

Graph Pooling: DIFFPOOL

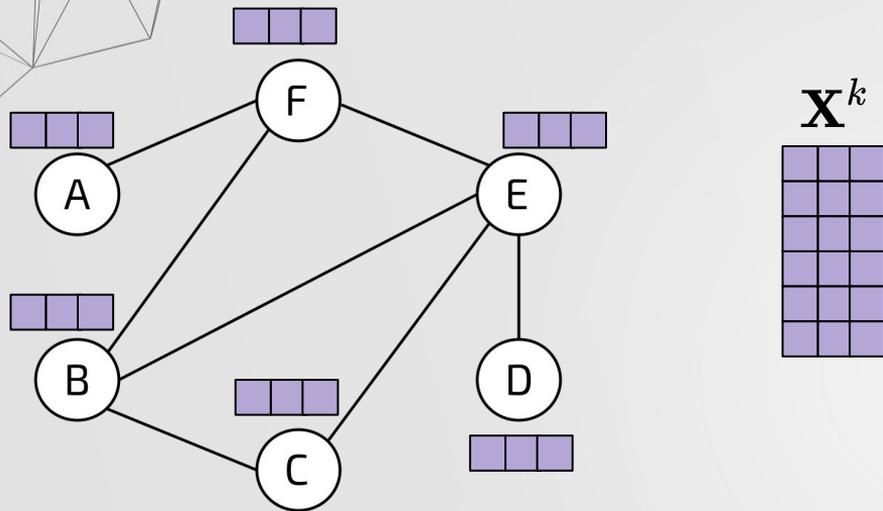
Giovanni Pellegrini^{1,2,3}

SML¹ Lab, University of Trento, Italy

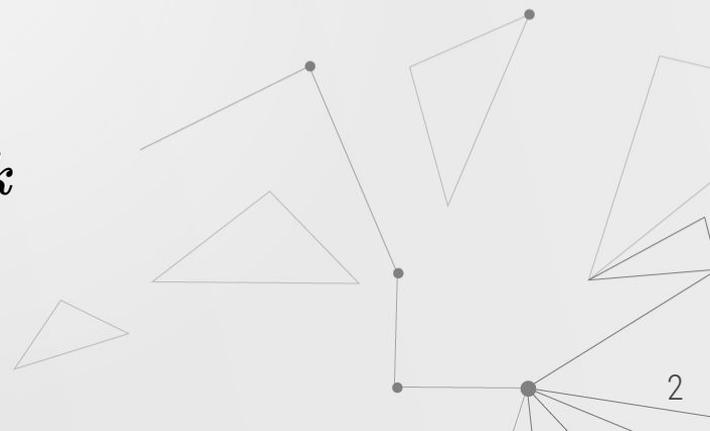
TIM²

EIT DIGITAL³

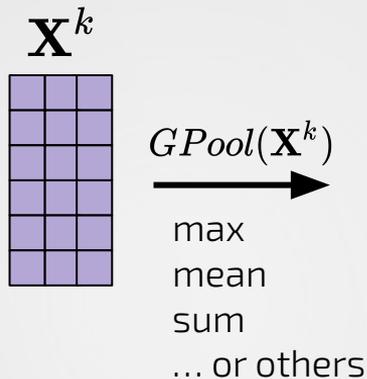
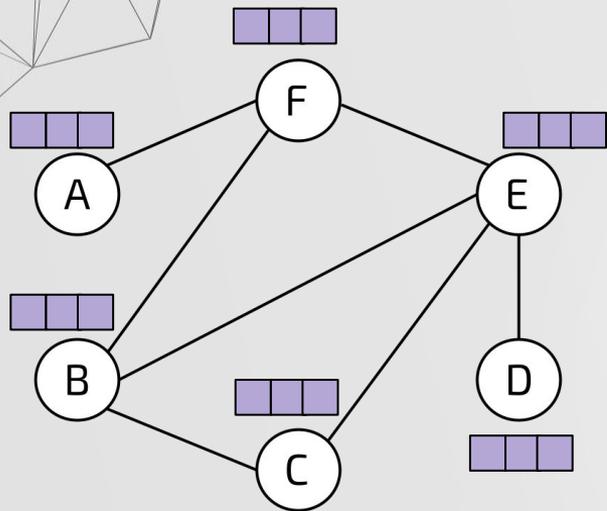
01 Graph prediction



$$\mathbf{X}^{t+1} = \text{GConv}(\mathbf{A}, \mathbf{X}^t, \mathbf{W}^t) \quad t = 1, \dots, k$$

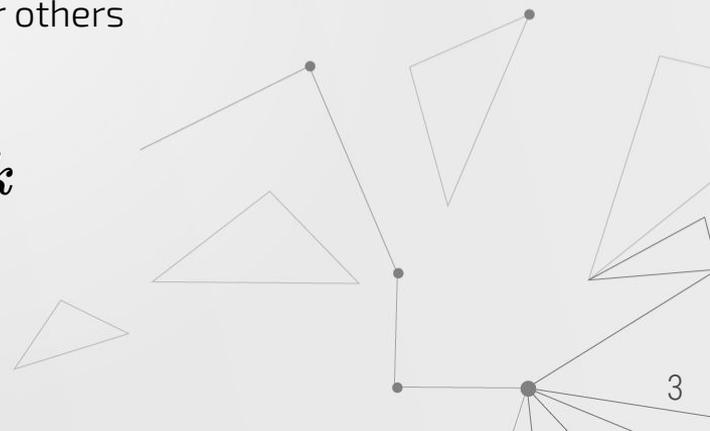


01 Graph prediction

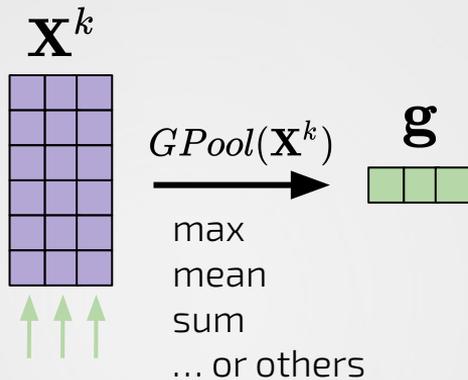
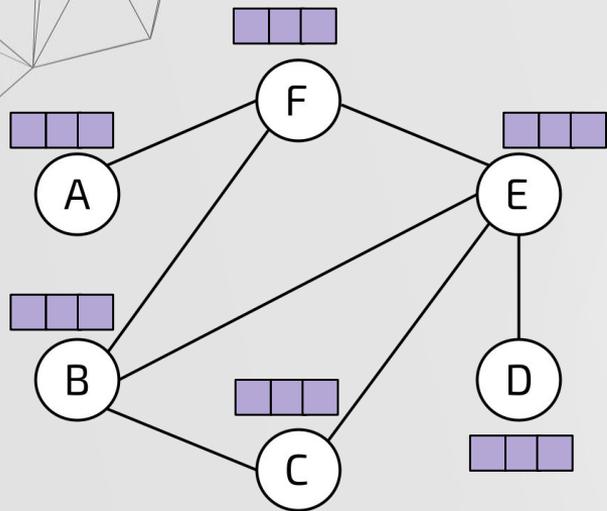


$$\mathbf{X}^{t+1} = GConv(\mathbf{A}, \mathbf{X}^t, \mathbf{W}^t) \quad t = 1, \dots, k$$

$GPool(\mathbf{X}^k) \rightarrow$ global pooling function



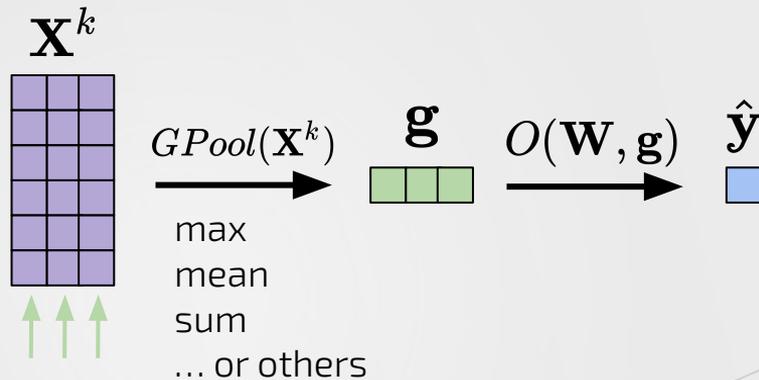
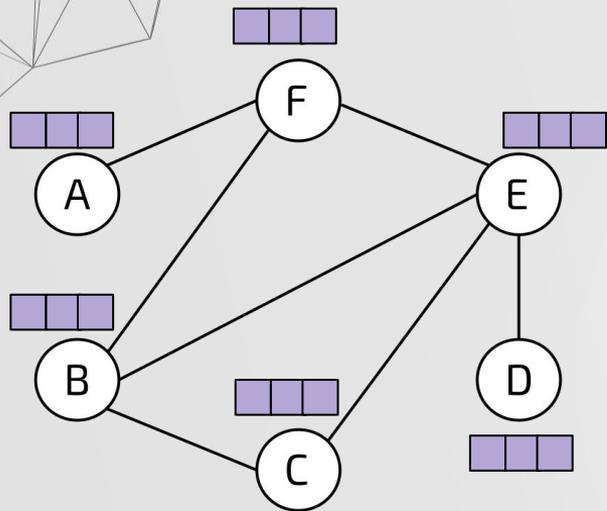
01 Graph prediction



$$\mathbf{X}^{t+1} = GConv(\mathbf{A}, \mathbf{X}^t, \mathbf{W}^t) \quad t = 1, \dots, k$$

$GPool(\mathbf{X}^k) \rightarrow$ global pooling function

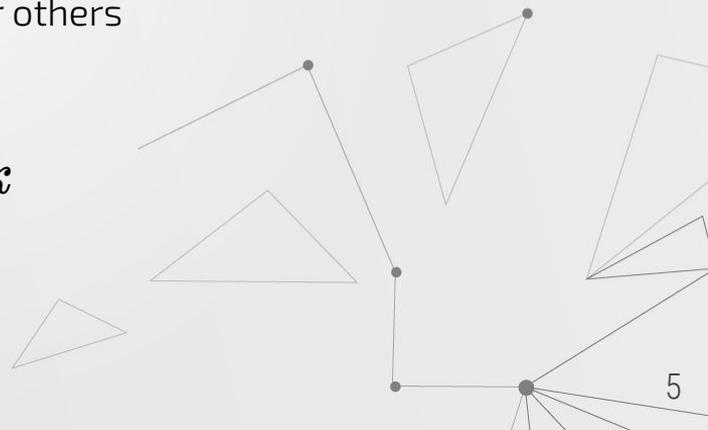
01 Graph prediction



$$\mathbf{X}^{t+1} = \text{GConv}(\mathbf{A}, \mathbf{X}^t, \mathbf{W}^t) \quad t = 1, \dots, k$$

$\text{GPool}(\mathbf{X}^k) \longrightarrow$ global pooling function

$O(\mathbf{W}, \mathbf{g}) \longrightarrow$ graph readout function



01 Graph prediction

Equivalent to having a virtual “supernode” in the last layer connected to all the nodes in the graph

Pooling the node embeddings all together do not capture hierarchical representations of the network



01 Graph prediction

Equivalent to having a virtual “supernode” in the last layer connected to all the nodes in the graph

Pooling the node embeddings all together do not capture hierarchical representations of the network

DIFFPOOL: hierarchical nodes pooling strategy



02 DIFFPOOL¹

DIFFerentiable **POOL**ing: Compute an hierarchical representation of the graph by aggregating “close” nodes

02 DIFFPOOL¹

DIFFerentiable POOLing: Compute an hierarchical representation of the graph by aggregating “close” nodes

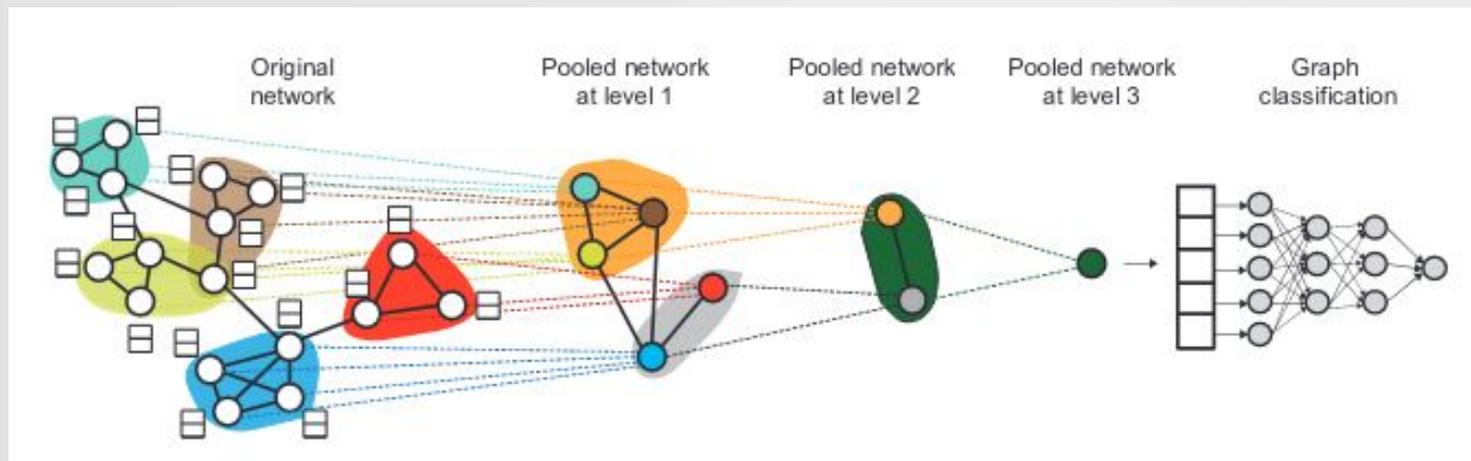
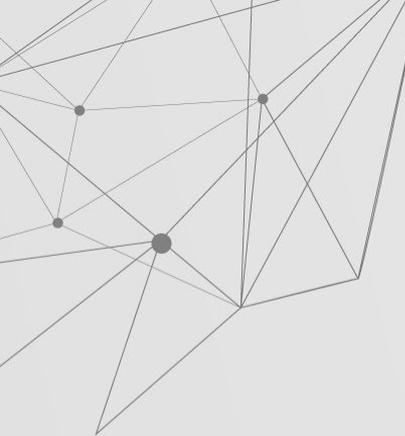


Image taken from the original publication.

Ying et al, *Hierarchical Graph Representation Learning with Differentiable Pooling*, NIPS '18

02 DIFFPOOL

Idea: stack several GNN and pooling layer on top of each other



\mathbf{A}^0

\mathbf{X}^0

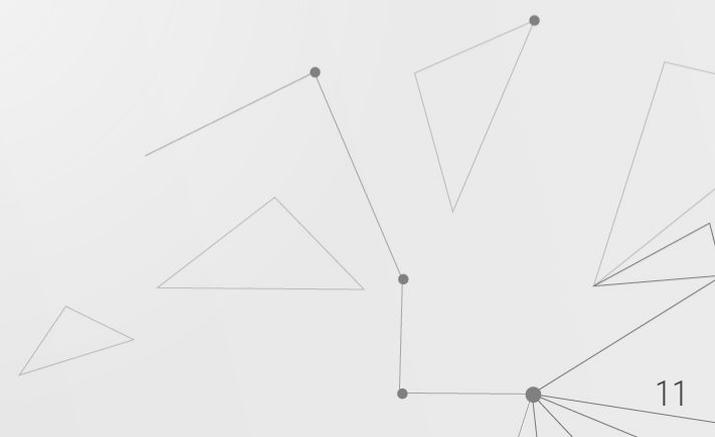


02 DIFFPOOL


$$G = (\mathbf{A}, \mathbf{X})$$

adjacency
matrix

node features
matrix



02 DIFFPOOL

$$G = (\mathbf{A}, \mathbf{X})$$

adjacency
matrix

node features
matrix

$$D = \{(\mathbf{G}_1, y_1), \dots, (\mathbf{G}_n, y_n)\}$$

02 DIFFPOOL


$$G = (\mathbf{A}, \mathbf{X})$$

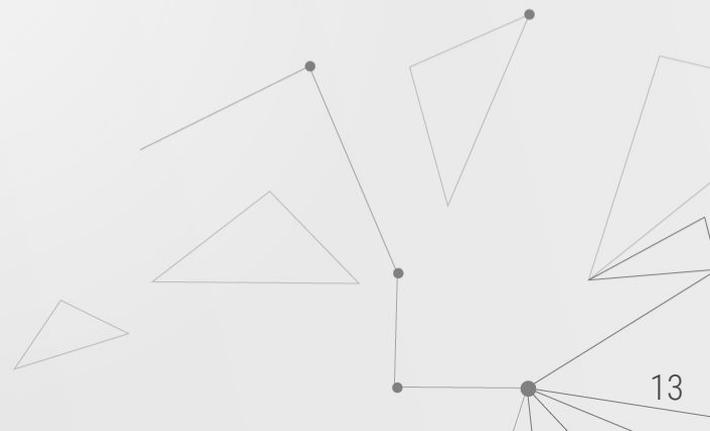
adjacency
matrix

node features
matrix

$$D = \{(\mathbf{G}_1, y_1), \dots, (\mathbf{G}_n, y_n)\}$$

$$\mathbf{X}^{t+1} = \text{GConv}(\mathbf{A}, \mathbf{X}^t, \mathbf{W}^t)$$

$$\mathbf{Z} = \text{GNN}(\mathbf{A}, \mathbf{X})$$



02 DIFFPOOL


$$G = (\mathbf{A}, \mathbf{X})$$

adjacency
matrix

node features
matrix

$$D = \{(\mathbf{G}_1, y_1), \dots, (\mathbf{G}_n, y_n)\}$$

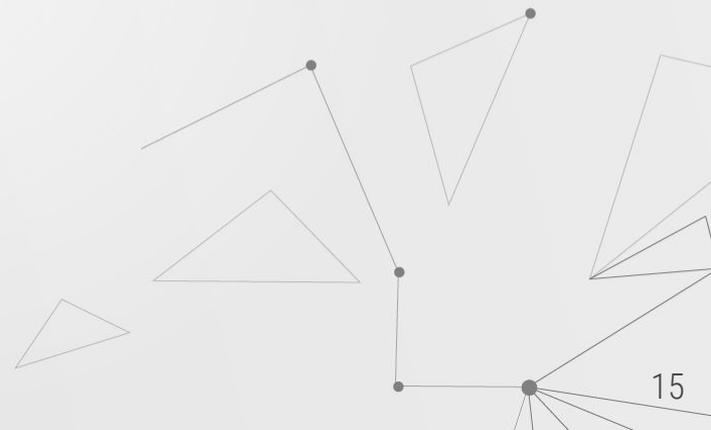
$$\mathbf{X}^{t+1} = \text{GConv}(\mathbf{A}, \mathbf{X}^t, \mathbf{W}^t)$$

$$\mathbf{Z} = \text{GNN}(\mathbf{A}, \mathbf{X})$$

arbitrary GNN that computes K iterations, $\mathbf{Z} = \mathbf{X}^K$

02 DIFFPOOL

Given \mathbf{A} and \mathbf{Z} , find a coarse representation of the graph



02 DIFFPOOL

Given \mathbf{A} and \mathbf{Z} , find a coarse representation of the graph

$$\mathbf{A} \in \mathbb{R}^{n \times n}$$

$$\mathbf{Z} \in \mathbb{R}^{n \times d}$$

02 DIFFPOOL

Given \mathbf{A} and \mathbf{Z} , find a coarse representation of the graph

$$\begin{array}{ccc} \mathbf{A} \in \mathbb{R}^{n \times n} & \longrightarrow & \mathbf{A}' \in \mathbb{R}^{m \times m} \\ \mathbf{Z} \in \mathbb{R}^{n \times d} & & \mathbf{Z}' \in \mathbb{R}^{m \times d} \end{array} \quad \text{with } m < n$$

02 DIFFPOOL

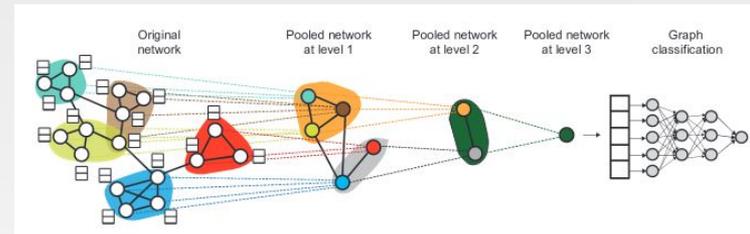
Given \mathbf{A} and \mathbf{Z} , find a coarse representation of the graph

$$\begin{array}{ccc} \mathbf{A} \in \mathbb{R}^{n \times n} & \longrightarrow & \mathbf{A}' \in \mathbb{R}^{m \times m} \\ \mathbf{Z} \in \mathbb{R}^{n \times d} & & \mathbf{Z}' \in \mathbb{R}^{m \times d} \end{array} \quad \text{with } m < n$$

Solution -> learn a cluster assignment for each node

03 Pooling in DIFFPOOL

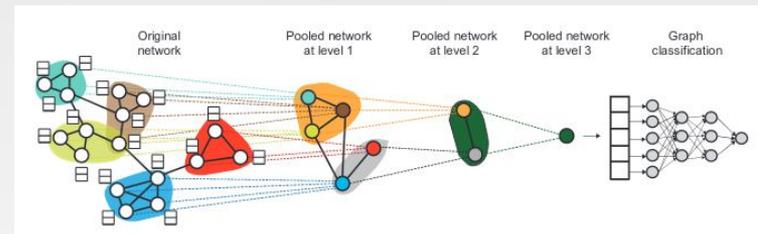
The final representation is obtained by coarsening the graph in L hierarchical steps. To do that, at each step a **cluster assignment matrix** is learned.



03 Pooling in DIFFPOOL

The final representation is obtained by coarsening the graph in L hierarchical steps. To do that, at each step a **cluster assignment matrix** is learned.

$$S^l \in R^{n^l \times n^{l+1}}$$



03 Pooling in DIFFPOOL

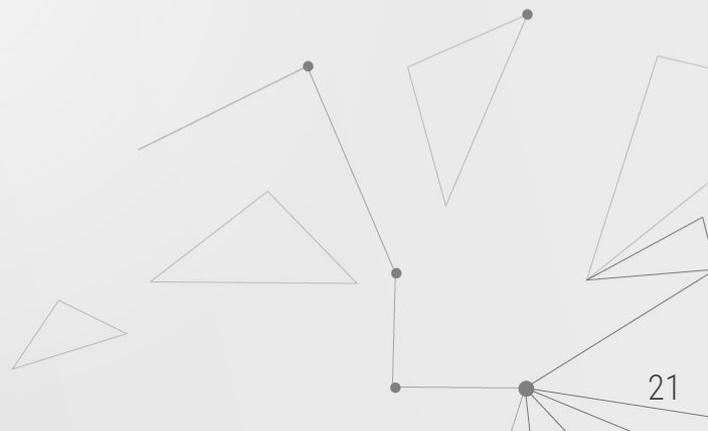
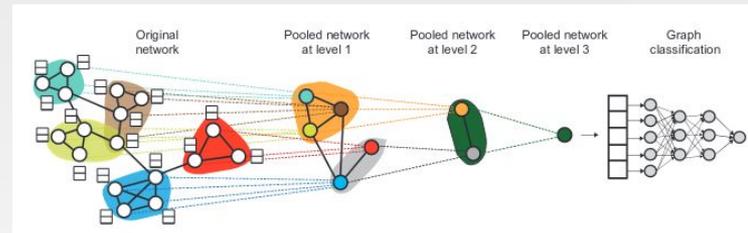
The final representation is obtained by coarsening the graph in L hierarchical steps. To do that, at each step a **cluster assignment matrix** is learned.

$$S^l \in R^{n^l \times n^{l+1}}$$

n^{l+1} = number of cluster at the next step

n^l = number of clusters at the current step

n^0 = number of nodes in the graph



03 Pooling in DIFFPOOL

The final representation is obtained by coarsening the graph in L hierarchical steps. To do that, at each step a **cluster assignment matrix** is learned.

$$\mathbf{S}^l \in \mathbb{R}^{n^l \times n^{l+1}}$$

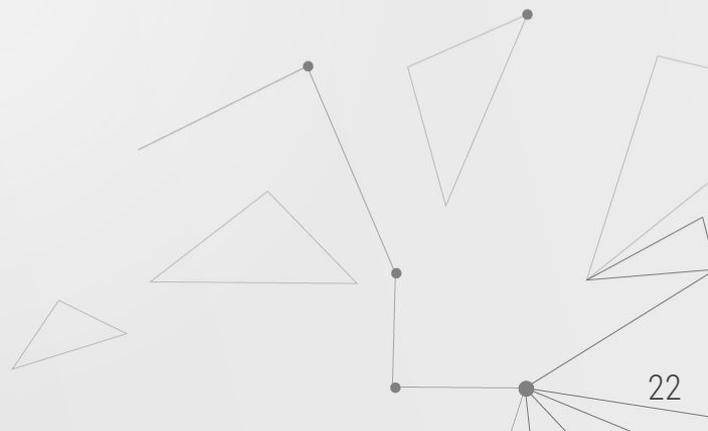
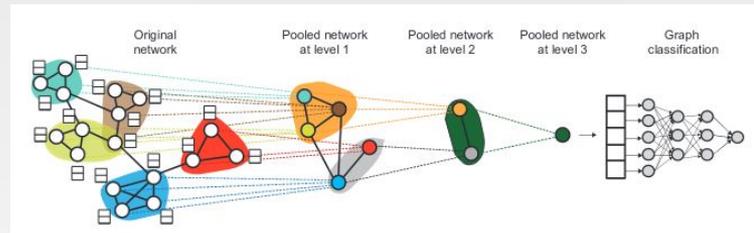
n^{l+1} = number of cluster at the next step

n^l = number of clusters at the current step

n^0 = number of nodes in the graph

$$\mathbf{X}^{l+1} = \mathbf{S}^{l^T} \mathbf{Z}^l \longrightarrow \mathbb{R}^{n^{l+1} \times d}$$

↗ transposed



03 Pooling in DIFFPOOL

The final representation is obtained by coarsening the graph in L hierarchical steps. To do that, at each step a **cluster assignment matrix** is learned.

$$\mathbf{S}^l \in \mathbb{R}^{n^l \times n^{l+1}}$$

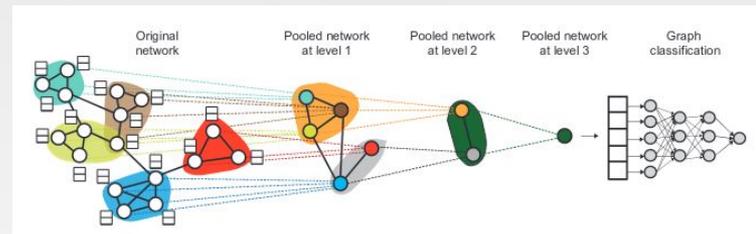
n^{l+1} = number of cluster at the next step

n^l = number of clusters at the current step

n^0 = number of nodes in the graph

$$\mathbf{X}^{l+1} = \mathbf{S}^{lT} \mathbf{Z}^l \longrightarrow \mathbb{R}^{n^{l+1} \times d}$$

$$\mathbf{A}^{l+1} = \mathbf{S}^{lT} \mathbf{A} \mathbf{S}^l \longrightarrow \mathbb{R}^{n^{l+1} \times n^{l+1}}$$



03 Pooling in DIFFPOOL

The final representation is obtained by coarsening the graph in L hierarchical steps. To do that, at each step a **cluster assignment matrix** is learned.

$$\mathbf{S}^l \in \mathbb{R}^{n^l \times n^{l+1}}$$

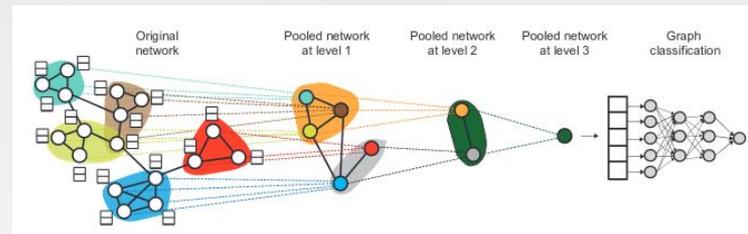
n^{l+1} = number of cluster at the next step

n^l = number of clusters at the current step

n^0 = number of nodes in the graph

$$\mathbf{X}^{l+1} = \mathbf{S}^{l^T} \mathbf{Z}^l \longrightarrow \mathbb{R}^{n^{l+1} \times d}$$

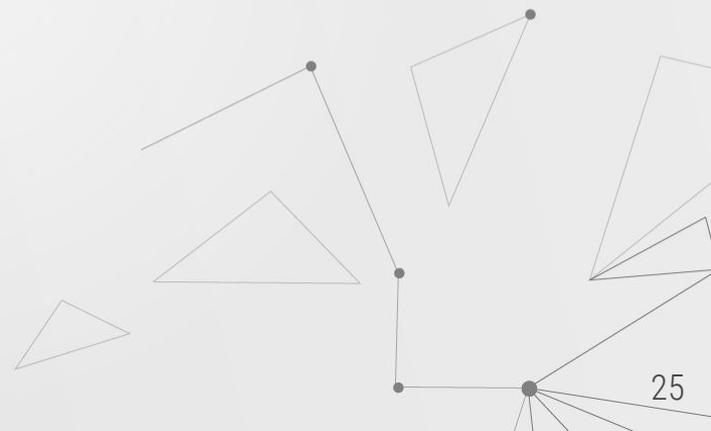
$$\mathbf{A}^{l+1} = \mathbf{S}^{l^T} \mathbf{A} \mathbf{S}^l \longrightarrow \mathbb{R}^{n^{l+1} \times n^{l+1}}$$



The number of clusters at each step is an hyperparameter!

04 Learning the assignment matrix

How is the cluster matrix learned?



04 Learning the assignment matrix

How is the cluster matrix learned?

$$\mathbf{Z}^l = \text{GNN}_{\text{emb}}^l(\mathbf{A}^l, \mathbf{X}^l)$$

04 Learning the assignment matrix

How is the cluster matrix learned?

$$\mathbf{Z}^l = \text{GNN}_{\text{emb}}^l(\mathbf{A}^l, \mathbf{X}^l)$$

$$\mathbf{S}^l = \text{softmax}(\text{GNN}_{\text{pool}}^l(\mathbf{A}^l, \mathbf{X}^l))$$

04 Learning the assignment matrix

How is the cluster matrix learned?

$$\mathbf{Z}^l = \text{GNN}_{\text{emb}}^l(\mathbf{A}^l, \mathbf{X}^l)$$

$$\mathbf{S}^l = \text{softmax}(\text{GNN}_{\text{pool}}^l(\mathbf{A}^l, \mathbf{X}^l))$$

$$\text{GNN}_{\text{pool}}^l \xrightarrow{\text{outputs}} \mathbf{R}^{n^l \times n^{l+1}}$$

Softmax is applied row-wise, it assigns probabilities to which clusters to belong to in the next step.

04 Prediction

How is the cluster matrix learned?

$$\mathbf{Z}^l = \text{GNN}_{\text{emb}}^l(\mathbf{A}^l, \mathbf{X}^l)$$

$$\mathbf{S}^l = \text{softmax}(\text{GNN}_{\text{pool}}^l(\mathbf{A}^l, \mathbf{X}^l))$$

differentiable!

$$\text{GNN}_{\text{pool}}^l \xrightarrow{\text{outputs}} \mathbf{R}^{n^l \times n^{l+1}}$$

Softmax is applied row-wise, it assigns probabilities to which clusters to belong to in the next step.

05 Final prediction

At the last step the graph is condensed in one single node (vector). The final prediction is the output of a MLP.

$$\hat{y} = \text{MLP}(\mathbf{Z}^L)$$

$$E = L(y, \hat{y})$$

05 Final prediction

At the last step the graph is condensed in one single node (vector). The final prediction is the output of a MLP.

$$\hat{y} = \text{MLP}(\mathbf{Z}^L)$$

$$E = L(y, \hat{y})$$

Let's switch to the notebook...