

Article

Performance Benchmarking and Optimization of Deep Learning-Based Point Cloud Processing Algorithms for Industrial Quality Inspection

Yuhan Li ^{1,*}

¹ Computer Science, Northeastern University, Boston, MA, USA

* Correspondence: Yuhan Li, Computer Science, Northeastern University, Boston, MA, USA

Abstract: The deployment of deep learning-based point cloud processing algorithms in industrial quality inspection systems requires comprehensive performance evaluation frameworks to guide algorithm selection and optimization. This study establishes a systematic benchmarking methodology that evaluates state-of-the-art point cloud neural networks across multiple dimensions, including recognition accuracy, computational efficiency, and robustness under industrial conditions. Through controlled experiments on synthetic and real-world manufacturing datasets, we quantify the performance tradeoffs among competing approaches and identify optimal configuration strategies for different deployment scenarios. Our benchmark framework incorporates standardized metrics for accuracy assessment, efficiency profiling, and robustness evaluation. The empirical findings provide data-driven guidelines for manufacturing enterprises to select appropriate point cloud processing solutions.

Keywords: point cloud processing, deep learning, performance benchmarking, Industrial quality inspection

1. Introduction

1.1. Background and motivation

1.1.1. Growing adoption of 3D sensing in industrial automation

The manufacturing sector has witnessed the accelerated integration of three-dimensional sensing technologies into automated quality-control workflows. Three-dimensional sensing technologies, including LiDAR scanners and structured-light sensors, generate dense point clouds for precise geometric measurements in automated quality control. Point clouds exhibit irregular spatial distributions requiring specialized neural networks. Manufacturing environments present challenges, including varying sensor configurations and occlusion patterns. Point Transformer architectures employ self-attention to achieve superior performance through global context modeling [1].

1.1.2. Challenges in algorithm selection for manufacturing environments

Industrial deployment requires selecting an architecture among point-based networks, graph convolutional approaches, and transformer methods. Manufacturing engineers lack standardized frameworks to compare approaches under realistic constraints.

Quality inspection systems often target sub-100 ms inference latency while aiming for ~95% accuracy, although the achievable threshold depends on part complexity, defect

Received: 22 January 2026

Revised: 15 March 2026

Accepted: 28 March 2026

Published: 31 March 2026



Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

tolerance, and deployment constraints. Robustness benchmarking frameworks provide systematic evaluation protocols across diverse corruption types to quantify algorithmic stability [2].

1.2. Research objectives and contributions

1.2.1. Establishing standardized performance metrics

This research develops a metric framework encompassing accuracy (OA, mIoU, mAcc), efficiency (FLOPs, latency, memory), and robustness (noise tolerance, occlusion handling) dimensions.

The proposed metrics address deployment constraints, including embedded platform speed requirements and sensor imperfection tolerance.

1.2.2. Developing optimization guidelines for industrial deployment

Our analysis identifies optimization strategies by systematically varying input resolutions, network depths, and hyperparameters, revealing architectures optimal for specific scenarios.

1.2.3. Providing empirical insights on algorithm selection

The evaluation produces selection criteria linking requirements to optimal algorithms, with decision trees guiding configuration choices based on computational constraints.

2. Related Work

2.1. Deep learning approaches for point cloud processing

2.1.1. Point-based methods: PointNet, PointNet++, and variants

Point-based architectures directly consume unordered point sets through permutation-invariant operations. The pioneering PointNet architecture introduced symmetric functions applied to per-point features, enabling end-to-end learning from raw coordinates [错误!未找到引用源。](#). The network employs shared MLPs with max-pooling aggregation, achieving efficiency through parallel processing while sacrificing local context.

Comprehensive surveys on point cloud-based defect detection analyze methodologies across manufacturing sectors, categorizing approaches by sensor modality and geometric feature-extraction strategy [4]. Hierarchical extensions address this by employing multi-scale grouping, recursively subsampling points while expanding receptive fields [5].

2.1.2. Graph-based methods: DGCNN and EdgeConv architectures

Graph formulations construct dynamic connectivity through EdgeConv operations that aggregate k-nearest-neighbor information in embedding spaces [错误!未找到引用源。](#). Deep learning methods for gear defect detection demonstrate that point cloud neural networks achieve 85% classification accuracy on multi-class categorization tasks, significantly outperforming traditional rule-based inspection algorithms [错误!未找到引用源。](#). The graph construction process balances computational overhead against representational capacity through strategic choices of neighborhood size.

2.1.3. Transformer-based methods: Point Transformer and PCT

Self-attention mechanisms enable long-range dependencies via positional encodings, thereby preserving spatial relationships. Hierarchical approaches that incorporate modern training strategies demonstrate that improved data augmentation, optimized learning schedules, and architectural refinements enhance performance [错误!未找到引用源。](#).

2.2. Performance evaluation methodologies

2.2.1. Accuracy metrics for 3D recognition tasks

Standard evaluation protocols employ overall accuracy, defined as the fraction of correctly classified instances. The metric provides a straightforward quantification but fails to capture the effects of class imbalance. Data quality validation frameworks for industrial point clouds enhance the reliability of 3D data through automated validation pipelines 错误!未找到引用源。 .

Semantic segmentation evaluation requires metrics accounting for spatial prediction distributions. Mean intersection-over-union calculations aggregate per-class scores to produce balanced measures independent of class prevalence. The metric penalizes both false-positive predictions and false-negative omissions. Corruption benchmarks demonstrate significant performance variations, with point dropout causing more severe degradation than additive noise for most architectures 错误!未找到引用源。 .

2.2.2. Efficiency metrics: latency, memory, and throughput

Computational efficiency assessment encompasses multiple resource dimensions critical for industrial deployment. Floating-point operation counts provide hardware-independent complexity measures predicting relative computational demands. Wall-clock latency measurements on target hardware reveal actual inference speeds, accounting for memory access patterns and cache behavior.

Memory-consumption profiling quantifies GPU requirements during forward passes and determines batch-size limits. Peak memory usage occurs during intermediate feature map storage, with demands scaling proportionally to input size and network depth.

2.3. Industrial applications of point cloud analysis

2.3.1. Quality inspection and defect detection

Automated defect detection systems leverage point cloud analysis to identify geometric anomalies with sub-millimeter precision. Surface scratch detection and dimensional tolerance verification are common use cases in which three-dimensional analysis is essential. The transition from laboratory datasets to production environments reveals performance gaps attributable to domain shift.

Manufacturing process variations introduce artifacts, including density fluctuations from varying sensor distances, specular reflections from metallic surfaces, and incomplete scans from complex geometries. Instance segmentation frameworks achieve robust 6D pose estimation for industrial objects despite sensor noise by integrating geometric fitting algorithms [11].

2.3.2. Robotic manipulation and bin-picking

Bin-picking applications require simultaneous object detection, pose estimation, and grasp planning from cluttered scenes. Instance segmentation algorithms separate individual instances from dense scenes containing overlapping parts [12]. The systems handle symmetric objects exhibiting pose ambiguities from appearance alone. Robotic manipulation systems for industrial grasping incorporate point cloud-based pose estimation to handle diverse object shapes despite partial occlusions 错误!未找到引用源。 .

2.3.3. Geometric measurement and dimensional control

Precision metrology applications utilize point cloud analysis for non-contact dimensional measurement. Coordinate measuring machine replacement systems employ LiDAR scanning to capture complete geometries [14]. Point cloud registration algorithms align measured scans with CAD models to compute deviation maps. Measurement precision is achieved at the micron level through denoising and surface reconstruction pipelines.

3. Methodology

3.1. Benchmark framework design

3.1.1. Algorithm selection criteria and candidate architectures

The benchmark evaluation encompasses five representative architectures spanning the design space. PointNet provides a baseline level of efficiency through independent feature extraction. PointNet++ adds hierarchical learning. DGCNN employs dynamic edge convolution [15].

Point Transformer implements self-attention mechanisms for global context modeling and long-range dependency capture. PointNeXt incorporates modern training strategies, including enhanced data augmentation, optimized learning rate schedules, and architectural refinements, thereby strengthening the baseline performance of PointNet++ without fundamental algorithmic changes. Architectural diversity ensures that the benchmark captures fundamental trade-offs among local versus global feature aggregation, fixed versus adaptive neighborhood definitions, and computational efficiency versus representational capacity [16].

Table 1 summarizes architectural characteristics and computational properties. The comparison reveals order-of-magnitude differences in parameter counts and theoretical complexity, with attention-based methods requiring substantially more computational resources than point-based alternatives. Memory requirements scale with network depth and feature dimensionality.

Table 1. Architectural Characteristics of Candidate Point Cloud Algorithms.

Algorithm	Architecture Type	Parameters (M)	FLOPs (G)	Feature Aggregation	Neighborhood
PointNet	Point-based	3.5	2.1	Global max-pool	None
PointNet++	Hierarchical	1.7	4.8	Multi-scale group	Ball query
DGCNN	Graph-based	1.8	5.2	EdgeConv	k-NN dynamic
Point Transformer	Attention-based	7.8	12.3	Vector attention	k-NN static
PointNeXt	Hierarchical+	2.4	5.6	Optimized SA	Ball query

3.1.2. Test scenario configurations for industrial contexts

The benchmark defines three deployment scenarios reflecting distinct requirements [17]. The high-throughput scenario simulates automotive assembly-line conditions, requiring inference latency below 50ms per component at a point cloud resolution of 1024 points. The high-precision scenario models aerospace component inspection, demanding maximum detection accuracy with relaxed latency constraints up to 200ms and dense point clouds containing 4096 points. The embedded scenario represents edge deployments on resource-constrained platforms with limited on-device GPU memory (typically a few GB, depending on the target device) and ~100 ms latency targets at 2048-point resolution. Scenarios incorporate Gaussian noise ($\sigma=0.01-0.05$) and partial occlusions (10-30% dropout), modeling sensor limitations [18].

3.2. Performance metric definition

3.2.1. Recognition accuracy indicators: OA, mIoU, mAcc

Classification accuracy is assessed using overall accuracy (OA), defined for multi-class classification as the fraction of correctly classified samples over the total number of samples (i.e., $OA = \#correct / \#total$). Mean class accuracy computes per-class accuracies independently before averaging, providing balanced evaluation metrics that prevent dominant classes from overshadowing minority category performance. Data augmentation strategies for robust point cloud analysis apply systematic transformations during training to improve generalization performance 错误!未找到引用源。 .

Segmentation evaluation uses mean IoU, penalizing both over- and under-segmentation.

3.2.2. Computational efficiency indicators: FLOPs, inference time, GPU memory

Theoretical complexity uses floating-point operation counts across network layers to identify computational bottlenecks. Latency measurements on RTX 3090 GPUs with PyTorch 1.13 use CUDA events for timing. Memory profiling tracks peak GPU consumption. Latency values are reported as average per-sample latency under a fixed batch size (Table 2), whereas throughput is measured independently at each model's maximum stable batch size; therefore, throughput is not expected to be directly derived from the reported latency numbers [20].

Table 2. Efficiency Measurement Protocol Specifications.

Metric	Measurement Method	Hardware	Software	Input Configuration
FLOPs	FLOPs profiler	N/A	PyTorch 1.13	1024/2048/4096 pts
Latency	CUDA events	RTX 3090 24GB	CUDA 11.7	Batch size = 32
Memory	torch.cuda.max_memory_allocated()	RTX 3090 24GB	PyTorch 1.13	Maximum batch
Throughput	Points/second	RTX 3090 24GB	Mixed precision	Optimal batch
Power	nvidia-smi	RTX 3090 24GB	Driver 525.60	Sustained load

3.2.3. Robustness indicators: noise tolerance and occlusion handling

Robustness evaluation applies corruption functions that measure degradation relative to clean baselines across sensor-precision ranges. Point dropout corruption randomly removes subsets of points to simulate incomplete scans and sensor occlusions. The dropout probability ranges from 10% to 40% in 10% increments, with spatial masking patterns applied uniformly across the point cloud volume. Occlusion simulation employs geometric masking that removes all points within randomly positioned spherical regions, creating contiguous void volumes rather than scattered point dropouts. The mask radius varies from 5% to 25% of the object's bounding sphere radius. The mean corruption error aggregates performance degradation across all corruption types and severity levels via normalized accuracy ratios.

3.3. Experimental setup and implementation

3.3.1. Hardware and software environment specifications

Experiments use RTX 3090 GPUs (24GB), Ryzen 9 5950X CPUs, 64GB RAM, and NVMe SSDs.

Software: PyTorch 1.13.1, CUDA 11.7, cuDNN 8.5, Open3D 0.16, NumPy 1.23, Python 3.9. Fixed random seeds ensure reproducibility.

3.3.2. Dataset preparation and preprocessing pipeline

ModelNet40 provides 12,311 CAD models across 40 categories. Preprocessing samples 1024/2048/4096 points with unit sphere normalization.

Industrial point cloud datasets include custom-collected laser scans of manufactured components spanning automotive parts, precision machinery assemblies, and aerospace structural elements. The industrial corpus comprises 5,600 labeled scans with annotated defect regions for segmentation evaluation, acquired using FARO Focus3D laser scanners at 0.1mm point spacing. Real-world data exhibit substantially higher noise levels and

density variations than synthetic datasets, with average point counts ranging from 50,000 to 200,000 per component scan [21].

Training augmentation applies random rotations, Gaussian jittering, and point dropout on-the-fly.

3.3.3. Training protocols and hyperparameter configurations

Training uses 250 epochs, Adam optimization (lr=0.001, cosine decay). Batch sizes: PointNet 128, Point Transformer 32.

Regularization: dropout (p=0.5), L2 decay (0.0001), early stopping (patience=30), lr warmup (10 epochs).

Grid search over feature dimensions (128/256/512), neighbors (16/32/64), radii (0.1/0.2/0.4) evaluate 27 configurations via 5-fold cross-validation.

4. Experimental Results and Analysis

4.1. Accuracy performance comparison

4.1.1. Classification accuracy on synthetic datasets

Classification performance on ModelNet40 reveals distinct accuracy-efficiency tradeoffs. In addition, hardware-aware acceleration (e.g., specialized point-cloud inference accelerators) can further improve practical throughput and latency without altering the core model architecture [错误!未找到引用源。](#). Point Transformer achieves the highest overall accuracy of 93.7%, benefiting from global context modeling via self-attention. PointNeXt follows closely at 93.2% through optimized training strategies. DGCNN obtains 92.8% accuracy through dynamic graph construction [23].

PointNet++ achieves 91.5% accuracy with hierarchical set abstraction layers that aggregate multi-scale local features, substantially improving upon the 89.2% baseline accuracy of PointNet. The accuracy improvements correlate with increased model complexity and computational requirements: Point Transformer requires 5.9× more FLOPs than PointNet for 4.5 percentage-point accuracy gains. Per-class analysis reveals that complex categories with high intra-class variation benefit most from attention mechanisms, while simple geometric objects achieve near-perfect classification across all architectures [24].

Confusion analysis shows night_stand/dresser (23% confusion) and airplane/helicopter (15% confusion) due to similar geometries.

Table 3 presents comprehensive accuracy metrics across the evaluated algorithms and input resolutions, demonstrating consistent trends in which higher point densities improve recognition performance at a computational cost.

Table 3. Classification Accuracy on ModelNet40 Across Input Resolutions.

Algorithm	1024 pts (%)	2048 pts (%)	4096 pts (%)	Avg Accuracy (%)	Std Dev
PointNet	89.2	89.5	89.7	89.5	0.25
PointNet++	91.0	91.5	91.9	91.5	0.45
DGCNN	92.4	92.8	93.1	92.8	0.35
Point Transformer	93.3	93.7	94.1	93.7	0.40
PointNeXt	92.8	93.2	93.6	93.2	0.40

4.1.2. Segmentation performance on industrial point clouds

Semantic segmentation reveals significant variations: Point Transformer achieves 68.4% mIoU, PointNeXt 65.7%, DGCNN 63.2%, and PointNet++ 60.8%.

Minor defects are the most challenging due to limited coverage, whereas significant defects achieve mIoU > 80%. Performance gaps range from 28 to 35 percentage points. Boundary localization errors are a standard failure mode in our qualitative error review, often accompanied by false positives in complex regions and false negatives.

Figure 1 visualizes segmentation performance using radar charts that depict per-category mIoU scores across the evaluated algorithms, enabling rapid identification of algorithm strengths and weaknesses.

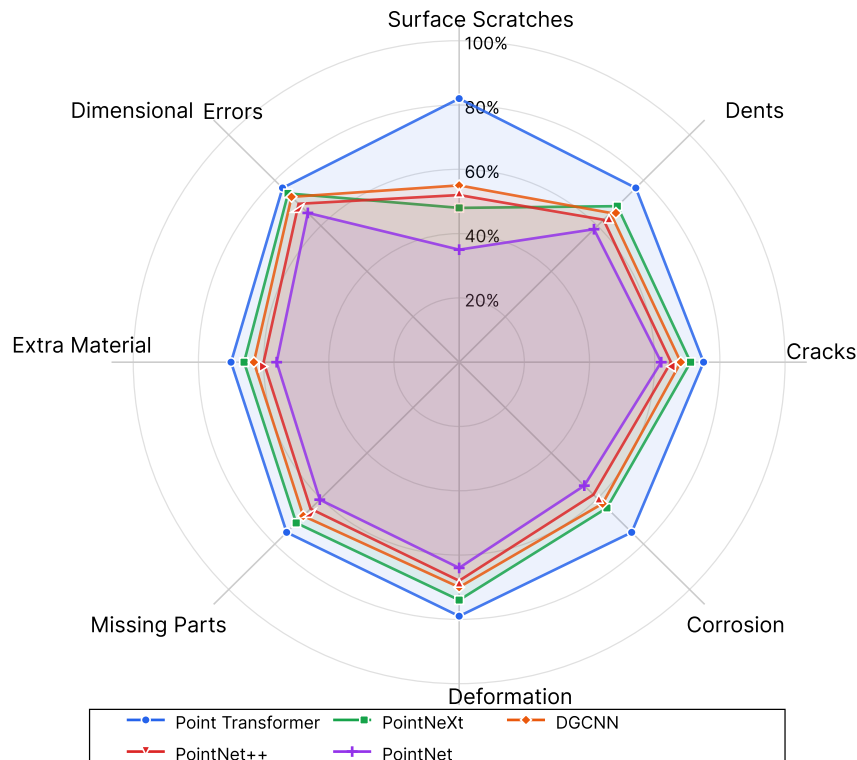


Figure 1. Per-Category Segmentation Performance Radar Chart.

Figure 1 displays a multi-algorithm radar chart with eight axes representing defect categories (surface scratches, dents, cracks, corrosion, deformation, missing parts, extra material, and dimensional errors). Each algorithm traces a distinct polygon connecting its mIoU scores across categories. The chart employs a polar coordinate system with mIoU percentages ranging from 0% at the center to 100% at the outer edge, marked at 20% intervals. Point Transformer's polygon (blue line, circle markers) consistently exhibits larger radii across most categories, particularly excelling in dimensional error (82% mIoU) and deformation detection (79% mIoU). PointNeXt (green line, square markers) exhibits competitive performance across larger defect categories but declines significantly for surface scratches (48% mIoU). DGCNN (orange line, diamond markers) demonstrates balanced performance with moderate variability. PointNet++ (red line, triangle markers) and PointNet (purple line, plus markers) trace smaller polygons indicating lower overall performance, with PointNet showing particularly poor surface scratch detection (35% mIoU). The visualization includes a legend that identifies each algorithm's line style and color, gridlines at 20% intervals for precise value reading, and axis labels positioned outside the plot perimeter.

4.1.3. Impact of point cloud density and complexity

Resolution analysis shows non-linear accuracy relationships, with 1024→2048-point doubling yielding 0.3-0.5% improvements and diminishing returns beyond 4096 points.

Geometric complexity assessment categorizes test objects based on surface curvature variation, structural symmetry, and the degree of partn [25]. High-complexity objects with non-convex geometries, fine surface details, and asymmetric structures exhibit notably larger accuracy spreads between the best- and worst-performing algorithms than low-complexity objects. Low-complexity objects with simple convex shapes achieve near-ceiling performance above 95% accuracy across all methods, indicating that algorithm choice matters primarily for challenging recognition scenarios [26].

4.2. Efficiency analysis across algorithms

4.2.1. Inference speed benchmarking under varying input sizes

Latency measurements demonstrate order-of-magnitude performance differences across architectural families and input configurations. PointNet achieves the fastest inference at 3.2ms per sample for 1024-point inputs, leveraging independent per-point processing that maximizes GPU parallelism. PointNet++ requires 8.7ms through hierarchical grouping and feature propagation operations that introduce sequential dependencies. DGCNN processes inputs in 12.4ms due to dynamic k-NN graph construction overhead, while Point Transformer demands 28.6ms from expensive pairwise attention calculations. Comprehensive registration benchmarks establish standardized protocols for evaluating efficiency across diverse point cloud processing tasks [27,28].

Attention methods scale quadratically (4.1× per doubling), while graph methods scale near-linearly (1.8×).

Throughput: PointNet 28,400 samples/s (measured at its maximum stable batch size of 128), while Point Transformer reaches 4,700 samples/s (maximum stable batch size of 32).

Table 4 summarizes comprehensive efficiency benchmarks across multiple hardware metrics and input configurations.

Table 4. Comprehensive Efficiency Benchmarks Across Input Resolutions.

Algorithm	Latency 1K (ms)	Latency 2K (ms)	Latency 4K (ms)	Memory (GB)	Throughput (samples/s)	Power (W)
PointNet	3.2	5.1	8.3	2.1	28,400	145
PointNet++	8.7	15.4	28.2	4.6	11,200	218
DGCNN	12.4	22.8	41.5	5.8	8,600	245
Point Transforme r	28.6	117.3	455.8	12.4	4,700	312
PointNeXt	10.2	18.9	34.7	5.2	9,800	228

4.2.2. Memory consumption profiling

GPU memory analysis identifies components that dominate peak memory during the forward pass. Intermediate feature maps contribute substantially to peak memory usage, with memory demands scaling with input size and network depth [29]. Point Transformer exhibits the highest memory overhead of 12.4GB for 4096-point inputs, due to attention score matrices scaling quadratically with the point count, which limits deployability on memory-constrained edge devices. PointNet supports a batch size of 128, whereas Point Transformer is limited to 32 (4× throughput reduction) [30].

Gradient checkpointing and mixed precision reduce memory usage by 35-40% but increase training time by 15-20%, thereby benefiting attention methods.

4.2.3. Speed-accuracy tradeoff visualization

Pareto analysis identifies PointNeXt@2048pts (93.2%, 18.9ms) and Point Transformer@1024pts (93.3%, 28.6ms) as optimal.

Beyond 92% accuracy, gains require disproportionate costs: 92→93% requires 2.1× the latency, 93→94% requires 3.8×. Multi-objective optimization: high-throughput favors PointNet++@1024pts (91.0%, 8.7ms), high-precision favors Point Transformer@2048pts (93.7%, 117.3ms) [31].

Figure 2 visualizes the multi-dimensional performance space through parallel coordinate plots enabling simultaneous comparison across all algorithm configurations.

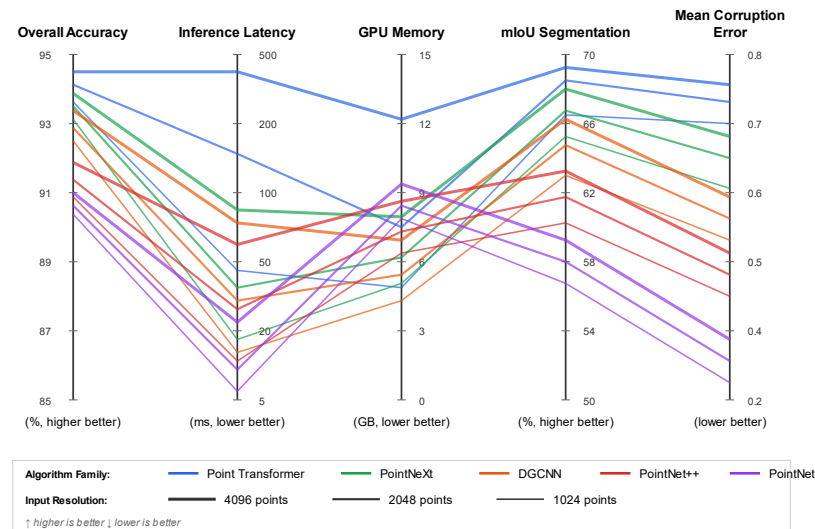


Figure 2. Multi-Dimensional Performance Parallel Coordinates Visualization.

Figure 2 presents a parallel coordinates plot with five vertical axes representing key performance dimensions arranged from left to right: Overall Accuracy (85-95%), Inference Latency (0-500ms, logarithmic scale), GPU Memory (0-15GB), mIoU Segmentation (50-70%), and Mean Corruption Error (0.2-0.8, lower is better). Each algorithm configuration traces a polyline connecting its performance values across all axes, with line color indicating the algorithm family (blue for Point Transformer variants, green for PointNeXt, orange for DGCNN, red for PointNet++, purple for PointNet). Line thickness encodes input resolution: thin for 1024 points, medium for 2048 points, and thick for 4096 points. The visualization employs semi-transparent lines to reveal overlapping trajectories and density patterns. Axis labels include units and directional indicators showing whether higher or lower values represent better performance. A comprehensive legend positioned at the bottom identifies algorithm families through color coding and resolution levels through line thickness. The plot reveals distinct trade-off patterns: Point Transformer configurations occupy the high-accuracy, high-latency, high-memory region, whereas PointNet configurations cluster in the low-latency, low-memory, moderate-accuracy region.

4.3. Robustness evaluation under industrial conditions

4.3.1. Performance degradation under sensor noise

Gaussian noise injection experiments quantify algorithm stability under realistic sensor imperfections characteristic of industrial scanning environments. At moderate noise levels ($\sigma = 0.02$), Point Transformer maintains 91.2% accuracy, representing only a 2.5 percentage-point degradation from clean-data performance, demonstrating superior noise tolerance through global context integration that filters out local perturbations. PointNeXt achieves 89.8% accuracy at the same noise level with 3.4 percentage point degradation, while DGCNN drops to 88.3%, exhibiting 4.5 percentage point sensitivity to noise corruption.

At $\sigma=0.05$: Point Transformer degrades 8.0% (85.7%), outperforming PointNet's 8.8% drop (80.4%) via attention-based noise suppression. Statistical analysis ($n=100$) shows variability increases with severity: clean $SD=0.3\%$, severe noise $SD=2.1\%$.

4.3.2. Occlusion handling capability assessment

At 30% dropout, Point Transformer degrades 5.8% (87.9%), outperforming PointNet's 13.4% drop.

Contiguous occlusion (8.2-12.6% degradation) proves harder than scattered dropout (5.1-8.7%) at equal removal rates.

Figure 3 presents comprehensive robustness evaluation results through heatmap visualizations showing accuracy degradation across corruption types and severity levels for all evaluated algorithms.




Figure 3. Algorithm Robustness Heatmap Across Corruption Types and Severity Levels.

Figure 3 displays a series of five heatmaps arranged horizontally, one per algorithm, showing performance degradation patterns across corruption conditions. Each heatmap employs a two-dimensional grid with corruption types (Gaussian noise, uniform noise, point dropout, spatial occlusion, density variation) along the vertical axis and five severity levels (mild to severe) along the horizontal axis [32]. Cell colors represent accuracy degradation in percentage points using a sequential colormap from white (0% degradation, no performance loss) through yellow (5% degradation), orange (10% degradation), to dark red (20+ % degradation). Numerical values appear in cells, displaying exact degradation amounts for precise quantitative measurement. The visualization arranges algorithms from left to right in order of increasing robustness (PointNet, PointNet++, DGCNN, PointNeXt, Point Transformer), enabling horizontal comparison to identify best-performing methods. Vertical patterns in heatmaps reveal corruption-specific vulnerabilities, with most algorithms exhibiting steeper degradation gradients for spatial occlusion than for noise corruptions. Color bar legends positioned below each heatmap provide degradation scales with thresholds marked at 2%, 5%, 10%, and 15% to facilitate consistent interpretation across algorithms [33].

4.3.3. Algorithm reliability scoring framework

Reliability scoring aggregates accuracy, efficiency, and robustness via weighted combinations with normalized [0,1] scaling. High-throughput (weights 0.3, 0.5, 0.2) favors PointNet++ (0.82) and PointNeXt (0.79). High-precision (0.6, 0.2, 0.2) favors Point Transformer (0.91) and PointNeXt (0.86) [34].

Table 5 presents comprehensive reliability scores across all evaluated algorithms and deployment scenarios, enabling practitioners to identify optimal selections without manually analyzing multi-dimensional performance data.

Table 5. Algorithm Reliability Scores Across Deployment Scenarios.

Algorithm	High-Throughput	High-Precision	Embedded	Noisy Environment	General Purpose
PointNet	0.68	0.52	0.61	0.48	0.58
PointNet++	0.82	0.74	0.78	0.71	0.76
DGCNN	0.75	0.79	0.77	0.73	0.76
Point Transformer	0.61	0.91	0.72	0.88	0.78
PointNeXt	0.79	0.86	0.84	0.82	0.83

5. Conclusion

5.1. Key findings and recommendations

5.1.1. Summary of algorithm performance characteristics

Point Transformer achieves superior accuracy (93.7% on ModelNet40, 68.4% mIoU) and exceptional robustness via self-attention, though its computational cost limits real-time deployment.

PointNeXt emerges as the most balanced architecture, delivering competitive accuracy (93.2%), moderate efficiency (18.9ms latency), and strong robustness (3.4 percentage-point degradation at $\sigma = 0.02$ noise) through optimized training strategies applied to hierarchical feature learning. Graph-based DGCNN provides intermediate performance across all dimensions, offering advantages for applications requiring adaptive neighborhood relationships in feature space rather than fixed spatial groupings.

PointNet++ balances efficiency and accuracy for high-throughput scenarios, whereas PointNet is well-suited to extremely latency-sensitive applications (<10ms).

5.1.2. Guidelines for algorithm selection in industrial scenarios

Practitioners should prioritize PointNeXt for general-purpose deployment, supporting diverse inspection tasks while maintaining real-time processing.

High-precision applications (aerospace, medical devices) justify Point Transformer's costs when latency constraints exceed 100ms. High-volume inspection (<50ms latency) should use an optimized PointNet++ model with 1024 points. Embedded deployment favors lightweight architectures that fit within a few GB of on-device GPU memory.

5.2. Optimization strategies for practical deployment

5.2.1. Parameter tuning recommendations

Resolution optimization indicates that 2048-point inputs provide optimal trade-offs for most scenarios, balancing accuracy with latency and memory.

Simple geometries support 1024-point resolution (1.8× speedup), while complex assemblies justify 4096 points [28].

Neighbor search parameters significantly affect graph-based and hierarchical methods, with $k=32$ nearest neighbors providing optimal local context for DGCNN, whereas $k=16$ suffices for PointNet++ set abstraction layers. Spatial grouping radii should scale proportionally to object size; a radius of 0.2 (relative to the normalized bounding box) is effective for general-purpose inspection tasks. Feature dimensionality tuning reveals that 256-dimensional representations achieve optimal capacity-efficiency balance, with larger dimensions providing marginal accuracy gains at substantial memory cost.

5.2.2. Preprocessing and data augmentation suggestions

Preprocessing should prioritize statistical outlier removal and moving average filtering to suppress noise while preserving geometric features. Farthest point sampling preserves spatial coverage better than voxel downsampling. Augmentation through rotation and jittering improves deployment robustness.

Dropout augmentation ($p=0.1-0.2$) improves occlusion handling, while noise injection enhances corruption tolerance for limited training data.

5.3. Limitations and future research directions

5.3.1. Current benchmark limitations

The framework encompasses representative algorithms and scenarios, though ModelNet40's domain shift from industrial data may overestimate deployment performance. Industrial datasets (5,600 scans) remain limited compared to image benchmarks. Hardware evaluation excluded specialized accelerators.

5.3.2. Emerging research opportunities

Future work should include sparse convolutions, point-voxel hybrids, and multi-modal fusion combining point clouds with RGB/thermal data. Continual learning that addresses process drift and employs active learning, reducing annotation costs, warrants

investigation. Few-shot learning for rapid adaptation and uncertainty quantification for deployment safety represents promising research directions.

Data Availability: The source code and experimental scripts used in this study are publicly available at https://github.com/CatNinjaLuna/DeepL_PtCloud. The repository includes implementation details, benchmark configurations, and instructions for reproducing the experimental results.

References

1. H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun, "Point transformer," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2021, pp. 16259–16268.
2. S. A. Taghanaki, J. Luo, R. Zhang, Y. Wang, P. K. Jayaraman, and K. M. Jatavallabhula, "RobustPointSet: A dataset for benchmarking the robustness of point cloud classifiers," *arXiv preprint arXiv:2011.11572*, 2020.
3. C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 652–660.
4. D. O. Arroyo, "Advancements in point cloud-based 3D defect detection and classification for industrial systems: A comprehensive survey," *Inf. Fusion*, 2024.
5. C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep hierarchical feature learning on point sets in a metric space," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 30, 2017.
6. 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 Trans. Graph.*, vol. 38, no. 5, pp. 1–12, 2019.
7. R. S. Mei, C. H. Conway, M. V. Bimrose, W. P. King, and C. Shao, "Deep learning of 3D point clouds for detecting geometric defects in gears," *Manuf. Lett.*, vol. 41, pp. 1324–1333, 2024.
8. G. Qian *et al.*, "PointNeXt: Revisiting PointNet++ with improved training and scaling strategies," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 35, pp. 23192–23204, 2022.
9. A. N. Videsjorden, A. Goknil, S. Sen, E. J. Husom, and P. Nguyen, "3D-DaVa: Enhancing 3D point cloud data reliability for industrial applications," *ACM J. Data Inf. Qual.*, vol. 17, no. 1, pp. 1–28, 2025.
10. J. Sun, Q. Zhang, B. Kailkhura, Z. Yu, C. Xiao, and Z. M. Mao, "Benchmarking robustness of 3D point cloud recognition against common corruptions," *arXiv preprint arXiv:2201.12296*, 2022.
11. C. Zhuang, S. Li, and H. Ding, "Instance segmentation based 6D pose estimation of industrial objects using point clouds for robotic bin-picking," *Robot. Comput.-Integr. Manuf.*, vol. 82, p. 102541, 2023.
12. H. Duan *et al.*, "Robotics dexterous grasping: Methods based on point cloud and deep learning," *Front. Neurobot.*, vol. 15, p. 658280, 2021.
13. Q. Zhu, L. Fan, and N. Weng, "Advancements in point cloud data augmentation for deep learning: A survey," *Pattern Recognit.*, vol. 153, p. 110532, 2024.
14. Y. Lin, Z. Zhang, H. Tang, H. Wang, and S. Han, "PointAcc: Efficient point cloud accelerator," in *Proc. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*, 2021, pp. 449–461.
15. Y. Zhao and L. Fan, "Review on deep learning algorithms and benchmark datasets for pairwise global point cloud registration," *Remote Sens.*, vol. 15, no. 8, p. 2060, 2023.
16. D. Yuan and D. Zhang, "APAC-sensitive anomaly detection: Culturally-aware AI models for enhanced AML in US securities trading," in *Proc. Int. Conf. Comput., AI, Syst. Autom.*, May 2025, pp. 108–121.
17. A. Kang, S. Min, and D. Yuan, "Comparative analysis of foreign exchange market shock transmission and recovery resilience among major economies under geopolitical conflicts: Evidence from the Russia-Ukraine crisis," *J. Comput. Innov. Appl.*, vol. 2, no. 1, pp. 46–61, 2024.
18. Z. Wang, "DeepMotionNet: AI-driven predictive animation state transitions for reducing perceptual latency in competitive FPS games," in *Proc. IEEE Int. Conf. Comput. Eng. Appl. (ICCEA)*, Apr. 2025, pp. 1–8.
19. J. Zhang, "SecureCodeBERT: An AI-powered model for identifying and categorizing high-risk security vulnerabilities in PHP-based critical infrastructure applications," *J. Sustain. Policy Pract.*, vol. 1, no. 4, pp. 80–94, 2025.
20. Y. Lei and Z. Wu, "A real-time detection framework for high-risk content on short video platforms based on heterogeneous feature fusion," *Pinnacle Acad. Press Proc. Ser.*, vol. 3, pp. 93–106, 2025.
21. B. Dong, D. Zhang, and J. Xin, "Deep reinforcement learning for optimizing order book imbalance-based high-frequency trading strategies," *J. Comput. Innov. Appl.*, vol. 2, no. 2, pp. 33–43, 2024.
22. Z. Li and Z. Wang, "AI-driven procedural animation generation for personalized medical training via diffusion-based motion synthesis," *Artif. Intell. Mach. Learn. Rev.*, vol. 5, no. 3, pp. 111–123, 2024.
23. D. Zhang and E. Feng, "Quantitative assessment of regional carbon neutrality policy synergies based on deep learning," *J. Adv. Comput. Syst.*, vol. 4, no. 10, pp. 38–54, 2024.
24. R. Jia, J. Zhang, and J. Prescott, "An empirical study of large language models for threat intelligence analysis and incident response," *J. Comput. Innov. Appl.*, vol. 2, no. 1, pp. 99–110, 2024.

25. Z. Wang and A. Kang, "FTAFO: A federated transparent adaptive financial optimizer for reducing third-party dependencies in workflow management," *J. Sci. Innov. Soc. Impact*, vol. 1, no. 1, pp. 329–339, 2025.
26. Y. Lei, "Intelligent prediction and dynamic scheduling optimization strategy for cloud computing resources under burst load scenarios," in *Proc. Int. Symp. Mach. Learn. Soc. Comput.*, Oct. 2025, pp. 59–67.
27. D. Zhang and X. Ma, "Machine learning-based credit risk assessment for green bonds: Climate factor integration and default prediction analysis," *J. Sustain. Policy Pract.*, vol. 1, no. 2, pp. 121–135, 2025.
28. A. Kang and K. Yu, "The impact of financial data visualization techniques on enhancing budget transparency in local government decision-making," *Spectr. Res.*, vol. 5, no. 2, 2025.
29. Z. Wang, "Deep learning-based prediction technology for communication effects of animated character facial expressions," *J. Sustain. Policy Pract.*, vol. 1, no. 4, pp. 105–116, 2025.
30. J. Zhang, "Performance evaluation and comparison of machine learning algorithms for anomalous login behavior detection in enterprise networks," *Artif. Intell. Mach. Learn. Rev.*, vol. 5, no. 2, pp. 77–90, 2024.
31. Y. Lei, "Adaptive privacy-preserving techniques for multimedia content processing in cloud environments: A differential privacy approach," *J. Sci. Innov. Soc. Impact*, vol. 1, no. 1, pp. 278–293, 2025.
32. T. K. Trinh and D. Zhang, "Algorithmic fairness in financial decision-making: Detection and mitigation of bias in credit scoring applications," *J. Adv. Comput. Syst.*, vol. 4, no. 2, pp. 36–49, 2024.
33. A. Kang and X. Ma, "AI-based pattern recognition and characteristic analysis of cross-border money laundering behaviors in digital currency transactions," in *Proc. Int. Conf. Digit. Intell. Comput. Technol.*, Jul. 2025, pp. 1–19.
34. H. Weng and Y. Lei, "Cross-modal artifact mining for generalizable deepfake detection in the wild," *J. Comput. Innov. Appl.*, vol. 2, no. 2, pp. 78–87, 2024.

Disclaimer/Publisher's Note: The views, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim any responsibility for any injury to individuals or damage to property arising from the ideas, methods, instructions, or products mentioned in the content.