

# Graph Neural Networks

## Introductory lecture

Advanced ML using NN  
spring 2021

# Introduction to Graph Neural Networks

Mats Granath  
24/3 2021

(GNN)

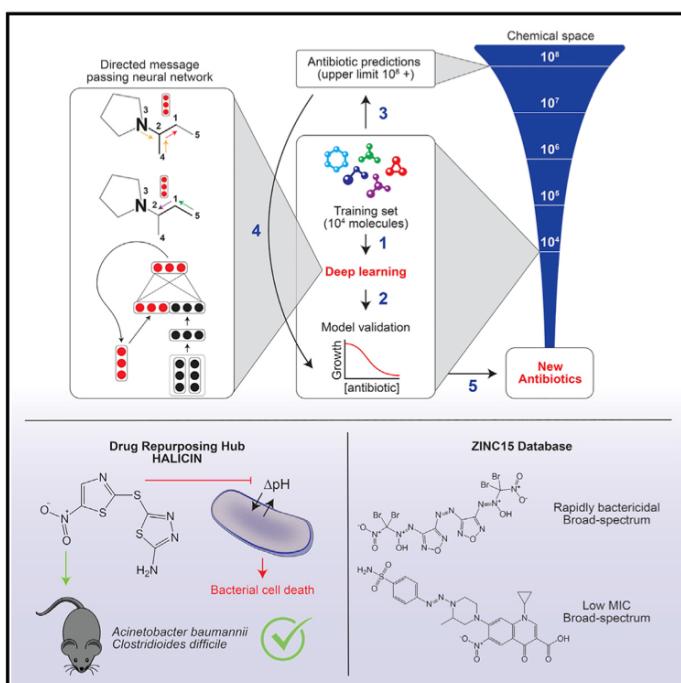
Cell

Recent GNN application

Article

## A Deep Learning Approach to Antibiotic Discovery

### Graphical Abstract



### Authors

Jonathan M. Stokes, Kevin Yang,  
Kyle Swanson, ..., Tommi S. Jaakkola,  
Regina Barzilay, James J. Collins

### Correspondence

regina@csail.mit.edu (R.B.),  
jimjc@mit.edu (J.J.C.)

### In Brief

A trained deep neural network predicts antibiotic activity in molecules that are structurally different from known antibiotics, among which Halicin exhibits efficacy against broad-spectrum bacterial infections in mice.

### Highlights

- A deep learning model is trained to predict antibiotics based on structure
- Halicin is predicted as an antibacterial molecule from the Drug Repurposing Hub
- Halicin shows broad-spectrum antibiotic activities in mice
- More antibiotics with distinct structures are predicted from the ZINC15 database

Stokes et al., 2020, Cell 180, 688–702  
February 20, 2020 © 2020 Elsevier Inc.  
<https://doi.org/10.1016/j.cell.2020.01.021>

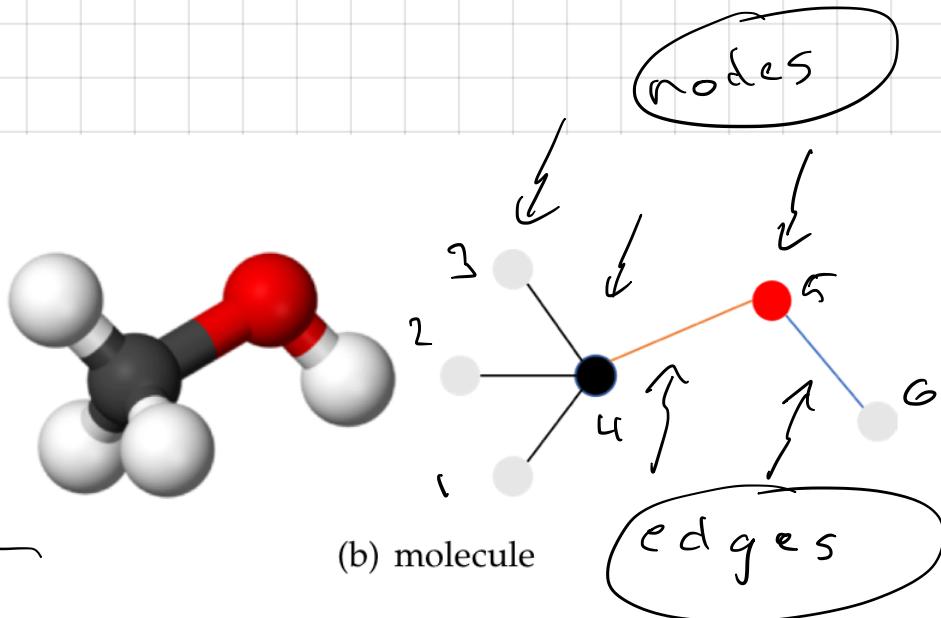
### Outline

- ~ what is a graph, how describe
- ~ Graph learning tasks
- ~ GNN vs. CNN
- ~ GNN layers
- ~ Homework

• Watch the YouTube videos for more general discussion!

What is a graph?

It's a collection of nodes and edges  
 $i \in \{1, \dots, n\}$        $\{e_{i \rightarrow j}\}$



$n = 6$

Structure given by Adjacency matrix

$$A_{ij} = 1 \text{ if edge } i \rightarrow j$$

$$A_{ij} = 0 \text{ if no edge}$$

We will consider undirected graphs  $A_{ji} = A_{ij}$

$$A = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

↙ Symmetric

But there will  
be more  
information  
than the  
structure

Graph Neural Networks:  
A Review of Methods and Applications

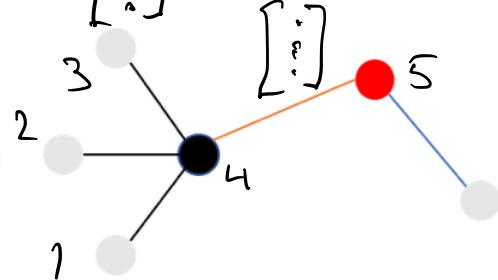
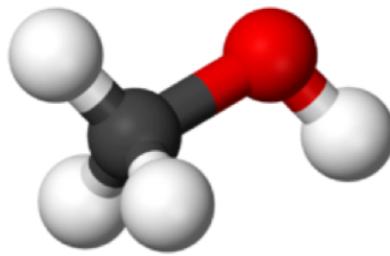
Feature vectors on nodes and edges

$$\vec{e}_{ij}$$

feature vector

$$\begin{bmatrix} ? \\ ? \\ ? \\ ? \\ ? \end{bmatrix}$$

$$\vec{x}_i$$

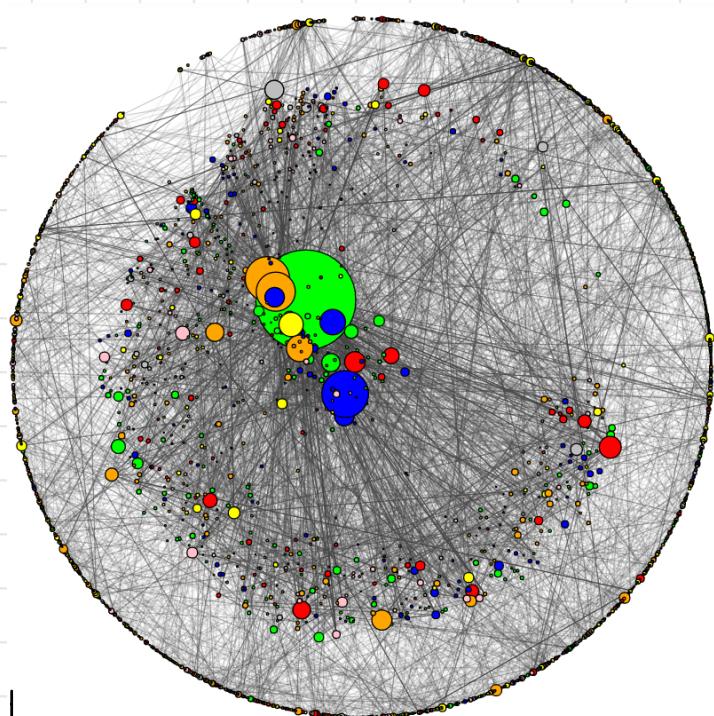


(b) molecule

Ex.  $e_{45}$  = distance 4 to 5 (scalar, not vector)

$x_3$  = label of type of molecule

Example Cora dataset



Citation network of papers.  
nodes

2708 nodes

1433 dim.

feature  
vector

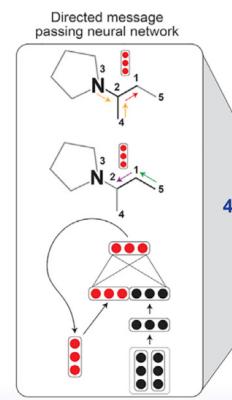
Based on  
keywords

7 classes  
of papers

# Machine learning tasks on graphs

## Graph classification

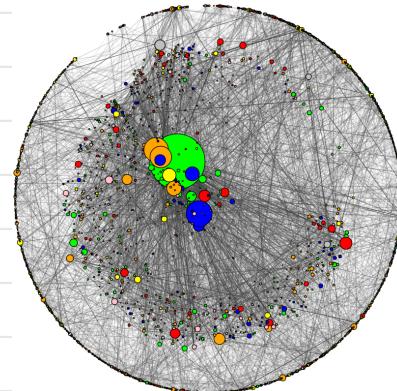
E.g. Does the molecule have antibiotic properties or not  $\rightarrow$



many graphs

## Node classification

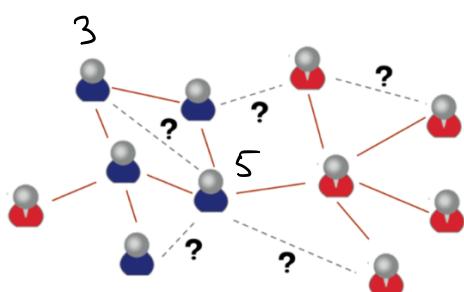
E.g. is the paper in class A,B,C, etc.?



single graph

## Link prediction

E.g. Is there a link (edge) between node 3 and 5?



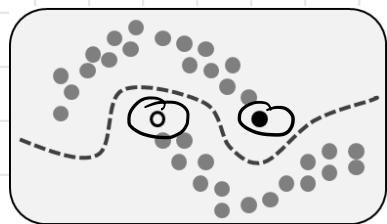
## Supervised learning:

Graph classification is typically supervised:  
set of labeled graphs, use for training.

## Semi-supervised learning:

Ex: two labeled data points,

many unlabeled  
use clustering to extract info. from unlabeled data points.



Node prediction is typically of this type.

Ex. Cora 140 labeled nodes  
out of 2708 nodes

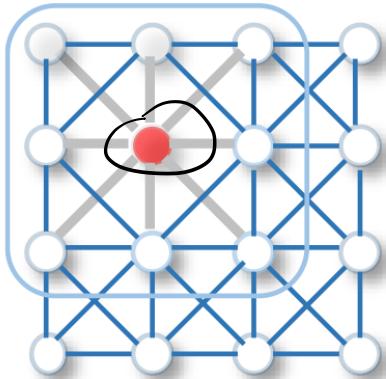
Use info about how papers are related (citation graph)

E.g. homophilous graph: adjacent nodes more likely to share properties.  
Unsupervised:

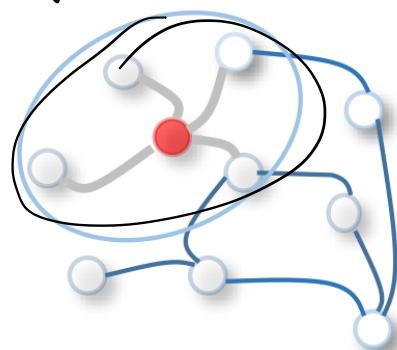
GNN do clustering even when not trained.

# Convolutions on graphs

CNN filter



GNN filter



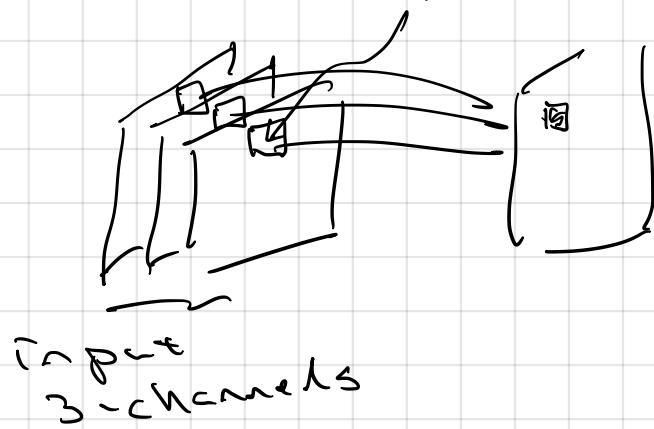
A Comprehensive Survey on Graph Neural Networks

Zonghan Wu, Shirui Pan, Member, IEEE, Fengwen Chen, Guodong Long,  
Chengqi Zhang, Senior Member, IEEE, Philip S. Yu, Fellow, IEEE

Standard CNN work on grids, regular structure: each "node" has fixed number of neighbors

filters

has fixed number of neighbors



different filters; weights to all pixels

3x3 filter has 9 trainable weights

How to set up similar object on graph?

Clearly: We want to use adjacency matrix  $A_{ij}$

Problem: not fixed number of adjacent nodes, how many weights?

# Graph convolutional layers, some examples

## Basic graph convolution

add "self-loops" to  $\hat{A}$

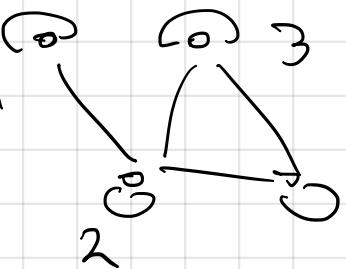
$$\tilde{A} = A + \mathbb{I}$$

$$\tilde{A} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & n & \dots & \dots \end{pmatrix}$$

its degree matrix

$$\tilde{D}: \text{diagonal } \tilde{D}_{ii} = \sum_j \tilde{A}_{ij} = \tilde{d}_i$$

$$D_{ij} = 0 \quad i \neq j$$



- Convolution takes a graph to an isomorphic graph, same structure, i.e. same  $A_{ij}$ .

Gives new node feature vectors

$$\vec{x}_0 \rightarrow \vec{x}'_0$$

dimension  $d'$

make feature matrix  $n \times d'$

$$\vec{x}' = \sigma \left( \underbrace{\tilde{D}^{-1/2} \tilde{A} \tilde{D}^{-1/2}}_{n \times n} \vec{x} \underbrace{W}_{d' \times d} \right)$$

$\vec{x}$  weight matrix  $d \times d'$

non-linearity  
ReLU, tanh.

Same weights on all nodes!

equivalently

$$\vec{x}_i^1 = \sigma \left( \frac{\sum_j \tilde{A}_{ij} \vec{x}_j}{\sqrt{\deg_i}} \right) \vec{w}^T \vec{x}_i$$

linear transform to  $d'$  dim vector

just like a single dense NN layer

collect from adjacent neighbors

(collects information isotropically from neighbors.  
"one-hop" per layer

Warning! Too many layers may smear out information over the graph.

No edge info used except for A

- Simple variation that uses an edge weight  $c_{ij}$  (scalar edge feature)

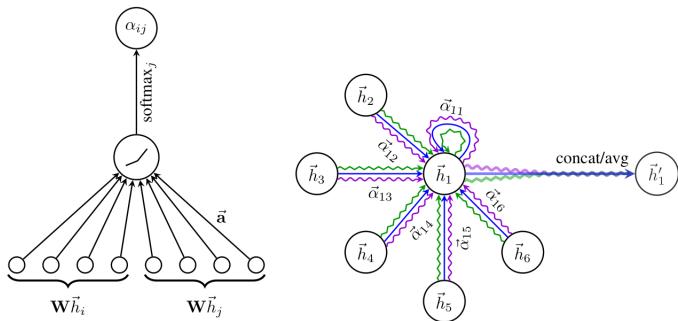
$$\vec{x}_i^1 = \sigma \left( \vec{w}_1 \vec{x}_i + \vec{w}_2 \left\{ \begin{array}{l} c_{ij} \vec{x}_j \\ \uparrow \end{array} \right\} \right)$$

"heavy" edges weighed more

# Graph Attention layers

The basic layer is isotropic, all adjacent nodes are treated the same.

We may want to be more selective.



## GRAPH ATTENTION NETWORKS

Petar Veličković\*  
Department of Computer Science and Technology  
University of Cambridge  
petar.velickovic@cst.cam.ac.uk

Guillem Cucurull\*  
Centre de Visió per Computador, UAB  
gcucurull@gmail.com

Arantxa Casanova\*  
Centre de Visió per Computador, UAB  
ar.casanova.8@gmail.com

Adriana Romero  
Montréal Institute for Learning Algorithms  
adriana.romero.soriano@umontreal.ca

Pietro Liò  
Department of Computer Science and Technology  
University of Cambridge  
pietro.liò@cst.cam.ac.uk

Yoshua Bengio  
Montréal Institute for Learning Algorithms  
yoshua.umontreal@gmail.com

Attention gives selective feedback from adjacent nodes.

- $\vec{x}'_i = \odot \left( \sum_j \alpha_{ij} \vec{w} \vec{x}_j \right)$

$\vec{x}'_i$        $\vec{x}_j$        $\vec{w}$        $d'$   
 $d'$        $d'$        $d'$        $d$   
 attention coefficients       $d' \times d$  weight matrix

- $\alpha_{ij} = \text{softmax}(\varepsilon_{ij}) = \frac{e^{\varepsilon_{ij}}}{\sum_k e^{\varepsilon_{ik}}}$

- $\varepsilon_{ij} = \vec{a}^\top \left( \vec{w}_{\vec{x}_i} || \vec{w}_{\vec{x}_j} \right) \hat{A}_{ij}$

$\vec{a}$        $\vec{w}_{\vec{x}_i}, \vec{w}_{\vec{x}_j}$        $\hat{A}_{ij}$   
 attention vector      concatenation  
 $d'$        $d' \times d'$        $K$   
 dim.  $2d'$        $d'$        $d'$   
 trainable weights       $d' \times \begin{bmatrix} \vec{w}_{\vec{x}_i} \\ \vec{w}_{\vec{x}_j} \end{bmatrix}$        $2d'$

- One can also have multiple attention heads  $w^k, q^k, k=1..K$

# Pooling layers

- Convolutional layers: graph  $\rightarrow$  some graph  
 (new features)
- How do we get output for graph classification?  $\leftarrow$  single label
  - We may want to decrease or unify graph size.

Use pooling layers: Pool nodes together

- Global pooling  $\leftrightarrow$  graph  $\rightarrow$  1 node

E.g. global mean pool

$$\vec{x}^1 = \frac{1}{n} \sum_{i=0}^n \vec{x}_i$$

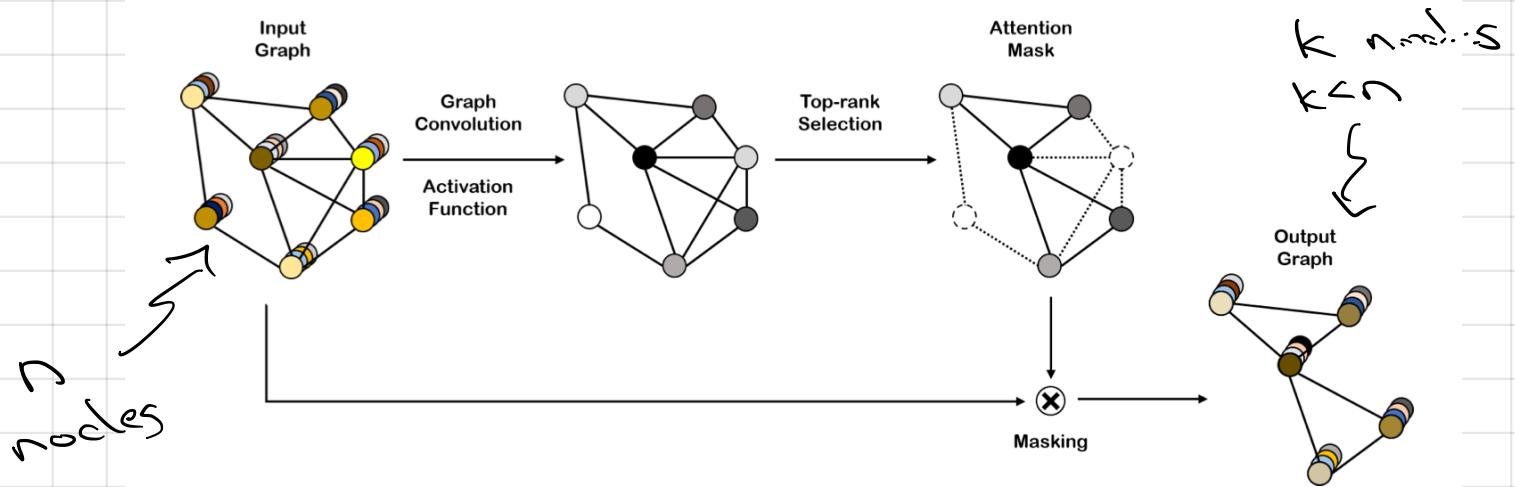
$\xrightarrow{\text{MLP}} \xrightarrow{\text{classifier}}$

single feature vector

Example

## Self-Attention Graph Pooling

Junhyun Lee \*<sup>1</sup> Inyeop Lee \*<sup>1</sup> Jaewoo Kang<sup>1</sup>



Select the  $k$  most important nodes

Self-attention  $\mathbb{Z}_i = \sigma \left( \sum_j D_{ij} \vec{A}_{ij} \frac{\vec{X}_j}{\sqrt{\sum_j D_{ij}}} \right)$

scalar value of each node

$\vec{A}$  attention vector

Keep  $k$  nodes with highest  $\mathbb{Z}$ -value

$$\vec{D}_{ij} \rightarrow \vec{A}_{\text{mask}(ij)}$$

$n \times n$

$k \times k$

$$\vec{X}'_i = \vec{X}_i \vec{Z}_i$$

Bottom line: Put together conv. layers + pooling layers depending on the task.

Beware of overdoing.

# Homework A

- Walk through

## Graph Neural Networks, Assignment A

**Course:** Advanced machine learning using neural networks, TIF360/FYM360

**Contact:** Mats Granath, mats.granath@physics.gu.se

David Fitzek, davidfi@chalmers.se

- Recommended GNN library for the h.w.

The screenshot shows a web browser window with the URL [pytorch-geometric.readthedocs.io](https://pytorch-geometric.readthedocs.io). The page displays the PyTorch Geometric Documentation. The left sidebar has a dark background with white text, listing sections like 'NOTES' (Installation, Introduction by Example, Creating Message Passing Networks, Creating Your Own Datasets, Advanced Mini-Batching, Memory-Efficient Aggregations, TorchScript Support, Colab Notebooks, External Resources) and 'PACKAGE REFERENCE' (torch\_geometric, torch\_geometric.nn, torch\_geometric.data, torch\_geometric.datasets, torch\_geometric.transforms, torch\_geometric.utils, torch\_geometric.io). The main content area has a light background with black text, featuring the title 'PYTORCH GEOMETRIC DOCUMENTATION'. It describes PyTorch Geometric as a geometric deep learning extension library for PyTorch, consisting of various methods for deep learning on graphs and other irregular structures. Below this, there are sections for 'Notes' (a bulleted list of the same items as the sidebar), 'Package Reference' (a bulleted list of package names), and 'Indices and Tables' (a bulleted list of 'Index' and 'Module Index'). At the bottom, there is copyright information ('© Copyright 2021, Matthias Fey Revision c711bb22.'), a note about the build ('Built with Sphinx using a theme provided by Read the Docs.'), and a small logo of a stylized brain or atom.

