Supplementary Materials for CP^3: Channel Pruning Plug-in for Point Cloud Network

In the supplementary material, we provide more experimental comparisons on object detection in Sec. A and segmentation tasks in Sec. B , and we showcase the effectiveness of the proposed knowledge recycling module in Sec. C, we also analyzed the pruning rates of different layers in Sec. D.

A. More Experimental Results on 3D Object Detection

To further illustrate the effectiveness of our proposed CP^3, we incorporated CP^3 with pruning methods HRank [17] and CHIP [40] to prune VoteNet [31] on ScanNetV2 [4] and SUN RGB-D [39] for 3D object detection.

Table 8. Comparisons of object detection performance on the ScanNetV2 dataset. The baseline PNN model is VoteNet.

Method	mAP@0.25	mAP@0.50	Params. (K)	GFLOPs (↓%)
Baseline	62.34	40.82	641.92	5.78 (–)
HRank	61.04	37.99	249.82	2.46 (57.4)
HRank+CP^3	61.66	39.25	239.43	2.44 (57.8)
HRank	59.46	35.98	178.16	1.87 (67.7)
HRank+CP^3	60.51	39.15	169.87	1.80 (68.9)
CHIP	62.17	41.37	247.72	2.49 (56.9)
CHIP+CP^3	62.33	41.49	245.45	2.45 (57.6)
CHIP	60.86	39.94	176.88	1.89 (67.3)
CHIP+CP^3	61.55	40.43	172.78	1.87 (67.6)

ScanNetv2 Tab. 8 shows the comparison results of directly applying advanced pruning methods (HRank, CHIP) and implementation them with CP^3. Overall, CP^3 consistently improved the performance of existing advanced CNN pruning methods under different pruning rates. For instance, in the case of applying HRank with 67.7% FLOPs reduction, by incorporating CP^3, the mAP@0.50 increased 3.17% (35.95% vs. 39.15%) while achieving 1.2% more FLOPs reduction (67.7% vs. 68.9%).

SUN RGB-D We reported the comparison results on the SUN RGB-D dataset in Tab. 9. For both HRank and CHIP, the implementation with CP^3 achieved higher accuracy performance with higher FLOPs reduction, similar to our observations on other tasks and datasets.

B. More Experimental Results on Semantic Segmentation

To investigate the generality of our work, we extended the comparisons on semantic segmentation. We conducted the experiment on the S3DIS [1] dataset of PointNeXt-S and PointNeXt-XL, and two advanced pruning methods are evaluated.

Table 9. Comparisons of object detection performance on the SUN RGB-D dataset. The baseline PNN model is VoteNet.

Method	mAP@0.25	mAP@0.50	Params. (K)	GFLOPs (↓%)
Baseline	59.78	35.77	641.92	5.78 (–)
HRank	59.22	34.26	249.82	2.46 (57.4)
HRank+CP^3	60.21	34.96	245.32	2.44 (57.8)
HRank	57.68	31.30	178.88	1.87 (67.7)
HRank+CP^3	59.22	33.18	176.03	1.85 (68.0)
CHIP	59.54	35.74	248.31	2.49 (56.9)
CHIP+CP^3	59.88	35.84	242.12	2.43 (58.0)
CHIP	58.63	35.07	176.23	1.89 (67.3)
CHIP+CP^3	59.13	35.32	172.02	1.87 (67.6)

Table 10. Comparisons of semantic segmentation performance on the S3DIS dataset (evaluated in Area-5) with PointNeXt-S [35].

Method	PointNeXt-S				
	OA	mAcc	mIoU	Params. (M)	GFLOPs (↓ %)
Baseline	88.20	70.70	64.20	0.80	3.60 (–)
HRank	85.89	67.27	60.49	0.33	1.53 (57.5)
HRank+CP^3	86.18	67.65	61.04	0.31	1.52 (57.8)
HRank	84.92	65.37	58.73	0.17	0.77 (78.6)
HRank+CP^3	85.12	67.48	60.12	0.16	0.74 (79.4)
CHIP	84.45	67.72	61.16	0.32	1.53 (57.5)
CHIP+CP^3	84.53	70.52	63.62	0.33	1.48 (58.9)
CHIP	84.39	67.29	60.63	0.15	0.74 (79.4)
CHIP+CP^3	85.04	69.02	61.45	0.16	0.73 (79.7)

PointNext-S Compared to other PointNeXt zoos, besides the fewer parameters, PointNeXt-S is designed with no InvResMLP blocks and is a simpler network architecture. The comparison result in Tab. 10 showed that the consistent outperformance of CP^3 compared to directly using HRank and CHIP **without** CP^3. In the case of applying CHIP with 57.5% FLOPs reduction, by incorporating CP^3, the mAcc increases 2.8 % while achieving 1.4 % more FLOPs reduction. The result indicated that even for a simpler network architecture, the accuracy performance degradation still occurred when directly implementation of CNN pruning methods, verifying the necessity of the implementation with CP^3.

PointNext-XL We further investigated on the more complex and larger network PointNext-XL. Similar observations on the improvement by CP^3 can be found in Tab. 11. For instance, our approach can achieve 69.8 % and 69.9 %

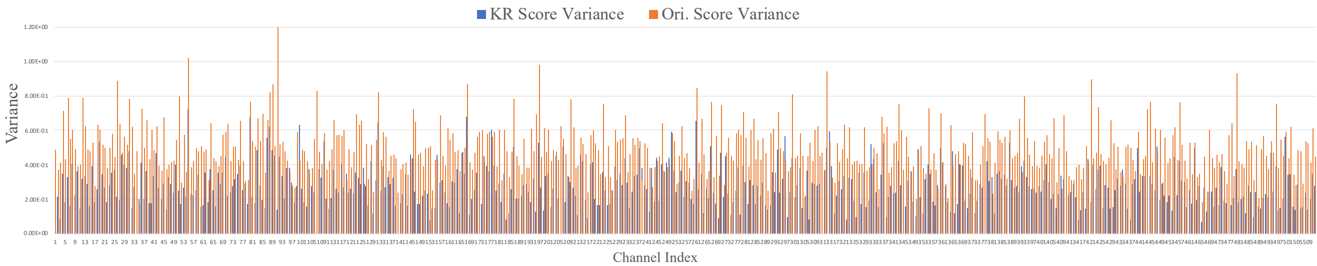

Figure 3. The variance of the importance score of each channel in the 10-th layer of PointNeXt-S. The x-axis represents channel indices, and the y-axis represents the variances of each channel importance scores, which are calculated by 2,468 input samples.

Table 11. Comparisons of semantic segmentation performance on the S3DIS dataset (evaluated in Area-5) with PointNeXt-XL [35].

Method	PointNeXt-XL				
	OA	**mAcc**	**mIoU**	**Params. (M)**	**GFLOPs (↓ %)**
Baseline	91.00	77.20	71.10	41.60	84.80 (–)
HRank	90.02	74.50	68.22	13.03	26.71 (68.5)
HRank+CP^3	90.80	75.85	69.98	12.57	25.78 (69.6)
HRank	89.79	74.40	68.09	8.80	18.05 (78.7)
HRank+CP^3	90.48	74.45	68.41	8.42	17.37 (79.5)
CHIP	89.98	75.43	69.21	17.54	35.97 (57.6)
CHIP+CP^3	90.57	75.90	69.40	17.00	34.68 (59.1)
CHIP	89.15	74.54	68.07	8.42	17.37 (79.5)
CHIP+CP^3	90.03	74.61	68.26	8.05	16.62 (80.4)

storage and computation reductions, respectively, with a 1.3 % and 1.7% accuracy increase for mAcc and mIoU over the baseline model.

Table 12. Comparisons on the SemanticKITTI with RandLA-Net.

Method	mIoU	Params. (M)	FLOPs (%)
Baseline (RandLA-Net)	50.30	0.95	100
CHIP	49.12	0.20	21.7
CHIP+CP^3	50.21	0.18	19.8

Outdoor Experiments CP^3 focuses on point-based networks (PNNs), and by following prevailing PNN works, we have experimented on the popular large-scale datasets such as ScanObjectNN and S3DIS and achieved promising results in the paper. To further illustrate the validity of our method, we experimented on a outdoor dataset (SemanticKITTI) with RandLA-Net and CHIP. The results are shown in Tab. 12.

C. Exploration on Knowledge Recycling

In this section, we took a deeper look into the Knowledge Recycling (KR) module. To verify the effectiveness of KR, we performed pruning methods and statistically analyzed the positive effect of KR. We took the PointNeXt-S with dataset ModelNet40 as an example. We performed the comparison between directly implementing HRank and HRank with CP^3. And we focus on the KR scores generated by the KR module (with discarded points) and HRank scores **without** KR module. We calculated the variances of scores to justify the robustness of CP^3. Fig. 3 shows the statistics comparison results on the 10-th layer with 512 channels on 2468 test meshed CAD models. Among 512 channels in the 10-th layer, the percentage is 93.2% in the case of the KR score variances lower than the original score variances. For instance, in the case of the 25-th channel, the KR score variance is much lower than the original score variance (0.21 vs. 0.89). Similar results can be found on other layers. These results verify that the KR module enabled the channel importance calculation to be more stable and robust.

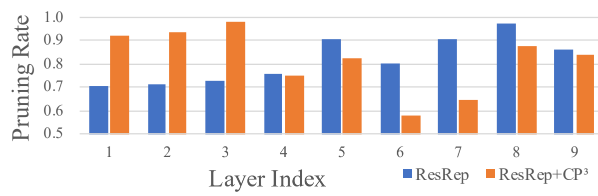

Figure 4. Different layer pruning rate.

D. Analysis of the layer-wise redundancy

We take the pruned PointNet++ on ScanObjectNN as an example and show the pruning rates for each layer in Fig. 4. Pruning with CP^3 eliminates more redundancy on shallow layers and less on 6th and 7th layers. CP^3 on ResRep achieved higher accuracy (84.80% vs 83.79%) with higher FLOPs reduction (84.8% vs 83.0%), indicating pruning with CP^3 effectively identifies the redundancy in point-based networks.