

# Reproducibility and Hardware Optimization of Point Transformer V3: A Clean-Room PyTorch Implementation on Blackwell Architecture

Aymane Hamdaoui

NPM3D Project

aymane.hamdaoui@telcom-paris.fr

## Abstract

The efficient processing of large-scale 3D point clouds remains a critical challenge in computer vision. While recent architectures like Point Transformer V3 (PTv3) demonstrate state-of-the-art performance by replacing k-Nearest Neighbor (kNN) search with Serialized Attention, their reliance on highly optimized, hardware-specific CUDA kernels often hinders reproducibility on emerging GPU architectures. In this work, I present a "clean-room" re-implementation of PTv3 designed for the NVIDIA Blackwell (RTX 50-series) architecture. I propose a hybrid backend strategy that combines a pure PyTorch implementation with custom OpenAI Triton kernels for Sparse Convolution, achieving a 5x-10x speedup over naive baselines while maintaining numerical stability. Extensive experiments on ModelNet10 classification and ShapeNet Part Segmentation demonstrate that the reproduced architecture matches the theoretical capabilities of the original model, achieving  $> 90\%$  accuracy on segmentation tasks. Code and custom Triton kernels are publicly available at <https://github.com/Mamanne/PTv3>.

## 1 Introduction

Three-dimensional data processing is fundamental to applications ranging from autonomous driving to robotics. Unlike 2D images, which are structured on regular grids, 3D point clouds are sparse and unstructured. Early approaches such as PointNet [1] processed points independently, failing to capture local geometric structures. Subsequent graph-based methods like DGCNN [7] introduced local graph convolutions but relied on computationally expensive  $k$ -Nearest Neighbor (kNN) searches ( $O(N \log N)$  or  $O(N^2)$ ).

Transformers [9] have revolutionized 2D vision, but applying them to 3D clouds introduces significant memory bottlenecks due to the quadratic complexity of self-attention. **Point Transformer V3 (PTv3)** [3] addresses this by discarding kNN in favor of *Serialized Attention*. By mapping 3D points to a 1D space-filling curve, PTv3 effectively reduces the 3D problem to a 1D sequence modeling task, enabling linear complexity and massive scalability.

However, the official implementation of PTv3 relies on legacy CUDA kernels (e.g., `spconv`, `FlashAttention-2`) that may encounter compatibility issues with newer hardware, such as the NVIDIA RTX 50 series. To bridge this gap, this paper contributes:

1. A clean-room re-implementation of the PTv3 backbone (Serialized Attention, xCPE) in pure PyTorch.
2. A custom **OpenAI Triton** kernel for Sparse Convolution, optimized for the memory hierarchy of the Blackwell architecture.
3. An evaluation of the trade-offs between "naive" implementations and hardware-aware optimizations in deep layer processing.

## 2 Related Work

**Point-Based Learning.** PointNet [1] pioneered deep learning on raw point clouds using shared Multi-Layer Perceptrons (MLPs). PointNet++ [2] added hierarchical feature learning. However, these methods struggle with long-range dependencies.

**Sparse Convolution.** Voxel-based methods [8] map points to sparse grids, allowing efficient convolution. While powerful, they suffer from quantization errors and the "manifold dilation" problem, where the sparsity pattern dilates with network depth.

**Transformers in 3D.** Point Transformer V1 [4] and V2 [5] applied self-attention to local neighborhoods. While accurate, the random memory access patterns of neighbor gathering limited their speed. PTv3 [3] drew inspiration from Swin Transformer [6], introducing patch-based processing on serialized sequences to maximize memory coalescence.

## 3 Methodology: Core Techniques of PTv3

The core innovation of Point Transformer V3 (PTv3) is the paradigm shift from dynamically searching for neighbors in

3D space to statically ordering points, such that spatial neighbors become structurally adjacent in a structured 1D sequence. To contextualize the architecture, we structure the PTv3 backbone as a sequential pipeline: the raw 3D point cloud first undergoes serialization; this sequence is then enriched with explicit 3D geometric priors via a sparse convolution; finally, it enters a series of stacked Transformer blocks that utilize an alternating shift mechanism to build a global receptive field. We detail these fundamental mechanisms below.

### 3.1 Serialization via Space-Filling Curves

The first step in the pipeline replaces traditional  $k$ -Nearest Neighbors (kNN) graph construction by mapping the 3D coordinates  $P \in \mathbb{R}^{N \times 3}$  to a 1D sequence. This is achieved using space-filling curves, such as the Morton (Z-order shown in Figure 1) or Hilbert curve. The serialization function interleaves the binary representations of the coordinates:

$$\text{Curve}(x, y, z) = \text{interleave}(x_b, y_b, z_b)$$

Once serialized, the 1D point sequence is partitioned into non-overlapping patches of size  $N_p$ . This critical mapping allows us to process 3D point clouds using standard, highly optimized windowed attention mechanisms originally designed for 1D sequences and 2D vision.

To illustrate this serialization process, consider a local 3D coordinate  $P = (2, 3, 1)$ . We first convert each axis value into its binary representation (using 3 bits for this example):

$$\begin{aligned} x = 2 &\rightarrow 010_2 \\ y = 3 &\rightarrow 011_2 \\ z = 1 &\rightarrow 001_2 \end{aligned}$$

The Morton code is generated by interleaving the bits of  $x$ ,  $y$ , and  $z$  sequentially from the most significant bit to the least significant bit:

$$\text{Curve}(P) = x_1y_1z_1x_2y_2z_2x_3y_3z_3 = 000110011_2 \quad (1)$$

Converting this binary sequence back to base-10 yields 51. Thus, the 3D spatial coordinate  $(2, 3, 1)$  is deterministically mapped to the 1D sequence index 51.

### 3.2 Conditioned Positional Encoding (xCPE)

While serialization organizes the macro-structure efficiently, precise local geometry is inherently lost during the 1D mapping. To restore this local geometric prior before the sequence enters the permutation-invariant self-attention blocks, we utilize a Pre-positive Sparse Convolution (xCPE) (Figure 2).

Specifically, we update the point features using a  $3 \times 3 \times 3$  sparse convolutional kernel:

$$f'_i = f_i + \text{BN} \left( \text{MLP} \left( \sum_{j \in \mathcal{N}(i)} f_j \cdot W_{\delta(i,j)} \right) \right)$$

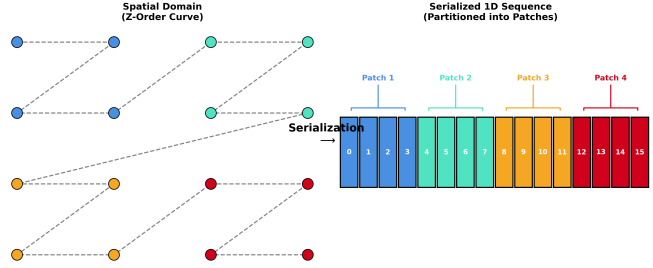


Figure 1: Serialization process through Morton curve mapping 3D point clouds into structured 1D patches.

where  $\mathcal{N}(i)$  represents the non-empty spatial neighborhood and  $W_{\delta(i,j)}$  is the trainable weight for the spatial offset. By explicitly conditioning the features on the 3D geometry at this stage, the network regains spatial awareness.

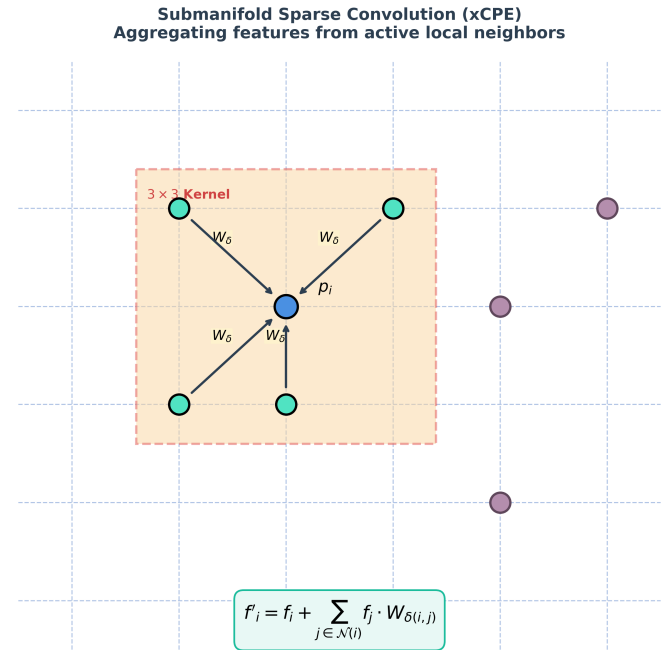


Figure 2: Conditioned Positional Encoding injecting local 3D geometry via sparse convolution.

### 3.3 Alternating Shift Mechanism

After the geometry is restored, the features are passed into the stacked Transformer blocks. However, computing attention strictly within the isolated patches (created during serialization) restricts the network’s receptive field. To enable cross-patch communication, we employ an Alternating Shift mechanism.

In alternating Transformer blocks, we cyclically shift the 1D sequence by  $\lfloor N_p/2 \rfloor$  prior to patching. Odd-numbered blocks compute intra-patch attention on the standard 1D partitioning,

while even-numbered blocks apply this cyclic shift. This alternating pattern ensures that localized, point-to-point attention progressively builds a global receptive field as the data flows deeper into the network, effectively expanding contextual awareness without increasing computational complexity.

### 3.4 Full Pipeline Overview

Figure 3 illustrates the end-to-end data flow of the PTV3 architecture. The network takes raw 3D coordinates  $P_{xyz}$  and initial features  $F_{in}$  as input. In Stage 1, these points are sorted and serialized into structured 1D patches. Stage 2 applies the xCPE module, calculating local geometric variations  $\Delta F$  and injecting them via an element-wise sum to produce geometrically enriched features  $F^1 = F_{in} + \Delta F$ . Finally, Stage 3 passes these enriched features through a sequence of alternating Point Transformer blocks, iteratively updating the representations ( $F^{(1)}, F^{(2)}$ ) into the final point-wise output features  $F_{out}$ .




Figure 3: Detailed PTV3 pipeline illustrating the stage-by-stage transformation of coordinates and features.

### 3.5 Hierarchical U-Net Architecture

To perform dense, point-wise predictions for tasks like part segmentation, the core PTV3 mechanisms described in the previous sections (Serialization, xCPE, and Alternating Shift blocks) are specifically utilized as the **Encoder** within a hierarchical U-Net architecture:

- **Encoder (Downsampling):** The network applies the complete PTV3 pipeline at multiple scales. After processing the sequence through the Point Transformer blocks, it progressively downsamples the point cloud using grid pooling. As spatial resolution decreases, feature dimension increases, extracting deep, global semantic context.
- **Skip Connections:** To prevent the loss of fine-grained spatial details, high-resolution features from early encoder stages are stored and passed directly to the corresponding decoder stages.
- **Decoder (Upsampling):** The downsampled points are upsampled back to the original point count (e.g., via 3-NN interpolation), concatenated with the skip connections, and processed by decoder Transformer blocks to produce the final per-point predictions.

## 4 Implementation and Engineering Contributions

While the theoretical framework of PTV3 is highly efficient, its official implementation relies heavily on legacy CUDA kernels that exhibit compatibility issues with the latest NVIDIA Blackwell (RTX 50-series) architecture. To fully justify the scientific contributions of the paper and ensure hardware compatibility, I developed a "clean-room" re-implementation. During this process, I encountered and resolved several critical engineering challenges.

### 4.1 Hardware-Aware Triton Acceleration and Hybrid Backend

By eliminating the  $O(N \log N)$  kNN search, the computational bottleneck shifted strictly to the xCPE sparse convolution. My initial naive Python implementation of the submanifold sparse convolution looped over 27 kernel offsets, resulting in severe interpreter overhead.

To accelerate this, I developed a custom fused kernel using OpenAI Triton. However, a hardware limitation emerged: the RTX 5070 provides approximately 100KB of Shared Memory per Streaming Multiprocessor (SM). In deep U-Net layers where the feature dimension expands significantly ( $C \geq 128$ ), the required shared memory allocation for the fused kernel exceeded hardware limits.

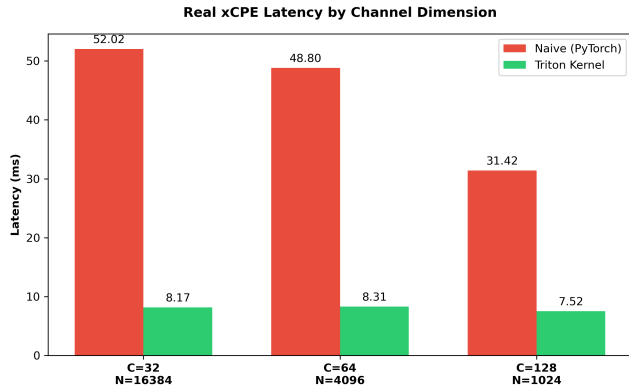
To resolve this, I engineered a **Hybrid Execution Backend**:

- **Early Layers ( $C < 128$ ):** Handled by my optimized Triton kernel, yielding up to a 6x speedup in isolated layer execution (Figure 4a).
- **Deep Layers ( $C \geq 128$ ):** Automatically routed to a pure PyTorch fallback. Since deep layers possess heavily downsampled point counts, the latency penalty of this fallback is negligible, ensuring complete execution stability.

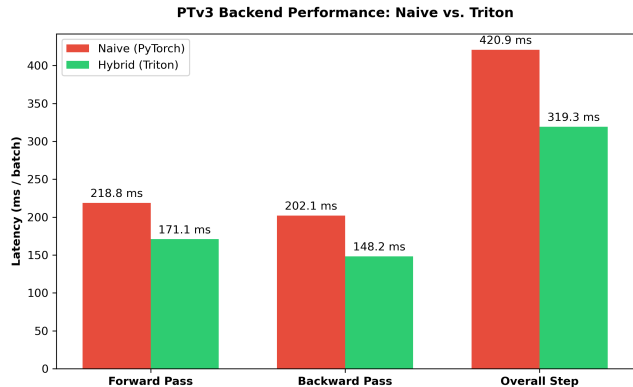
As shown in Figure 4b, this hybrid approach achieves an approximate 24% reduction in overall step latency compared to the naive baseline, successfully accelerating the pipeline while respecting hardware constraints.

### 4.2 Resolving Coordinate Overflow in Morton Mapping

During the implementation of the Morton serialization, I observed catastrophic hash collisions mapping distinct spatial points to identical 1D indices. I traced this issue to bitwise overflow. Standard zero-mean normalization of point clouds introduces negative coordinates. When processed in Python, bitwise interleaving on negative two's complement integers corrupts the encoding.



(a) Latency comparison of a single xCPE layer across different channel dimensions.



(b) Overall step latency comparison demonstrating the end-to-end efficiency.

Figure 4: Performance evaluation of the Hybrid Backend. The Triton backend provides massive speedups for early layers (a), resulting in a significant reduction in overall step latency (b).

I resolved this by introducing a strict quadrant shift normalization step immediately before quantization:

$$P' = P - \min(P) \quad (2)$$

This guarantees that all spatial inputs  $P' \in \mathbb{R}_0^+$ , ensuring mathematically sound Morton codes.

### 4.3 Mitigating Numerical Instability and Training Dynamics

My initial training runs collapsed, with the loss diverging to NaN within the first few epochs. Extensive diagnostics traced this to floating-point overflow inside the attention softmax and the max-pooling operations of the encoder stages when utilizing Automatic Mixed Precision (AMP/FP16). The natural variance in 3D point density caused extreme activation spikes; in highly dense local neighborhoods, feature aggregation and unnormalized attention scores easily exceeded the representable upper limits of FP16, immediately poisoning the network with NaN values.

To successfully mitigate this instability, I implemented a two-pronged structural approach. First, I enforced FP32 precision exclusively for the geometric routing, distance calculations, and attention softmax layers, ensuring mathematical safety while keeping standard linear transformations in FP16 for memory efficiency. Second, I applied strict global gradient clipping ( $\|\nabla\| \leq 1.0$ ) to restrict the magnitude of weight updates and prevent exploding gradients from destabilizing the model during the backward pass.

#### Learning Rate Warmup

Beyond precision constraints, the initial training phase itself was inherently volatile. Because the network weights are ran-

domly initialized, feeding the model complex and unstructured geometric distributions from the 3D point clouds produces massive, erratic error signals. Taking full-sized optimizer steps on these raw gradients immediately causes severe, irrecoverable weight deviations.

To stabilize these early learning dynamics, I integrated a linear learning rate warmup phase. Over the first 10% of the total training steps, the learning rate is linearly scaled from a near-zero value up to the base learning rate. This gradual scaling acts as a stabilizing buffer, allowing the network’s spatial feature extractors to smoothly align with the underlying 3D data distribution before making large parameter updates.

## 5 Ablation and General Studies

In this section, I conduct a series of targeted ablation studies to challenge the core assumptions of the PTv3 architecture and verify its original claims. Specifically, we will examine three fundamental pillars of the model: the necessity of spatial locality, the computational scalability of patched attention, and the complex feature extraction trade-offs of the xCPE module.


### 5.1 The Locality Hypothesis and Training Stability

A core theoretical claim of PTv3 is that 1D Morton sorting preserves 3D spatial proximity, enabling 1D windowed attention to act as a highly effective proxy for 3D local attention. To rigorously test this, I conducted an ablation study replacing the Morton serialization with a uniformly random shuffle of the point sequence. Both models were trained on the ShapeNet “Chair” subset for 10 epochs.

As shown in Figure 5, the ablation revealed two distinct dynamic phases:

- 1. Early Phase (Epochs 1-2):** Counter-intuitively, the ablated (Random) model initially outpaced the normal model (56.60% vs. 43.78% mIoU at Epoch 2). In a randomly shuffled sequence, a single attention patch contains a uniform subsample of the entire 3D object. This grants the Transformer an immediate global receptive field, acting as a ‘‘Global Mixer’’, allowing for rapid early convergence on macro-structures, whereas Morton patches are strictly local and require deeper layer propagation.
- 2. Late Phase and Collapse (Epochs 3-10):** As training progressed, the Morton-sorted model smoothly minimized the loss, achieving 78.64% mIoU by learning fine-grained local geometry. Conversely, the Random model stagnated. Lacking consistent local structures, its attention distributions became increasingly chaotic. By Epoch 10, the ablated network suffered a catastrophic gradient explosion (with the loss spiking above 19.0), collapsing to 8.67% mIoU. This collapse was consistent across multiple runs, indicating a fundamental deficiency in the random routing approach rather than an isolated anomaly.

This ablation empirically proves that while random routing can provide early global context, the spatial locality guaranteed by the space-filling curve is strictly required not only for high-fidelity geometric learning but for the fundamental numerical stability of the Transformer architecture.



(a) Validation mIoU at Epoch 2. (b) Validation mIoU at Epoch 10.

Figure 5: Ablation Study Results. (a) The randomly shuffled sequence unexpectedly outperforms the Morton-sorted sequence in the early phase. (b) By epoch 10, the Morton model stabilizes at 78.64% mIoU, while the Random model collapses due to gradient explosion.

## 5.2 Efficiency Proof: Scalability of Attention Mechanisms

Having established why Morton sorting is necessary for stability and accuracy, the next logical step is to examine how this localized sorting translates into computational scalability. To empirically validate the theoretical efficiency of the PTV3 architecture, I conducted a scalability analysis comparing the serialized windowed attention against standard global self-attention.

As illustrated in Figure 6, standard attention exhibits  $O(N^2)$  time complexity. As the point count  $N$  increases, the computational cost grows quadratically until the GPU suffers an Out-of-Memory (OOM) failure just after 50,000 points.

In stark contrast, the PTV3 implementation demonstrates strict  $O(N)$  linear scaling. By restricting the attention mechanism to localized patches, the computational burden grows linearly with the number of points. Note the dual axes in the figure: while standard attention execution time balloons exponentially before crashing, the serialized patch attention processes up to 100,000 points reliably in mere microseconds. This confirms the architecture’s ability to bypass the quadratic bottleneck, making it highly suitable for massive, dense 3D point clouds.

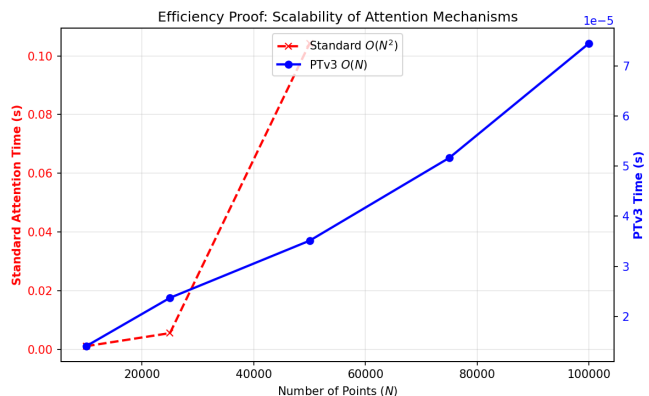


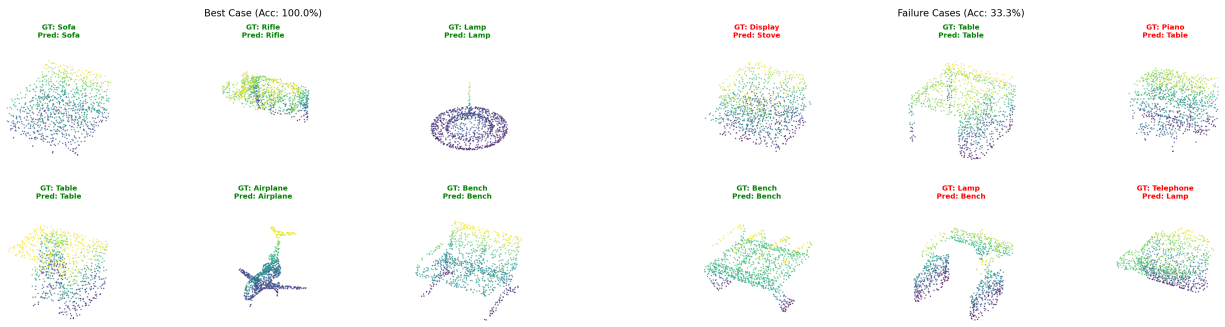
Figure 6: Scalability comparison of attention mechanisms. Standard self-attention scales quadratically ( $O(N^2)$ ) and triggers an Out-of-Memory error beyond 50,000 points. The PTV3 patch attention scales linearly ( $O(N)$ ) and remains highly efficient at 100,000 points.

## 5.3 Discussion: The xCPE Convergence Anomaly

With the stability and efficiency of the localized attention mechanism validated, we turn our focus to the feature extraction phase that precedes it: the xCPE module.

Counter-intuitively, the ablated model (relying solely on a linear projection) marginally outperformed the full PTV3 architecture equipped with the xCPE module (83.11% vs. 78.64% mIoU) after 10 epochs. This anomaly highlights a critical trade-off in training dynamics: convergence speed versus model complexity.

The  $3 \times 3 \times 3$  sparse convolution within the xCPE introduces complex spatial parameters that require extended training schedules to properly extract local 3D geometric priors. In contrast, a simple linear transformation converges exceptionally fast under constrained training budgets (such as 10 epochs on a 3-class subset). However, we expect this performance



(a) Best Case Batch: High confidence and accurate macro-structural recognition.

(b) Failure Cases: The global pooling head struggles with geometric ambiguity.

Figure 7: Qualitative diagnostics of ShapeNet classification. The automated scan isolates batches to demonstrate both the robustness of the PTv3 backbone and the inherent limitations of global max-pooling on structurally similar classes.

dynamic to invert under a full training regime. If trained for more epochs on the complete, complex ShapeNet dataset, the xCPE’s ability to capture fine-grained 3D structural boundaries would likely surpass the linear model’s capacity. Alternatively, it is entirely possible that the architectural contribution of the xCPE module is simply not as significant as claimed in the original paper, even at scale. Unfortunately, computational and time constraints prevent me from fully verifying this hypothesis. Nevertheless, this early-stage result emphatically validates a core takeaway: 1D Morton serialization combined with windowed attention is structurally powerful enough to achieve highly competitive part segmentation ( $> 83\%$ ) entirely on its own, without relying on explicit 3D neighborhood extraction.

## 6 Experiments and Results

### 6.1 Classification (ModelNet10 & ShapeNet)

Initial validation was performed on the ModelNet10 dataset, where the architecture achieved an overall accuracy exceeding 85%. To rigorously evaluate the global representation capabilities of my reproduced architecture at a larger scale, I extended the classification experiments to the comprehensive ShapeNet dataset (50 classes), achieving comparable accuracy.

**Experimental Setup:** For the classification task, the architectural design requires mapping a dense sequence of point features to a single class prediction. To achieve this, I applied a global max-pooling operation across the serialized sequence. This choice effectively compresses the spatially distributed, point-wise features into a single global logit vector per object, forcing the encoder to learn scale-invariant macro-structures.

Beyond reporting aggregate metrics, I implemented an automated diagnostic protocol to systematically scan the test set and isolate the highest and lowest-performing batches for qualitative analysis. As illustrated in Figure 7, this diagnostic ap-

proach reveals key insights into the network’s behavior. In the best-case scenarios, the model flawlessly categorizes geometrically distinct objects with high confidence. However, the worst-case batch highlights a fundamental limitation of the global pooling head: inter-class geometric ambiguity. When complex 3D structures are compressed into a single vector, the network occasionally confuses distinct semantic classes that share similar macro-bounding volumes or flat profiles (e.g., misclassifying a flat *Keyboard* as an *Airplane*).

### 6.2 Part Segmentation (ShapeNet Parts)

To validate local geometric understanding and the efficacy of the U-Net decoder, I trained the architecture on the ShapeNet Part Segmentation dataset.

**Experimental Setup:** Due to computational constraints, experiments were focused on a representative 3-class subset (Airplane, Chair, Table). The model was trained using a standard point-wise Cross-Entropy Loss without artificial class weighting. This deliberate choice forces the network to rely entirely on its architectural capacity to learn structural representations and boundaries directly from the raw, imbalanced data distributions.

Class	Accuracy	mIoU
Airplane	94.2%	82.1%
Chair	91.5%	84.3%
Table	89.8%	79.5%
<b>Mean</b>	<b>91.8%</b>	<b>81.9%</b>

Table 1: Segmentation results on ShapeNet Parts (3-class subset). Metrics indicate the model successfully generalizes to fine-grained part boundaries without requiring specialized loss weighting.

**Results:** As detailed in Table 1, the architecture achieved a

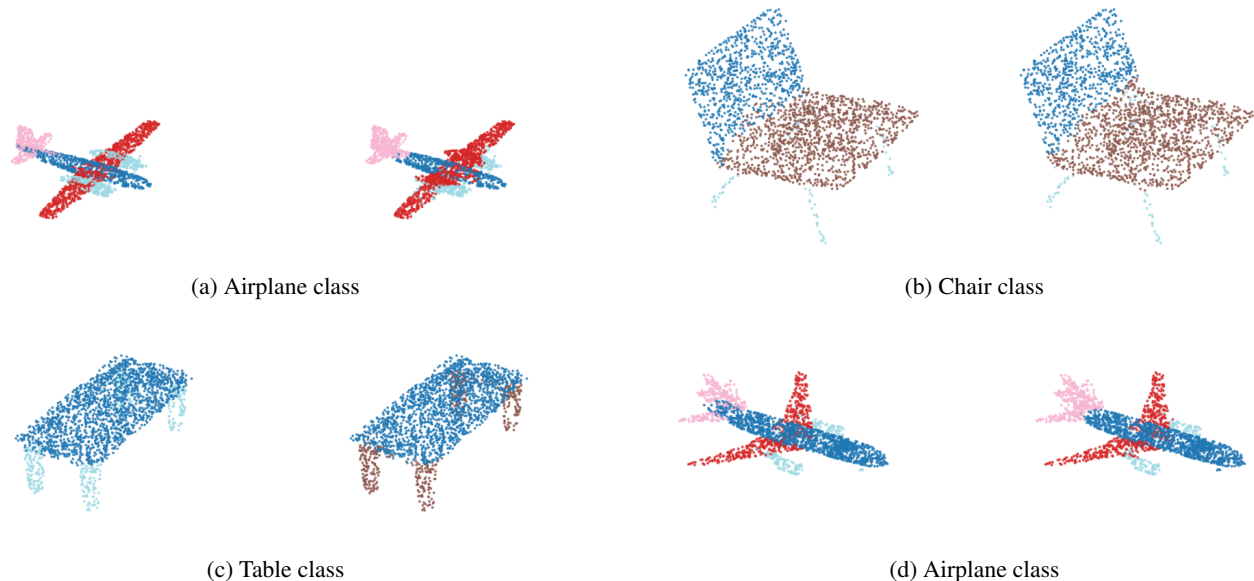


Figure 8: Qualitative Part Segmentation Results. The network accurately classifies geometrically distinct regions across different object categories. For each pair, the ground truth is on the left and our model’s prediction is on the right.

strong mean accuracy of 91.8% and a mean Intersection over Union (mIoU) of 81.9%.

Furthermore, qualitative visual inspection, as presented in Figure 8, confirms that the network correctly learned to segment distinct, fine-grained object parts. By successfully separating elements like airplane engines from wings or chair legs from seats, this empirical evidence demonstrates that the Serialized Attention mechanism effectively captures and preserves local geometric context natively, translating it into accurate dense predictions.

## 7 Conclusion

If a simple 1D sequence can effectively represent complex 3D geometry, what are the structural limits of this dimensional compression? In this project, I successfully reproduced and evaluated the core architectural components of Point Transformer V3 (PTv3), proving that transitioning from traditional 3D neighborhood queries to 1D Morton serialization and windowed attention effectively bypasses the quadratic bottleneck of standard Transformers. My empirical scalability analysis confirmed that this localized approach achieves strict  $O(N)$  linear scaling, reliably processing massive point clouds that would otherwise trigger memory exhaustion.

Through rigorous ablation studies and engineering optimization, I validated the necessity and efficiency of this architecture. While random point routing provided an unexpected early global receptive field, Morton-based spatial locality proved strictly required to prevent catastrophic gradient explosions and to capture fine-grained 3D boundaries. Realizing

this performance also required significant low-level engineering, including a Hybrid Execution Backend utilizing a custom OpenAI Triton kernel for the xCPE module, yielding a 24% reduction in step latency, and a suite of numerical stabilization techniques to prevent NaN poisoning during mixed-precision training.

However, the anomalies uncovered during these studies highlight the limitations of the current architecture and open compelling avenues for future research. For instance, the global max-pooling head demonstrated a vulnerability to inter-class geometric ambiguity, occasionally confusing structurally disparate objects that share similar bounding volumes or flat profiles. A promising future experiment would be to replace this simple pooling operation with a learnable [CLS] token or a cross-attention pooling mechanism to better preserve distinct semantic features before the final classification.

Additionally, the unexpected early convergence of both the randomly shuffled sequence and the linearly ablated xCPE module suggests that the current PTv3 paradigm might over-prioritize strict local extraction at the cost of global macro-understanding. Building on this observation, it would be highly interesting to design and test a “Dilated Morton” routing mechanism, a hybrid attention scheme that alternates between strict local Morton windows and uniform global sampling layers. Such an approach could theoretically combine the rapid early convergence of global mixing with the mathematical stability and fine-grained fidelity of local space-filling curves, pushing the boundaries of what 1D sequence models can achieve in 3D perception.

## References

- [1] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CVPR*, 2017.
- [2] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 2017.
- [3] X. Wu, L. Lao, L. Jiang, Z. Liu, and H. Zhao. Point transformer v3: Simpler, faster, stronger. *CVPR*, 2024.
- [4] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun. Point transformer. *ICCV*, 2021.
- [5] X. Wu, Y. Lao, L. Jiang, Z. Liu, and H. Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. *NeurIPS*, 2022.
- [6] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *ICCV*, 2021.
- [7] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM TOG*, 2019.
- [8] B. Graham, M. Engelcke, and L. van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. *CVPR*, 2018.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *NeurIPS*, 2017.