StruMamba3D represents a more robust and efficient Mamba-based framework for point cloud structural modeling, achieving the highest performance among all single-scale models.

Few-shot Learning on ModelNet40. To evaluate the effectiveness of the proposed methods with limited fine-tuning data, we conduct few-shot classification experiments on the ModelNet40 dataset. Following the standard procedure [25, 42], we perform 10 independent experiments for each setting and report the mean accuracy and standard deviation. As shown in Tab. 3, our method demonstrates a significant improvement over Mamba-based methods, outperforming PointMamba [21] by 0.6% in 5-way-10-shot and 0.5% in 10-way-10-shot settings, and surpassing Mamba3D [15] by 1.1% and 0.9%, respectively. These results indicate that our method can efficiently adapt to downstream tasks with longer input sequences, even under the constraint of limited training data.

4.3. Ablation Study

To explore the architectural design, we perform ablation studies on the PB-T50-RS split of ScanObjectNN dataset, as well as ModelNet40 and ShapeNetPart datasets.

Effect of Designed Modules. As shown in Tab. 4, we conduct an ablation study to evaluate the effectiveness of two core modules: the Structural SSM block and the sequence length-adaptive strategy. The baseline model is built upon the standard Mamba [9] block, following the same

Table 4. **Effect of the designed modules.** ‘SLAS’ indicates the sequence length-adaptive strategy used in pre-training phase.

| Structural SSM Block | SLAS | Overall Accuracy | | mIoU _c |
|----------------------|--------------|------------------|--------------|-------------------|
| | | ScanNN | MN40 | SNPart |
| \times | \times | 87.23 | 91.86 | 81.56 |
| \checkmark | \times | 92.09 | 94.45 | 84.49 |
| \checkmark | \checkmark | 92.75 | 95.06 | 84.96 |

Table 5. **Ablation on Structural SSM Block.** The baseline uses the standard Mamba block to replace the Structural SSM Block.

| Method | Overall Accuracy | | mIoU _c |
|------------------------------|------------------|--------------|-------------------|
| | ScanNN | MN40 | SNPart |
| baseline | 88.24 | 92.50 | 82.08 |
| w/ Structural SSM | 91.78 | 93.92 | 84.15 |
| w/ Lightweight Conv. for h | 92.22 | 94.65 | 84.62 |
| w/ Lightweight Conv. for x | 92.40 | 94.81 | 84.77 |
| w/ Bidirectional Scanning | 92.75 | 95.06 | 84.96 |

architecture as PointMAE [25]. Compared to this baseline, the Structural SSM block significantly enhances model performance, yielding improvements of 4.86% on ScanObjectNN, 2.59% on ModelNet40, and 2.93% on ShapeNetPart, demonstrating its effectiveness in modeling spatial dependencies and capturing richer structural information in point clouds. Furthermore, integrating the sequence length-adaptive strategy further enhances performance by improving the generalization to varying input lengths, thereby ensuring optimal performance across downstream tasks.

Effect of Structural SSM Block. As shown in Tab. 5, we further validate the effectiveness of individual components within the Structural SSM block. Starting from the baseline model with the standard Mamba block, we incrementally incorporate key modules from the Structural SSM block. The most significant performance gains stem from the Structural SSM and the lightweight convolution designed for spatial states, demonstrating that modeling structural information through hidden states in SSM is crucial. Moreover, the lightweight convolution for input points and the bidirectional scanning mechanism also enhance the performance across downstream tasks.

Effect of Structural SSM. Structural SSM is a crucial component of our method, as it effectively captures the structural information of point clouds through spatial states. To further analyze the impact of spatial state initialization and state-wise update strategy, we conduct ablation studies as shown in Tab. 6. Although spatial state initialization provides the states with positional information, the improvement is minimal, as Eq. (2) does not consider spatial relationships between input points and states. When integrating the state-wise update strategy, the model achieves a significant performance boost, with improvements of 1.42% on ScanObjectNN, 0.94% on ModelNet40, and 1.09% on ShapeNetPart. The results indicate that incorporating spatial relationships to control state information updates and propagation within the SSM is essential for effectively mod-

Table 6. **Ablation on Structural SSM.** ‘SS’ represents the use of spatial states in the SSM. ‘ $\phi(x)$ ’ denotes the SSM parameters depending on the point feature, while ‘MLP(ΔP)’ represents the SSM parameters depending on the spatial relationship.

| $\phi(x)$ | SS | MLP(ΔP) | Overall Accuracy | | mIoU _c |
|--------------|--------------|-------------------|------------------|--------------|-------------------|
| | | | ScanNN | MN40 | SNPart |
| \checkmark | \times | \times | 90.94 | 93.84 | 83.62 |
| \checkmark | \checkmark | \times | 91.33 | 94.12 | 83.87 |
| \checkmark | \checkmark | \checkmark | 92.75 | 95.06 | 84.96 |
| \times | \checkmark | \checkmark | 91.12 | 94.25 | 84.02 |

Table 7. **Ablation on sequence length-adaptive strategy.** ‘ASUM’ indicates the adaptive state update mechanism, and ‘ λ ’ indicates the weight of the spatial state consistency loss \mathcal{L}_{ssc} .

| ASUM | \mathcal{L}_{ssc} | Overall Accuracy | | |
|--------------|---------------------|------------------|--------------|--------------|
| | | ScanNN | MN40 | SNPart |
| \times | \times | 92.09 | 94.45 | 84.49 |
| \checkmark | \times | 92.26 | 94.57 | 84.56 |
| \checkmark | $\lambda = 1$ | 92.47 | 94.94 | 84.81 |
| \checkmark | $\lambda = 2$ | 92.75 | 95.06 | 84.96 |
| \checkmark | $\lambda = 5$ | 92.57 | 94.89 | 84.77 |
| \times | $\lambda = 2$ | 92.16 | 94.65 | 84.56 |

eling structural information. Moreover, we observe that relying solely on spatial relationships to generate SSM parameters results in performance degradation. This is likely due to the varying importance of point features, which cannot be fully captured by spatial position alone.

Effect of Sequence Length-Adaptive Strategy. The sequence length-adaptive strategy consists of two key components: the adaptive state update mechanism and the spatial state consistency loss. Tab. 7 presents an ablation study on these components. When applied individually, both the adaptive state update mechanism and the spatial state consistency loss yield only marginal performance gains. However, when used together, they lead to a significant improvement, demonstrating their complementary effects. The adaptive state update mechanism ensures a consistent total sampling time across different input lengths, while the spatial state consistency loss ensures the consistency of spatial state features. In addition, we perform an ablation study on the weight λ of \mathcal{L}_{ssc} and observe that the model achieves optimal performance when λ is set to 2.

5. Conclusion

We identify two key issues limiting Mamba potential: 3D point adjacency distortion during SSM processing and inadequate long-sequence memory retention in downstream tasks. To address these issues, we propose StruMamba3D with two core components. The structural SSM block maintains spatial dependencies during state selection and propagation, while the sequence length-adaptive strategy reduces the sensitivity of Mamba-based models to varying input lengths. Our method achieves substantial improvements across four downstream tasks, with ablation studies confirming the effectiveness of each module.

6. Acknowledgment

This work is supported by the National Defense Basic Scientific Research program (JCKY2022911B002).

References

- [1] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 6, 11
- [2] Binjie Chen, Yunzhou Xia, Yu Zang, Cheng Wang, and Jonathan Li. Decoupled local aggregation for point cloud learning, 2023. 6, 7
- [3] Guangyan Chen, Meiling Wang, Yi Yang, Kai Yu, Li Yuan, and Yufeng Yue. Pointgpt: Auto-regressively generative pre-training from point clouds. *arXiv preprint arXiv:2305.11487*, 2023. 3, 6, 7, 12, 13
- [4] Siheng Chen, Baoan Liu, Chen Feng, Carlos Vallespi-Gonzalez, and Carl Wellington. 3d point cloud processing and learning for autonomous driving: Impacting map creation, localization, and perception. *IEEE Signal Processing Magazine*, 38(1):68–86, 2020. 1
- [5] Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*, 2024. 3
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pages 4171–4186. Association for Computational Linguistics, 2019. 2, 3
- [7] Runpei Dong, Zekun Qi, Linfeng Zhang, Junbo Zhang, Jianjian Sun, Zheng Ge, Li Yi, and Kaisheng Ma. Autoencoders as cross-modal teachers: Can pretrained 2d image transformers help 3d representation learning? In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. 6, 7
- [8] Deng-Ping Fan, Ge-Peng Ji, Tao Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Pranet: Parallel reverse attention network for polyp segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 263–273. Springer, 2020. 6
- [9] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 2, 3, 7, 14
- [10] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021.
- [11] Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. Combining recurrent, convolutional, and continuous-time models with linear state space layers. *Advances in neural information processing systems*, 34:572–585, 2021. 3
- [12] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7:187–199, 2021. 2, 3
- [13] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3d point clouds: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(12):4338–4364, 2020. 1
- [14] Ziyu Guo, Renrui Zhang, Longtian Qiu, Xianzhi Li, and Pheng-Ann Heng. Joint-mae: 2d-3d joint masked autoencoders for 3d point cloud pre-training. *arXiv preprint arXiv:2302.14007*, 2023. 6
- [15] Xu Han, Yuan Tang, Zhaoxuan Wang, and Xianzhi Li. Mamba3d: Enhancing local features for 3d point cloud analysis via state space model. In *Proceedings of the 32nd ACM International Conference on Multimedia*, page 4995–5004. Association for Computing Machinery, 2024. 2, 3, 6, 7
- [16] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 2, 3
- [17] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 1960. 3
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 11, 12
- [19] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems*, 31, 2018. 1, 2, 3, 6
- [20] Ying Li, Lingfei Ma, Zilong Zhong, Fei Liu, Michael A Chapman, Dongpu Cao, and Jonathan Li. Deep learning for lidar point clouds in autonomous driving: A review. *IEEE Transactions on Neural Networks and Learning Systems*, 32(8):3412–3432, 2020. 1
- [21] Dingkang Liang, Xin Zhou, Wei Xu, Xingkui Zhu, Zhikang Zou, Xiaoqing Ye, Xiao Tan, and Xiang Bai. Pointmamba: A simple state space model for point cloud analysis. In *Advances in Neural Information Processing Systems*, 2024. 2, 3, 6, 7, 12, 13
- [22] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. In *European Conference on Computer Vision*, pages 657–675. Springer, 2022. 2, 3, 6, 7, 13
- [23] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 11, 12
- [24] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. *arXiv preprint arXiv:2202.07123*, 2022. 3, 6
- [25] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *European conference on computer vision*, pages 604–621. Springer, 2022. 2, 3, 5, 6, 7, 8, 11
- [26] François Pomerleau, Francis Colas, Roland Siegwart, et al. A review of point cloud registration algorithms for mobile robotics. *Foundations and Trends® in Robotics*, 4(1):1–104, 2015. 1

- [27] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 1, 2, 3, 4, 6, 7
- [28] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 1, 2, 3, 6, 7, 11
- [29] Zekun Qi, Runpei Dong, Guofan Fan, Zheng Ge, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. In *International Conference on Machine Learning (ICML)*, 2023. 3, 6
- [30] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *Advances in neural information processing systems*, 35:23192–23204, 2022. 1, 3, 6
- [31] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6411–6420, 2019. 1, 3
- [32] Chuxin Wang, Jiacheng Deng, Jianfeng He, Tianzhu Zhang, Zhe Zhang, and Yongdong Zhang. Long-short range adaptive transformer with dynamic sampling for 3d object detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(12):7616–7629, 2023. 1
- [33] Chuxin Wang, Yixin Zha, Jianfeng He, Wenfei Yang, and Tianzhu Zhang. Rethinking masked representation learning for 3d point cloud understanding. *IEEE Transactions on Image Processing*, 2024. 3
- [34] ChuXin Wang, Wenfei Yang, Xiang Liu, and Tianzhu Zhang. State space model meets transformer: A new paradigm for 3d object detection. In *The Thirteenth International Conference on Learning Representations*, 2025. 3
- [35] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (tog)*, 38(5):1–12, 2019. 2, 3, 5, 6
- [36] Chengzhi Wu, Junwei Zheng, Julius Pfommer, and Jürgen Beyerer. Attention-based point cloud edge sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5333–5343, 2023. 7
- [37] Xiaoyang Wu, Yixing Lao, Li Jiang, Xihui Liu, and Hengshuang Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. *Advances in Neural Information Processing Systems*, 35:33330–33342, 2022. 3
- [38] Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong He, and Hengshuang Zhao. Point transformer v3: Simpler faster stronger. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4840–4851, 2024. 3
- [39] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 7
- [40] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 574–591. Springer, 2020. 3
- [41] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM Transactions on Graphics (ToG)*, 35(6):1–12, 2016. 6, 7, 12
- [42] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19313–19322, 2022. 2, 3, 6, 7, 11, 13
- [43] Yixin Zha, Chuxin Wang, Wenfei Yang, and Tianzhu Zhang. Exploring semantic masked autoencoder for self-supervised point cloud understanding. *arXiv preprint arXiv:2506.21957*, 2025.
- [44] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Point-m2ae: multi-scale masked autoencoders for hierarchical point cloud pre-training. *Advances in neural information processing systems*, 35:27061–27074, 2022. 2, 3, 6, 7
- [45] Renrui Zhang, Liuhui Wang, Yu Qiao, Peng Gao, and Hongsheng Li. Learning 3d representations from 2d pre-trained models via image-to-point masked autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21769–21780, 2023. 3
- [46] Renrui Zhang, Liuhui Wang, Yu Qiao, Peng Gao, and Hongsheng Li. Learning 3d representations from 2d pre-trained models via image-to-point masked autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21769–21780, 2023. 6
- [47] Tao Zhang, Xiangtai Li, Haobo Yuan, Shunping Ji, and Shuicheng Yan. Point cloud mamba: Point cloud learning via state space model. *arXiv preprint arXiv:2403.00762*, 2024. 2, 3, 6, 7
- [48] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10252–10263, 2021. 3
- [49] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021. 3