

MambaTron: Efficient Cross-Modal Point Cloud Enhancement using Aggregate Selective State Space Modeling

Sai Tarun Inaganti
Robotics Institute, University of Minnesota
inaga015@umn.edu

Gennady Petrenko
Homothereum
gennady@homothereum.org

Abstract

Point cloud enhancement is the process of generating a high-quality point cloud from an incomplete input. This is done by filling in the missing details from a reference like the ground truth via regression, for example. In addition to unimodal image and point cloud reconstruction, we focus on the task of view-guided point cloud completion, where we gather the missing information from an image, which represents a view of the point cloud and use it to generate the output point cloud. With the recent research efforts surrounding state-space models, originally in natural language processing and now in 2D and 3D vision, Mamba has shown promising results as an efficient alternative to the self-attention mechanism. However, there is limited research towards employing Mamba for cross-attention between the image and the input point cloud, which is crucial in multi-modal problems. In this paper, we introduce MambaTron, a Mamba-Transformer cell that serves as a building block for our network which is capable of unimodal and cross-modal reconstruction which includes view-guided point cloud completion. We explore the benefits of Mamba’s long-sequence efficiency coupled with the Transformer’s excellent analytical capabilities through MambaTron. This approach is one of the first attempts to implement a Mamba-based analogue of cross-attention, especially in computer vision. Our model demonstrates a degree of performance comparable to the current state-of-the-art techniques while using a fraction of the computation resources.

1. Introduction

In recent years, the point-cloud format has garnered attention in the 3D computer vision research community due to properties like packaging efficiency and applications in various domains like scene reconstruction and understanding, robotics, geographical information systems and autonomous driving. The point cloud is generally an un-

ordered collection of data-points with each point represented by its 3D location coordinates (x, y, z) . Its space complexity is only dependent on the number of points contained, unlike that of the other formats like volumetric grids for which it is cubic. Pointnet [30] pioneered a class of point cloud processing techniques that operated directly on the input instead of converting it into a voxel grid.

One practical problem with the raw point cloud data collected from 3D scanners like LiDAR is that they tend to be sparse, noisy and incomplete due to factors like occlusion, hardware limitations and optical phenomena like reflections and transparency. The data needs to be enhanced with density and completeness before it can be usable for downstream tasks like classification and segmentation.

Point cloud completion [33] is a category of enhancement where the missing data in an incomplete point cloud is recovered. In order to advance the research methodologies behind point cloud enhancement and completion, datasets like PCN [45], KITTI [10], ModelNet [40], ShapeNet [3] and ShapeNet-ViPC [49] were introduced, among which were CAD models or raw data directly collected from hardware like LiDAR and laser scanners. Metrics like Chamfer Distance [5], Earth Mover’s Distance [5, 34] and Jensen-Shannon Divergence [1] were considered to measure the performance of point cloud completion techniques [7]. Parametrized completion techniques were introduced which primarily depended on the 3D prior knowledge of the input point cloud alone with no other auxiliary data. This approach was susceptible to accuracy errors, for example, in situations where two incomplete point clouds which belonged to different classes, but most of the difference is in the missing data that needs to be estimated. ViPC [49] introduced a multimodal processing technique where the missing information is inferred and enforced from a reference view or image which significantly aided in closing in the accuracy compared to unimodal methods.

The transformer [37] and recently the state-space model [13, 29] architectures have contributed significantly to the field of natural language processing. They have been adapted for 2D [4, 52] and 3D [24, 28, 39, 44, 48] computer

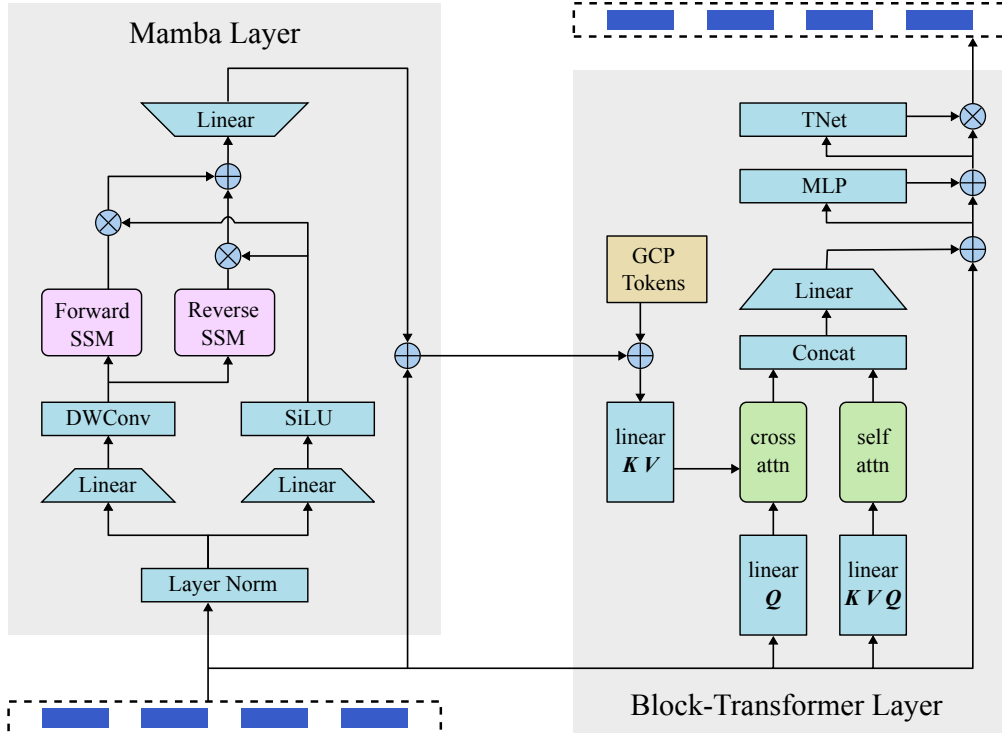


Figure 1. Illustration of the fully-bidirectional MambaTron cell. Input embeddings block is at the bottom-left in addition to the output block at the top right. The Mamba layer processes the entire input sequence at once while the Block-Transformer layer operates on a block of tokens at a time. We set the block size to 4 as illustrated above. The context state embeddings output by the Mamba layer are added to the Geometric Center Position (GCP) tokens which are the key-point coordinates while being fed to the Block-Transformer.

vision tasks with impressive results. Block-based transformers [6, 20, 22] have tackled the problem of long token sequences by operating on blocks or windows of tokens at any given time. Our research adapts these natural language techniques towards the task of view-guided point cloud completion. In this paper, we detail our contributions, for which the summaries are given below.

1. We propose the fully bidirectional MambaTron block, which consists of a Mamba State Space Model (SSM) layer for long-range global contextualization and an attention block for short-range neighborhood attention. This block is responsible for generating the piecewise contextualized local embeddings for both the auxiliary view and the partial point cloud.
2. We present our end-to-end neural network model for point cloud reconstruction that optionally takes a reference view as input and use it to train the MambaTron cells. We use the MambaTron cell to encode the multimodal input (point cloud OR image) and also model the relationship between the two inputs (point cloud AND image).
3. The Mamba layer is sensitive to the order in which the input tokens are presented and so we describe our

Adjacency-Preserving Reordering (APR) technique to optimize MambaTron’s performance while maintaining geometric adjacency among neighboring tokens.

4. We demonstrate the performance on the ShapeNet-ViPC dataset, compare it with the SOTA image-assisted point cloud completion methods and show how our MambaTron-based network keeps up with state-of-the-art techniques with a fraction of resource utilization.

2. Related Work

2.1. Point Cloud Analysis and Completion

Studies were conducted on deep-learning parametric methods surrounding point cloud analysis [7] and enhancement [33, 47] pioneered by PointNet’s [30] method to directly process 3D points and PCN’s [45] ability to directly operate on point clouds without any assumptions about any of the global characteristics like structural information. PointNet++ [31] was an encoder model that popularized the Farthest Point Sampling (FPS) method to downsample 3D points which helped the model learn local features which both the point-level and global features missed and pro-

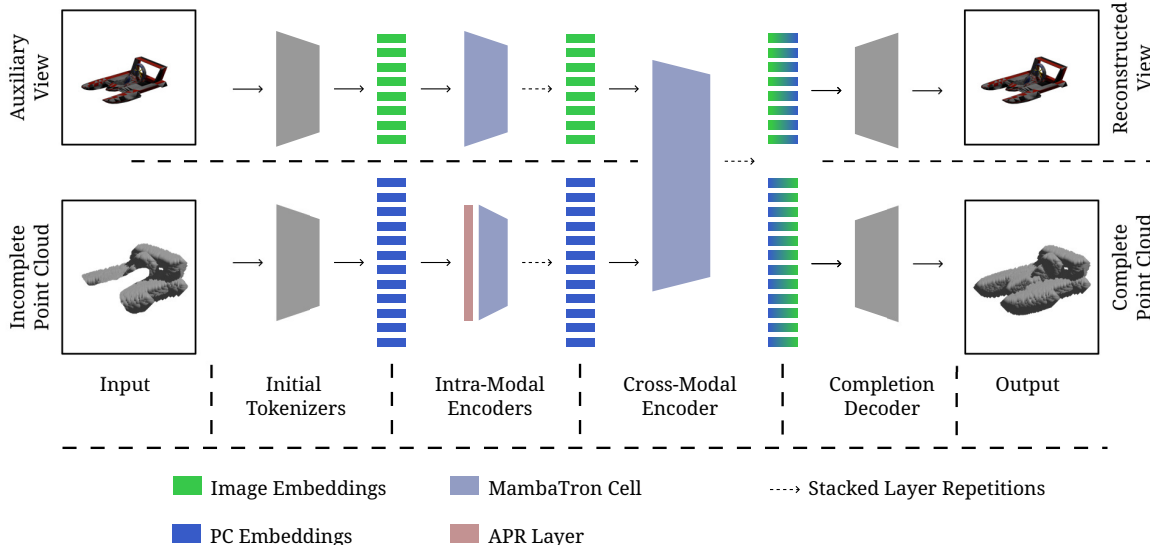


Figure 2. An overview of our MambaTron-based model for the View-Guided Point Cloud Completion Task.

moted compactness in the input. FPS was adapted for its effectiveness by several other encoders [24,26,28,39,44,48] and also in end-to-end completion techniques [23,41]. Various architectural solutions were repurposed from the domain of natural language and speech processing like transformers [37,44] and state space models [13,24,48]. 2D vision also had some ideas to offer, like masked autoencoders [18,28], GANs [11] and DDPMs [19,25].

On top of these ideas, which mostly described the fundamental structure and operations of the encoder/decoder layers, various point cloud enhancement techniques were recently introduced to tackle problems inherent to point clouds like surface properties, sparsity etc. GCN-based methods [38,50] use an encoder that relies on a graph that connects the input points as vertices, even in the feature space. FoldingNet [43] and AtlasNet [12] deform a 2D texture grid onto a 3D surface in a series of folding operations. TopNet [35] and SnowflakeNet [41] implement specialized decoder modules to upsample/grow a point cloud in a tree pattern. Mamba-based enhancement methods [23,51] have leveraged Mamba’s ability to efficiently process long 3D point sequences. Unimodal point cloud enhancement methods like the ones mentioned above have only the input priors to rely on and are more prone to performance degradation in completion tasks, compared to denoising and upsampling tasks, unless the input is constrained by class, geometry etc. These models are bound to be fully supervised as they rely on the availability of the ground truth point clouds.

2.2. View-Guided Point Cloud Completion

In order to combat the problem faced by point cloud completion tasks, the ViPC [49] architecture proposed using an auxiliary input like an image which was a view of

the object in the input data that contained the missing information. This spanned the class of view-assisted methods [2,42,47,49,53] that leveraged multimodal information fusion to generate a fulfilled point cloud. With the help of auxiliary views, the models can limit their reliance on the availability of ground truths by self-supervision. XMFnet [2] is a transformer-based model that utilizes cross-attention to fuse the multimodal features to form the completed output. EGInet [42] improves on top of XMFnet by adapting and incorporating style and content losses [9] during the feature fusion process. These models are primarily attention based which have a quadratic computational complexity $O(L^2)$ with respect to the input token sequence length L . Our MambaTron-based completion network aims to reduce this to a near-linear theoretical complexity. Our network design is based on the cross-modal feature transfer mechanism of Joint-MAE [17] and the point cloud groupwise-decoder designs of AXform [46] and Joint-MAE.

2.3. State Space Models

2.3.1 History

The Structured State Space Sequence (S4) model [15] is the result of applying the principles of control theory in combination with sequential deep learning techniques like RNNs to model languages. It is the first attempt at efficiently processing long sequence data, surpassing its predecessor LSSL [16] by a wide margin in state computation efficiency, both of which utilize the HiPPO matrix [14] for state representation. These works spawned further SSM variants which made incremental progress towards modeling performance. Mamba [13] discards the linear time invariance (LTI) of its HiPPO matrices and introduces se-

lection as a way to efficiently compress and package the context, which enabled SSMs to outperform transformer-based models while maintaining subquadratic computational complexity.

Following its success in language modeling, Mamba was adapted for computer vision tasks. ViM [52] adapts SSMs for image processing tasks as an analogue to the ViT [4]. Point Mamba [24] orders point cloud patches using Hilbert/Trans-Hilbert ordering operations before feeding them to a Mamba layer for downstream tasks. PCM [48] stacks Mamba encoders and uses a Geometric Affine Module (GAM) [26], uses code-based reordering and coupled encoder-decoder pairs for segmentation tasks. PoinTramba [39] uses a transformer layer to generate patch features and a learning-based ordering technique to optimize Mamba’s performance. 3dmambacomplete [23] is an end-to-end point cloud completion technique that implements grid-folding in the decoder.

Despite the advantages of SSMs, studies [8, 27] have shown the limitations on the extent by which SSMs outperform transformers. MambaTron aims to overcome this by placing Mamba with a Block-Transformer to carry over the benefits of both the architectures while minimizing the bottlenecks.

2.3.2 Analytical Study

Quantitatively, S4 [15] introduces the Structured SSM which implements a linear control system with parameters $(\Delta, \mathbf{A}, \mathbf{B}, \mathbf{C})$ in two stages. This system is a continuous sequence-to-sequence transformation $x(t) \rightarrow y(t)$ that maintains a latent state $h(t)$ and is formulated below:

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

Discretization. The first stage implements the continuous-to-discrete parameter transformation $(\Delta, \mathbf{A}, \mathbf{B}) \rightarrow (\bar{\mathbf{A}}, \bar{\mathbf{B}})$ using a discretization rule (f_A, f_B) like the zero-order hold (ZOH).

$$\bar{\mathbf{A}} = f_A(\Delta, \mathbf{A}), \bar{\mathbf{B}} = f_B(\Delta, \mathbf{A}, \mathbf{B}) \quad (2)$$

Computation. The second stage computes the discrete sequence-to-sequence transformation $x_t \rightarrow y_t$ with state h_t using the new set of parameters $(\bar{\mathbf{A}}, \bar{\mathbf{B}}, \mathbf{C})$.

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

Here, $x_t \in \mathbb{R}, y_t \in \mathbb{R}, h_t \in \mathbb{R}^N, \bar{\mathbf{A}} \in \mathbb{R}^{N \times N}, \bar{\mathbf{B}} \in \mathbb{R}^{N \times 1}, \mathbf{A} \in \mathbb{R}^{N \times N}, \mathbf{B} \in \mathbb{R}^{N \times 1}, \mathbf{C} \in \mathbb{R}^{N \times 1}$ where N is the length of the latent state vector h .

S4 models are Linear Time Invariant (LTI), meaning both the continuous parameters $(\Delta, \mathbf{A}, \mathbf{B}, \mathbf{C})$ and discrete parameters $(\bar{\mathbf{A}}, \bar{\mathbf{B}})$ are fixed for all time steps. Mamba/S6 introduces the selective scanning mechanism and turns the

parameters $\Delta, \mathbf{B}, \mathbf{C}$ into time-variant functions of x_t which have a length dimension L .

2.4. Block-based Transformers

Language Transformers have mostly replaced general RNNs like LSTMs due to their parallel processing capabilities and attention mechanism which enables them to handle long input sequences without vanishing gradients. But their quadratic complexity quickly makes training transformer models on long input sequences expensive and unstable. To handle the complexity problem, various block/window-based architectures were proposed which limit the size of the attention matrix by introducing variations of state/context vectors. Block-Recurrent Transformer (BRecT) [22] cells, inspired by LSTM cells, employ a sliding window to maintain a state vector like in an RNN while limiting a dimension of the attention matrix by the window size $W \ll L$, bringing the complexity down from $\mathcal{O}(L^2)$ of the vanilla transformer to $\mathcal{O}(LW)$. Block-transformers [20] operate on blocks of tokens, and employ a block-decoder followed by a token-level decoder, giving the same complexity. Block-State Transformers (BST) [6] replace the recurrent cell in BRecTs with an SSM and processes each block in parallel instead of utilizing a sliding window, further bringing down the complexity to $\mathcal{O}(W^2) + \mathcal{O}(L \log L)$. The MambaTron cell is based on the BRecT cell and benefits from subquadratic complexity like BSTs.

3. Qualitative Method Description

3.1. Problem Statement

Our ultimate goal is to train a unified network that is capable of both unimodal reconstruction and cross-modal point cloud completion. We say reconstruction (reconstructed output is the same as the input) for unimodal inputs and not completion, which is beyond the scope of this work. We define multiple conditions that must be fulfilled in order to achieve this.

1. If an incomplete point cloud and a complete reference image are given, the network generates embeddings that represent the complete image and also a complete version of the point cloud.
2. If only a complete point cloud is given, the model must generate embeddings which can be decoded back to the input. This case is meant for downstream tasks and not enhancement.

We design a model that utilizes MambaTron cells to generate embeddings for images and point clouds while incorporating missing information from the other modality when

available. We later observe that the two conditions are complementary and help with the performance of the task corresponding to the other condition.

3.2. MambaTron

The word "MambaTron" is a portmanteau of three terms: Mamba, Transformer and Perceptron. The MambaTron cell takes advantage of the selective scanning mechanism of S6 and its expressive power to compute the context states and embeddings from the input sequence. It consists of two layers, one after the other. The cell and its inner workings are depicted in Figure 1.

Mamba Layer. Given an input sequence of length L , this layer processes the input bidirectionally so that every element in the input contributes to the context corresponding to a given input embedding. It outputs the context state vectors corresponding to each input embedding. The context sequence inherits its length L from the input sequence.

Block-Transformer Layer. This layer divides the input sequence into blocks of size $W \ll L$ and processes them blockwise. For any given block, this layer's computation only depends on the the block's input and context, not on any other block's data, which means that they can be processed in parallel. This layer outputs the context-imbued feature embeddings of length W per block. The output blocks are then concatenated back to give the output sequence of length L .

3.3. Adjacency-Preserving Reordering

In the point cloud encoding research works, there is a general consensus that the encoder must be able to capture three different types of embeddings to a fair extent: individual positional embeddings, the local neighborhood group embeddings and the global embeddings. Farthest Point Sampling (FPS) [31] to obtain a set of key points followed by a neighborhood selection algorithm like the K Nearest Neighbors (KNN) technique to divide the point cloud into groups of the same size, collectively abbreviated as FPS+KNN is a popular technique whose output serves as the input to the point cloud tokenizers that gives us the positional embeddings of the key points and also the collective embeddings of each neighborhood group. Depending on the tokenizer, pooling the positional embeddings can give us the global features. The feature embeddings of the point cloud groups are ordered before being sent to the Mamba encoder layer in any Mamba-based point cloud encoder.

As witnessed in some of the previous works [24, 39, 48] involving point cloud encoding tasks surrounding SSMs, Mamba's performance in computer vision tasks is greatly influenced by the ordering of the input tokens to the SSM layer. The previous works introduce ordering techniques to combat this which is described in Section 2.3.1. The problem with these techniques is that they are either too rigid

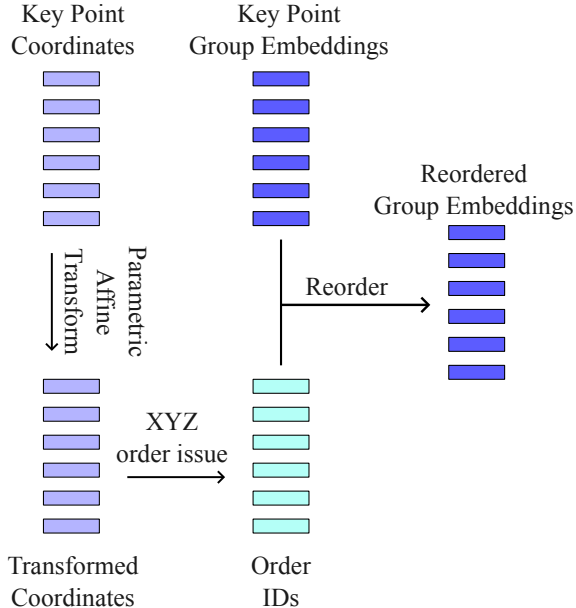


Figure 3. The Adjacency-Preserving Reordering (APR) scheme

(like PointMamba's Hilbert ordering) and apply the same order for every layer or are too flexible (like PoinTramba's BIO ordering). PCM's code-based XYZ ordering technique allows for flexibility while maintaining a level of rigidity, i.e. adjacent groups in the input sequence are adjacent geometrically, which is what ViM [52] goes for in 2D image encoding. However, PCM's technique is still too rigid as any point cloud is only privy to six different styles of ordering based on the permutations of the X, Y and Z directions. We argue that introducing more adjacency-respecting orders helps optimize Mamba's performance further. To achieve this, we introduce the Adjacency Preserving Reordering (APR) technique.

In the technique, we take the sparse point cloud containing the key points from the FPS algorithm and run it through a Transformation Network (TNet) which rotates and/or mirrors it while preserving the structure. Then we run the XYZ ordering scheme on the new point cloud. Since the number of possible affine transformations is infinite, there are an infinite ways to order the point cloud groups while maintaining the adjacency, only limited by the finite number of key points in the point cloud. Since a TNet is involved, this is a parameter-based ordering technique like BIO.

3.4. Model

Our model pipeline is a series of four sections, as shown in Figure 2.

1. Initial Tokenizer
2. Intra-modal Encoder

Table 1. Quantitative results on the original **ShapeNet-ViPC** dataset for the **Known** categories.

Methods	Avg	Airplane	Cabinet	Car	Chair	Lamp	Sofa	Table	Watercraft
<i>Mean CD per point $\times 10^{-3}$ (lower is better)</i>									
ViPC [49]	3.308	1.760	4.558	3.183	2.476	2.867	4.481	4.990	2.197
CSDN [53]	2.570	1.251	3.670	2.977	2.835	2.554	3.240	2.575	1.742
XMFnet [2]	1.443	0.572	1.980	1.754	1.403	1.810	1.702	1.386	0.945
EGInet [42]	1.211	0.534	1.921	1.655	1.204	0.776	1.552	1.227	0.802
Ours	1.199	0.537	1.905	1.620	1.184	0.774	1.561	1.234	0.779
<i>Mean F-Score @ 0.001 (higher is better)</i>									
ViPC [49]	0.591	0.803	0.451	0.512	0.529	0.706	0.434	0.594	0.730
CSDN [53]	0.695	0.862	0.548	0.560	0.669	0.761	0.557	0.729	0.782
XMFnet [2]	0.796	0.961	0.662	0.691	0.809	0.792	0.723	0.830	0.901
EGInet [42]	0.836	0.969	0.693	0.723	0.847	0.919	0.756	0.857	0.927
Ours	0.839	0.969	0.697	0.728	0.851	0.920	0.764	0.856	0.931

3. Cross-modal Encoder (for cross-modal instances only)

4. Reconstruction Decoder

In this paper, we focus on the design of sections 2 and 3 of which the MambaTron cell is the building block.

4. Quantitative Model Execution

We are presented with a point cloud $(N, 3)$ with or without an auxiliary image (H, W) , the end-to-end neural network model from Figure 2 receives as input and outputs a reconstructed version of the input point cloud.

4.1. Forward Pass

4.1.1 Initial tokenizer

We borrow PointMamba’s technique to extract the initial tokens which gives us the key point cloud $(n_p, 3)$ and the corresponding position embeddings (n_p, C) and group embeddings (n_p, C) . We reference ViM for the 2D counterpart, which gives us the image patch embeddings (n_i, C) .

4.1.2 Intra-Modal Encoder

This is a serially stacked MambaTron encoder which takes in either the image or the point cloud tokens that are output by the tokenizer and output their feature embeddings imbued with the global context. For the image input, we prepend a special $\langle I \rangle$ token and append a $\langle STOP \rangle$ token. For the point cloud input, we first reorder the tokens using APR from section 3.3, then replace $\langle I \rangle$ with $\langle P \rangle$ along with adding the GCP tokens to the context vectors as shown in Figure 1, which are basically the key point positional embeddings. The output feature embeddings (n_i, C) or (n_p, C) with the special tokens removed are imbued with global context by this encoder.

4.1.3 Cross-Modal Encoder

This encoder is used to extract context from one modality and impart it to the other. The image and point cloud group embeddings are arranged in the following order to form the input.

1. $\langle I \rangle$ token
2. image embeddings
3. $\langle P \rangle$ token
4. key point group embeddings
5. $\langle STOP \rangle$ token

Similar to the intra-modal encoder, the group context vectors are supplied with the GCP tokens at the block-transformer layer. The outputs (n_i, C) and (n_p, C) with the special tokens removed now contain the relationships of each image patch or point cloud group with every token from the other modality.

4.1.4 Decoder

The decoder takes the embeddings (n_i, C) or (n_p, C) from the intra-modal or the cross-modal encoder, depending on the number of inputs and reconstructs the image or the point cloud with the same dimension as the input. Due to the modularity of our method, the decoder is selected based on its ability to process image embeddings if needed and is task-specific which will be described later in section 5.

4.2. Training

Our training method is a two-stage process.

Table 2. Quantitative results on the **ShapeNet-ViPC** dataset for the **Novel** categories.

Methods	Avg		Bench		Monitor		Speaker		Phone	
	CD	F	CD	F	CD	F	CD	F	CD	F
ViPC [49]	4.601	0.498	3.091	0.654	4.419	0.491	7.674	0.313	3.219	0.535
CSDN [53]	3.656	0.631	1.834	0.798	4.115	0.598	5.690	0.485	2.985	0.644
XMFnet [2]	2.671	0.710	1.278	0.862	2.806	0.677	4.823	0.556	1.779	0.748
EGInet [42]	2.354	0.750	1.047	0.902	2.513	0.716	4.282	0.591	1.575	0.792
Ours	2.333	0.761	1.041	0.909	2.499	0.723	4.268	0.599	1.566	0.797

Unimodal stage. The model is trained on complete point cloud inputs. Even though this is called the unimodal training stage, we actually derive the top-down image projection of the input point cloud similar to Joint-MAE [17] and pass it along as well, making this a cross-modal task. We also partially mask the inputs before the decoder stage to promote robustness. This stage teaches the model to develop an idea of how complete point clouds look like and how they correlate to images.

Cross-modal stage. After the first stage, the model is further trained with an incomplete point cloud and an auxiliary view as inputs and a completed point cloud as the output which is compared to the ground truth point cloud for regression.

We first look at the various types of training loss that we consider for regression in both the stages.

Chamfer Distance. This loss consolidates the reconstructed point cloud to the ground truth.

$$\mathcal{L}_{CD} = \sum_{x \in P_1} \min_{y \in P_2} \|x - y\|_2^2 + \sum_{y \in P_2} \min_{x \in P_1} \|y - x\|_2^2 \quad (4)$$

Style Loss. The style loss [9, 21, 42, 53] regression ensures that the model is learning to apply the image features to the point cloud and vice-versa. The style loss is defined as:

$$\mathcal{L}_{style} = \frac{(G(F_i) - G(F'_p))^2 + (G(F_p) - G(F'_i))^2}{n_P \times C} \quad (5)$$

where F_p and F_i are the intra-modal features, whereas F'_p and F'_i are the cross-modal features. G is the gram matrix function $G(F) = F^T \cdot F$.

Projection Loss. Similar to Joint-MAE’s [17] reconstruction loss, we project the reconstructed point cloud onto a 2D plane by the same view of the input projection and compare it with the initial projection (unimodal stage only).

$$\mathcal{L}_{proj} = MSE(I_{proj}(P_{output}), I) \quad (6)$$

2D Loss. This loss compares the final image to the auxiliary view in the cross-modal case and the input projection.

$$\mathcal{L}_{2D} = MSE(I_{output}, I) \quad (7)$$

Table 3. ScanObjectNN dataset, accuracy in %

Method	Params. (M)	FLOPs (G)	PB.T50_RS ↑
Point-MAE [28]	22.1	4.8	85.18
Joint-MAE [17]	-	-	86.07
PointMamba [24]	12.3	3.1	89.31
PCM [48]	34.2	45.0	88.10
Ours	13.6	3.9	90.17
RECON [32]	43.6	5.3	90.63

For the unimodal training stage, the training loss is below.

$$\mathcal{L}_{uni} = \mathcal{L}_{CD} + \mathcal{L}_{2D} + \mathcal{L}_{proj} \quad (8)$$

For the cross-modal stage, regression is performed using the following loss.

$$\mathcal{L}_{cross} = \mathcal{L}_{CD} + \mathcal{L}_{2D} + \mathcal{L}_{style} \quad (9)$$

5. Experiments and Results

5.1. Cross-Modal Point Cloud Completion

We first train our model with a Joint-MAE [17] based decoder at the unimodal stage using complete point clouds from the ShapeNet55 dataset which consists of images belonging to one of 55 categories, then train on the ShapeNet-ViPC dataset with an 80-20 train-test split which contains objects with noisy and noiseless views from 13 categories at the cross-modal stage. We set MambaTron’s window size to 4. For the evaluation, we use the l2-normalized CD and the F-score metrics similar to the other state-of-the-art techniques for this task. On average, our model outperforms the latest state-of-the-art technique with only 3.92 M parameters compared to EGInet’s 9.03 M and XMFnet’s 9.57 M. Our model captures the objects with more complex features and limited samples more effectively compared to EGInet, especially on novel categories. The results are presented in Tables 1 and 2.

5.2. Unimodal Point Cloud Analysis

Though the main focus of this paper is view-guided point cloud completion, we also test the efficacy of our model




Figure 4. GPU usage comparison

on point cloud analysis tasks by pretraining. We train our model with an AXform-based [46] decoder at the unimodal stage similar to the previous task on the ScanObjectNN [36] dataset, but this time with an 80-20 split. On the same 80% training data, we replace the decoder with a classification head and train a second time. The window size is the same as before. We observe comparable results comparable to the SOTA methods on average while on a much smaller computational footprint, giving the accuracies of $94.97 \pm 0.02\%$ on OBJ_BG, $93.15 \pm 0.02\%$ on OBJ_ONLY and $90.17 \pm 0.02\%$ on PB_T50_RS, which are the evaluation metrics used.

We also test the computational efficiency [24] of our MambaTron-based model where we map the GPU usage to the length of the input sequence to show how MambaTron’s subquadratic complexity is superior to that of the transformer-based models, which is shown on Figure 4.

5.3. Ablation Studies

We study the effects of various design decisions like the cell level components, network level components, the multiple stages of training and the losses.

Effect of the Block-Transformer. A primary question would be the purpose of the Block-Transformer in the MambaTron cell. The Mamba layer of the MambaTron block focuses on the relationships among the input elements and not the elements themselves. But our design allows for smoothly swapping out the block, in which case the output embeddings are replaced by the context state vectors. We tested the model with the block-transformer disabled and found that the model performed worse than even the baseline ViPC model on the ShapeNet-ViPC dataset at its size with the novel categories being affected the most. For the pretraining task, the performance dropped below that of PointMamba. The MambaTron cell converged quicker than its bare Mamba counterpart, yet the performance still tanked on the novel categories.

Effect of sharing the intra-modal encoder. In our model, a common stacked MambaTron encoder accepts tokens from either the image or the point cloud and adds context to each token. Using a separate intra-modal encoder for each encoder increased the number of parameters while the

Table 4. Ablation study on ShapeNet-ViPC

w/o	CD (avg)	
	known	novel
-	1.199	2.333
block-tr	1.754	3.217
shared intra	1.317	2.625
cross-modal	1.363	2.699
style/proj loss	1.322	2.676

slowing down the training as the loss function takes longer to converge. The model performed worse on more complex and novel categories.

Effect of the cross-modal encoder. Since the shared intra-modal encoder already processes both the image and the point cloud, we tested the model without a cross-modal encoder or the style loss. The projection loss in the unimodal training stage stays intact.

Effect of the style and projection losses. This time, we trained the model without the mentioned losses contributing to the regression, while the cross-modal encoder is intact, to show the effect of the image on the encoder.

The ablation performance is recorded and the average CD for the ShapeNet-ViPC dataset are given in Table 4 for the known and novel categories.

6. Conclusion and Future Work

In this paper, we have described the MambaTron cell and built a modular network model using cell for the tasks of view-guided point cloud completion and unimodal point cloud pretraining which can be used for downstream tasks like classification. Our model performs on par with or surpasses the current state-of-the-art models for both the tasks with a limited footprint while being quicker, owing to MambaTron’s fast inference coupled with expressive power. This version of the MambaTron is limited to outputting embeddings which are constrained to the input’s dimensions. Further research can remove this constraint, thus helping generalize the cell’s applications. At the network level, improvements can be applied to further integrate the MambaTron cell into the model instead of simply utilizing it as a building block for a conventional feed-forward network. Future works can focus on using the cell for unguided point cloud completion in addition to other enhancement tasks like upsampling and denoising. We hope that our work helps other researchers in point cloud and multimodality research with applications like scene-awareness and path-planning in fields like Robotics.

Acknowledgment. This work was supported by the Minnesota Robotics Institute (MnRI).

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 40–49. PMLR, 10–15 Jul 2018. 1
- [2] Emanuele Aiello, Diego Valsesia, and Enrico Magli. Cross-modal learning for image-guided point cloud shape completion. In *Advances in Neural Information Processing Systems*, 2022. 3, 6, 7
- [3] Angel X. Chang, Thomas A. Funkhouser, Leonidas J. Guibas, Pat Hanrahan, Qi-Xing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. Shapenet: An information-rich 3d model repository. *CoRR*, abs/1512.03012, 2015. 1
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. 1, 4
- [5] Haoqiang Fan, Hao Su, and Leonidas J. Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1
- [6] Mahan Fathi, Jonathan Pilault, Orhan Firat, Christopher Pal, Pierre-Luc Bacon, and Ross Goroshin. Block-state transformers, 2023. 2, 4
- [7] Ben Fei, Weidong Yang, Wenming Chen, Zhijun Li, Yikang Li, Tao Ma, Xing Hu, and Lipeng Ma. Comprehensive review of deep learning-based 3d point cloud completion processing and analysis, 2022. 1, 2
- [8] Daniel Y Fu, Tri Dao, Khaled Kamal Saab, Armin W Thomas, Atri Rudra, and Christopher Re. Hungry hungry hippos: Towards language modeling with state space models. In *The Eleventh International Conference on Learning Representations*, 2023. 4
- [9] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. A neural algorithm of artistic style, 2015. 3, 7
- [10] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. 1
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 3
- [12] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [13] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2024. 1, 3
- [14] Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1474–1487. Curran Associates, Inc., 2020. 3
- [15] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022. 3, 4
- [16] 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. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 572–585. Curran Associates, Inc., 2021. 3
- [17] 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. In Edith Elkind, editor, *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 791–799. International Joint Conferences on Artificial Intelligence Organization, 8 2023. Main Track. 3, 7
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2022. 3
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. 3
- [20] Namgyu Ho, Sangmin Bae, Taehyeon Kim, Hyunjik Jo, Yireun Kim, Tal Schuster, Adam Fisch, James Thorne, and Se-Young Yun. Block transformer: Global-to-local language modeling for fast inference. *arXiv preprint arXiv:2406.02657*, 2024. 2, 4
- [21] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 7
- [22] DeLesley Hutchins, Imanol Schlag, Yuhuai Wu, Ethan Dyer, and Behnam Neyshabur. Block-recurrent transformers, 2022. 2, 4
- [23] Yixuan Li, Weidong Yang, and Ben Fei. 3dmambacomplete: Exploring structured state space model for point cloud completion, 2024. 3, 4
- [24] Dingkan 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. 1, 3, 4, 5, 7, 8
- [25] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2837–2845, June 2021. 3

- [26] 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, 2022. 3, 4
- [27] William Merrill, Jackson Petty, and Ashish Sabharwal. The illusion of state in state-space models. In *Forty-first International Conference on Machine Learning*, 4
- [28] Yatian Pang, Wenxiao Wang, Francis E. H. Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, page 604–621, Berlin, Heidelberg, 2022. Springer-Verlag. 1, 3, 7
- [29] Badri Narayana Patro and Vijay Srinivas Agneeswaran. Mamba-360: Survey of state space models as transformer alternative for long sequence modelling: Methods, applications, and challenges, 2024. 1
- [30] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CoRR*, abs/1612.00593, 2016. 1, 2
- [31] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: deep hierarchical feature learning on point sets in a metric space. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 5105–5114, Red Hook, NY, USA, 2017. Curran Associates Inc. 2, 5
- [32] 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 Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *International Conference on Machine Learning, ICML 2023, 23–29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 28223–28243. PMLR, 2023. 7
- [33] Siwen Qian, Junhao Yu, Ziming Nie, Muze Wang, Sijia Feng, Pei An, and Jiaqi Yang. Deep learning for 3d point cloud enhancement: A survey. 2024. 1, 2
- [34] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. The earth mover’s distance as a metric for image retrieval. *Int. J. Comput. Vision*, 40(2):99–121, Nov. 2000. 1
- [35] Lyne P. Tchammi, Vineet Kosaraju, Hamid Rezatofighi, Ian Reid, and Silvio Savarese. Topnet: Structural point cloud decoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3
- [36] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 8
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017. 1, 3
- [38] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. 38(5), Oct. 2019. 3
- [39] Zicheng Wang, Zhenghao Chen, Yiming Wu, Zhen Zhao, Luping Zhou, and Dong Xu. Pointtramba: A hybrid transformer-mamba framework for point cloud analysis, 2024. 1, 3, 4, 5
- [40] 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 (CVPR)*, June 2015. 1
- [41] Peng Xiang, Xin Wen, Yu-Shen Liu, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Zhizhong Han. Snowflakenet: Point cloud completion by snowflake point deconvolution with skip-transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5499–5509, October 2021. 3
- [42] Hang Xu, Chen Long, Wenxiao Zhang, Yuan Liu, Zhen Cao, Zhen Dong, and Bisheng Yang. Explicitly guided information interaction network for cross-modal point cloud completion, 2024. 3, 6, 7
- [43] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3
- [44] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3
- [45] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *2018 International Conference on 3D Vision (3DV)*, pages 728–737, 2018. 1, 2
- [46] Kaiyi Zhang, Ximing Yang, Yuan Wu, and Cheng Jin. Attention-based transformation from latent features to point clouds. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(3):3291–3299, Jun. 2022. 3, 8
- [47] Kun Zhang, Ao Zhang, Xiaohong Wang, and Weisong Li. Deep-learning-based point cloud completion methods: A review. *Graphical Models*, 136:101233, 2024. 2, 3
- [48] Tao Zhang, Haobo Yuan, Lu Qi, Jiangning Zhang, Qianyu Zhou, Shunping Ji, Shuicheng Yan, and Xiangtai Li. Point cloud mamba: Point cloud learning via state space model, 2024. 1, 3, 4, 5, 7
- [49] Xuancheng Zhang, Yutong Feng, Siqi Li, Changqing Zou, Hai Wan, Xibin Zhao, Yandong Guo, and Yue Gao. View-guided point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15890–15899, 2021. 1, 3, 6, 7
- [50] Yingxue Zhang and Michael Rabbat. A graph-cnn for 3d point cloud classification. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6279–6283, 2018. 3
- [51] Qingyuan Zhou, Weidong Yang, Ben Fei, Jingyi Xu, Rui Zhang, Keyi Liu, Yeqi Luo, and Ying He. 3dmambaipf: A state space model for iterative point cloud filtering via differentiable rendering, 2024. 3
- [52] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient

visual representation learning with bidirectional state space model. In *Forty-first International Conference on Machine Learning*. 1, 4, 5

- [53] Zhe Zhu, Liangliang Nan, Haoran Xie, Honghua Chen, Jun Wang, Mingqiang Wei, and Jing Qin. Csdn: Cross-modal shape-transfer dual-refinement network for point cloud completion. *IEEE Transactions on Visualization and Computer Graphics*, 30(7):3545–3563, 2024. 3, 6, 7

— Supplementary Material —

MambaTron: Efficient Cross-Modal Point Cloud Enhancement using Aggregate Selective State Space Modeling

Sai Tarun Inaganti
Robotics Institute, University of Minnesota
inaga015@umn.edu

Gennady Petrenko
Homothereum
gennady@homothereum.org

Table 1. Effects of APR

APR	affine	CD (avg)	
		known	novel
✗	✗	1.354	2.737
✓	✗	1.317	2.611
✓	✓	1.199	2.333

Table 2. ModelNet40 dataset

Method	OA (%)
PCM [6]	93.4
PointMamba [2]	93.6
Point-MAE [3]	93.8
Joint-MAE [1]	94.0
Ours	94.3

1. Effects of the APR scheme

The Adjacency Preserving Reordering (APR) scheme is space-filling, similar to the Hilbert [2] and XYZ [6] ordering techniques. Additionally, the scheme theoretically provides an infinite ways to order the points, only limited by the number of points, thanks to a learnable affine transformation layer. We look at the effects of each component through a further ablation study on the ShapeNet-ViPC dataset in Table 1.

2. Performance statistics on additional datasets

The task of cross-modal point cloud completion is a recent idea introduced [7] along with the ShapeNet-ViPC dataset, with no equivalent dataset that is just as comprehensive, i.e. contains class-mapped pointclouds with corresponding reference view images. For unimodal analysis, in addition to the real-world 3D object classification on the ScanObjectNN dataset, we perform experiments that we describe below with the same encoder configuration as that of the classification task on ScanObjectNN. The training data is subject to random rotation and slight scaling operations at each epoch. For comparison purposes, we take the results reported in the other research works.

Classification on ModelNet40. This [4] is a synthetic CAD dataset with 40 classes that contains a total of 12,311 noiseless pre-aligned models. The classification head is based on that of PointMamba [2]. We report the Overall Accuracy (OA) in percentage in Table 2.

Table 3. ShapeNetPart dataset

Method	mIoU (%)	
	class	instance
Point-MAE [3]	84.2	86.1
Joint-MAE [1]	85.4	86.3
PointMamba [2]	84.4	86.2
PCM [6]	87.0	85.3
Ours	87.4	86.9

Segmentation on ShapeNetPart. We perform the task of part segmentation on the ShapeNetPart [5] dataset which contains 16,881 shapes from 16 categories, with 50 segmentation parts in total. We report the mean Intersection over Union (mIoU) percentages at both the class level and the instance level in Table 3. Our model demonstrates superior performance by a good margin, especially on the instance level.

References

- [1] 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. In Edith Elkind, editor, *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 791–799. International Joint Conferences on Artificial Intelligence Organization, 8 2023. Main Track. 1

- [2] Ding kang 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. [1](#)
- [3] Yatian Pang, Wenxiao Wang, Francis E. H. Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, page 604–621, Berlin, Heidelberg, 2022. Springer-Verlag. [1](#)
- [4] 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 (CVPR)*, June 2015. [1](#)
- [5] 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. 35(6), Dec. 2016. [1](#)
- [6] Tao Zhang, Haobo Yuan, Lu Qi, Jiangning Zhang, Qianyu Zhou, Shunping Ji, Shuicheng Yan, and Xiangtai Li. Point cloud mamba: Point cloud learning via state space model, 2024. [1](#)
- [7] Xuancheng Zhang, Yutong Feng, Siqi Li, Changqing Zou, Hai Wan, Xibin Zhao, Yandong Guo, and Yue Gao. View-guided point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15890–15899, 2021. [1](#)