Directional Graph Networks with Hard Weight Assignments

Miguel Dominguez and Raymond Ptucha (from Rochester Institute of Technology)

Introduction

- Point cloud networks emulate traditional CNNs
- Map learnable weight matrices to points in a neighborhood
- This mapping may not be 1:1
- Other works create soft mapping between every weight matrix and every neighbor: $O(K^2)$
- A hard mapping $O(K)$ performs better

Method

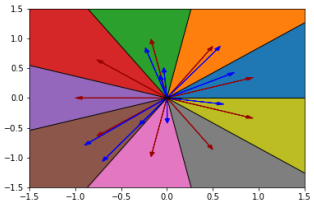

- Map each weight matrix to angular (directional) regions on a sphere
- Generate KNN graph
- Select one edge per weight matrix in each neighborhood which is either in angular region or as close as possible
- Perform weighted sum on the chosen edges

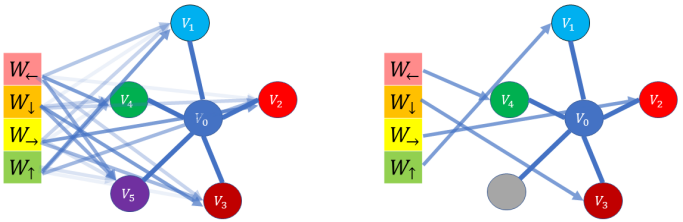

a) **b)**

- Left: $O(K^2)$ Soft assignment filters
- Right: $O(K)$ Hard assignment filter

Results

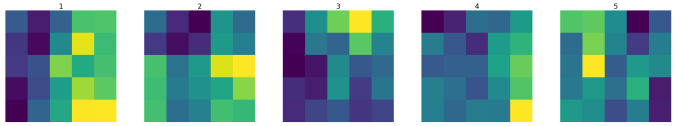

- Above: Visualization of 1st layer filters
- Below: Overall results on Benchmarks

Network	MN40 (1k)	ShapeNet	SHREC'15
PointNet++	91.9%	85.1	94.1%
DGCNN	92.9%	85.2	-
SpiderCNN	92.4%	85.3	95.8%
PointCNN	92.2%	**86.1**	-
HDGN (Ours)	**93.9%**	85.4	**100%**

Ablations

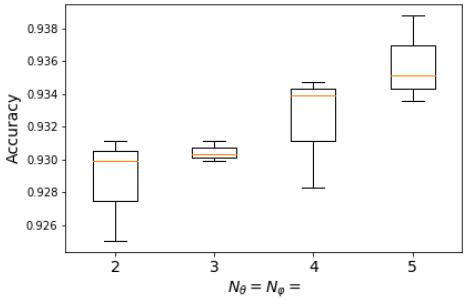

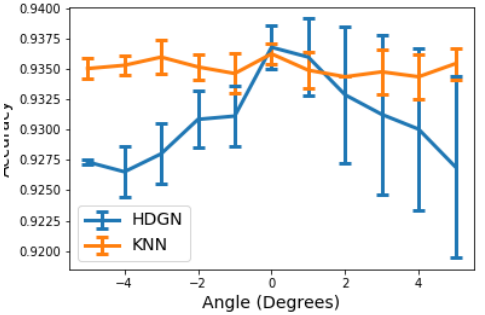

- Top: More directional weight matrices improves accuracy
- Bottom: Directional weight matrices are sensitive to angle if not trained on rotational variation

Conclusion

- First step in reducing complexity of point cloud networks that use graphs
- Future work will address the large overhead to store graph information