Graph Neural Networks

Analysis of complex interconnected data

Comp 599: Network Science, Fall 2022



Neural Networks - Short Intro - Linear Model

J

• Linear regression: $f(x) = w^{\top} x = \sum_{d} w_d x_d$ $x = [1, x_1, \dots, x_D]^{\top}$ $w = [w_0, w_1, \dots, w_D]^{\top}$

Model: linear combination of features and weights Learning: find the weights that minimize a cost function Cost: sum of losses per individual point



$$X = \begin{bmatrix} x^{(1)^{\top}} \\ x^{(2)^{\top}} \\ \vdots \\ x^{(N)^{\top}} \end{bmatrix} = \begin{bmatrix} x_1^{(1)}, & x_2^{(1)}, & \cdots, & x_D^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{(N)}, & x_2^{(N)}, & \cdots, & x_D^{(N)} \end{bmatrix} \text{ one instance} \in \mathbb{R}^{N \times D}$$
one feature

$$V(w) = rac{1}{2} \sum_{n=1}^{N} \left(y^{(n)} - w^{ op} x^{(n)}
ight)^2$$

 $w^* = rgmin_w J(w)$

Neural Networks - Short Intro - Beyond Linear

- Linear regression: assume input of one dimension
- More expressive: use nonlinear bases:
 - Transform the input with nonlinearities then apply linear model

Example: blue dots are the datapoints, generated from blue ground truth curve with some random noise. We can get a very close nonlinear fit, green line, with this linear regression model when using 10 nonlinear Gaussian bases, equally spaced



curve-fitting using nonlinear Gaussian bases

$$f(x;w) = \sum_{m} w_m \phi_m(x)$$

f(x; w) = wx

Neural Networks - Short Intro - Adaptive Bases

- Linear regression: assume input of one dimension
- More expressive: use nonlinear bases:
 - Transform the input with fixed nonlinearities then apply linear model
- Neural networks use **adaptive** nonlinear bases
 - Transform the input with learnable nonlinearities then apply linear model



$$J(x;w) = wx$$

C/

$$f(x;w) = \sum_{m} w_m \phi_m(x)$$

$$f(x; w, \mu) = \sum_{m} w_{m} \phi_{m}(x; \mu_{m})$$

Neural Networks - Short Intro - Adaptive Bases

- Linear regression: assume input of one dimension
- More expressive: use nonlinear bases:
 - Transform the input with fixed nonlinearities then apply linear model
- Neural networks use **adaptive** nonlinear bases
 - Transform the input with learnable nonlinearities then apply linear model





$$f(x;w) = wx$$

C/

$$f(x;w) = \sum_{m} w_m \phi_m(x)$$

$$f(x; w, \mu) = \sum_{m} w_{m} \phi_{m}(x; \mu_{m})$$

° (_____

Neural Networks - Short Intro - MLP

- Neural networks use **adaptive** nonlinear bases
 - Learning the (weights of) nonlinear bases

$$\hat{y}_c = g\left(\sum_m W_{c,m}hig(\sum_d V_{m,d}x_dig)
ight)$$

nonlinearity, activation function: we have different choices

```
more compressed form
```

 $\hat{y} = f(x; W, V) = g(Wh(Vx))$

• The most common non-linearity leaky ReLU $h(x) = \max(0, x) + \gamma \min(0, x)$





Neural Networks - Short Intro - Deep Networks

- Neural networks use **adaptive** nonlinear bases •
 - Example: 2 layers MLP 0
- **Deep** networks: neural networks with many layers
 - Stack/compose layers of adaptive nonlinear bases Ο
 - # layers is called the depth of the network Ο
 - each layer can be fully connected (dense) or sparse Ο
 - empirically, increasing the depth is often more effective than Ο increasing the width (#parameters per layer) Not the case for GNNs
 - parameters may be shared across units (e.g., in conv-nets) Ο





output of one layer is input to the next



Deep Networks for Graphs?

Can we feed an adjacency matrix to this?

Not the best choice: Order $\mathcal{O}(n)$ parameters, also Sensitive to order of nodes

What if it is for Graph Classification, flatten the matrix into a vector?

$$z = MLP(A[1] \bigoplus A[2] \bigoplus \dots A[n])$$

concatenation

What is the size of D? Input dimension? n^2

Still same issues, depends on the arbitrary ordering of nodes that we used in the adjacency matrix



Permutation Invariance & Equivariance

function f that takes an adjacency matrix A as input should be:

• Permutation Invariance

$$f(PAP^{\top}) = f(A)$$

or

• Permutation Equivariance

 $f(PAP^{\top}) = Pf(A)$

where P is a permutation matrix that reorders nodes



Since changing order of nodes in the adjacency does not change the graph

Graph Neural Networks

• GNNs consist of multiple permutation equivariant / invariant functions



<u>ک</u>

<u>e</u>

Graph Neural Networks

Use the local neighbourhood similar to convolution on images: use local neighbourhood similar to CNN for images





GNN model uses a message passing framework



See <u>https://petar-v.com/talks/GNN-Wednesday.pdf</u>



Graph Neural Networks: Neural Message Passing

vector messages are exchanged between nodes and updated using neural networks



Graph Neural Networks: Neural Message Passing

vector messages are exchanged between nodes and updated using neural networks



UPDATE and AGGREGATE are arbitrary differentiable functions

AGGREGATE function takes a set as input, GNNs defined in this way are permutation equivariant by design.

initial embeddings at k = 0 are set to the input features for all the nodes, if no feature can be identity: a one-hot indicator feature, which uniquely identifies that node

Basic GNN

vector messages are exchanged between nodes and updated using neural networks TARGET NODE

a hidden embedding
$$\mathbf{h}_{u}^{(k)}$$

$$\mathbf{h}_{u}^{(k+1)} = \text{UPDATE}^{(k)} \left(\mathbf{h}_{u}^{(k)}, \text{AGGREGATE}^{(k)}(\{\mathbf{h}_{v}^{(k)}, \forall v \in \mathcal{N}(u)\}) \right)$$

$$= \text{UPDATE}^{(k)} \left(\mathbf{h}_{u}^{(k)}, \mathbf{m}_{\mathcal{N}(u)}^{(k)} \right),$$

$$\mathbf{m}_{\mathcal{N}(u)} = \sum_{v \in \mathcal{N}(u)} \mathbf{h}_{v},$$

$$\text{UPDATE}(\mathbf{h}_{u}, \mathbf{m}_{\mathcal{N}(u)}) = \sigma \left(\mathbf{W}_{\text{self}} \mathbf{h}_{u} + \mathbf{W}_{\text{neigh}} \sum_{v \in \mathcal{N}(u)} \mathbf{h}_{v}^{(k-1)} + \mathbf{b}^{(k)} \right)$$

$$\text{trainable parameter matrices In matrix form:}$$

$$\mathbf{H}^{(t)} = \sigma \left(\mathbf{AH}^{(k-1)} \mathbf{W}_{\text{neigh}}^{(k)} + \mathbf{H}^{(k-1)} \mathbf{W}_{\text{self}}^{(k)} \right)$$

Intuition

If we have:

$$A_{ij} = \begin{cases} 1, & i \text{ links to } j \\ 0, & \text{otherwise} \end{cases} \quad X_{ik} = \begin{cases} 1, & i \text{ has } k \\ 0, & \text{otherwise} \end{cases}$$

Then simple matrix multiplication of A and X, AX, gives us the number of neighbors of a particular attribute / type for each node, i.e.

- kth column of AX shows the number of type k neighbors for all nodes,
 - e.g., number of 'male' friends each person has.
- ith row of AX shows the number of neighbors node i for all types,
 - e.g., number of friends 'smith' has of each type, say male and female

For graph neural networks, after the first iteration (k = 1), every node embedding contains a information from its 1-hop neighbourhood

$$\mathbf{H}^{(t)} = \sigma \left(\mathbf{A} \mathbf{H}^{(k-1)} \mathbf{W}_{\text{neigh}}^{(k)} + \mathbf{H}^{(k-1)} \mathbf{W}_{\text{self}}^{(k)} \right)$$

GCN (Kipf & Welling, ICLR'17)





 $H^{l+1} = \phi(AH^l W^l)$

Multi-layer Graph Convolutional Network (GCN) with first-order filters.

From <u>https://petar-v.com/talks/GNN-</u> <u>Wednesday.pdf</u> https://www.youtube.com/watch?v=uF53xsT7mjc

30

° (

Hidden laver



GCN (Kipf & Welling, ICLR'17): has self-loop and symmetric normalization

$$\mathbf{m}_{\mathcal{N}(u)} = \sum_{v \in \mathcal{N}(u)} \frac{\mathbf{h}_v}{\sqrt{|\mathcal{N}(u)||\mathcal{N}(v)|}}$$



Multi-layer Graph Convolutional Network (GCN) with first-order filters.

ReLU

Hidden layer

From <u>https://petar-v.com/talks/GNN-Wednesday.pdf</u> <u>https://www.youtube.com/watch?v=uF53xsT7mjc</u>

Input

GCN (Kipf & Welling, ICLR'17)



Multi-layer Graph Convolutional Network (GCN) with first-order filters.



 $H^{l+1} = \phi(\hat{D}^{-\frac{1}{2}}\hat{A}\hat{D}^{-\frac{1}{2}}H^{l}W^{l})$ $h_i^{l+1} = \phi(\sum_{j \in \mathcal{N}(i)} \frac{1}{c_{ij}} h_j^l W^l)$

From <u>https://petar-v.com/talks/GNN-</u> <u>Wednesday.pdf</u> <u>https://www.youtube.com/watch?v=uF53xsT7mjc</u>

GCN (Kipf & Welling, ICLR'17)







3-layer with random weights (untrained!)

From <u>https://petar-v.com/talks/GNN-</u> <u>Wednesday.pdf</u> <u>https://www.youtube.com/watch?v=uF53xsT7mjc</u>

° (* 1997)

GCN (Kipf & Welling, ICLR'17)





Iteration $0 \Rightarrow$ Gets better as we learn the weights

From https://arxiv.org/pdf/1609.02907.pdf

° (201

GCN (Kipf & Welling, ICLR'17)





More layers do not help due to over smoothing

From https://arxiv.org/pdf/1609.02907.pdf

මා

کی ا

Attentional GNN

GAT (Veličković et al., ICLR'18)

compute scalar value in each edge







Transductive			
Method	Cora	Citeseer	Pubmed
MLP	55.1%	46.5%	71.4%
ManiReg (Belkin et al., 2006)	59.5%	60.1%	70.7%
SemiEmb (Weston et al., 2012)	59.0%	59.6%	71.7%
LP (Zhu et al., 2003)	68.0%	45.3%	63.0%
DeepWalk (Perozzi et al., 2014)	67.2%	43.2%	65.3%
ICA (Lu & Getoor, 2003)	75.1%	69.1%	73.9%
Planetoid (Yang et al., 2016)	75.7%	64.7%	77.2%
Chebyshev (Defferrard et al., 2016)	81.2%	69.8%	74.4%
GCN (Kipf & Welling, 2017)	81.5%	70.3%	79.0%
MoNet (Monti et al., 2016)	$81.7\pm0.5\%$	—	$78.8\pm0.3\%$
GCN-64*	$81.4\pm0.5\%$	$70.9\pm0.5\%$	$\textbf{79.0} \pm 0.3\%$
GAT (ours)	$\textbf{83.0}\pm0.7\%$	$\textbf{72.5} \pm 0.7\%$	$\textbf{79.0} \pm 0.3\%$

From <u>https://petar-v.com/talks/GNN-</u> <u>Wednesday.pdf</u>

° (* 1



Overview: Link Prediction

Goal: Learn f(G, v) and f(G, u) for accurate link prediction



کی ا

UNIVERSIT

Classical Supervised Learning Pipeline



1

وي ال



1

وي ال



Weisfeiler-Lehman (WL-1) Algorithm

- Weisfeiler-Lehman Algorithm (WL-1)
 - a.k.a. Color Refinement Algorithm
- Recursive algorithm to determine if two graphs are isomorphic
 - Works by recursively coloring nodes
 - Valid isomorphism test for most graphs (Babai and Kucera, 1979)
 - Cai et al., 1992 shows counter-examples not distinguished by it
 - Earliest works in graph kernels use these colors as features (Shervashidze et al., 2011)

Comp 599: N

D 11 1



Shervashidze et al. 2011

Initialize: h_v = initial color of node v

k = 0

function WL-1-coloring(G): while vertex attributes change do: $k \leftarrow k + 1$ for all vertices $v \in G$ do

these $h_{k,v} \leftarrow hash(h_{k-1,v}, \{h_{k-1,u}: \forall u \in Neighbors(v)\})$ Return $\{h_{k,v}: \forall v \in G\}$ neighbors of node v

کی ا

Resources: Libraries and Datasets









graphneural.network



github.com/deepmind/graph_nets





github.com/deepmind/jraph

https://pytorch-geometric.readthedocs. io/en/latest/modules/datasets.html

github.com/graphdeeplearning/benchmarking-gnns

slides based on https://petar-v.com/talks/GNN-Wednesday.pdf



Graphlearning.io

27