

Cornell Bowers CIS


## Logistics

- HW2 is due today
  - One extra slip day because of huggingface maintenance
  - Run the cells in the coding assignment
  - Turn off GPUs after using
- Project Proposal is due March 7
- HW 3 will be released at the end of the week
- We won't release HW solutions
- Change in office hours
  - Zach's office hours will be 4-5 on Wednesday's (instead of 6-7pm)
  - Varsha's office hours will be 1-2pm Tuesday's (instead of 9-10am)


Cornell Bowers CIS

## What are Graphs?

Graphs are a general language for describing and analyzing entities with relations/interactions



Molecule




Knowledge Graph


Cornell Bowers CIS

## What are Graphs?

Graphs = Nodes + Edges



Molecule




Knowledge Graph

## Graph: Directed vs Undirected

How the edges link the nodes allows us to distinguish between undirected graphs vs directed graphs


Graph G with 3 nodes



Examples:

- Phone Calls
- Following on Twitter

What we will focus on now



Examples:


- Academic collaborations
- Friendships on Facebook

## Adjacency Matrix - A

A represents the edges in a given graph

$A_{ij} = 1$  if an edge exists between nodes  $i$  and  $j$ , else 0

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




## Degree Matrix - D

D is a diagonal matrix, where each diagonal entry represents the degree of each node in a given graph


$D_{i,i} = \text{degree}(i)$

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

$$D = \begin{bmatrix} 3 & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \end{bmatrix}$$




## GraphML vs NLP vs CV



The cat sat on the mat.



VS.




➡ No spatial locality (unlike grids)

➡ No rank ordering or fixed reference point


## Why Do We Care About Learning on Graphs?

There are many different settings where we might care about learning on graphs:


- Graph classification
- Node classification
- Link prediction
- Community detection
- Graph embedding
- Graph generation



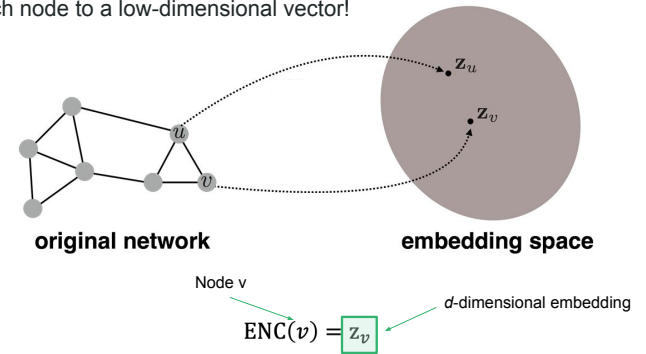
## Representation Learning > Feature Engineering




## Representation Learning + Classifier



Map each node to a low-dimensional vector!




### Example Graph mapped into 2 dimensions




(a) Input: Karate Graph

(b) Output: Representation

### Classify Embeddings!



### Brief Review: word2vec



(words close in sentences → close in embedding space)


List of sentences

“During the Battle of Endor, the Death Star II’s energy shield was destroyed...”  
W W W W W  
“In the third film, Anakin becomes Vader when...”  
“Samuel L. Jackson portrayed Mace Windu in the prequel trilogy...”

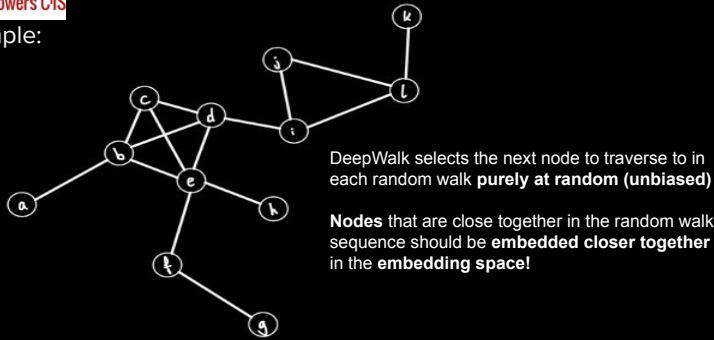
Can we do this on graphs?

### DeepWalk: word2vec For Graphs

This is **exactly** the same optimization as word2vec, but we instead optimize over **sequences of random walks on a graph**.



Example:



DeepWalk selects the next node to traverse to in each random walk **purely at random (unbiased)**


**Nodes** that are close together in the random walk sequence should be **embedded closer together** in the **embedding space!**

These are the "sentences" that we generate!

Random walk:

node2vec: The Introduction of Bias...

node2vec = DeepWalk + control over local vs global exploration (via two additional hyperparameters that we won't discuss in detail)



Breadth First Search (BFS) {s1, s2, s3} **Local** microscopic view

Depth First Search (DFS) {s4, s5, s6} **Global** macroscopic view

→ BFS


→ DFS

Homophily

Structural equivalence

*Grover and Leskovec., ACM SIGKDD, 2016*

# GRAPH NEURAL NETWORKS



Convolutional Layer in CNN

Translation-invariant

1 <sub>x1</sub>	1 <sub>x0</sub>	1 <sub>x3</sub>	0	0
0 <sub>x0</sub>	1 <sub>x1</sub>	1 <sub>x0</sub>	1	0
0 <sub>x1</sub>	0 <sub>x0</sub>	1 <sub>x3</sub>	1	1
0	0	1	1	0
0	1	1	0	0

Image

4		

Convolved Feature

How about for non-Euclidean data? Can we do something similar with graphs?


Can you perform convolutions on graphs? What would that look like? Will you run into any problems?

**Locality:** you can tell a lot about a particular **pixel** based on the properties of their neighbors

### Locality vs Homophily


**Homophily:** you can tell a lot about a particular **node** based on the properties of their neighbors

#### Image Convolutions




Generate next layer embedding vectors for each **pixel** in an input image by **aggregating the transformed feature vectors** of each of the pixel's neighbors

#### Graph Convolutions




Generate next layer embedding vectors for each **node** in an input graph by **aggregating the transformed feature vectors** of each of the node's neighbors

Let's look at a single layer of a graph convolution



**Thomas Kipf**  
PhD @ University of Amsterdam  
Currently: Research scientist @ Google Brain



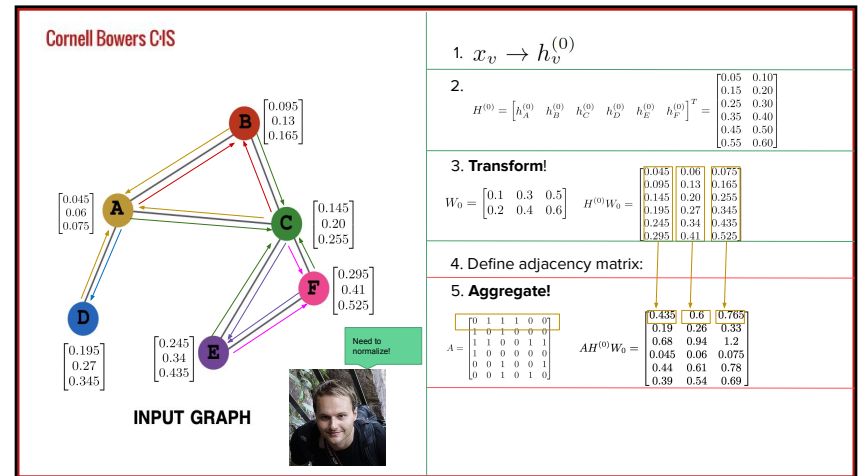
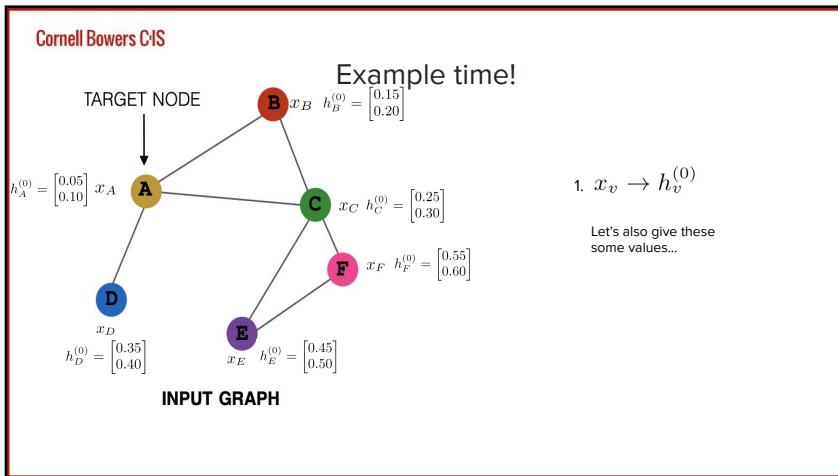
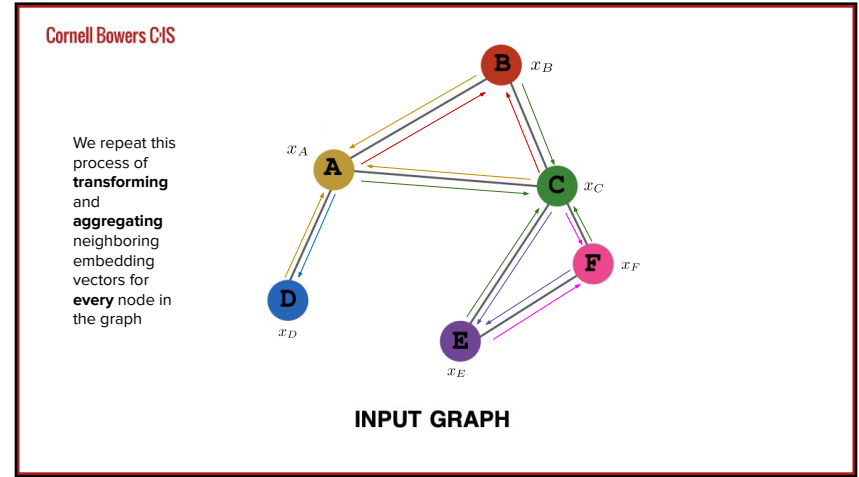
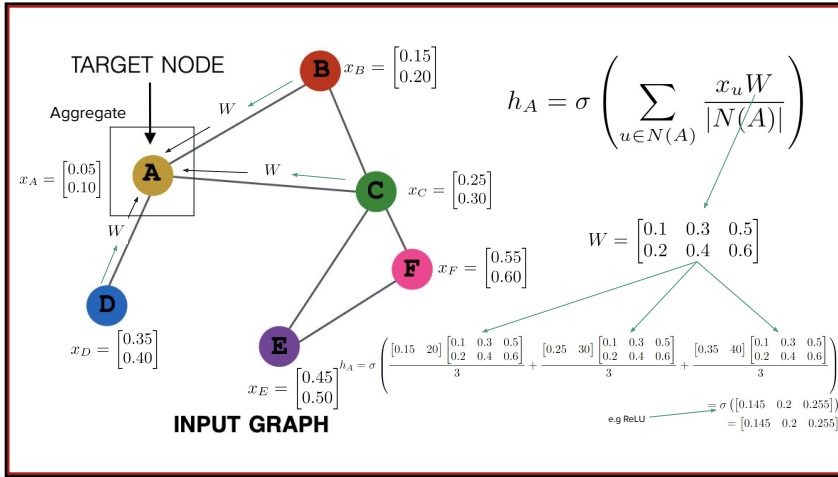
$$h_A = \sigma \left( \sum_{u \in N(A)} \frac{x_u W}{|N(A)|} \right)$$

Aggregate

Note: **Aggregation function MUST be permutation-invariant!**

- Mean()
- Sum()
- Max()

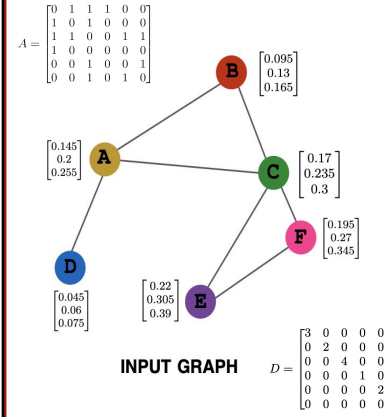
Let's choose Mean() for now...



Cornell Bowers CIS

You left multiply by  $D^{-1}$ . Compute  $D^{-1}$  and then think about why we multiply by  $D^{-1}$ .  
 (As a reminder, D is the diagonal matrix where each entry is a the degree of each vertex)

Cornell Bowers CIS



- $x_v \rightarrow h_v^{(0)}$
- $H^{(0)} = [h_A^{(0)} \ h_B^{(0)} \ h_C^{(0)} \ h_D^{(0)} \ h_E^{(0)} \ h_F^{(0)}]^T = \begin{bmatrix} 0.05 & 0.10 \\ 0.15 & 0.20 \\ 0.25 & 0.30 \\ 0.35 & 0.40 \\ 0.45 & 0.50 \\ 0.55 & 0.60 \end{bmatrix}$
- Transform!**  
 $W_0 = \begin{bmatrix} 0.1 & 0.3 & 0.5 \\ 0.2 & 0.4 & 0.6 \end{bmatrix}$   $H^{(0)}W_0 = \begin{bmatrix} 0.045 & 0.06 & 0.075 \\ 0.095 & 0.13 & 0.165 \\ 0.145 & 0.20 & 0.255 \\ 0.195 & 0.27 & 0.345 \\ 0.245 & 0.34 & 0.435 \\ 0.295 & 0.41 & 0.525 \end{bmatrix}$
- Define adjacency matrix:
- Aggregate!**  
 $AH^{(0)}W_0 = \begin{bmatrix} 0.435 & 0.6 & 0.765 \\ 0.19 & 0.26 & 0.33 \\ 0.68 & 0.94 & 1.2 \\ 0.045 & 0.06 & 0.075 \\ 0.44 & 0.61 & 0.78 \\ 0.39 & 0.54 & 0.69 \end{bmatrix}$
- Normalize
- Pass through non-linearity  
 $H^{(1)} = \sigma(D^{-1}AH^{(0)}W_0)$   
 $= ReLU(D^{-1}AH^{(0)}W_0)$   
 $= Z$

Cornell Bowers CIS

This is the **only** thing we **optimize!**

$$H^{(1)} = \sigma \left( \begin{array}{|c|c|} \hline & \\ \hline \end{array} D^{-1} AH^{(0)} W_0 \right)$$

2. Aggregate      1. Transform

Cornell Bowers CIS

This is the **only** thing we **optimize!**

$$Z = \begin{array}{|c|c|} \hline \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} & \frac{1}{2} X \Theta \\ \hline \end{array}$$

2. Aggregate      1. Transform

This is often the equation you will see for GCNs!



### GCN > Random Walks

$$H^{(1)} = \sigma(D^{-1}AH^{(0)}W_0)$$

Table 2: Summary of results in terms of classification accuracy (in percent).

Method	Citeseer	Cora	Pubmed	NELL
ManiReg [3]	60.1	59.5	70.7	21.8
SemiEmb [28]	59.6	59.0	71.1	26.7
LP [32]	45.3	68.0	63.0	26.5
DeepWalk [22]	43.2	67.2	65.3	58.1
ICA [18]	69.1	75.1	73.9	23.1
Planetoid* [29]	64.7 (26s)	75.7 (13s)	77.2 (25s)	61.9 (185s)
<b>GCN (this paper)</b>	<b>70.3 (7s)</b>	<b>81.5 (4s)</b>	<b>79.0 (38s)</b>	<b>66.0 (48s)</b>
GCN (rand. splits)	67.9 ± 0.5	80.1 ± 0.5	78.9 ± 0.7	58.4 ± 1.7

Just plug test nodes here!


With random walks, what we're optimizing are the **final embedding vectors** themselves, **not weights**...

...so for every **new/unseen node** that we're given (e.g. in a test set), we have to use SGD AGAIN to optimize their embeddings, which is **computationally expensive!**

**THIS IS A BIG REASON WHY WE USE A WEIGHT MATRIX!**

### What do we do with Z?

$$\begin{aligned} H^{(1)} &= \sigma(D^{-1}AH^{(0)}W_0) \\ &= ReLU(D^{-1}AH^{(0)}W_0) \\ &= Z \end{aligned}$$



Depends on the downstream prediction task:

- Feed Z into a **MLP + Softmax decoder** for **node-level classification/regression**
- For **graph-level predictions** (e.g. classifying an entire graph), can **concat/sum/mean all vectors in Z**, and then feed this long vector into a **MLP**
  - Just like in CNNs!

### Stacking GCN Layers

$$H^{(1)} = \sigma(D^{-1}AH^{(0)}W_0)$$

Input to the next layer

Note: new weight matrix! Weight matrices in GNNs are **layer-specific**.

$$H^{(2)} = \sigma(D^{-1}AH^{(1)}W_1)$$

But D and A never change!

### Stacking GCN Layers

Final GCN update rules:


Node-level update rule: 
$$h_v^{l+1} = \sigma \left( \sum_{u \in N(v)} \frac{h_u^l W_l}{|N(v)|} \right)$$

Graph-level update rule: 
$$H^{(l+1)} = \sigma(D^{-1}AH^{(l)}W_l)$$

Let's just keep adding more layers, right?

**BIG problem!**

## Stacking GCN Layers



$$H^{(2)} = \sigma(D^{-1}AH^{(1)}W_1)$$

In order to calculate A's  $h_A^2$  vector, we need to calculate  $h_u^1$  for each  $u$  in  $\text{Neighbors}(A)$

$H^2$  looks at neighbors' neighbors' neighbors, etc...this becomes MASSIVE on large graphs

$$h_A^2 = \sigma \left( \sum_{u \in N(A)} \frac{h_u^1 W_1}{|N(A)|} \right)$$



In order to calculate each node  $u$ 's  $h_u^1$  vector, we need to calculate  $h_v^0$  for each  $v$  in  $\text{Neighbors}(u)$


$$h_u^1 = \sigma \left( \sum_{v \in N(u)} \frac{h_v^0 W_0}{|N(u)|} \right)$$

Just to calculate  $h_A^2$ , we need to look at A's neighbors' neighbors


Therefore, number of layer in graph neural networks is a **very important hyperparameter!**

## The over-smoothing problem


### Receptive field for 1-layer GNN



### Receptive field for 2-layer GNN



### Receptive field for 3-layer GNN



Receptive field: the set of all nodes that are used to calculate an  $l$ -th layer embedding vector for a node  $v$


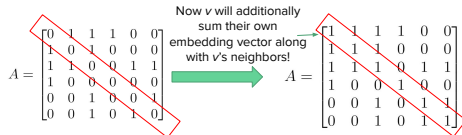
Here, we encounter the **over-smoothing problem**, where final-layer node embeddings (in  $Z$ ) become highly similar.

## GraphSAGE

2 **BIG** problems with GCNs:

**Problem 1:**  $h_v^{l+1}$  doesn't aggregate  $h_v^l$   $h_v^{l+1} = \sigma \left( \sum_{u \in N(v)} \frac{h_u^l W_l}{|N(v)|} \right)$

**Solution 1:** Add **self-loops!**



**Jure Leskovec**  
Postdoc @ Cornell  
Currently: Professor @ Stanford  
Until very recently: Chief Scientist @ Pinterest  
Created **node2vec**

**Problem 2:** Just Mean()? How about the rest?

**Solution 2:** Make the aggregation function a **hyperparameter!**

## GraphSAGE > GCN

Table 1: Prediction results for the three datasets (micro-averaged F1 scores). Results for unsupervised and fully supervised GraphSAGE are shown. Analogous trends hold for macro-averaged scores.

Name	Citation		Reddit		PPI	
	Unsup. F1	Sup. F1	Unsup. F1	Sup. F1	Unsup. F1	Sup. F1
Random	0.206	0.206	0.043	0.042	0.396	0.396
Raw features	0.575	0.575	0.585	0.585	0.422	0.422
DeepWalk	0.565	0.565	0.324	0.324	—	—
DeepWalk + features	0.701	0.701	0.691	0.691	—	—
GraphSAGE-GCN	0.742	0.772	<b>0.908</b>	0.930	0.465	0.500
GraphSAGE-mean	0.778	0.820	0.897	0.950	0.486	0.598
GraphSAGE-LSTM	0.788	0.832	<b>0.907</b>	<b>0.954</b>	0.482	<b>0.612</b>
GraphSAGE-pool	<b>0.798</b>	<b>0.839</b>	0.892	0.948	<b>0.502</b>	0.600

## Simplifying GCNs

Remember this?

Graph-level update rule:

$$H^{(l+1)} = \boxed{\times} \boxed{D^{-1} A} H^{(l)} W_l$$

Define:  $S = D^{-1} A$ 

$$H^1 = S H^{(0)} W_0$$

Why does this work so well?

The strength of GNNs comes from their ability to **propagate node features**, not from non-linearities

Get rid of the non-linearities!

<https://arxiv.org/pdf/1902.07153.pdf>
Write an expression for  $H^{(l+1)}$ 

Table 2. Test accuracy (%) averaged over 10 runs on citation networks. <sup>†</sup>We remove the outliers (accuracy < 75/65/75%) when calculating their statistics due to high variance.

	Cora	Citeseer	Pubmed
<b>Numbers from literature:</b>			
GCN	81.5	70.3	79.0
GAT	83.0 ± 0.7	72.5 ± 0.7	79.0 ± 0.3
GLN	81.2 ± 0.1	70.9 ± 0.1	78.9 ± 0.1
AGNN	83.1 ± 0.1	71.7 ± 0.1	79.9 ± 0.1
LNet	79.5 ± 1.8	66.2 ± 1.9	78.3 ± 0.3
AdaLNet	80.4 ± 1.1	68.7 ± 1.0	78.1 ± 0.4
DeepWalk	70.7 ± 0.6	51.4 ± 0.5	76.8 ± 0.6
DGI	82.3 ± 0.6	71.8 ± 0.7	76.8 ± 0.6
<b>Our experiments:</b>			
GCN	81.4 ± 0.4	70.9 ± 0.5	79.0 ± 0.4
GAT	83.3 ± 0.7	72.6 ± 0.6	78.5 ± 0.3
FastGCN	79.8 ± 0.3	68.8 ± 0.6	77.4 ± 0.3
GIN	77.6 ± 1.1	66.1 ± 0.9	77.0 ± 1.2
LNet	80.2 ± 3.0 <sup>†</sup>	67.3 ± 0.5	78.3 ± 0.6 <sup>†</sup>
AdaLNet	81.9 ± 1.9 <sup>†</sup>	70.6 ± 0.8 <sup>†</sup>	77.8 ± 0.7 <sup>†</sup>
DGI	82.5 ± 0.7	71.6 ± 0.7	78.4 ± 0.7
SGC	81.0 ± 0.0	71.9 ± 0.1	78.9 ± 0.0

## Summary

- Learning on graphs: Classify nodes and entire graphs, predict links or detect communities and even generate graphs and their embeddings
- Feature Engineering 🚫 Representation Learning 😊
- Random Walks, DeepWalk + node2vec: word2vec on graphs, embed nearby nodes on the random walk closer together
- GCN: CNN on graphs, transform + aggregate neighbors. Homophily in GCNs similar to locality in CNNs.
- Over-smoothing problem: Can't stack too many layers