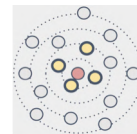


# DeepGCNs for Representation Learning on Graphs

Guohao Li





- ① **DeepGCNs: Can GCNs Go as Deep as CNNs? (ICCV'2019 Oral)**  
Guohao Li\*, Matthias Müller\*, Ali Thabet, Bernard Ghanem
- ② **DeepGCNs: Making GCNs Go as Deep as CNNs (arXiv'2019)**  
Guohao Li\*, Matthias Müller\*, Guocheng Qian, Itzel C. Delgadillo, Abdullellah Abualshour, Ali Thabet, Bernard Ghanem
- ③ **SGAS: Sequential Greedy Architecture Search (arXiv'2019)**  
Guohao Li\*, Guocheng Qian\*, Itzel C. Delgadillo\*, Matthias Müller, Ali Thabet, Bernard Ghanem



Guohao Li



Matthias Müller



Guocheng Qian



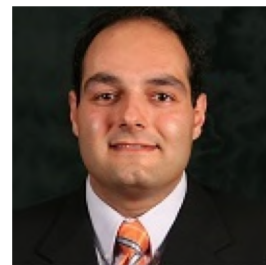
Itzel C. Delgadillo



Abdullellah  
Abualshour

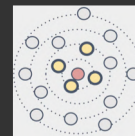


Ali Thabet

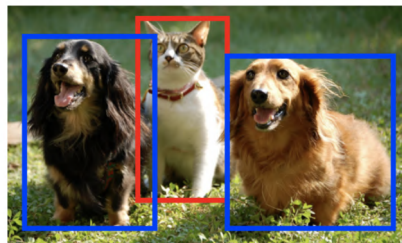


Bernard Ghanem

# Grid data vs. General graphs



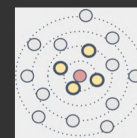
DeepGCNs.org



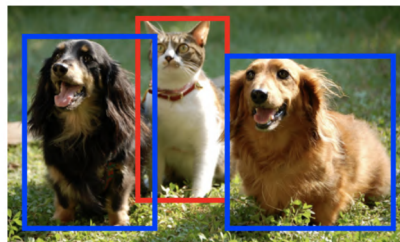
CAT, DOG

Grid Data :  
• Image

# Grid data vs. General graphs



DeepGCNs.org



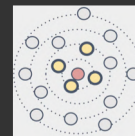
CAT, DOG



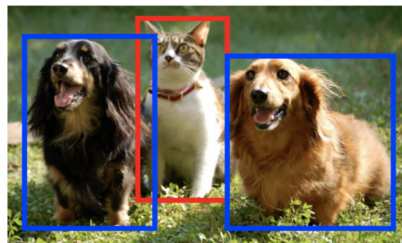
- Grid Data :
- Image
  - Video



# Grid data vs. General graphs



DeepGCNs.org



CAT, DOG

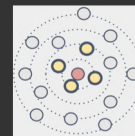


Grid Data :

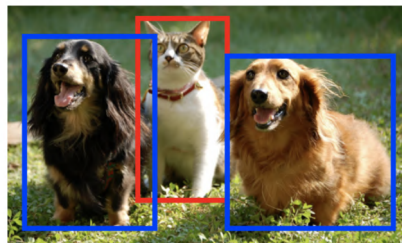
- Image
- Video
- Audio
- Text



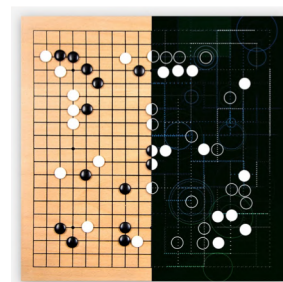
# Grid data vs. General graphs



DeepGCNs.org



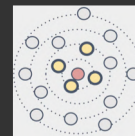
CAT, DOG



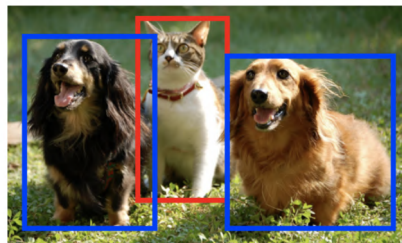
Grid Data :

- Image
- Video
- Audio
- Text
- Grid game (Go)
- ...

# Grid data vs. General graphs



DeepGCNs.org

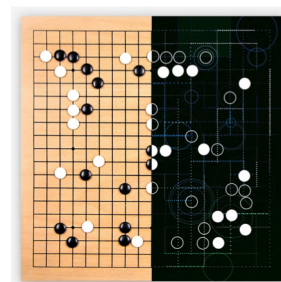


CAT, DOG

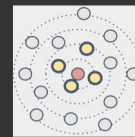


Grid Data :

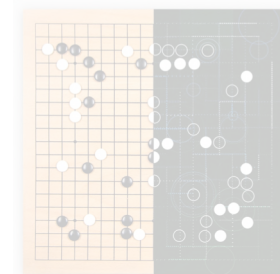
- Image
- Video
- Audio
- Text
- Grid game (Go)
- ...

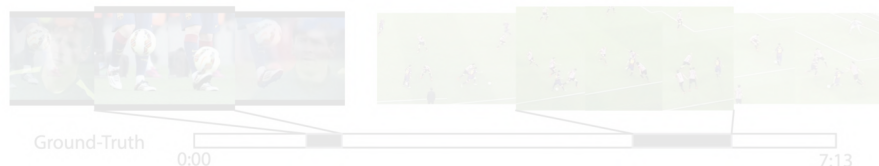
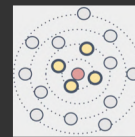


CNN works well



## Why do we need graph convolutional networks?





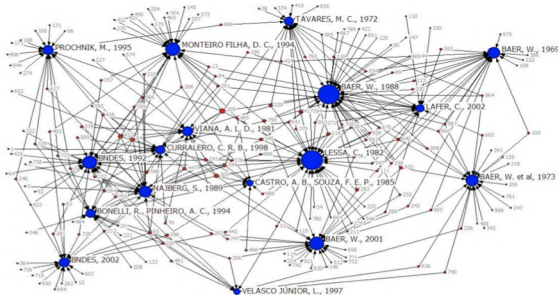
## Why we need graph convolutional networks?

## Tremendous non-grid graph structured data

Lots of real-world applications need to deal with **Non-Grid** data

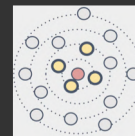


- Social Networks
- Citation Networks



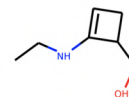
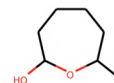
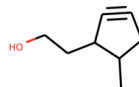
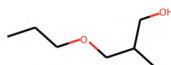


# Grid data vs. General graphs



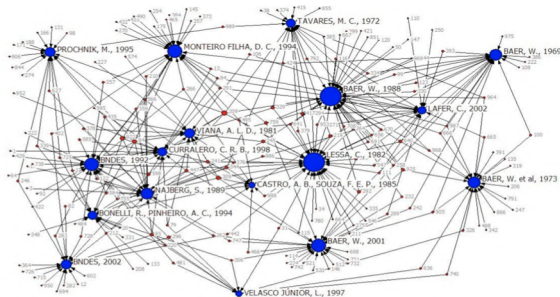
DeepGCNs.org

Lots of real-world applications need to deal with **Non-Grid** data

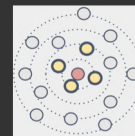


General Graphs :

- Social Networks
- Citation Networks
- Molecules

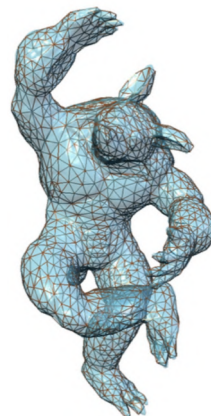
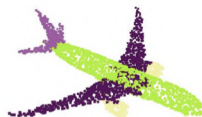
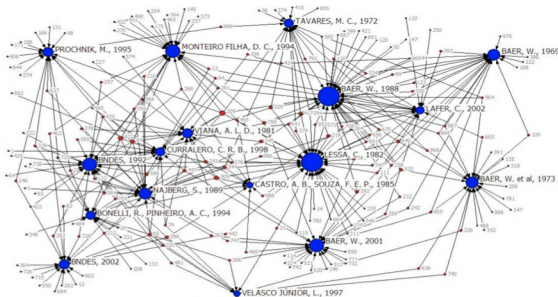
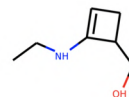
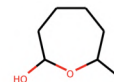
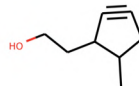
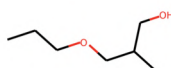


# Grid data vs. General graphs



DeepGCNs.org

Lots of real-world applications need to deal with **Non-Grid** data

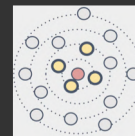


General Graphs :

- Social Networks
- Citation Networks
- Molecules
- Point Clouds
- 3D Meshes
- ...

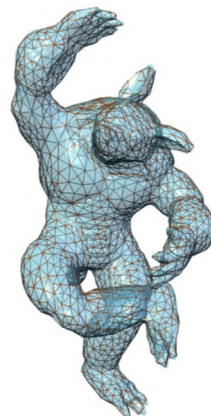
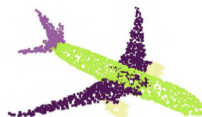
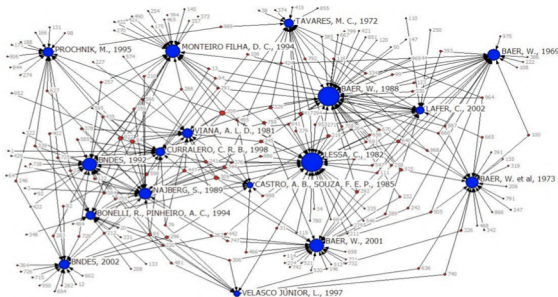
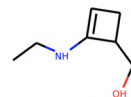
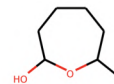
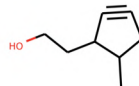
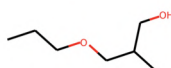


# Grid data vs. General graphs



DeepGCNs.org

Lots of real-world applications need to deal with **Non-Grid** data



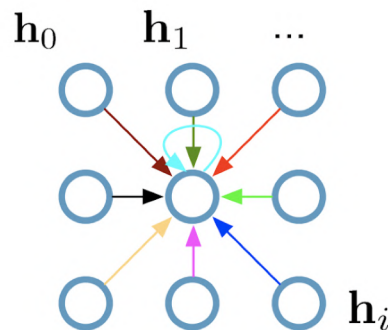
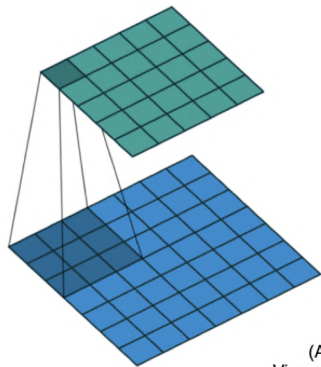
General Graphs :

- Social Networks
- Citation Networks
- Molecules
- Point Clouds
- 3D Meshes
- ...

CNN doesn't work  
**GCN** to rescue

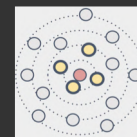


## Single CNN layer with 3x3 filter:

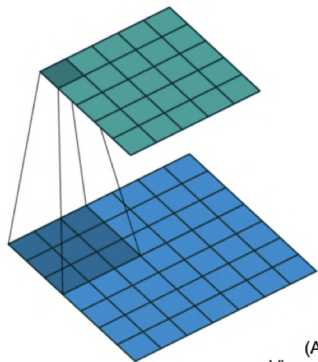


$\mathbf{h}_i \in \mathbb{R}^F$  are (hidden layer) activations of a pixel/node

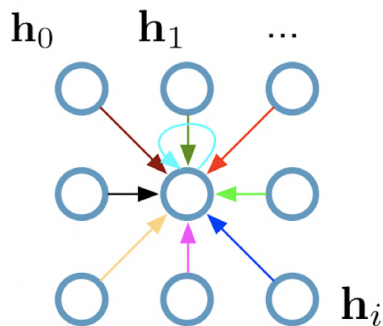
Slides by Thomas Kipf



## Single CNN layer with 3x3 filter:



(Animation by  
Vincent Dumoulin)

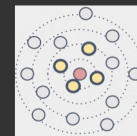


### Update for a single pixel:

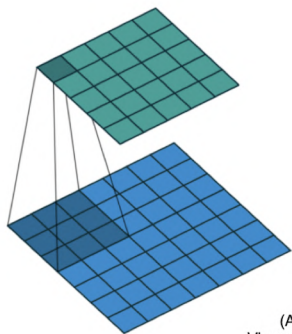
- Transform messages individually  $\mathbf{W}_i \mathbf{h}_i$
- Add everything up  $\sum_i \mathbf{W}_i \mathbf{h}_i$

$\mathbf{h}_i \in \mathbb{R}^F$  are (hidden layer) activations of a pixel/node

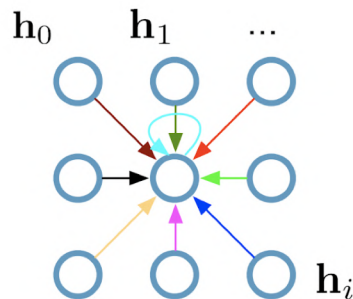
Slides by Thomas Kipf



## Single CNN layer with 3x3 filter:



(Animation by  
Vincent Dumoulin)



### Update for a single pixel:

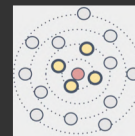
- Transform messages individually  $\mathbf{W}_i \mathbf{h}_i$
- Add everything up  $\sum_i \mathbf{W}_i \mathbf{h}_i$

$\mathbf{h}_i \in \mathbb{R}^F$  are (hidden layer) activations of a pixel/node

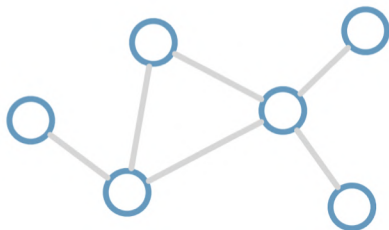
### Full update:

$$\mathbf{h}_4^{(l+1)} = \sigma \left( \mathbf{W}_0^{(l)} \mathbf{h}_0^{(l)} + \mathbf{W}_1^{(l)} \mathbf{h}_1^{(l)} + \dots + \mathbf{W}_8^{(l)} \mathbf{h}_8^{(l)} \right)$$

Slides by Thomas Kipf

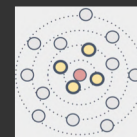


Consider this  
undirected graph:



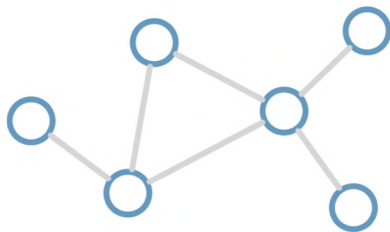
Slides by Thomas Kipf

# CNN vs. GCN - Introduction: GCN

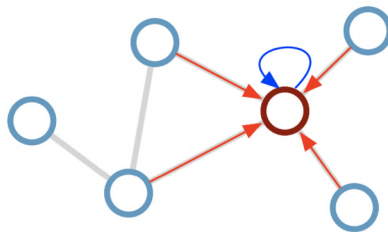


DeepGCNs.org

Consider this  
undirected graph:

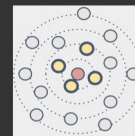


Calculate update  
for node in red:



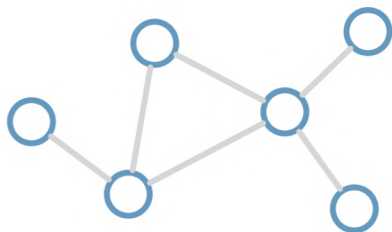
Slides by Thomas Kipf

# CNN vs. GCN - Introduction: GCN

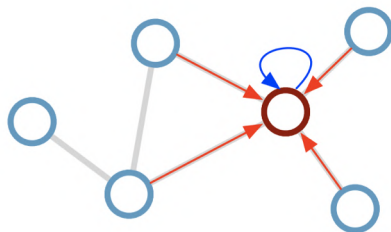


DeepGCNs.org

Consider this  
undirected graph:



Calculate update  
for node in red:



$\mathcal{N}_i$  : neighbor indices

$c_{ij}$  : norm. constant  
(fixed/trainable)

**Update rule:**

$$\mathbf{h}_i^{(l+1)} = \sigma \left( \mathbf{h}_i^{(l)} \mathbf{W}_0^{(l)} + \sum_{j \in \mathcal{N}_i} \frac{1}{c_{ij}} \mathbf{h}_j^{(l)} \mathbf{W}_1^{(l)} \right)$$

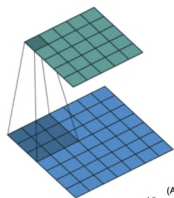
Slides by Thomas Kipf

# CNN vs. GCN - Comparison

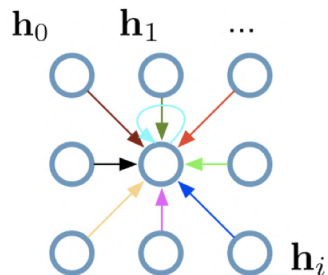


DeepGCNs.org

**Single CNN layer  
with 3x3 filter:**



(Animation by  
Vincent Dumoulin)



Convolutional Neural Network (CNN)

Slides by Thomas Kipf

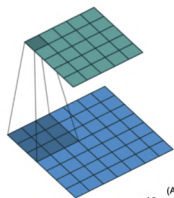


# CNN vs. GCN - Comparison

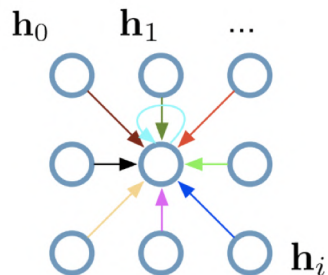


DeepGCNs.org

**Single CNN layer  
with 3x3 filter:**



(Animation by  
Vincent Dumoulin)



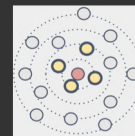
**Full update:**

$$\mathbf{h}_4^{(l+1)} = \sigma \left( \mathbf{w}_0^{(l)} \mathbf{h}_0^{(l)} + \mathbf{w}_1^{(l)} \mathbf{h}_1^{(l)} + \dots + \mathbf{w}_8^{(l)} \mathbf{h}_8^{(l)} \right)$$

Convolutional Neural Network (CNN)

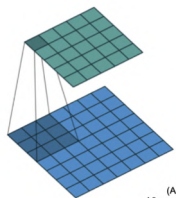
Slides by Thomas Kipf

# CNN vs. GCN - Comparison

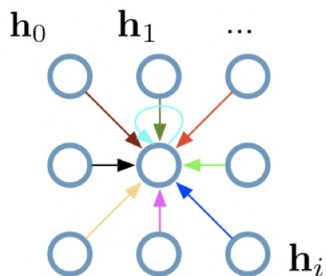


DeepGCNs.org

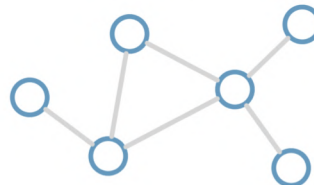
**Single CNN layer with 3x3 filter:**



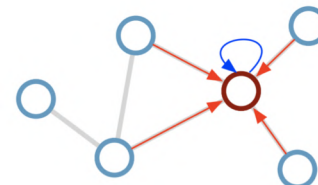
(Animation by Vincent Dumoulin)



Consider this undirected graph:



Calculate update for node in red:



**Full update:**

$$\mathbf{h}_4^{(l+1)} = \sigma \left( \mathbf{w}_0^{(l)} \mathbf{h}_0^{(l)} + \mathbf{w}_1^{(l)} \mathbf{h}_1^{(l)} + \dots + \mathbf{w}_8^{(l)} \mathbf{h}_8^{(l)} \right)$$

Convolutional Neural Network (CNN)

Graph Convolutional Network (GCN)

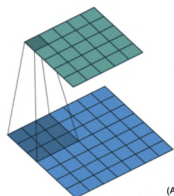
Slides by Thomas Kipf

# CNN vs. GCN - Comparison

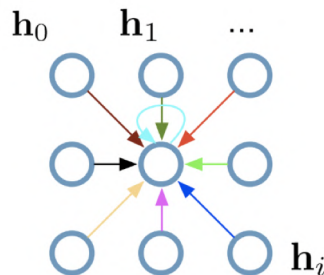


DeepGCNs.org

**Single CNN layer with 3x3 filter:**



(Animation by Vincent Dumoulin)

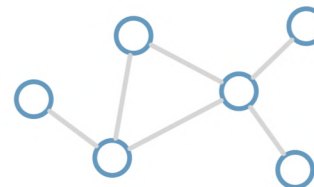


**Full update:**

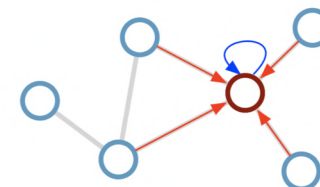
$$\mathbf{h}_i^{(l+1)} = \sigma \left( \mathbf{W}_0^{(l)} \mathbf{h}_0^{(l)} + \mathbf{W}_1^{(l)} \mathbf{h}_1^{(l)} + \dots + \mathbf{W}_8^{(l)} \mathbf{h}_8^{(l)} \right)$$

Convolutional Neural Network (CNN)

Consider this undirected graph:



Calculate update for node in red:

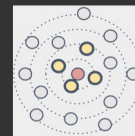


**Update rule:**

$$\mathbf{h}_i^{(l+1)} = \sigma \left( \mathbf{h}_i^{(l)} \mathbf{W}_0^{(l)} + \sum_{j \in \mathcal{N}_i} \frac{1}{c_{ij}} \mathbf{h}_j^{(l)} \mathbf{W}_1^{(l)} \right)$$

Graph Convolutional Network (GCN)

Slides by Thomas Kipf



Node  
Features

Neighbor's  
Features

$$\mathbf{x}_i^{(k)} = \gamma^{(k)} \left( \mathbf{x}_i^{(k-1)}, \bigoplus_{j \in \mathcal{N}(i)} \phi^{(k)} \left( \mathbf{x}_i^{(k-1)}, \mathbf{x}_j^{(k-1)}, \mathbf{e}_{i,j} \right) \right),$$

By <https://pytorch-geometric.readthedocs.io>

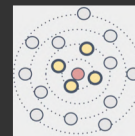
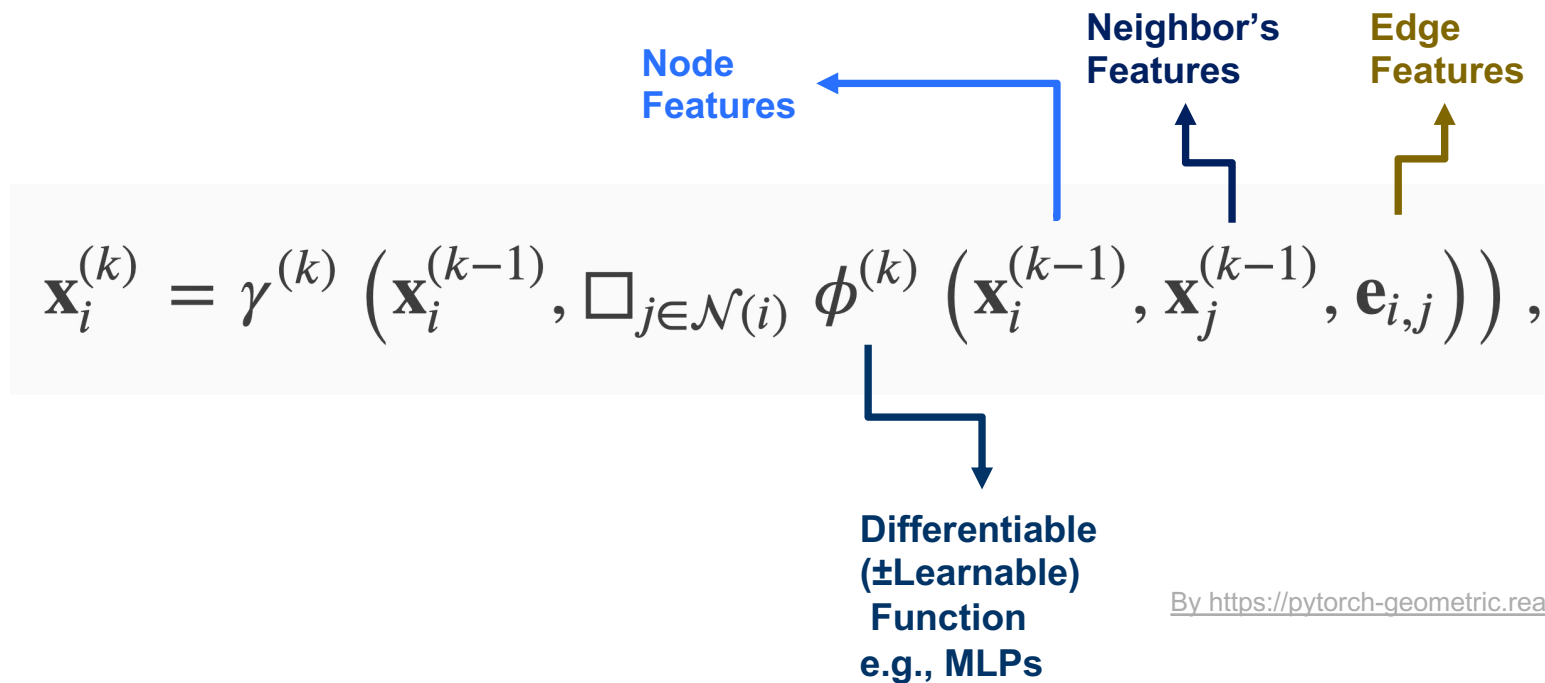
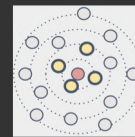


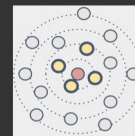
Diagram illustrating the message passing process in a Graph Convolutional Network (GCN). The diagram shows three inputs to the function  $\phi^{(k)}$ : Node Features (blue arrow), Neighbor's Features (blue arrow), and Edge Features (yellow arrow). The function  $\phi^{(k)}$  takes these inputs and produces the output  $\mathbf{x}_i^{(k)}$ .

$$\mathbf{x}_i^{(k)} = \gamma^{(k)} \left( \mathbf{x}_i^{(k-1)}, \bigoplus_{j \in \mathcal{N}(i)} \phi^{(k)} \left( \mathbf{x}_i^{(k-1)}, \mathbf{x}_j^{(k-1)}, \mathbf{e}_{i,j} \right) \right),$$

By <https://pytorch-geometric.readthedocs.io>



By <https://pytorch-geometric.readthedocs.io>



$$\mathbf{x}_i^{(k)} = \gamma^{(k)} \left( \mathbf{x}_i^{(k-1)}, \bigoplus_{j \in \mathcal{N}(i)} \phi^{(k)} \left( \mathbf{x}_i^{(k-1)}, \mathbf{x}_j^{(k-1)}, \mathbf{e}_{i,j} \right) \right),$$

**Node Features** (blue arrow pointing to  $\mathbf{x}_i^{(k)}$ )

**Neighbor's Features** (blue arrow pointing to  $\mathbf{x}_j^{(k-1)}$ )

**Edge Features** (yellow arrow pointing to  $\mathbf{e}_{i,j}$ )

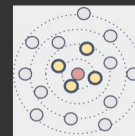
**Permutation Invariant Function**  
e.g., sum, mean or max (red arrow pointing to  $\bigoplus$ )

**Differentiable ( $\pm$ Learnable) Function**  
e.g., MLPs (blue arrow pointing to  $\phi^{(k)}$ )

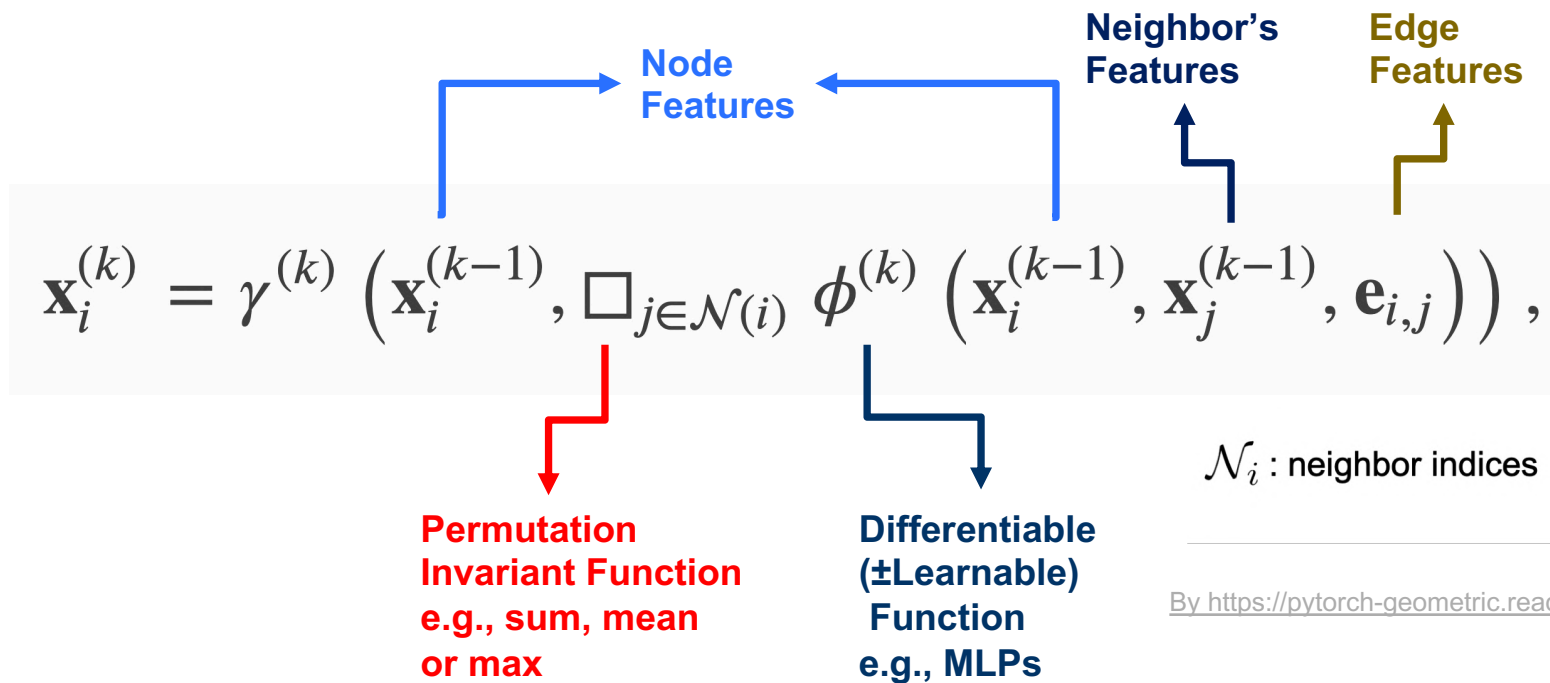
$\mathcal{N}_i$  : neighbor indices

By <https://pytorch-geometric.readthedocs.io>

# CNN vs. GCN - Message Passing



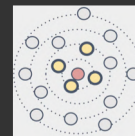
DeepGCNs.org



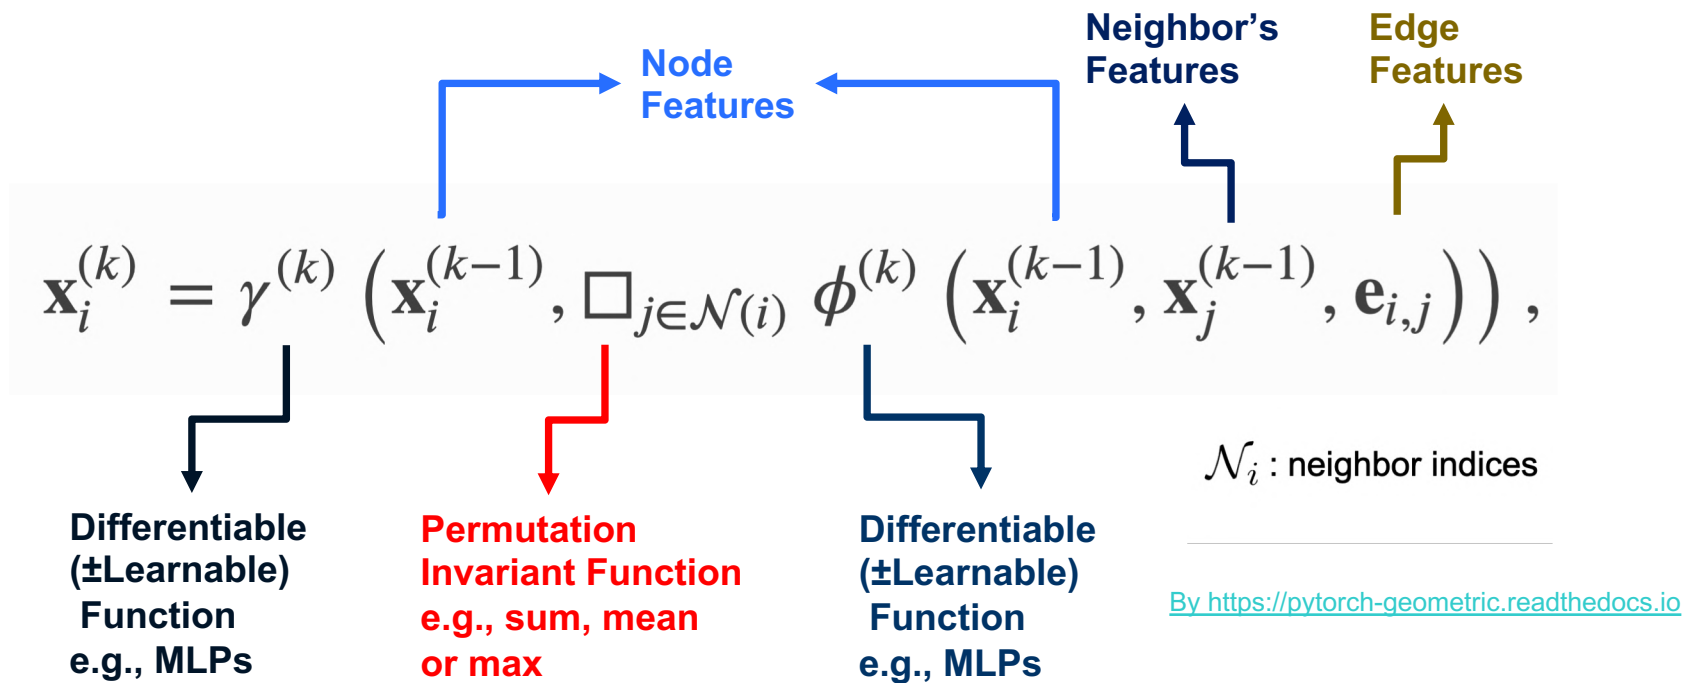
By <https://pytorch-geometric.readthedocs.io>

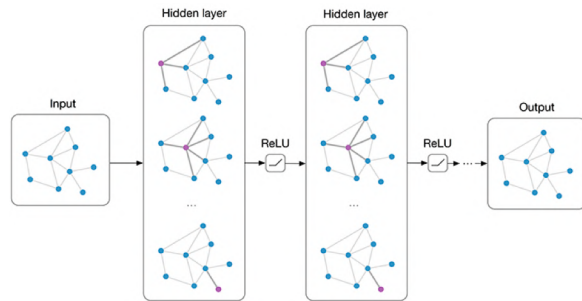


# CNN vs. GCN - Message Passing

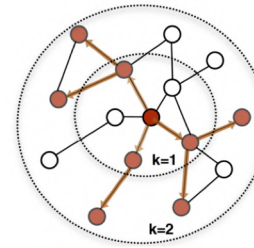


DeepGCNs.org

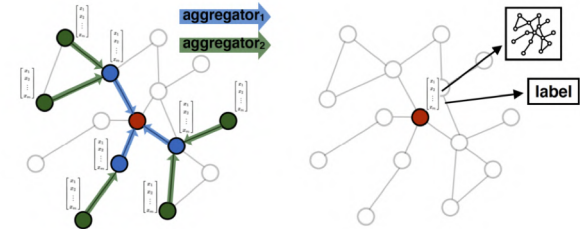




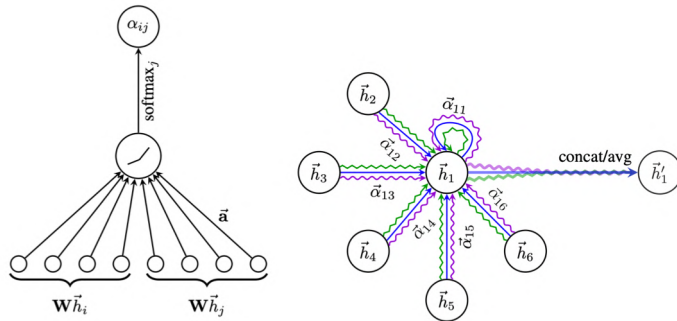
Kipf, T.N. and Welling, M., 2016. Semi-Supervised Classification with Graph Convolutional Networks.



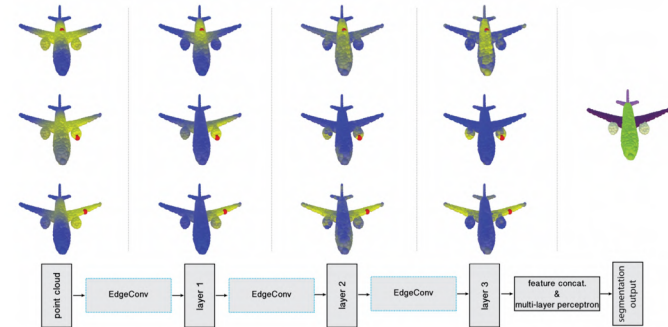
Hamilton, W.L., Ying, R. and Leskovec, J., 2017. Inductive Representation Learning on Large Graphs.



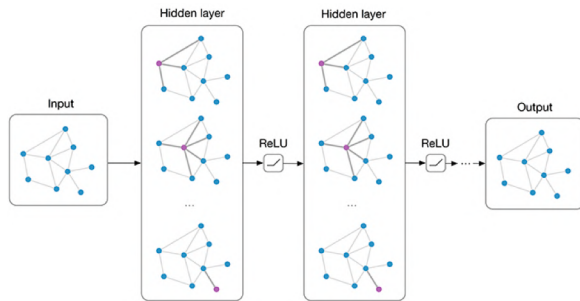
Most SOTA GCN models are no deeper than 3 or 4 layers.



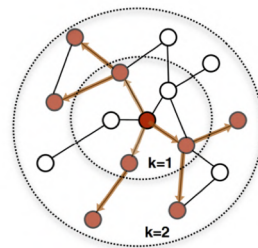
Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P. and Bengio, Y., 2018. Graph Attention Networks.



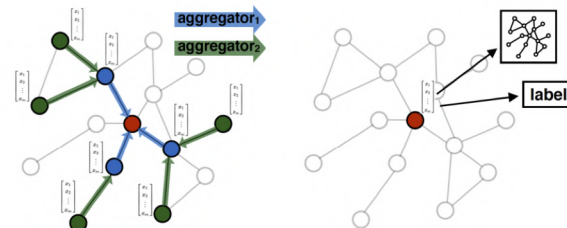
Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M. and Solomon, J.M., 2018. Dynamic Graph CNN for Learning on Point Clouds.



Kipf, T.N. and Welling, M., 2016. Semi-Supervised Classification with Graph Convolutional Networks.

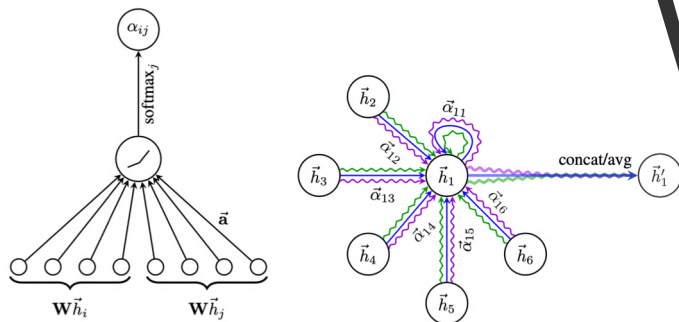


Hamilton, W.L., Ying, R. and Leskovec, J., 2017. Inductive Representation Learning on Large Graphs.

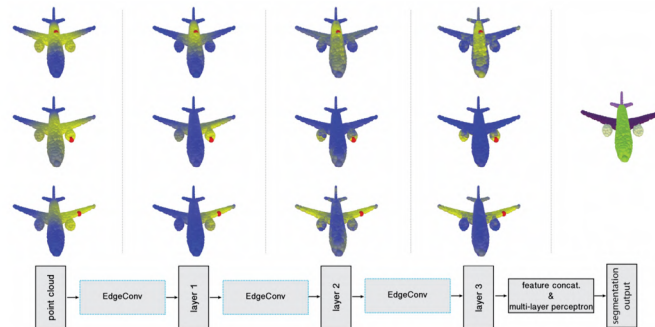


# Most SOTA GCN models are no deeper than 3 or 4 layers.

## Why?

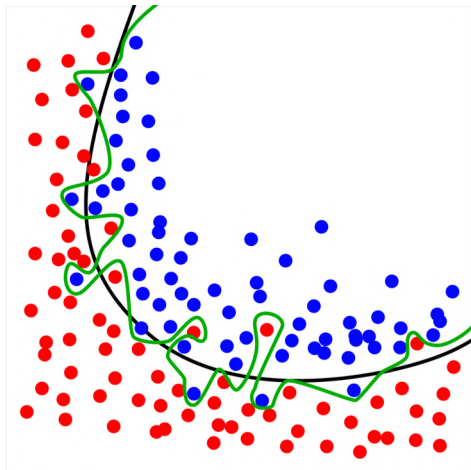


Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P. and Bengio, Y., 2018. Graph Attention Networks.

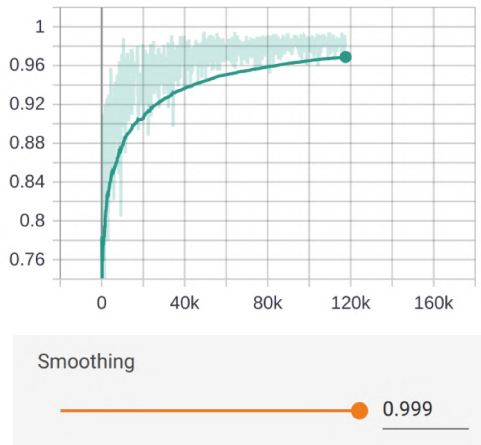


Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M. and Solomon, J.M., 2018. Dynamic Graph CNN for Learning on Point Clouds.

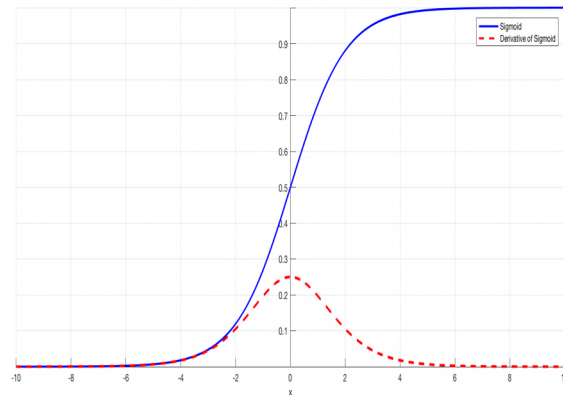
# Why GCNs are limited to shallow structures?



Over-fitting



Over-smoothing

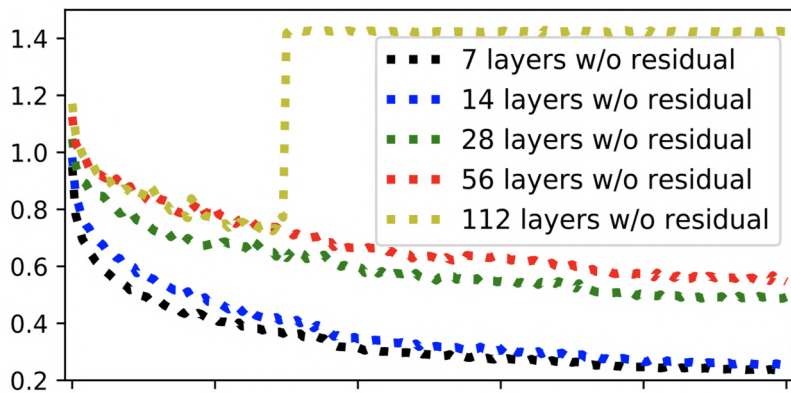


Vanishing Gradient

Figures from <https://towardsdatascience.com/the-vanishing-gradient-problem-69bf08b15484>

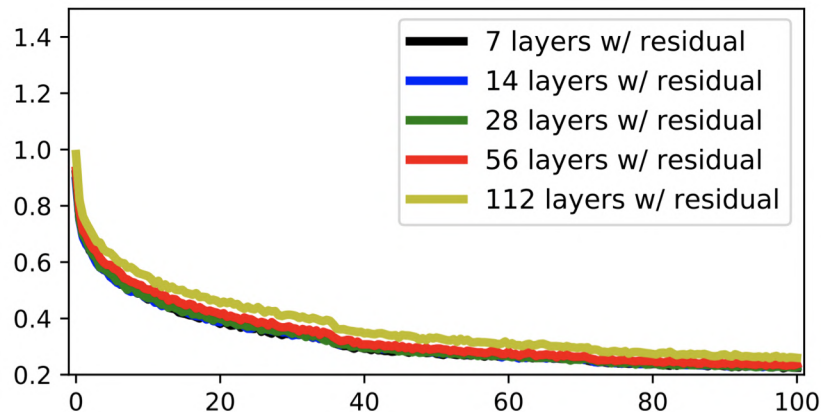
# Training Loss of GCNs with varying depth

Deeper GCNs don't converge well.



PlainGCNs

Even a 112-layer deep GCN converges well!!!!

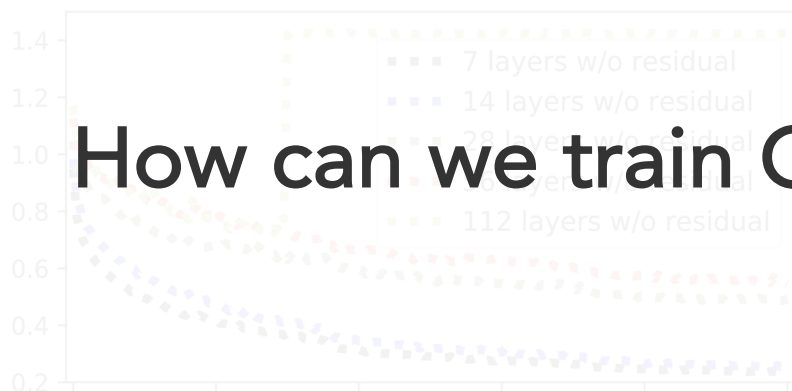


ResGCNs

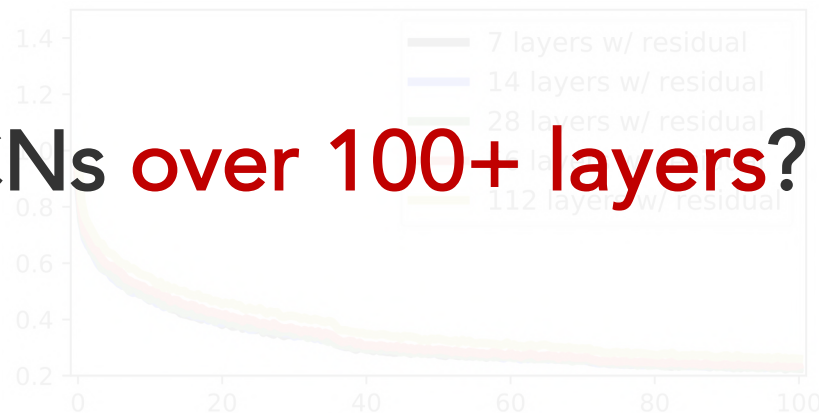
# Training Loss of GCNs with varying depth

Deeper GCNs don't converge well.

Even a 112-layer deep GCN converges well!!!



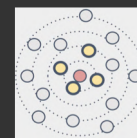
PlainGCNs



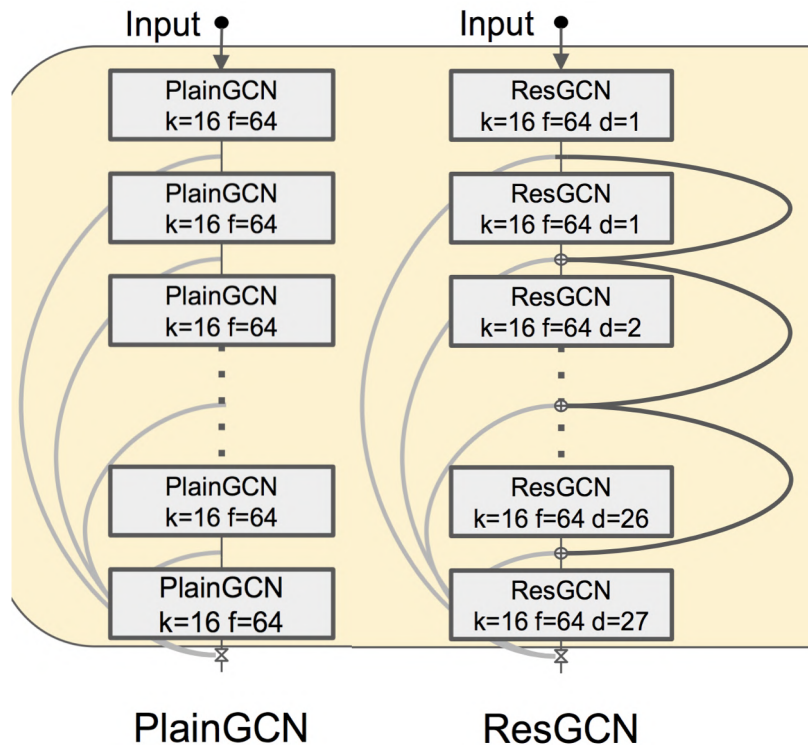
ResGCNs

How can we train GCNs **over 100+ layers?**

# Residual Graph Connections



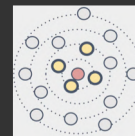
DeepGCNs.org



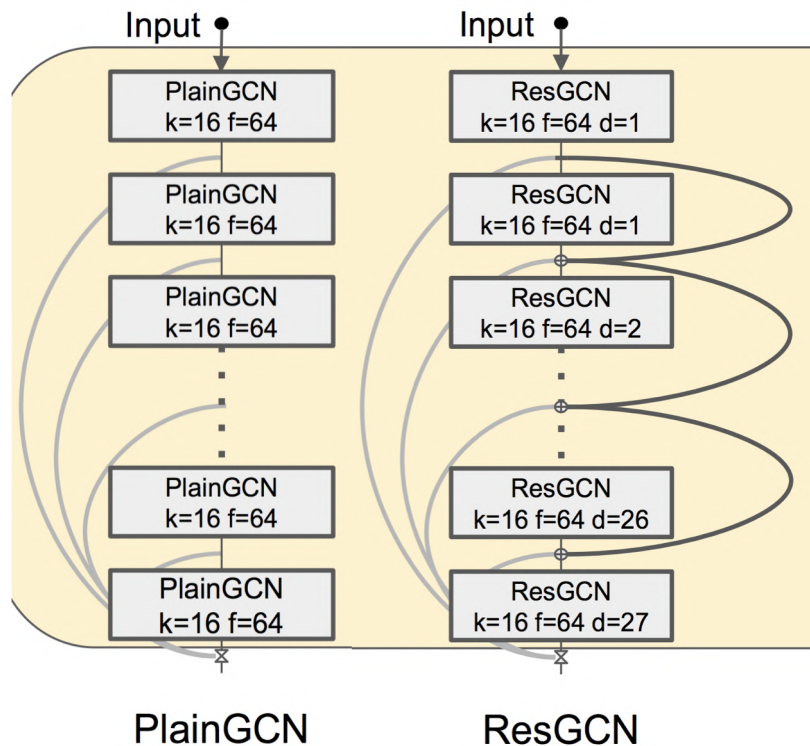
$$\begin{aligned}\mathcal{G}_{l+1} &= \mathcal{H}(\mathcal{G}_l, \mathcal{W}_l) \\ &= \mathcal{F}(\mathcal{G}_l, \mathcal{W}_l) + \mathcal{G}_l.\end{aligned}$$



# Residual Graph Connections



DeepGCNs.org



$$\begin{aligned}\mathcal{G}_{l+1} &= \mathcal{H}(\mathcal{G}_l, \mathcal{W}_l) \\ &= \mathcal{F}(\mathcal{G}_l, \mathcal{W}_l) + \mathcal{G}_l.\end{aligned}$$

An example: ResMRGCN

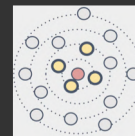
$$h_{\mathcal{N}^{(d)}(v_l)}^{res} = \max \left( \{h_{u_l} - h_{v_l} | u_l \in \mathcal{N}^{(d)}(v_l)\} \right), \quad \text{Aggregate}$$

$$h_{v_{l+1}}^{res} = \text{mlp} \left( \text{concat} \left( h_{v_l}, h_{\mathcal{N}^{(d)}(v_l)}^{res} \right) \right), \quad \text{Update}$$

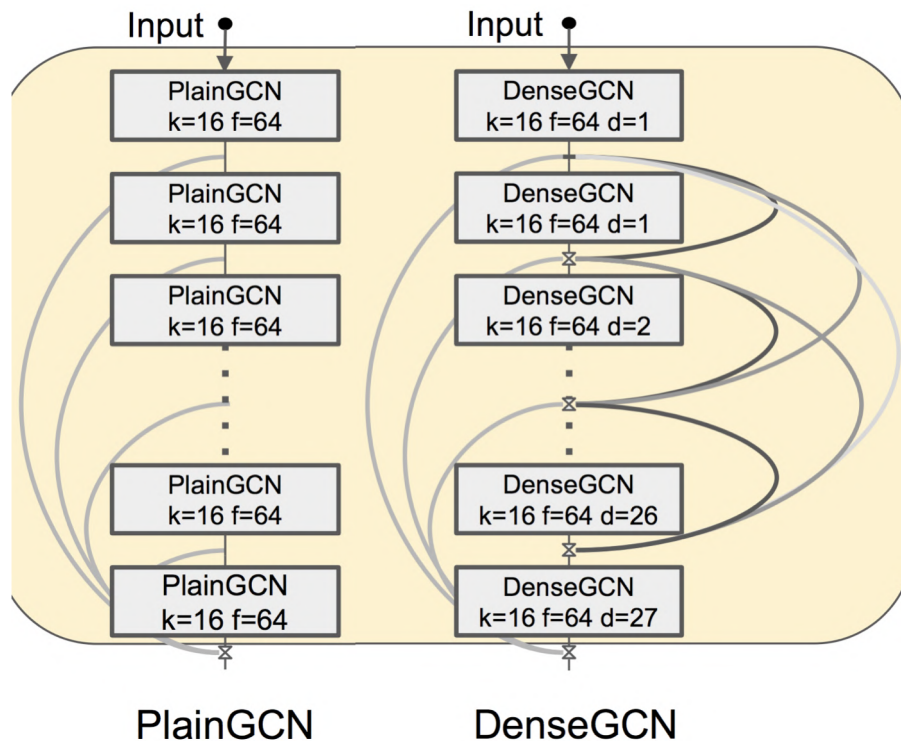
$$h_{v_{l+1}} = h_{v_{l+1}}^{res} + h_{v_l}. \quad \text{Skip connection}$$



# Dense Graph Connections

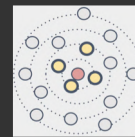


DeepGCNs.org

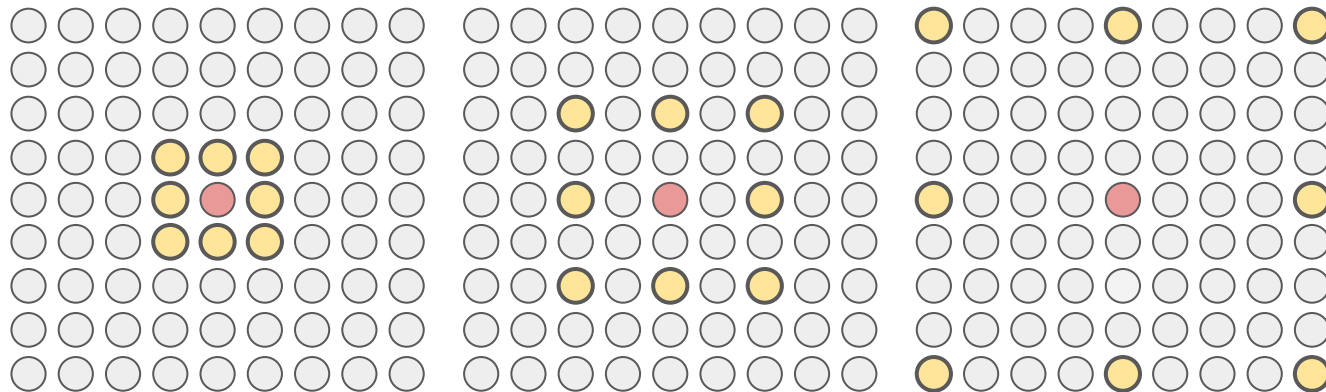


$$\begin{aligned}\mathcal{G}_{l+1} &= \mathcal{H}(\mathcal{G}_l, \mathcal{W}_l) \\ &= \mathcal{T}(\mathcal{F}(\mathcal{G}_l, \mathcal{W}_l), \mathcal{G}_l) \\ &= \mathcal{T}(\mathcal{F}(\mathcal{G}_l, \mathcal{W}_l), \dots, \mathcal{F}(\mathcal{G}_0, \mathcal{W}_0), \mathcal{G}_0).\end{aligned}$$

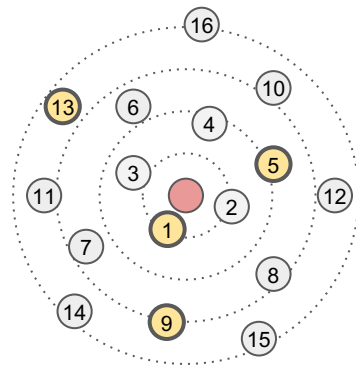
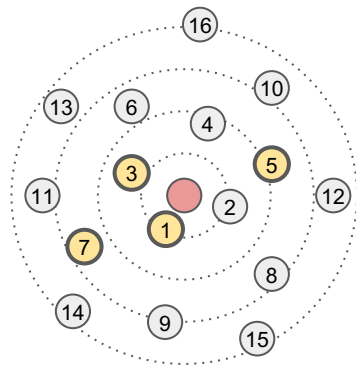
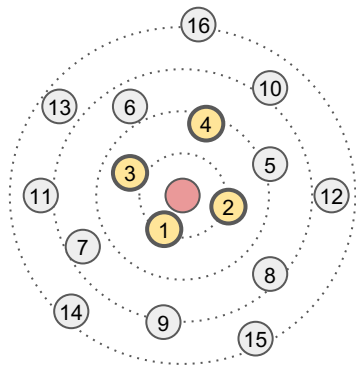
# Dilated Graph Convolutions



DeepGCNs.org

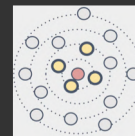


Dilated Convolution  
on a regular graph,  
e.g. 2D image

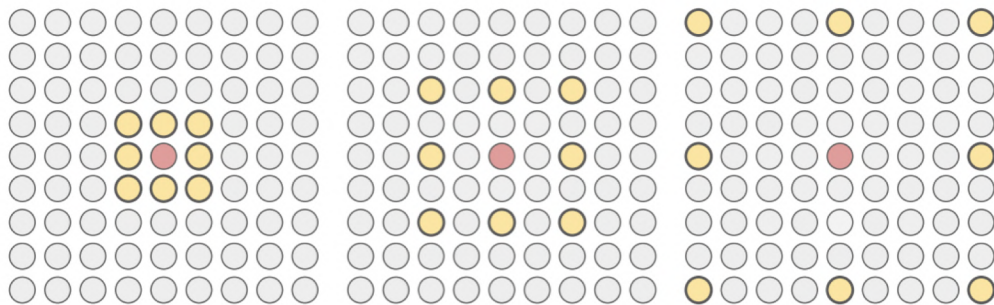


Dilated graph  
Convolution on an  
irregular graph, e.g.  
3D point cloud

# Dilated Graph Convolutions

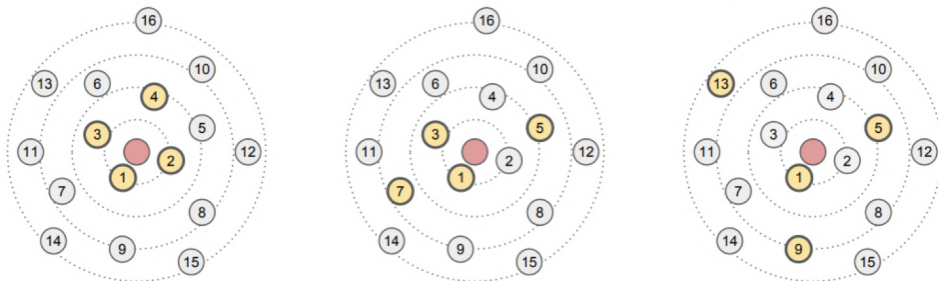


DeepGCNs.org

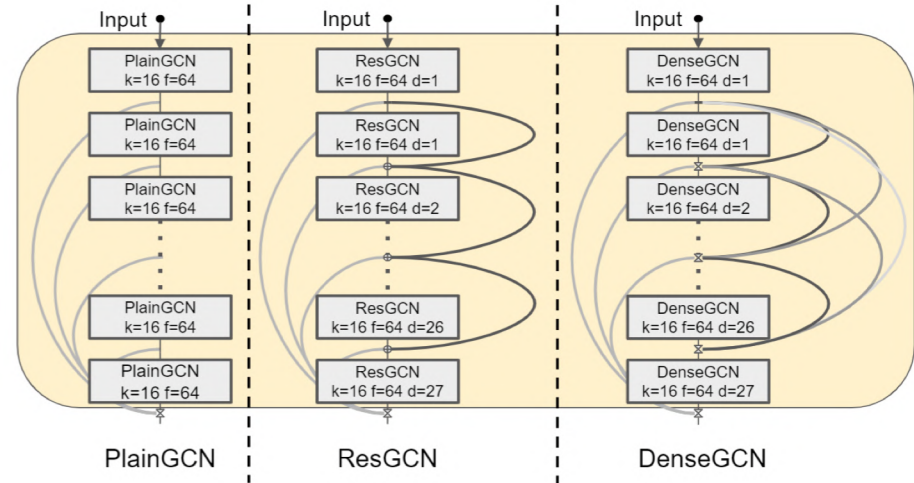
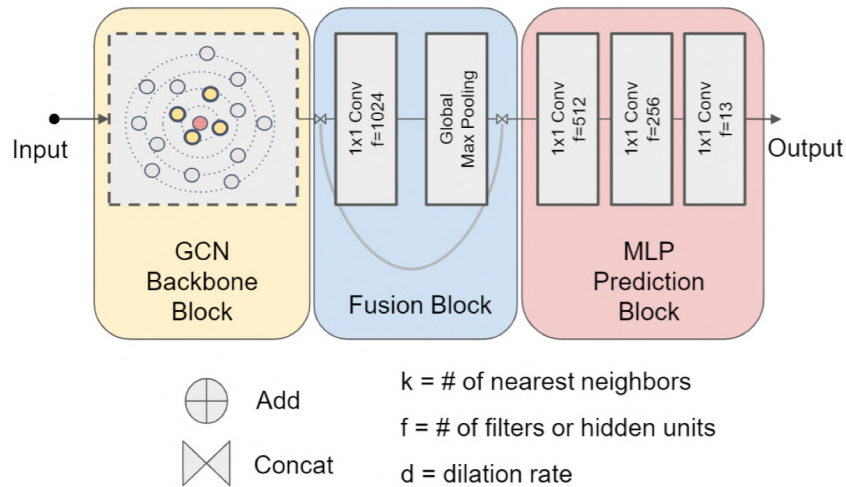


$$\mathcal{N}^{(d)}(v) = \{u_1, u_{1+d}, u_{1+2d}, \dots, u_{1+(k-1)d}\}.$$

$d$  = dilation rate

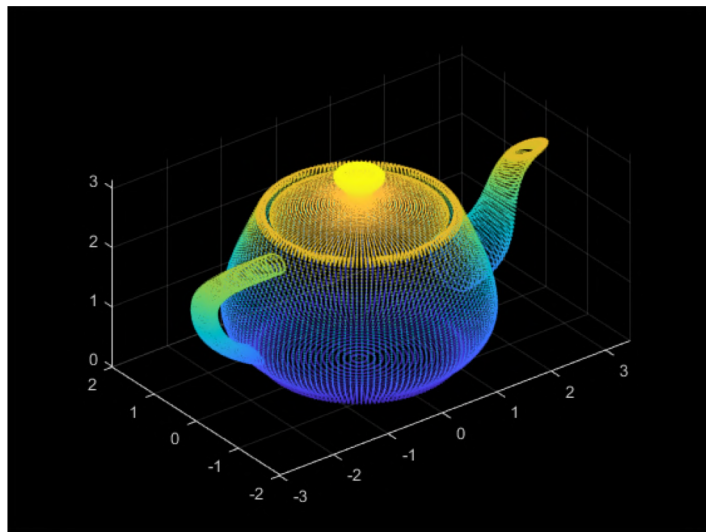


# Deep Graph Convolutional Networks (GCNs)



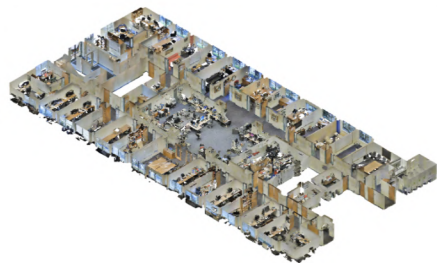
# Experiments

# Graph Learning on 3D Point Clouds

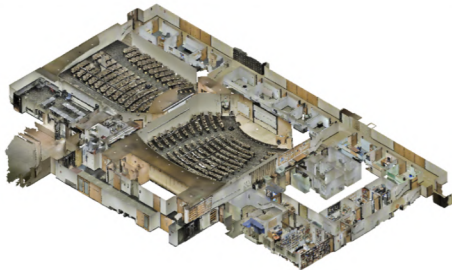


- Point clouds are unordered and irregular
- Represented by 3D coordinates and extra features such as color, surface normal, etc.
- We use k-NN to construct the directed dynamic edges between points at every GCN layer in the feature space.

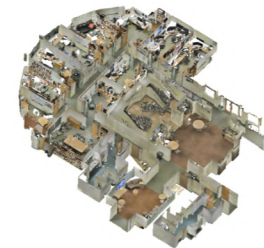
# Stanford 3D Large-Scale Indoor Spaces Dataset



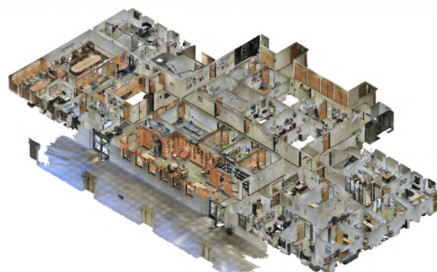
Area 1



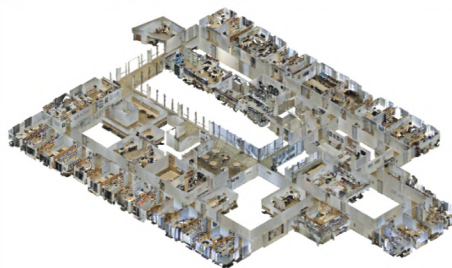
Area 2



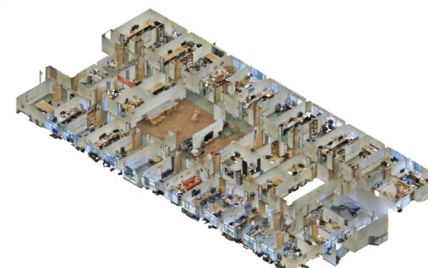
Area 3



Area 4



Area 5



Area 6

<http://buildingparser.stanford.edu/dataset.html>



## We outperform other SOTA in 9 out of 13 classes

Method	OA	mIOU	ceiling	floor	wall	beam	column	window	door	table	chair	sofa	bookcase	board	clutter
PointNet [29]	78.5	47.6	88.0	88.7	69.3	42.4	23.1	47.5	51.6	54.1	42.0	9.6	38.2	29.4	35.2
MS+CU [8]	79.2	47.8	88.6	<b>95.8</b>	67.3	36.9	24.9	48.6	52.3	51.9	45.1	10.6	36.8	24.7	37.5
G+RCU [8]	81.1	49.7	90.3	92.1	67.9	<b>44.7</b>	24.2	52.3	51.2	58.1	47.4	6.9	39.0	30.0	41.9
PointNet++ [31]	-	53.2	90.2	91.7	73.1	42.7	21.2	49.7	42.3	62.7	59.0	19.6	45.8	48.2	45.6
3DRNN+CF [51]	<b>86.9</b>	56.3	92.9	93.8	73.1	42.5	25.9	47.6	59.2	60.4	<b>66.7</b>	24.8	<b>57.0</b>	36.7	51.6
DGCNN [43]	84.1	56.1	-	-	-	-	-	-	-	-	-	-	-	-	-
<b>ResGCN-28 (Ours)</b>	<b>85.9</b>	<b>60.0</b>	<b>93.1</b>	95.3	<b>78.2</b>	33.9	<b>37.4</b>	<b>56.1</b>	<b>68.2</b>	<b>64.9</b>	61.0	<b>34.6</b>	51.5	<b>51.1</b>	<b>54.4</b>

Table 1. Comparison of ResGCN-28 with state-of-the-art.



Consistent improvements  
across all the classes.

Class	DGCNN [6]	ResGCN-28 ( <i>Ours</i> )
ceiling	92.7	<b>93.1</b>
floor	93.6	<b>95.3</b>
wall	77.5	<b>78.2</b>
beam	32.0	<b>33.9</b>
column	36.3	<b>37.4</b>
window	52.5	<b>56.1</b>
door	63.7	<b>68.2</b>
table	61.1	<b>64.9</b>
chair	60.2	<b>61.0</b>
sofa	20.5	<b>34.6</b>
bookcase	47.7	<b>51.5</b>
board	42.7	<b>51.1</b>
clutter	51.5	<b>54.4</b>
<b>mIOU</b>	56.3	<b>60.0</b>

Table 2. Comparison of ResGCN-28 with DGCNN\* (Our shallow baseline model).

\* We reproduced the results of DGCNN on all classes since the results across all classes were not provided in the DGCNN paper.

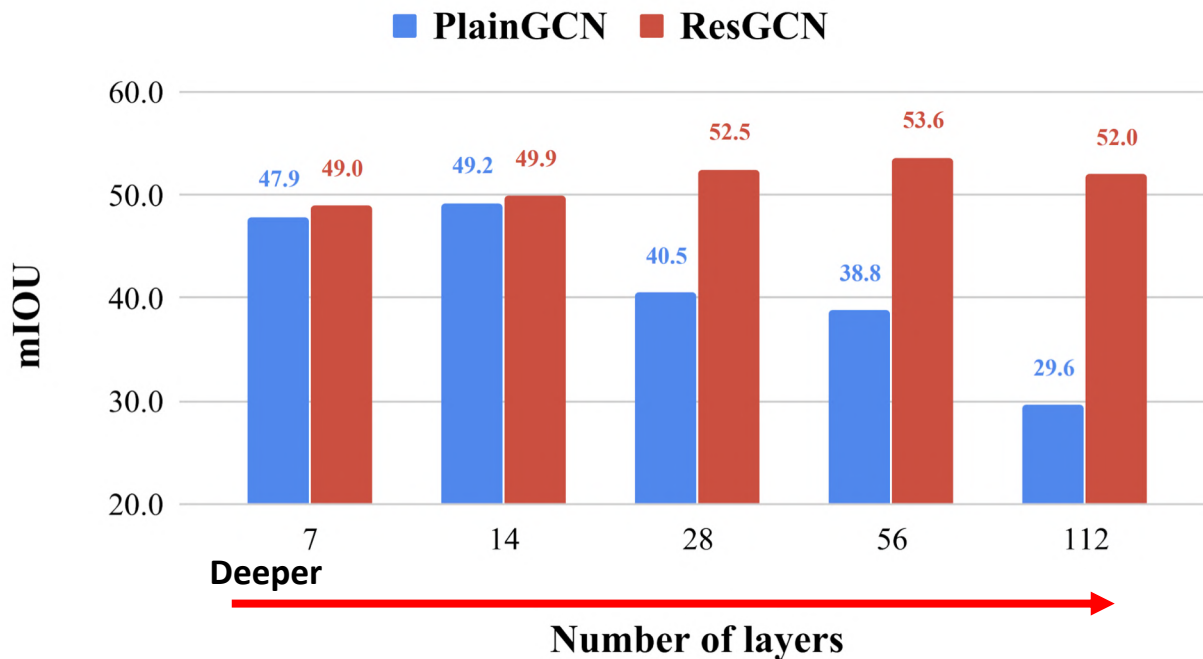
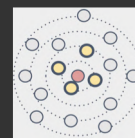
Consistent improvements  
across all the classes.

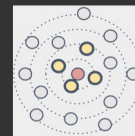
Class	DGCNN [6]	ResGCN-28 ( <i>Ours</i> )
ceiling	92.7	<b>93.1</b>
floor	93.6	<b>95.3</b>
wall	77.5	<b>78.2</b>
beam	32.0	<b>33.9</b>
column	36.3	<b>37.4</b>
window	52.5	<b>56.1</b>
door	63.7	<b>68.2</b>
table	61.1	<b>64.9</b>
chair	60.2	<b>61.0</b>
sofa	20.5	<b>34.6</b>
bookcase	47.7	<b>51.5</b>
board	42.7	<b>51.1</b>
clutter	51.5	<b>54.4</b>
<b>mIOU</b>	56.3	<b>60.0</b>

~ 4% boost in mIOU.

Table 2. Comparison of ResGCN-28 with DGCNN\* (Our shallow baseline model).

\* We reproduced the results of DGCNN on all classes since the results across all classes were not provided in the DGCNN paper.





## Visualizations on S3DIS

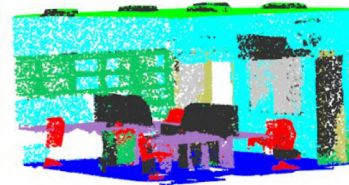
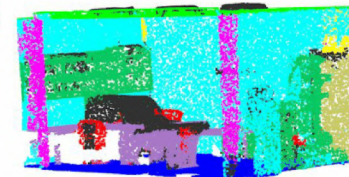
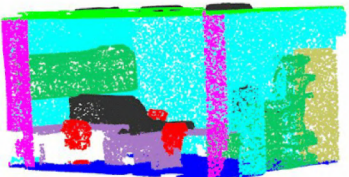
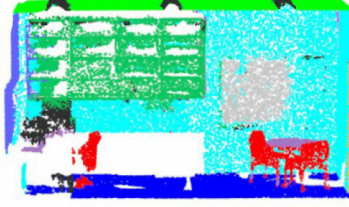
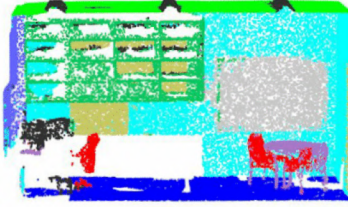
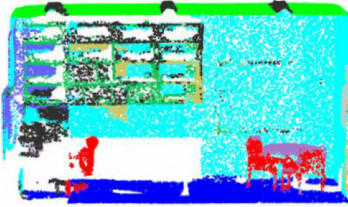
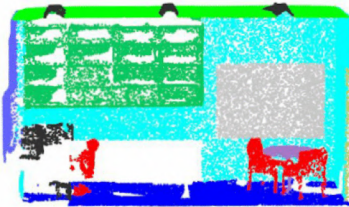
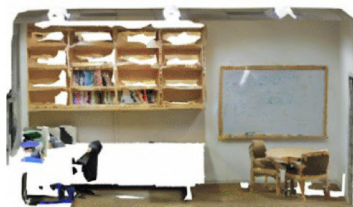
Original

Ground Truth

PlainGCN

ResGCN

DenseGCN



Ceiling

Floor

Wall

Beam

Column

Window

Door

Table

Chair

Sofa

Bookcase

Board

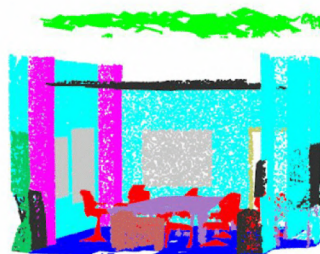
Clutter



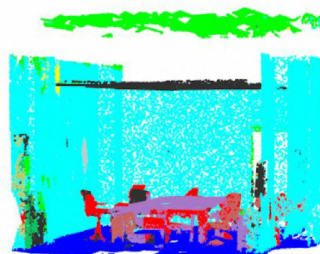
Original



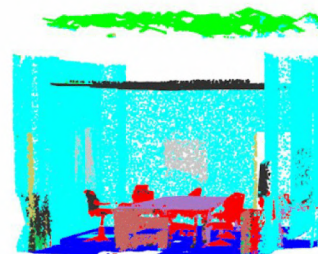
Ground Truth



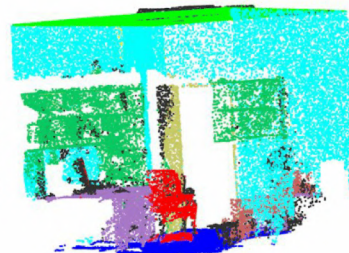
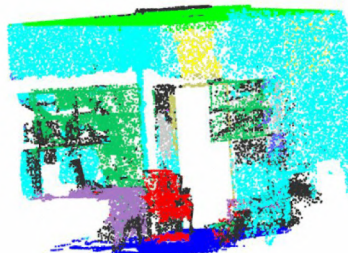
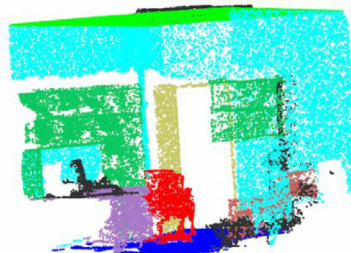
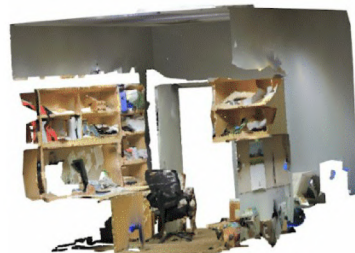
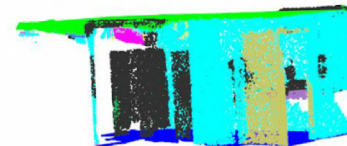
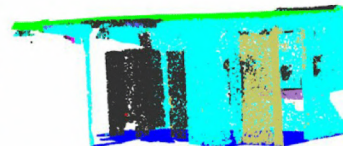
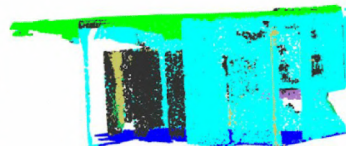
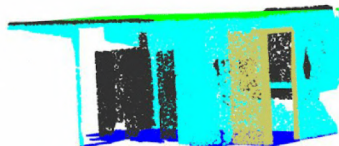
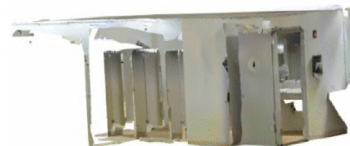
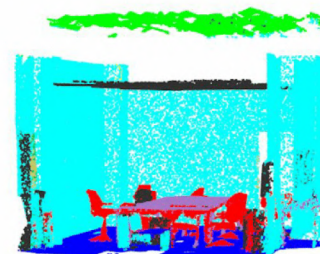
PlainGCN



ResGCN



DenseGCN



Ceiling

Floor

Wall

Beam

Column

Window

Door

Table

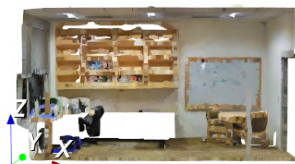
Chair

Sofa

Bookcase

Board

Clutter



Original



Ground Truth



ResGCN-28

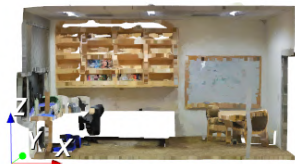


1/2x filters



1/4x filters

### Reduce Network Width



Original



Ground Truth



ResGCN-28

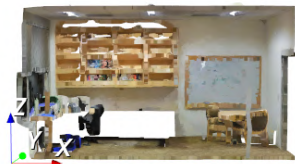


1/2x NNs



1/4x NNs

### Reduce Kernel Size



Original



Ground Truth



ResGCN-28



1/2x layers

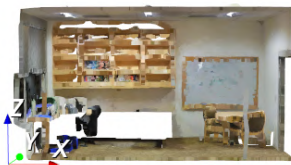


1/4x layers

### Reduce Network Depth







Original



Ground Truth



ResGCN-28

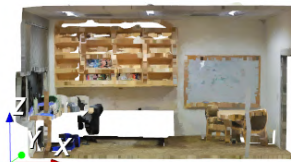


w/o stochastic



w/o dilation

No Dilation



Original



Ground Truth



ResGCN-28



ResGCN-28W



ResGCN-56

Wider

Deeper

Ceiling

Floor

Wall

Beam

Column

Window

Door

Table

Chair

Sofa

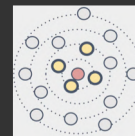
Bookcase

Board

Clutter







## GCN variants

- ResEdgeConv
- ResGraphSAGE
- ResGIN
- ResMRGCN

Model	mIoU	$\Delta$ mIoU	dynamic	connection	dilation	stochastic	# NNs	# filters	# layers
<i>ResEdgeConv-28</i>	<b>52.49</b>	0.00	✓	$\oplus$	✓	✓	16	64	28
<i>PlainGCN-28</i>	<b>40.31</b>	-12.18	✓				16	64	28
<i>ResGraphSAGE-28</i>	<b>49.20</b>	-3.29	✓	$\oplus$	✓	✓	16	64	28
<i>ResGraphSAGE-N-28</i>	<b>49.02</b>	-3.47	✓	$\oplus$	✓	✓	16	64	28
<i>ResGIN-<math>\epsilon</math>-28</i>	<b>42.81</b>	-9.68	✓	$\oplus$	✓	✓	16	64	28
<i>ResMRGCN-28</i>	<b>51.17</b>	-1.32	✓	$\oplus$	✓	✓	16	64	28

Table 4. Comparisons of Deep GCNs variants on area 5 of S3DIS.

$$h_{v_{l+1}}^{res} = \max \left( \text{mlp}(\{\text{concat}(h_{v_l}, h_{u_l} - h_{v_l}) | u_l \in \mathcal{N}^{(d)}(v_l)\}) \right),$$

$$h_{v_{l+1}} = h_{v_{l+1}}^{res} + h_{v_l}.$$

ResEdgeConv

$$h_{\mathcal{N}^{(d)}(v_l)}^{res} = \max \left( \{\text{mlp}(h_{u_l}) | u_l \in \mathcal{N}^{(d)}(v_l)\} \right),$$

$$h_{v_{l+1}}^{res} = \text{mlp} \left( \text{concat} \left( h_{v_l}, h_{\mathcal{N}^{(d)}(v_l)}^{res} \right) \right),$$

$$h_{v_{l+1}} = h_{v_{l+1}}^{res} + h_{v_l},$$

ResGraphSAGE

$$h_{v_{l+1}}^{res} = \text{mlp} \left( (1 + \epsilon) \cdot h_{v_l} + \text{sum}(\{h_{u_l} | u_l \in \mathcal{N}^{(d)}(v_l)\}) \right),$$

$$h_{v_{l+1}} = h_{v_{l+1}}^{res} + h_{v_l}.$$

ResGIN

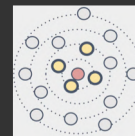
$$h_{\mathcal{N}^{(d)}(v_l)}^{res} = \max \left( \{h_{u_l} - h_{v_l} | u_l \in \mathcal{N}^{(d)}(v_l)\} \right),$$

$$h_{v_{l+1}}^{res} = \text{mlp} \left( \text{concat} \left( h_{v_l}, h_{\mathcal{N}^{(d)}(v_l)}^{res} \right) \right),$$

$$h_{v_{l+1}} = h_{v_{l+1}}^{res} + h_{v_l}.$$

ResMRGCN

# More Results



DeepGCNs.org



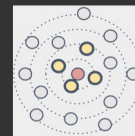
By John Morris.

Deeper

Wider

Number of filters	32	64	128	256
<i>PlainMRGCN-3</i>	95.84	97.60	98.58	99.13
<i>PlainMRGCN-7</i>	97.35	98.69	99.22	<b>99.38</b>
<i>PlainMRGCN-14</i>	97.55	99.02	99.31	99.34
<i>PlainMRGCN-28</i>	98.09	99.00	99.02	99.31
<i>PlainMRGCN-56</i>	92.70	97.43	97.31	97.61
<i>PlainMRGCN-112</i>	60.75	71.97	89.69	91.50
<i>ResMRGCN-3</i>	96.04	97.60	98.53	99.09
<i>ResMRGCN-7</i>	97.00	98.43	99.19	99.30
<i>ResMRGCN-14</i>	97.75	98.88	99.26	<b>99.38</b>
<i>ResMRGCN-28</i>	98.50	99.16	99.29	<b>99.41</b>
<i>ResMRGCN-56</i>	98.62	99.27	<b>99.36</b>	<b>99.40</b>
<i>ResMRGCN-112</i>	98.41	99.34	<b>99.38</b>	<b>99.39</b>
<i>DenseMRGCN-3</i>	95.96	97.85	98.66	99.11
<i>DenseMRGCN-7</i>	97.87	98.47	99.31	<b>99.36</b>
<i>DenseMRGCN-14</i>	98.93	99.00	99.01	<b>99.43</b>
<i>DenseMRGCN-28</i>	99.16	99.29	<b>99.42</b>	-
<i>DenseMRGCN-56</i>	99.22	-	-	-

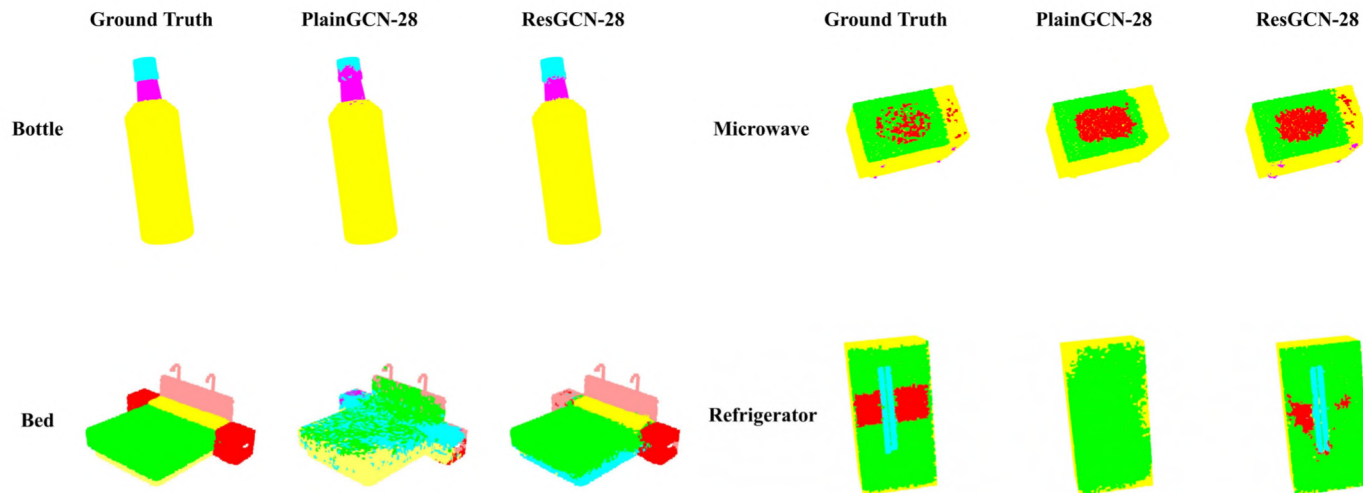
Table 5. Node classification of biological networks.



By John Morris.

Model	m-F1 score (%)
GraphSAGE [42]	61.20
GATConv [43]	97.30
VR-GCN [57]	97.80
GaAN [58]	98.71
GeniePath [59]	98.50
Cluster-GCN [56]	99.36
<b><i>ResMRGCN-28 (Ours)</i></b>	<b>99.41</b>
<b><i>DenseMRGCN-14 (Ours)</i></b>	<b>99.43</b>

Table 6. Comparison of DeepGCNs with state-of-the-art on PPI node classification.



Method	bed	bottle	chair	clock	dishw.	disp.	door	earph.	fauc.	knife	lamp	micro.	fridge	st. furn.	table	tr. can	vase
PointNet [33]	13.4	29.5	27.8	28.4	48.9	76.5	30.4	33.4	47.6	32.9	18.9	37.2	33.5	38.0	29.0	34.8	44.4
PointNet++ [34]	30.3	<b>41.4</b>	<b>39.2</b>	<b>41.6</b>	50.1	80.7	32.6	38.4	<b>52.4</b>	34.1	<b>25.3</b>	48.5	36.4	40.5	<b>33.9</b>	46.7	49.8
SpiderCNN [53]	<b>36.2</b>	32.2	30.0	24.8	50.0	80.1	30.5	37.2	44.1	22.2	19.6	43.9	39.1	<b>44.6</b>	20.1	42.4	32.4
<b>ResGCN-28 (Ours)</b>	35.2	36.8	33.8	32.6	<b>52.7</b>	<b>84.4*</b>	<b>42.5*</b>	<b>41.9</b>	49.7	<b>35.4*</b>	20.0	<b>54.3</b>	<b>46.1*</b>	42.5	14.8	<b>49.7</b>	<b>50.8</b>
PointCNN [54]	41.9	41.8	43.9	36.3	58.7	82.5	37.8	48.9	60.5	34.1	20.1	58.2	42.9	49.4	21.3	53.1	58.9

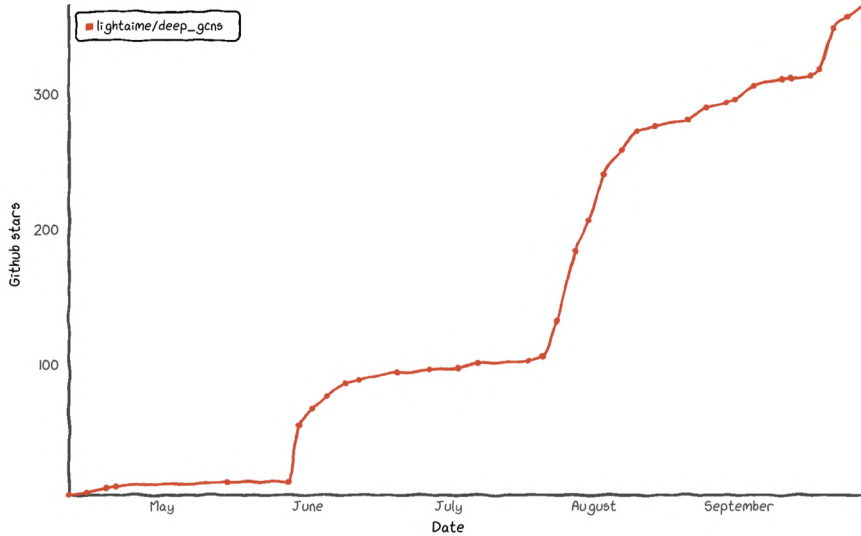
Table 7. Comparison of ResGCN-28 with other methods on PartNet Part Segmentation.

TensorFlow Repo

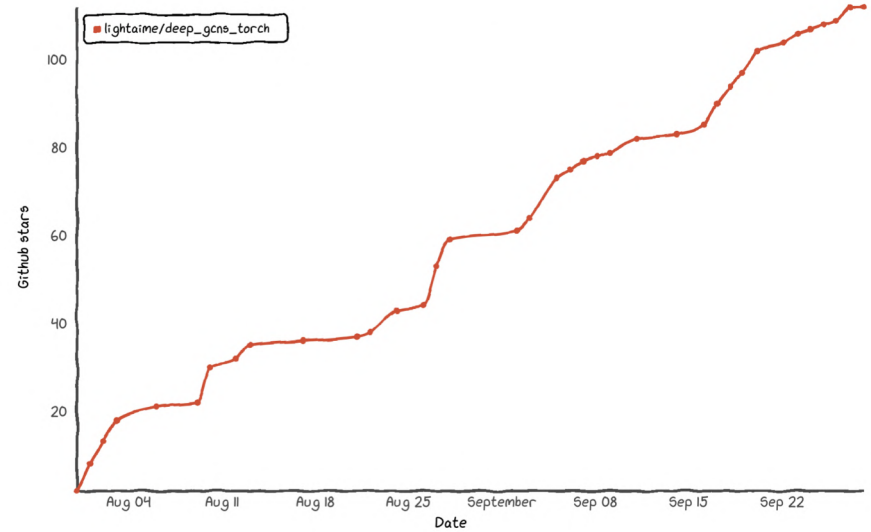
★ 750+ Stars

Pytorch Repo

Star history



Star history



<https://www.deepgcns.org>

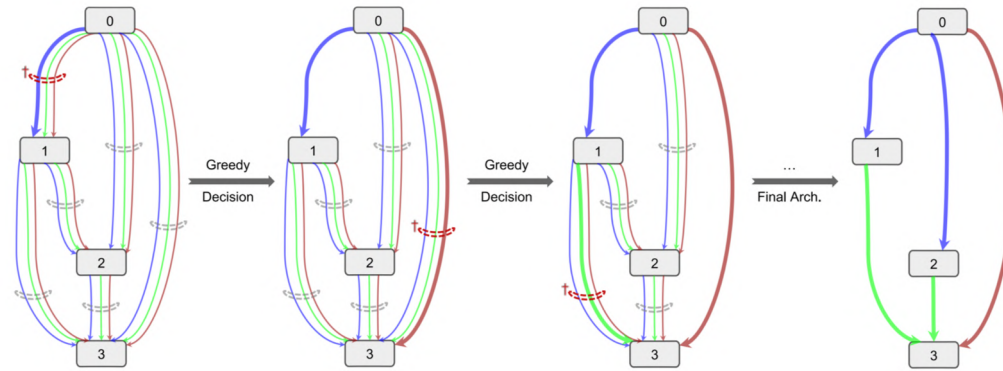
# Can we learn how to design Deep GCN architectures automatically?

Can we learn how to design Deep GCN architectures automatically?

**Neural Architecture Search!**



# SGAS: Sequential Greedy Architecture Search



SGAS: Sequential Greedy Architecture Search (arXiv 2019, Guohao Li et.al)

<https://sites.google.com/kaust.edu.sa/sgas>

# SGAS: Sequential Greedy Architecture Search

Aiming to alleviate this common issue, we introduce **sequential greedy architecture search** (SGAS), an efficient method for neural architecture search.

By dividing the search procedure into **sub-problems**, SGAS chooses and prunes candidate operations in a greedy fashion.

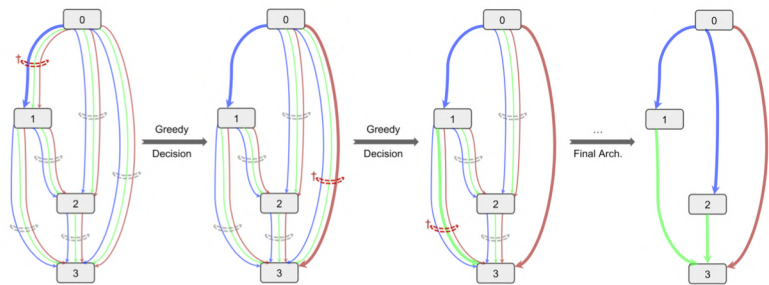


Figure 2. Illustration of Sequential Greedy Architecture Search.

# SGAS: Sequential Greedy Architecture Search

We apply SGAS to search architectures for **Convolutional Neural Networks** (CNN) and **Graph Convolutional Networks** (GCN).

Extensive experiments show that SGAS is able to find SOTA architectures with minimal computational cost for tasks such as:

- image classification,
- point cloud classification,
- node classification in protein-protein interaction graphs.

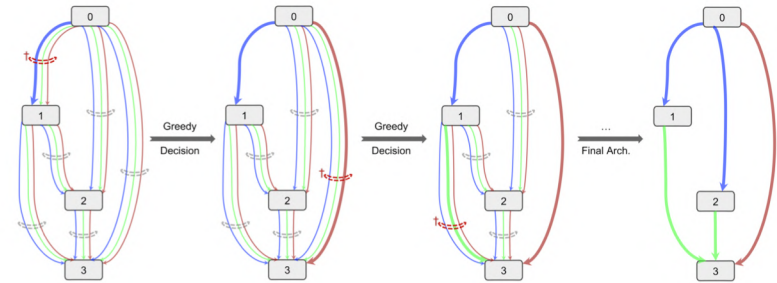


Figure 2. Illustration of Sequential Greedy Architecture Search.

# SGAS: Sequential Greedy Architecture Search

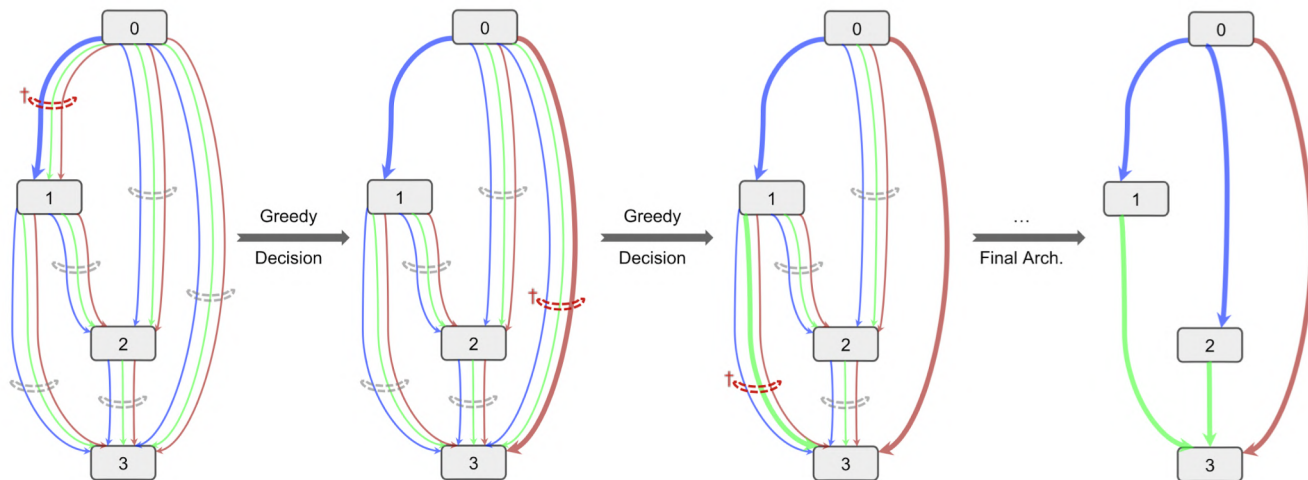


Figure 2. Illustration of Sequential Greedy Architecture Search.

# SGAS: Sequential Greedy Architecture Search

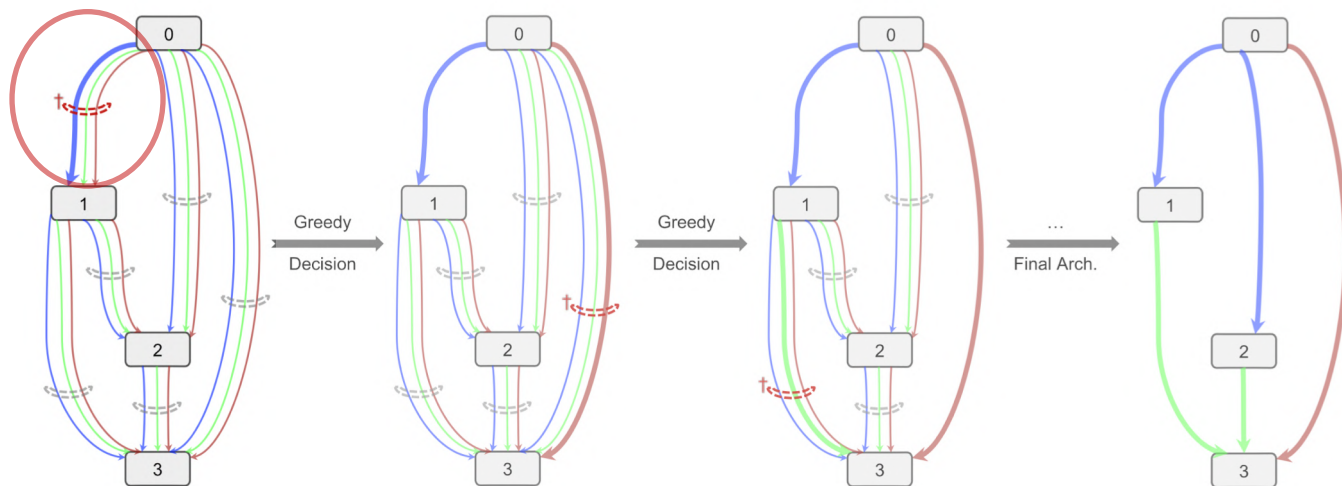


Figure 2. Illustration of Sequential Greedy Architecture Search.

- ① If a decision epoch, select an edge  $(i^\dagger, j^\dagger)$  based on the greedy *Selection Criterion*

# SGAS: Sequential Greedy Architecture Search

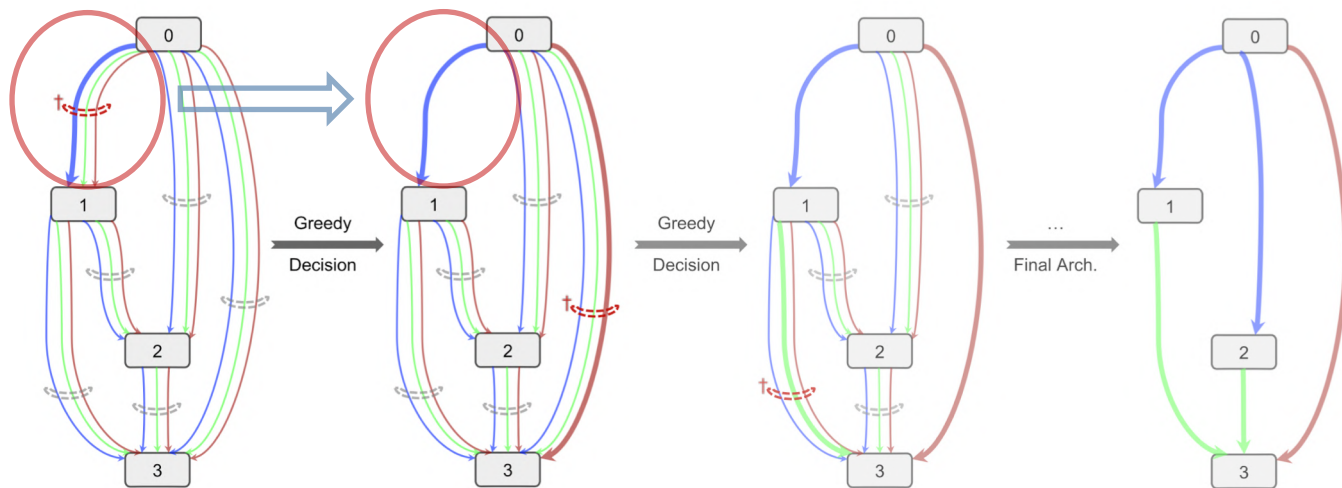


Figure 2. Illustration of Sequential Greedy Architecture Search.

- ① If a decision epoch, select an edge  $(i^\dagger, j^\dagger)$  based on the greedy *Selection Criterion*
- ② Determine the operation by replacing  $\bar{o}^{(i^\dagger, j^\dagger)}$  with  $o^{(i^\dagger, j^\dagger)} = \operatorname{argmax}_{o \in \mathcal{O}} \alpha_o^{(i^\dagger, j^\dagger)}$

# SGAS: Sequential Greedy Architecture Search

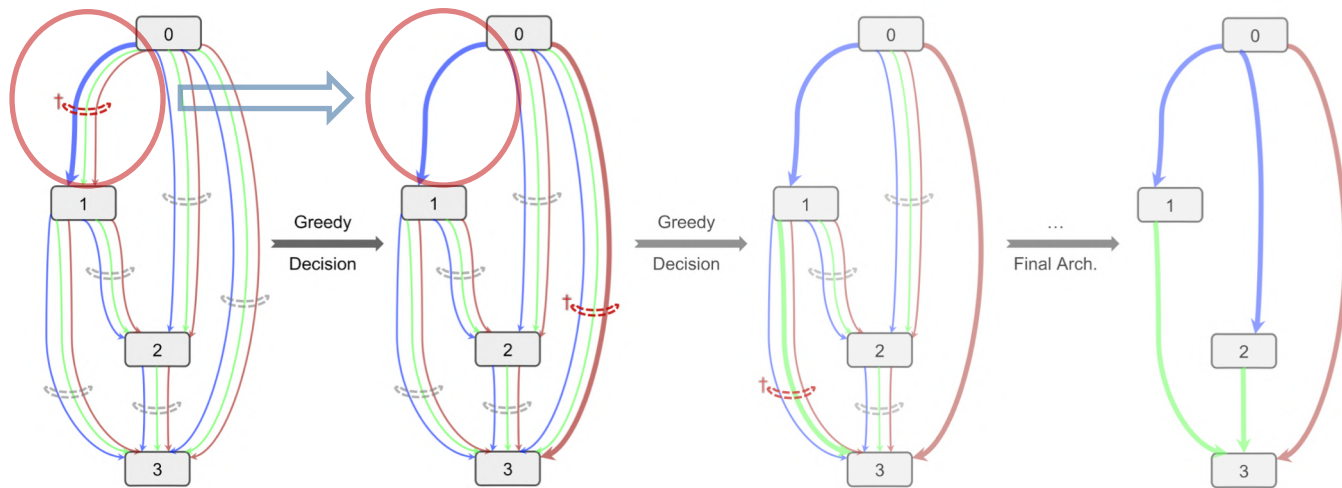


Figure 2. Illustration of Sequential Greedy Architecture Search.

- ① If a decision epoch, select an edge  $(i^\dagger, j^\dagger)$  based on the greedy *Selection Criterion*
- ② Determine the operation by replacing  $\bar{o}^{(i^\dagger, j^\dagger)}$  with  $o^{(i^\dagger, j^\dagger)} = \operatorname{argmax}_{o \in \mathcal{O}} \alpha_o^{(i^\dagger, j^\dagger)}$
- ③ Prune unchosen weights from  $\mathcal{W}$ , Remove  $\alpha^{(i^\dagger, j^\dagger)}$  from  $\mathcal{A}$

# SGAS: Sequential Greedy Architecture Search

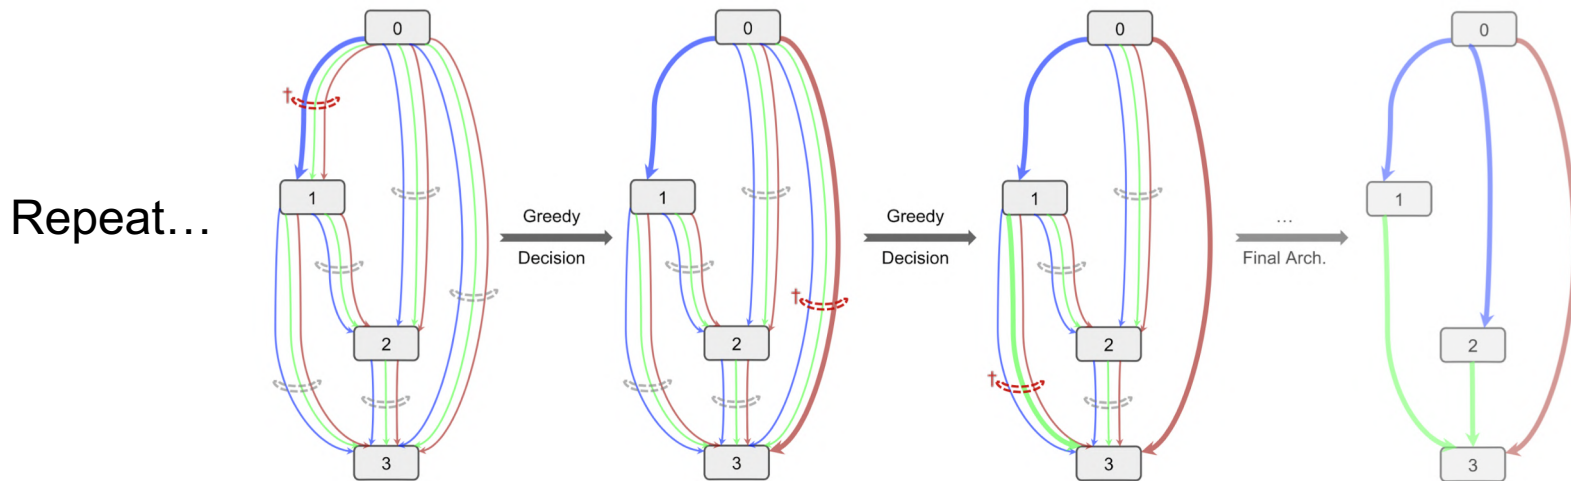


Figure 2. Illustration of Sequential Greedy Architecture Search.

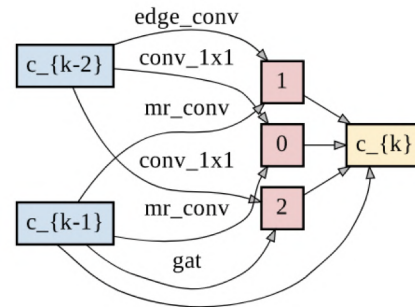
- ① If a decision epoch, select an edge  $(i^\dagger, j^\dagger)$  based on the greedy *Selection Criterion*
- ② Determine the operation by replacing  $\bar{o}^{(i^\dagger, j^\dagger)}$  with  $o^{(i^\dagger, j^\dagger)} = \operatorname{argmax}_{o \in \mathcal{O}} \alpha_o^{(i^\dagger, j^\dagger)}$
- ③ Prune unchosen weights from  $\mathcal{W}$ , Remove  $\alpha^{(i^\dagger, j^\dagger)}$  from  $\mathcal{A}$



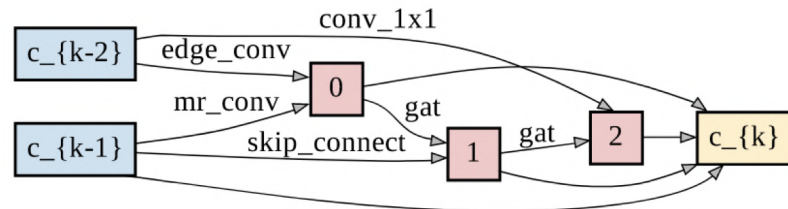
# Results – SGAS for GCN on ModelNet

Architecture	OA (%)	Params (M)	Search Cost (GPU-days)
3DmFV-Net [4]	91.6	45.77	manual
SpecGCN [54]	91.5	2.05	manual
PointNet++ [42]	90.7	1.48	manual
PCNN [3]	92.3	8.2	manual
PointCNN [31]	92.2	0.6	manual
DGCNN [55]	92.2	1.84	manual
KPConv [51]	92.9	14.3	manual
Random Search	92.65±0.33	8.77	random
SGAS (Cri.1 avg.)	92.69±0.20	8.78	0.19
SGAS (Cri.1 best)	92.87	8.63	0.19
SGAS (Cri.2 avg.)	92.93±0.19	8.87	0.19
SGAS (Cri.2 best)	<b>93.23</b>	8.49	0.19
SGAS (Cri.2 small best)	93.07	3.86	0.19

Table 3. Comparison with state-of-the-art architectures for 3D object classification on ModelNet40.



(a) Normal cell of the best model with SGAS (Cri. 1) on ModelNet

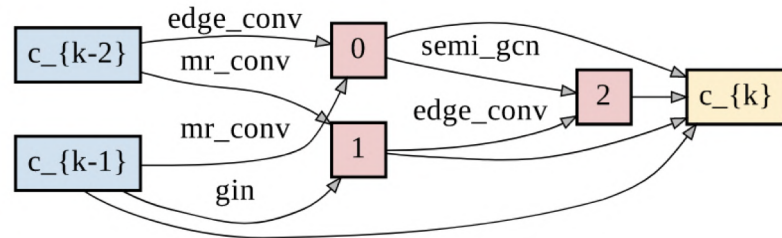


(b) Normal cell of the best model with SGAS (Cri. 2) on ModelNet

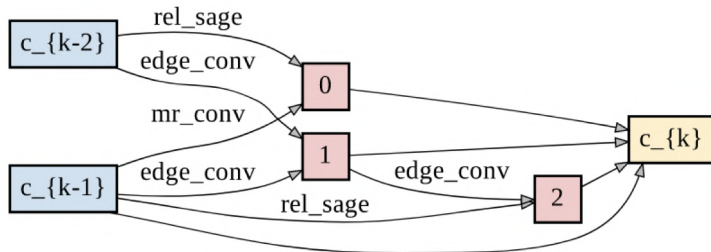
# Results – SGAS for GCN on PPI

Architecture	micro-F1 (%)	Params (M)	Search Cost (GPU-days)
GraphSAGE (LSTM) [14]	61.2	0.26	manual
GeniePath [30]	97.9	1.81	manual
GAT [44]	97.3±0.2	3.64	manual
DenseMRGCN-14 [23]	99.43	53.42	manual
ResMRGCN-28 [23]	99.41	14.76	manual
Random Search	99.36±0.04	23.70	random
SGAS (Cri.1 avg.)	99.38±0.17	25.01	0.003
SGAS (Cri.1 best)	<b>99.46</b>	23.18	0.003
SGAS (Cri.2 avg.)	99.40±0.09	25.93	0.003
SGAS (Cri.2 best)	<b>99.46</b>	29.73	0.003
SGAS (small)	98.89	0.40	0.003

Table 4. Comparison with state-of-the-art architectures for node classification on PPI.



(a) Normal cell of the best model with SGAS (Cri. 1) on PPI



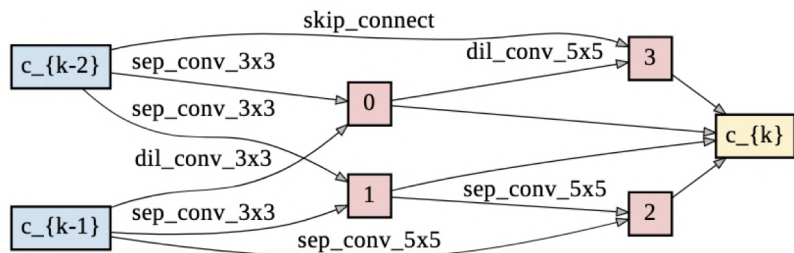
(b) Normal cell of the best model with SGAS (Cri. 2) on PPI

# Results – SGAS for CNN on CIFAR-10

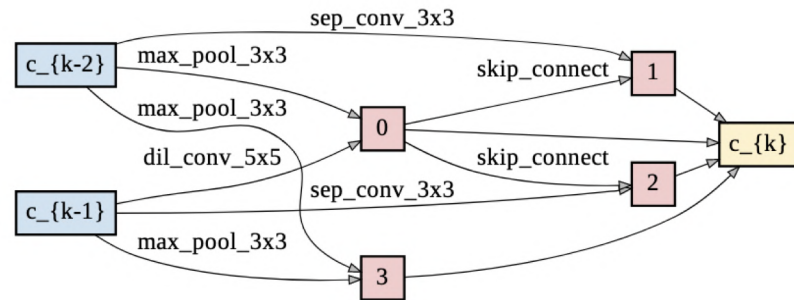
Architecture	Test Err. (%)	Params (M)	Search Cost (GPU-days)	Search Method
DenseNet-BC [18]	3.46	25.6	-	manual
NASNet-A [55]	2.65	3.3	1800	RL
AmoebaNet-A [36]	3.34±0.06	3.2	3150	evolution
AmoebaNet-B [36]	2.55±0.05	2.8	3150	evolution
Hier-Evolution [28]	3.75±0.12	15.7	300	evolution
PNAS [27]	3.41±0.09	3.2	225	SMBO
ENAS [34]	2.89	4.6	0.5	RL
NAONet-WS [31]	3.53	3.1	0.4	NAO
DARTS (1 <sup>st</sup> order) [29]	3.00±0.14	3.3	0.4	gradient
DARTS (2 <sup>nd</sup> order) [29]	2.76±0.09	3.3	1	gradient
SNAS (mild) [49]	2.98	2.9	1.5	gradient
ProxylessNAS [7]	2.08	-	4	gradient
P-DARTS [8]	2.5	3.4	0.3	gradient
BayesNAS [52]	2.81±0.04	3.4	0.2	gradient
PC-DARTS [50]	2.57±0.07	3.6	0.1	gradient
SGAS (Cri.1 avg.)	2.66±0.24*	3.7	0.25	gradient
SGAS (Cri.1 best)	2.39	3.8	0.25	gradient
SGAS (Cri.2 avg.)	2.67±0.21*	3.9	0.25	gradient
SGAS (Cri.2 best)	2.44	4.1	0.25	gradient

Table 1. Performance comparison with state-of-the-art image classifiers on CIFAR-10.

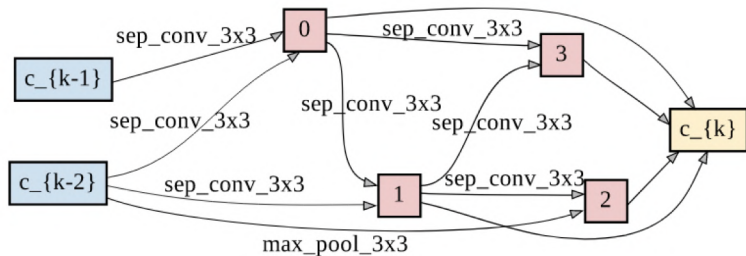
# Results – SGAS for CNN on CIFAR-10



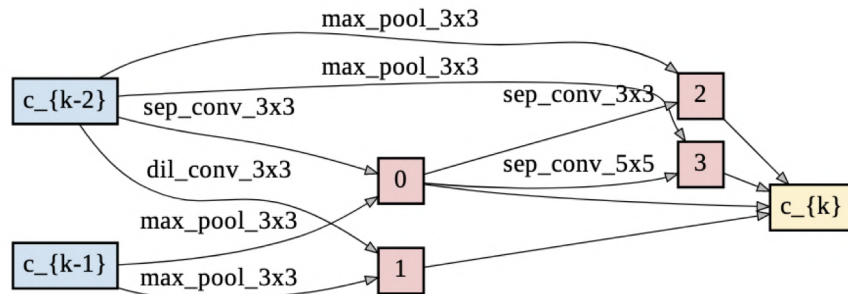
(a) Normal cell of the best model with SGAS (Cri. 1) on CIFAR-10



(b) Reduction cell of the best model with SGAS (Cri. 1) on CIFAR-10



(c) Normal cell of the best model with SGAS (Cri. 2) on CIFAR-10



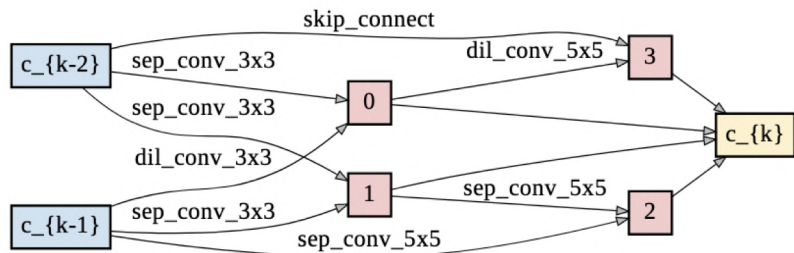
(d) Reduction cell of the best model with SGAS (Cri. 2) on CIFAR-10

# Results – SGAS for CNN on ImageNet

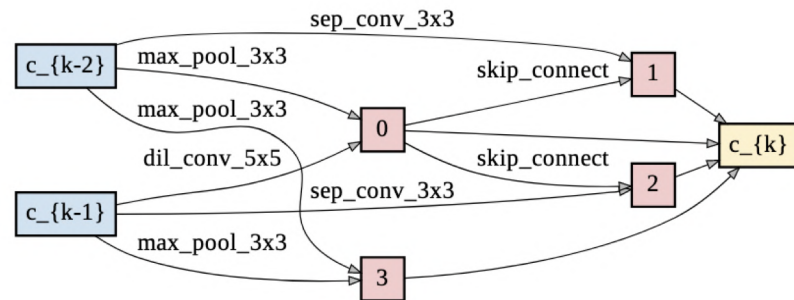
Architecture	Test Err. (%)		Params (M)	$\times +$ (M)	Search Cost (GPU-days)	Search Method
	top-1	top-5				
Inception-v1 [41]	30.2	10.1	6.6	1448	-	manual
MobileNet [16]	29.4	10.5	4.2	569	-	manual
ShuffleNet 2x (v1) [51]	26.4	10.2	$\sim 5$	524	-	manual
ShuffleNet 2x (v2) [32]	25.1	-	$\sim 5$	591	-	manual
NASNet-A [55]	26	8.4	5.3	564	1800	RL
NASNet-B [55]	27.2	8.7	5.3	488	1800	RL
NASNet-C [55]	27.5	9	4.9	558	1800	RL
AmoebaNet-A [36]	25.5	8	5.1	555	3150	evolution
AmoebaNet-B [36]	26	8.5	5.3	555	3150	evolution
AmoebaNet-C [36]	24.3	7.6	6.4	570	3150	evolution
PNAS [27]	25.8	8.1	5.1	588	225	SMBO
MnasNet-92 [42]	25.2	8	4.4	388	-	RL
DARTS (2 <sup>nd</sup> order) [29]	26.7	8.7	4.7	574	4.0	gradient
SNAS (mild) [49]	27.3	9.2	4.3	522	1.5	gradient
ProxylessNAS [7]	24.9	7.5	7.1	465	8.3	gradient
P-DARTS [8]	24.4	7.4	4.9	557	0.3	gradient
BayesNAS [52]	26.5	8.9	3.9	-	0.2	gradient
PC-DARTS [50]	25.1	7.8	5.3	586	0.1	gradient
SGAS (Cri.1 avg.)	24.4 $\pm$ 0.2	7.3 $\pm$ 0.1	5.3	579	0.25	gradient
SGAS (Cri.1 best)	24.2	7.2	5.3	585	0.25	gradient
SGAS (Cri.2 avg.)	24.4 $\pm$ 0.2	7.4 $\pm$ 0.1	5.4	597	0.25	gradient
SGAS (Cri.2 best)	<b>24.1</b>	7.3	5.4	598	0.25	gradient

Table 2. Performance comparison with state-of-the-art image classifiers on ImageNet.

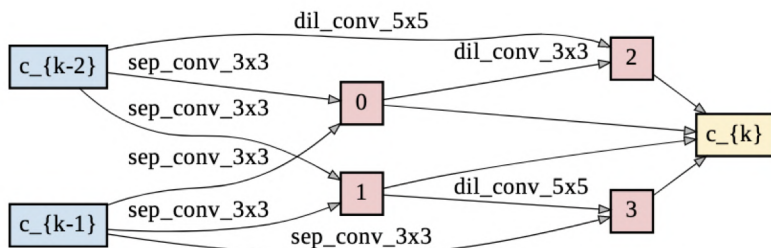
# Results – SGAS for CNN on ImageNet



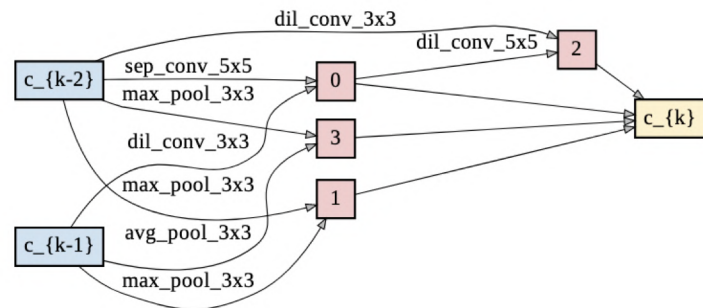
(a) Normal cell of the best model with SGAS (Cri. 1) on ImageNet



(b) Reduction cell of the best model with SGAS (Cri. 1) on ImageNet



(c) Normal cell of the best model with SGAS (Cri. 2) on ImageNet



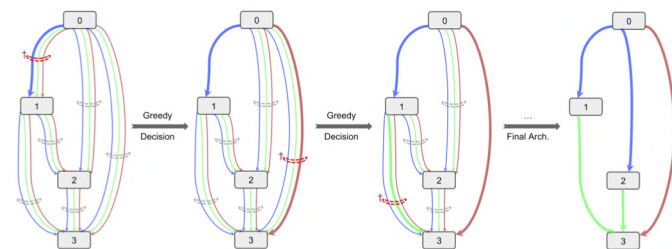
(d) Reduction cell of the best model with SGAS (Cri. 2) on ImageNet



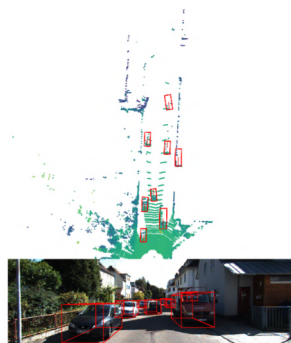
# Follow-up works



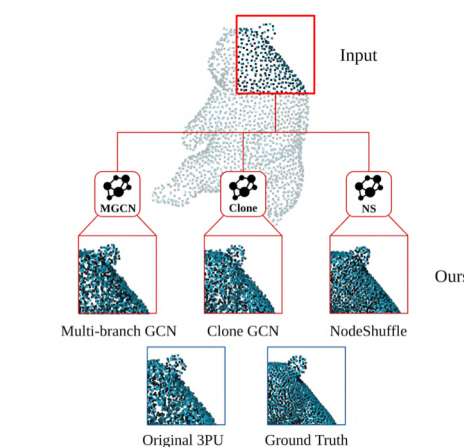
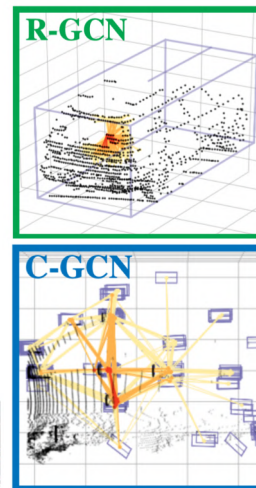
DeepGCNs.org



SGAS: Sequential Greedy Architecture Search. Guohao Li. et al.



PointRGCN: Graph Convolution Networks for 3D Vehicles Detection Refinement. Jesue Zarzar. et al.

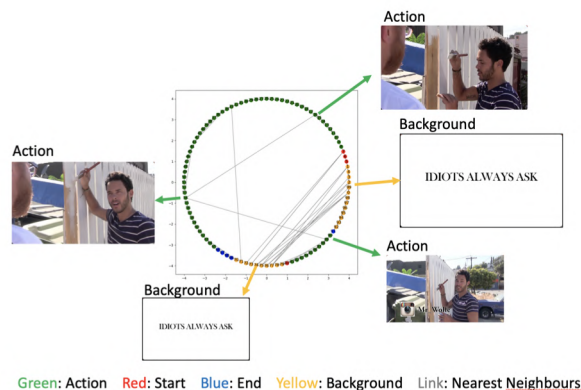


PU-GCN: Point Cloud Upsampling via Graph Convolutional Network. Guocheng Qian. et al.

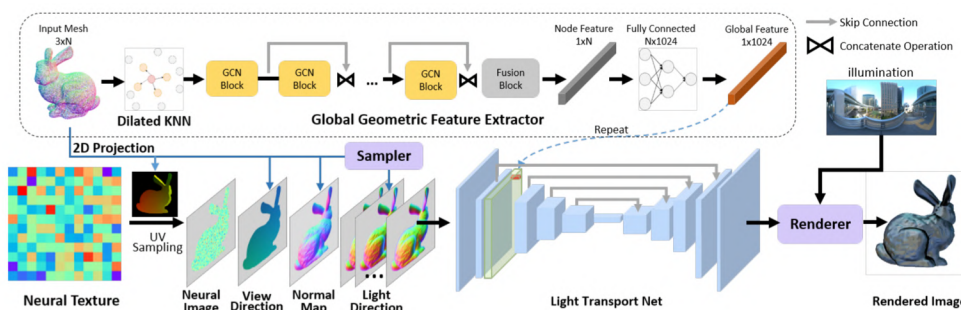
# Follow-up works



DeepGCNs.org

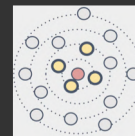


G-TAD: Sub-Graph Localization for Temporal Action Detection. Mengmeng xu. et al.

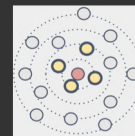


A Neural Rendering Framework for Free-Viewpoint Relighting. Zhang Chen. et al.

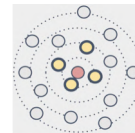




- Thomas Kipf: <http://tkipf.github.io/misc/SlidesCambridge.pdf>
- Stanford SNAP: <http://snap.stanford.edu/proj/embeddings-www/files/nrltutorial-part2-gnns.pdf>



- Thomas Kipf: <http://tkipf.github.io/misc/SlidesCambridge.pdf>
- Stanford SNAP: <http://snap.stanford.edu/proj/embeddings-www/files/nrltutorial-part2-gnns.pdf>
- Pytorch Geometric: <https://pytorch-geometric.readthedocs.io>
- Deep Graph Library: <https://www.dgl.ai/>
- TensorFlow Graphics: <https://github.com/tensorflow/graphics>



THANK YOU

# DeepGCNs for Representation Learning on Graphs

**DeepGCNs: Can GCNs Go as Deep as CNNs? (ICCV'2019 Oral)**

Guohao Li\*, Matthias Müller\*, Ali Thabet, Bernard Ghanem

**DeepGCNs: Making GCNs Go as Deep as CNNs (arXiv'2019)**

Guohao Li\*, Matthias Müller\*, Guocheng Qian, Itzel C. Delgadillo, Abdullellah Abualshour, Ali Thabet, Bernard Ghanem

**SGAS: Sequential Greedy Architecture Search (arXiv'2019)**

Guohao Li\*, Guocheng Qian\*, Itzel C. Delgadillo\*, Matthias Müller, Ali Thabet, Bernard Ghanem



KAUST

Contact: [guohao.li@kaust.edu.sa](mailto:guohao.li@kaust.edu.sa)

About me: <https://ghli.org>

VCC VISUAL  
COMPUTING  
CENTER

IVUL