FastPoint: Accelerating 3D Point Cloud Model Inference via Sample Point Distance Prediction

Donghyun Lee eudh1206@snu.ac.kr
Dawoon Jeong daun20211@snu.ac.kr
Jae W. Lee jaewlee@snu.ac.kr
Hongil Yoon hongilyoon@google.com

Department of CSE
Seoul National University

IEEE/CVF International Conference on Computer Vision (ICCV 2025), Honolulu, Hawaii

Deep Learning on 3D Point Clouds

- Since the release of PointNet and PointNet++, various 3D point cloud models have been proposed, enhancing the model performance and efficiency.
- These models all share similar model architecture, utilizing Set Abstraction and Feature Propagation layers as core components.

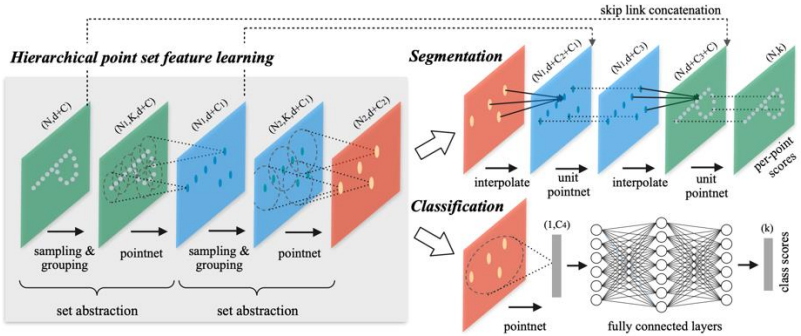

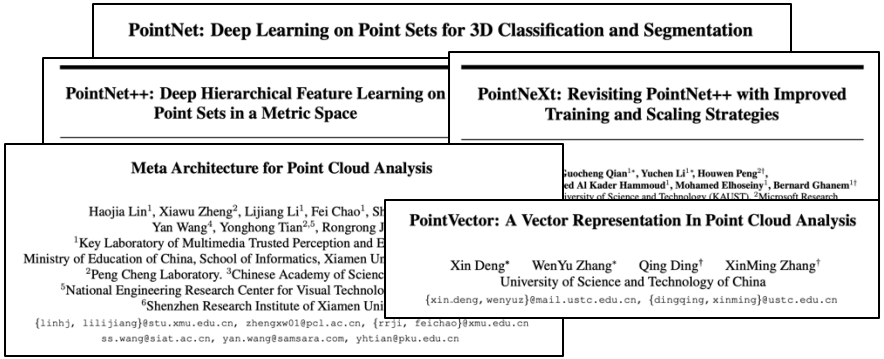

Figures from Qi, C.R. et al. "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space"

Challenges in 3D Point Cloud Model Inference

⚠ Challenges

Irregular nature of point cloud → ↑ computational challenges

Performance Bottleneck

Farthest Point Sampling (FPS): Takes up average **80.9%** of overall model execution time.

Neighbor Search: Takes up average **8.0%** of overall execution time.

Our Proposal

#1. Minimum Distance Prediction Sampling (MDPS): Accelerates FPS via **sample point distance prediction.**

#2. Redundancy Free Neighbor Search: Eliminates **redundant distance calculations.**

Farthest Point Sampling - Observations

Observation #1. As sampling proceeds, the distance between the sampled points decreases, and **the decreasing trend of minimum distance** displays a smooth curve.

Observation #2. We can accurately estimate this curve using **only a few initial FPS iterations.** This allows us to achieve **comparable sampling quality** with FPS with **low costs** by simply sampling points according to the estimated curve.

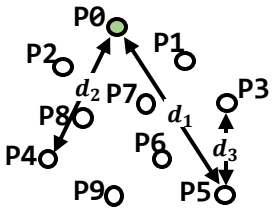

Farthest Point Sampling

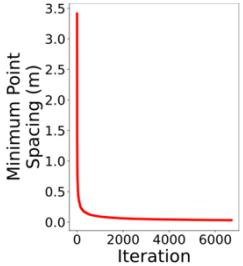

Minimum Distance Curve

Solution #1. Minimum Distance Prediction Sampling

- We propose ***Minimum Distance Prediction Sampling (MDPS).***

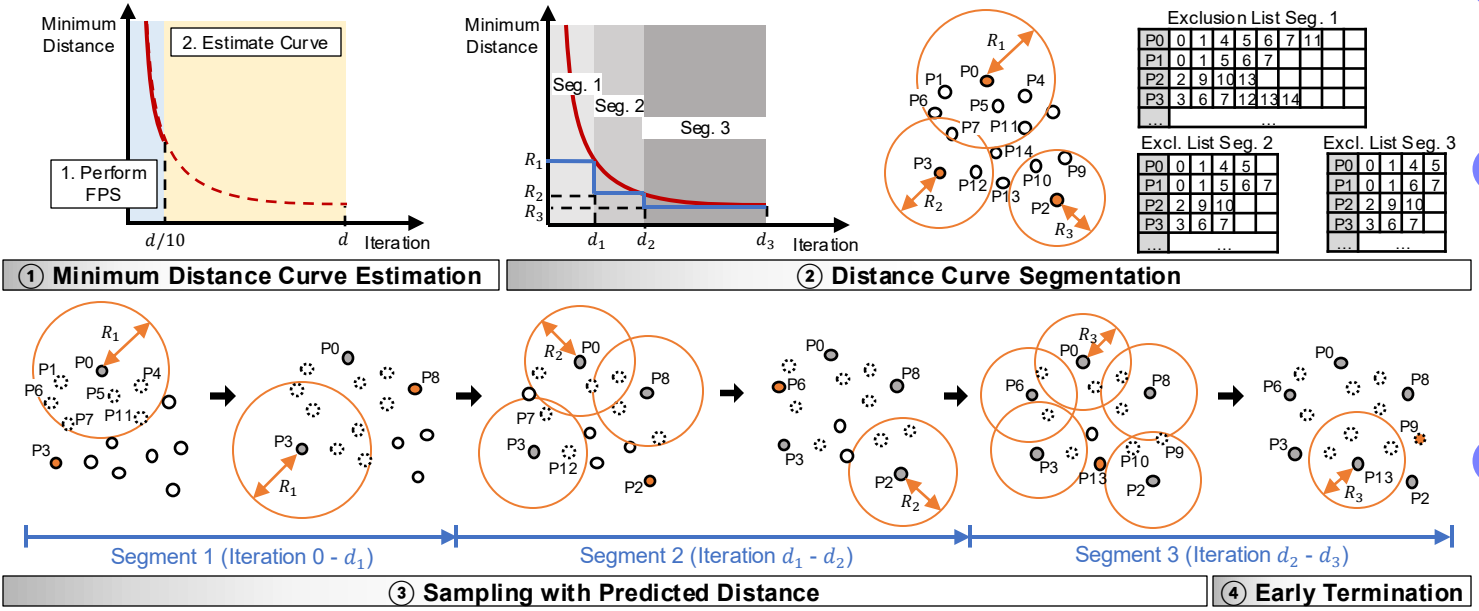

Needs only 1/10 FPS
Requires only 1/10 of the FPS iterations.

Fully Parallelizable
Distance calculations in exclusion list construction are fully parallelizable.

Fast Sampling Process
The use of exclusion list eliminates the need for distance calculations.

Neighbor Search - Observation

Observation #3. There are **redundant distance computations** across ball query, k-NN, and MDPS.

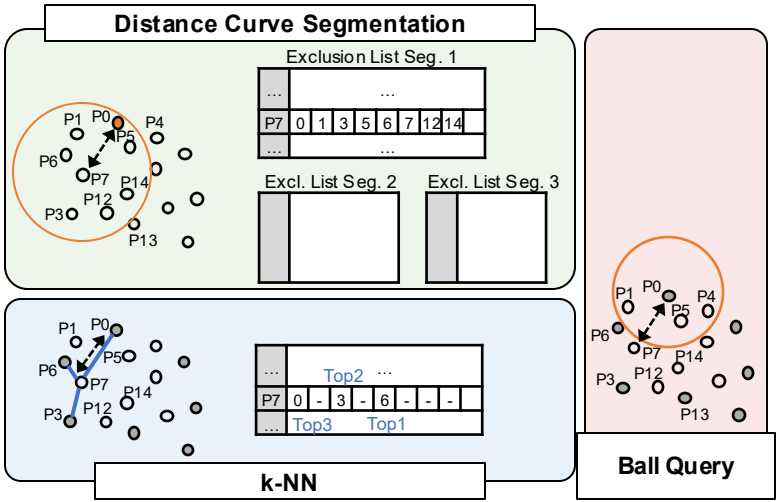

[MDPS & k-NN & Ball Query] Redundant distance calculation between P_7 and P_0.

Solution #2. Redundancy-Free Neighbor Search

- We propose ***Redundancy-Free Neighbor Search***, which eliminates redundant distance calculations in k-NN and Ball Query.

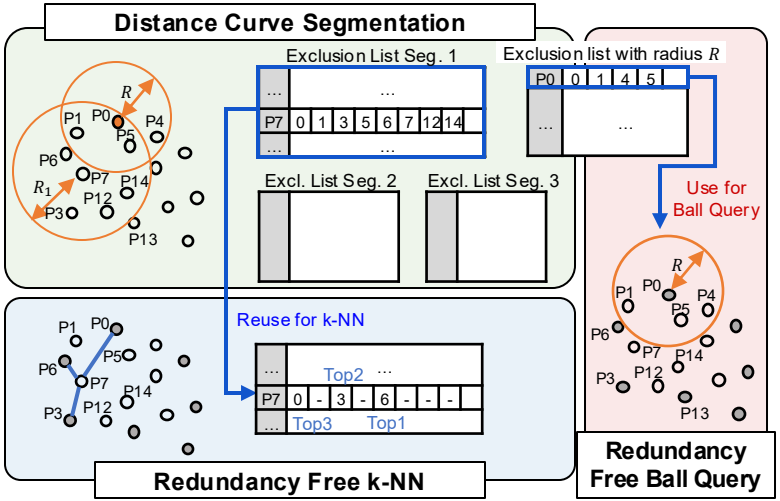

Redundancy-Free k-NN: Reuse exclusion list to reduce the search space while performing k-NN.

Redundancy-Free Ball Query: Add extra exclusion list while sampling and use this list for ball query. Adding extra list does not incur significant latency overhead.

Evaluation

Model	Dataset	mIoU (Diff. to Baseline FPS)			
		Baseline (FPS)	Random Sampling	Grid Sampling	MDPS
PV-L	S3DIS	71.33	68.21 (-3.12)	70.97 (-0.36)	71.37 (+0.04)
	ScanNet	70.70	67.66 (-3.04)	69.61 (-1.09)	70.63 (-0.07)
	Semantic KITTI	50.91	38.46 (-12.45)	50.62 (-0.29)	50.79 (-0.12)
PMB-L	S3DIS	69.72	64.83 (-4.89)	69.54 (-0.18)	69.74 (+0.02)
	ScanNet	70.86	59.94 (-10.92)	70.36 (-0.50)	70.89 (+0.03)
	Semantic KITTI	52.19	47.26 (-4.93)	52.27 (+0.08)	52.09 (-0.10)

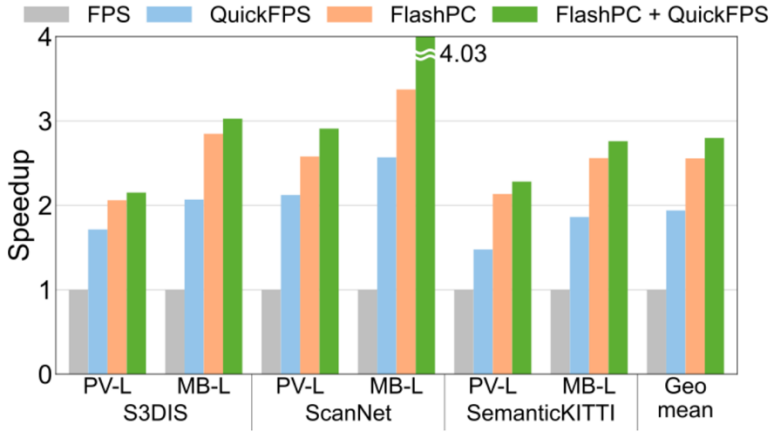

NVIDIA RTX 3090

Accuracy Max **0.12** mIoU loss, potential mIoU gain of **0.04**

Throughput **2.55x** end-to-end speedup

Please contact to the author or refer to the full paper for more details
http://arc.snu.ac.kr/pubs/iccv25_fastpoint.pdf

Accelerating 3D Point Cloud Model Inference via Sample Point Distance Prediction

Donghyun Lee eudh1206@snu.ac.kr

Dawoon Jeong daun20211@snu.ac.kr

Jae W. Lee jaewlee@snu.ac.kr

Hongil Yoon hongilyoon@google.com

Sourced code available at https://github.com/SNU-ARC/FastPoint

Department of CSE
Seoul National University

IEEE/CVF International Conference on Computer Vision (ICCV 2025), Honolulu, Hawaii