

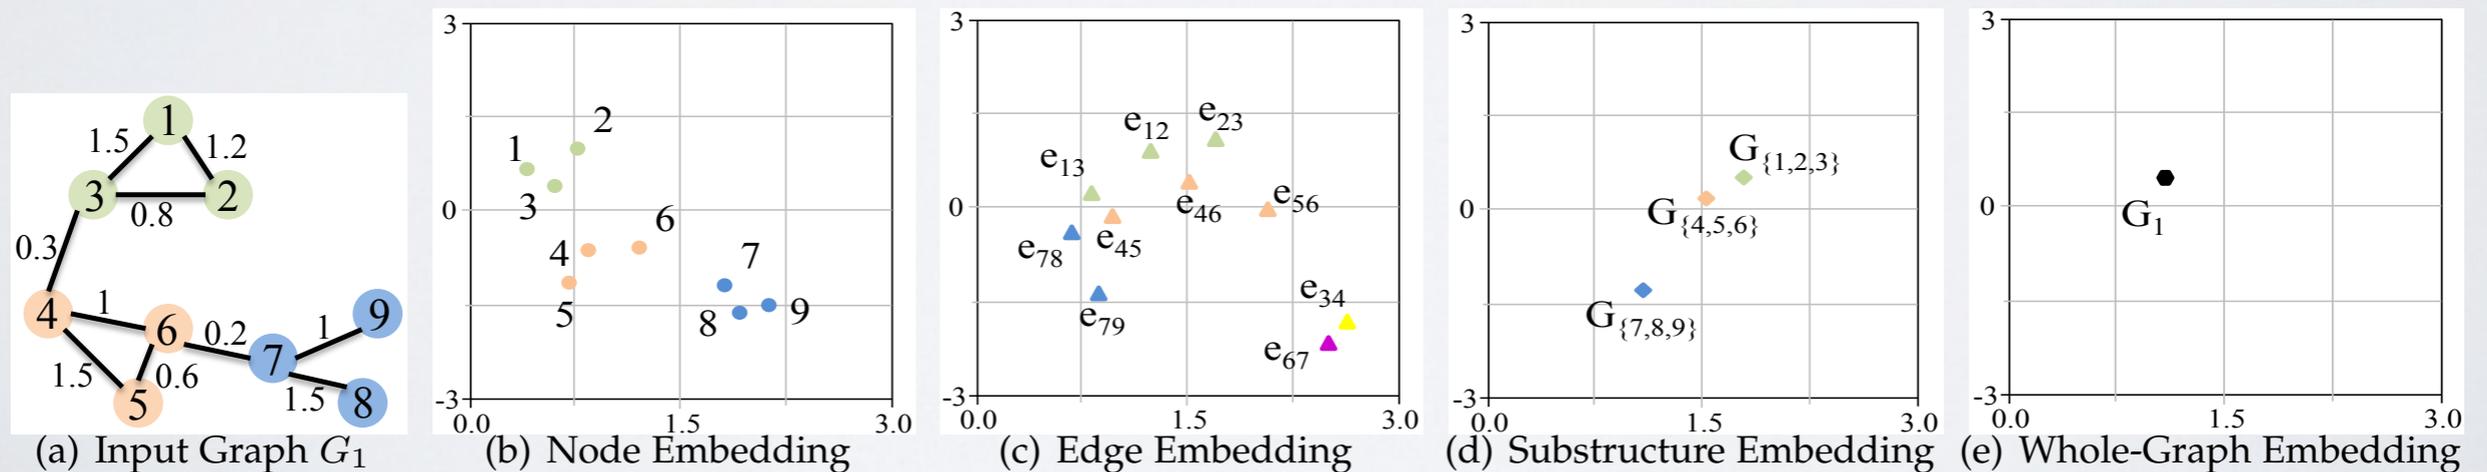
GRAPH/NODE EMBEDDING

Goyal, P., & Ferrara, E. (2018). Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*, 151, 78-94.

Cai, H., Zheng, V. W., & Chang, K. C. C. (2018). A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 30(9), 1616-1637.

VARIANT

- We can differentiate:
 - Node embedding
 - Edge Embedding
 - Substructure embedding
 - Whole graph Embedding
- In this course, only *node embedding* (often called graph embedding)



IN CONCRETE TERMS

- A graph is composed of
 - Nodes (possibly with labels)
 - Edges (possibly directed, weighted, with labels)
- A graph/node embedding technique in **d** dimensions will assign a vector of length **d** to each node, that will be useful for **what we want to do with the graph**.
 - It captures some aspect of the network structure
- A vector can be assigned to an edge (u,v) by combining vectors of u and v

WHAT TO DO WITH EMBEDDINGS?

- Two possible ways to use an embedding:
 - ▶ Unsupervised learning:
 - The *distance* between vectors in the embedding is used for *something*
 - ▶ Supervised learning:
 - Algorithm learn to predict *something* from the features in the embedding

WHAT CAN WE DO WITH
EMBEDDINGS ?

EMBEDDING TASKS

- Common tasks:
 - ▶ Link prediction (supervised)
 - ▶ Graph reconstruction (unsupervised link prediction ? / ad hoc)
 - ▶ Community detection (unsupervised)
 - ▶ Node classification (supervised community detection ?)
 - ▶ Role definition (Variant of node classification, can be unsupervised)
 - ▶ Visualisation (distances, like unsupervised)

OVERVIEW OF MOST POPULAR METHODS

MATRIX FACTORIZATION

LE: LAPLACIAN EIGENMAPS

- Introduced 2001

- Objective function:

$$y^* = \min \sum_{i \neq j} \|y_i - y_j\|^2 S_{ij}$$

- y^* : optimal embedding
- y_i : embedding of node i
- S_{ij} : similarity between nodes i and j (A , heuristic, ...)

- Minimize the product between **distance in the embedding** and **similarity in the graph**

- If nodes are *similar*, they must be *close* in the embedding

LE: LAPLACIAN EIGENMAPS

- $y^* = \min \sum_{i \neq j} \|y_i - y_j\|^2 S_{ij}$
- Can be written (with $S=A$) in matrix form as:
 - $\min y^T L y$
 - L : Laplacian
- To avoid trivial solution, we impose the constraint:
 - $y^T D y = I$
 - D : Degree matrix
- Solution: d eigenvectors of lowest eigenvalues of $D^{-1/2} L D^{-1/2}$

HOPE: HIGHER-ORDER PROXIMITY PRESERVED EMBEDDING

- Preserve a proximity matrix

- $y^* = \min \sum_{i,j} |S_{ij} - y_i y_j^T|$

- S can be the adjacency matrix, or number of common neighbors, Adamic Adar, etc.
- As similarity tends towards 0, embedding vectors must tend towards orthogonality (orthogonal vectors: $y_i y_j^T = 0$)

RANDOM WALKS BASED

DEEPWALK

- The first *Random Walk+Neural Networks* graph embedding method.
 - First of a long series
- Adaptation of **word2vec/skipgram** to graphs

SKIPGRAM

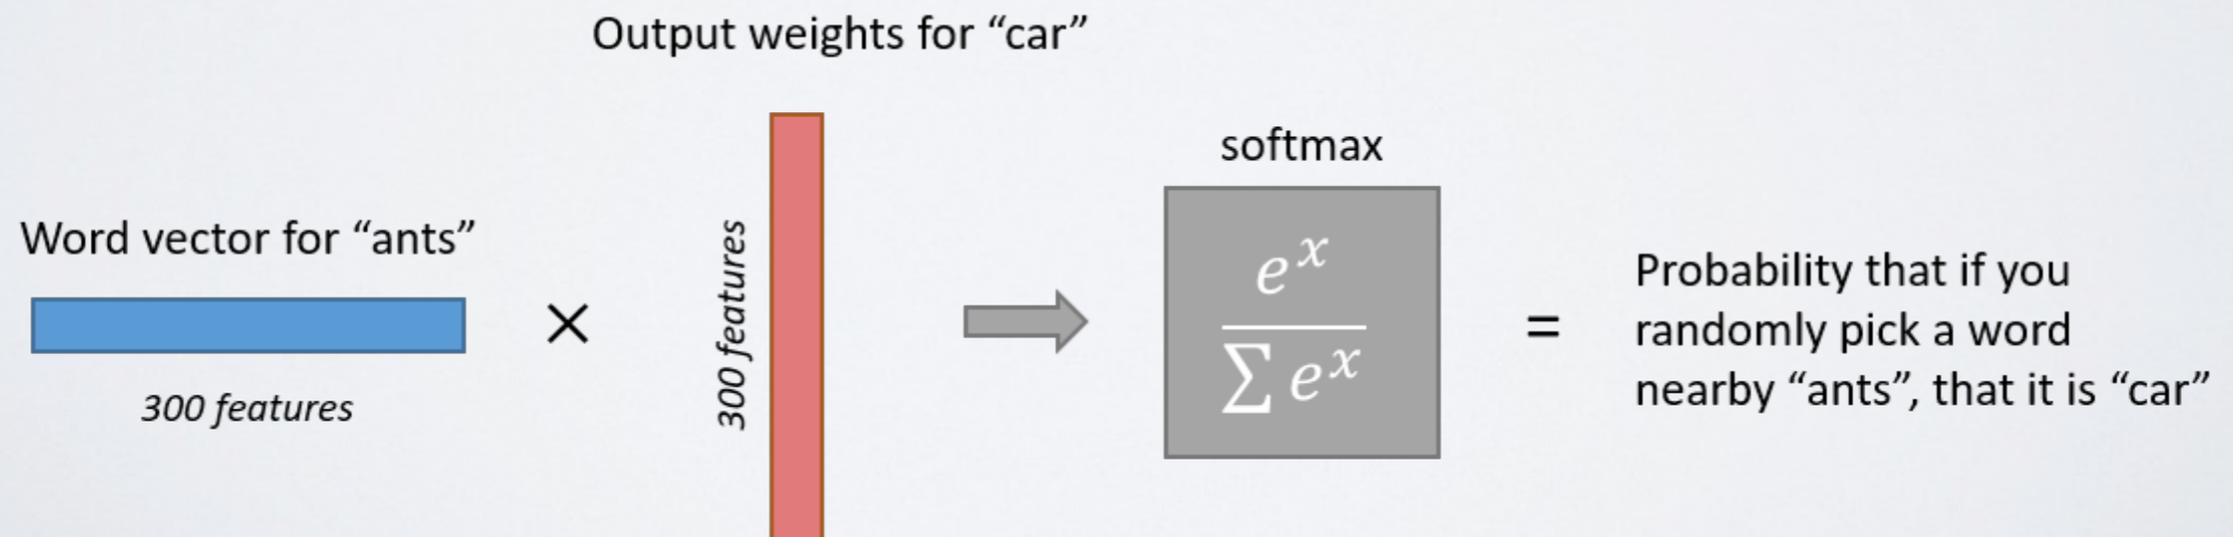
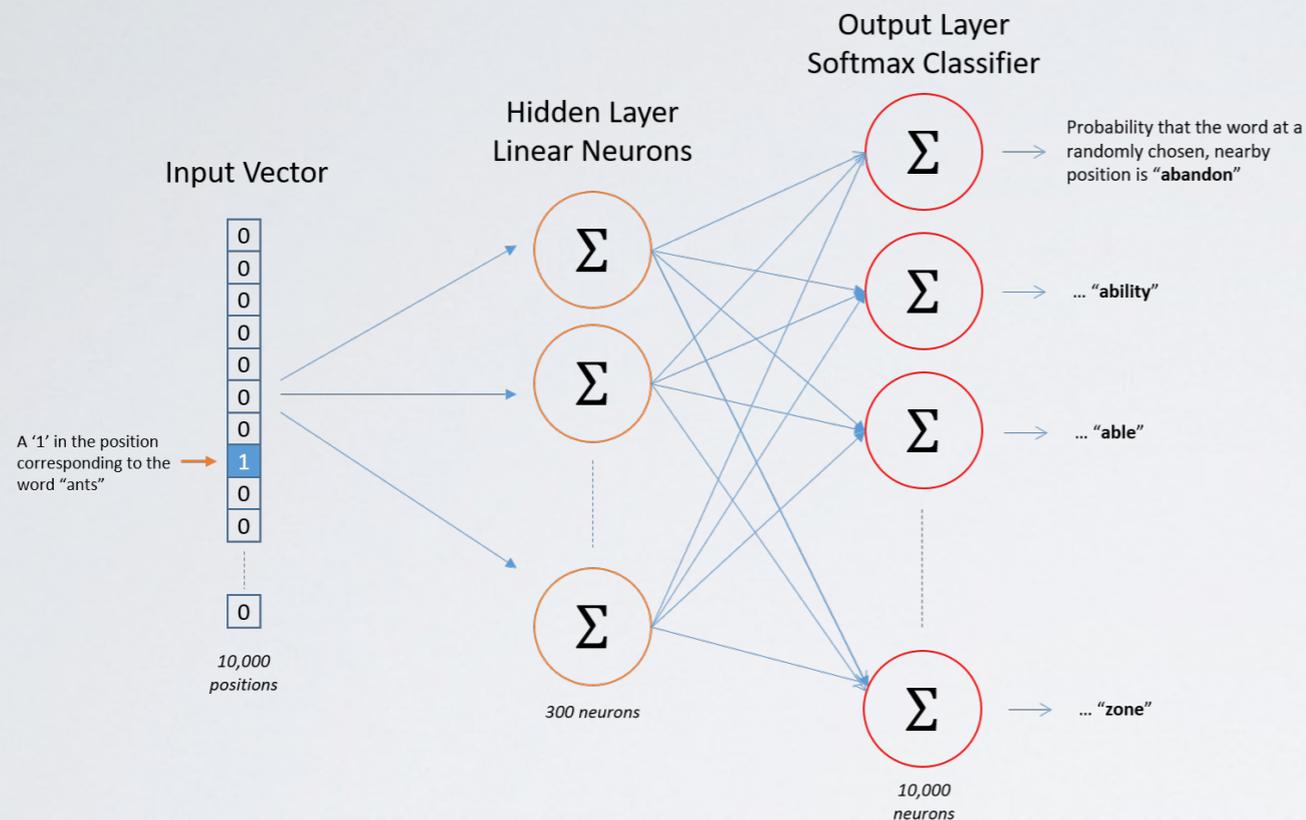
Word embedding

Corpus \Rightarrow Word = vectors

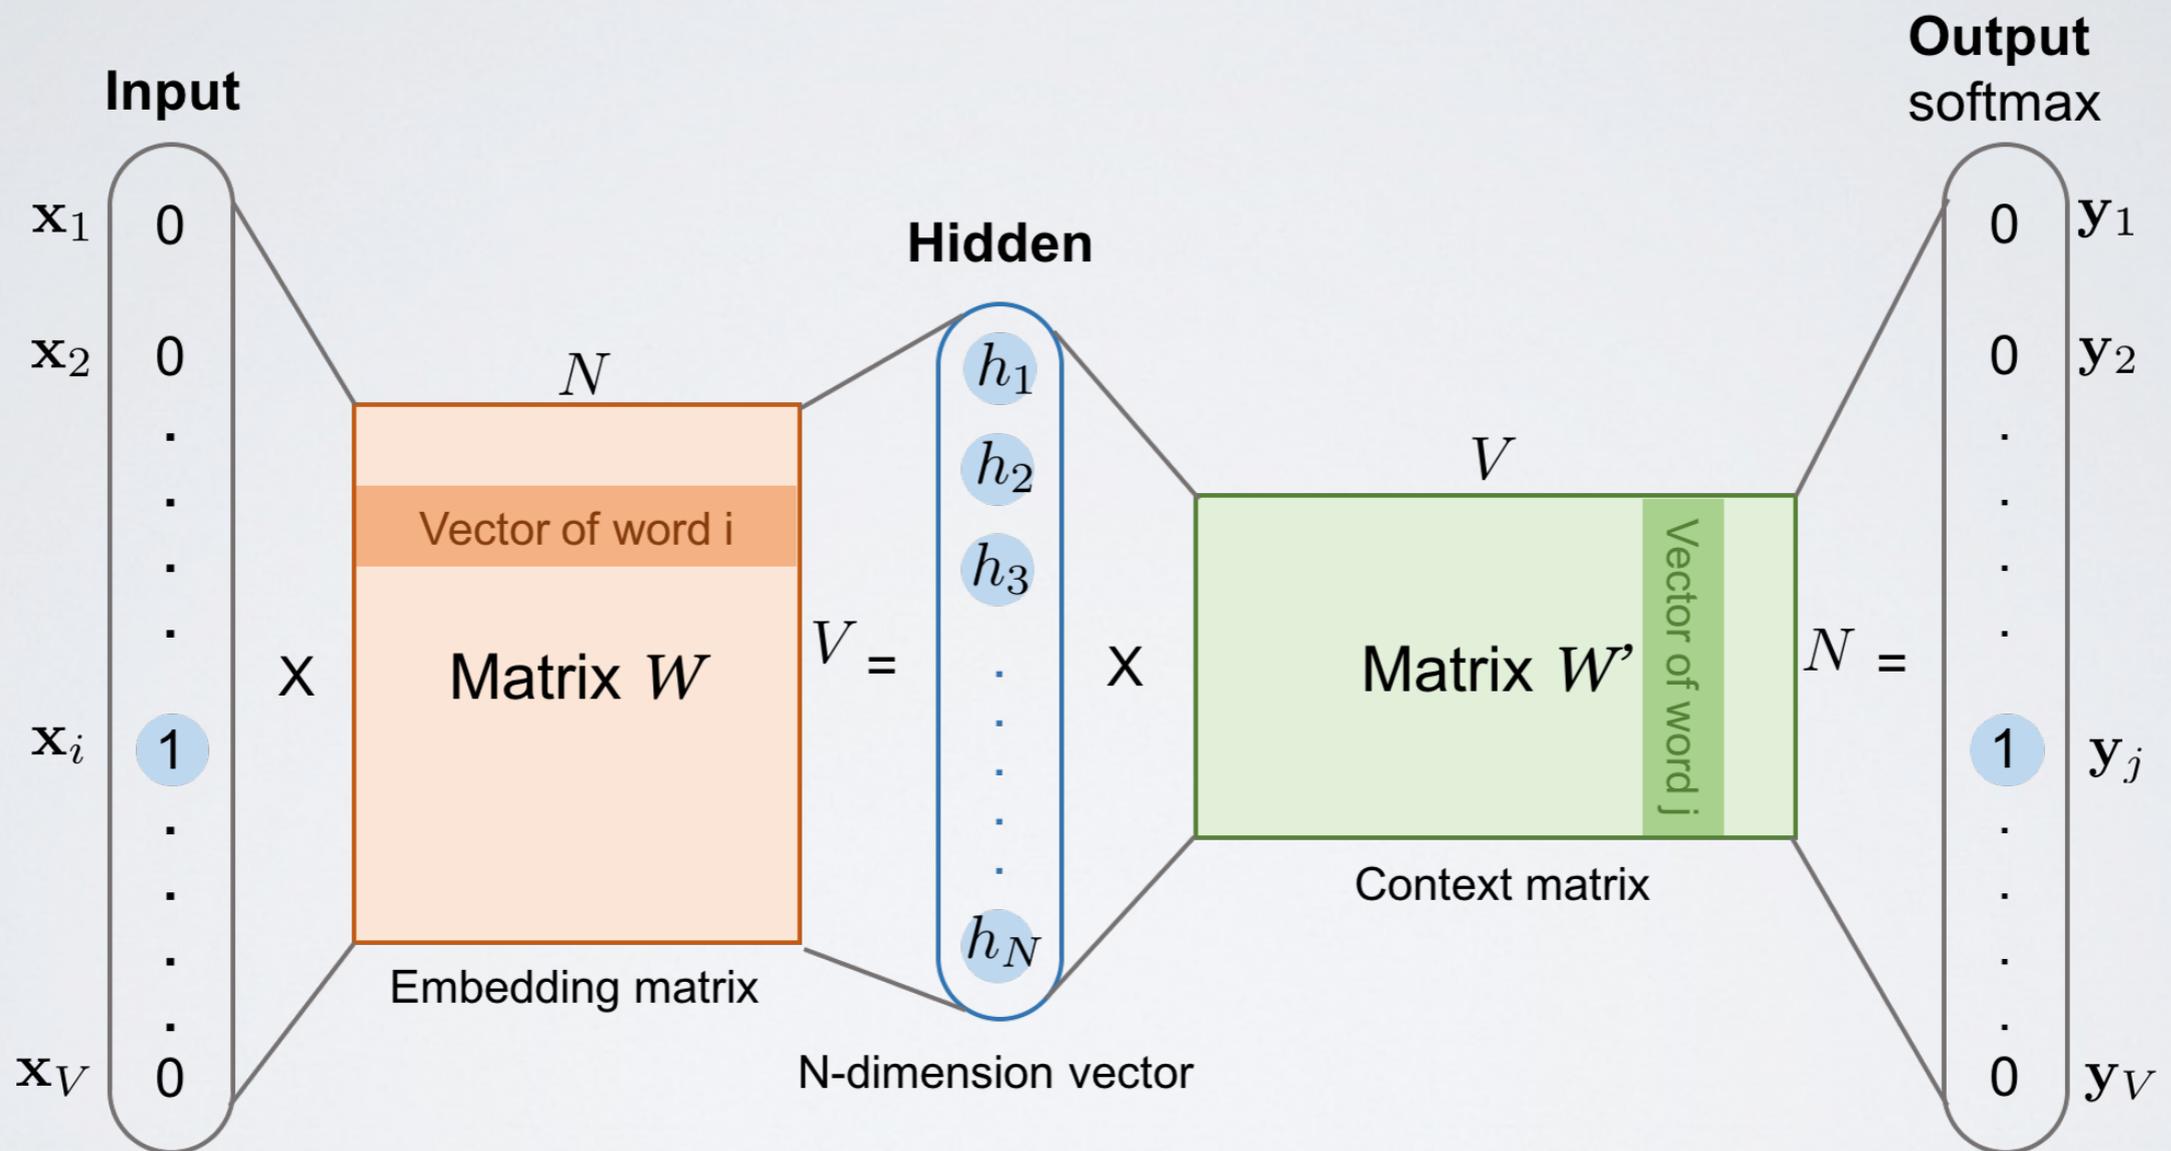
Similar embedding = similar **context**

Source Text	Training Samples
The quick brown fox jumps over the lazy dog. \rightarrow	(the, quick) (the, brown)
The quick brown fox jumps over the lazy dog. \rightarrow	(quick, the) (quick, brown) (quick, fox)
The quick brown fox jumps over the lazy dog. \rightarrow	(brown, the) (brown, quick) (brown, fox) (brown, jumps)
The quick brown fox jumps over the lazy dog. \rightarrow	(fox, quick) (fox, brown) (fox, jumps) (fox, over)

SKIPGRAM

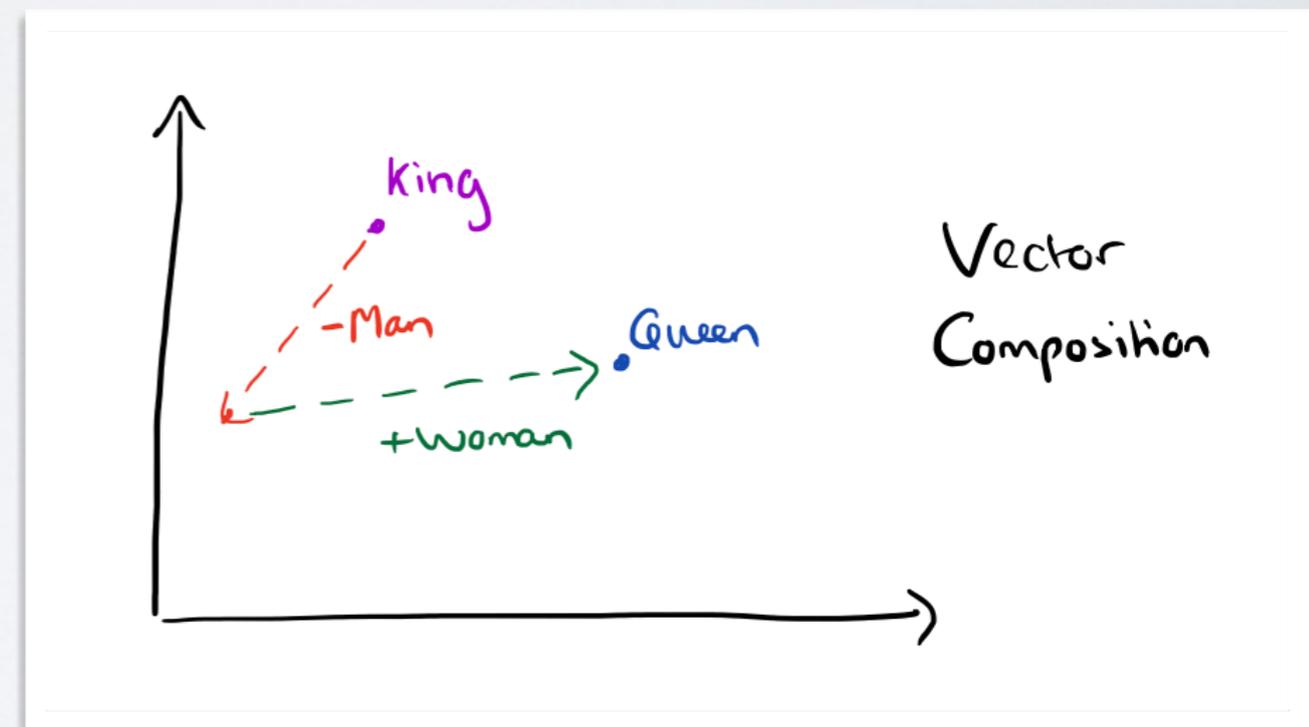
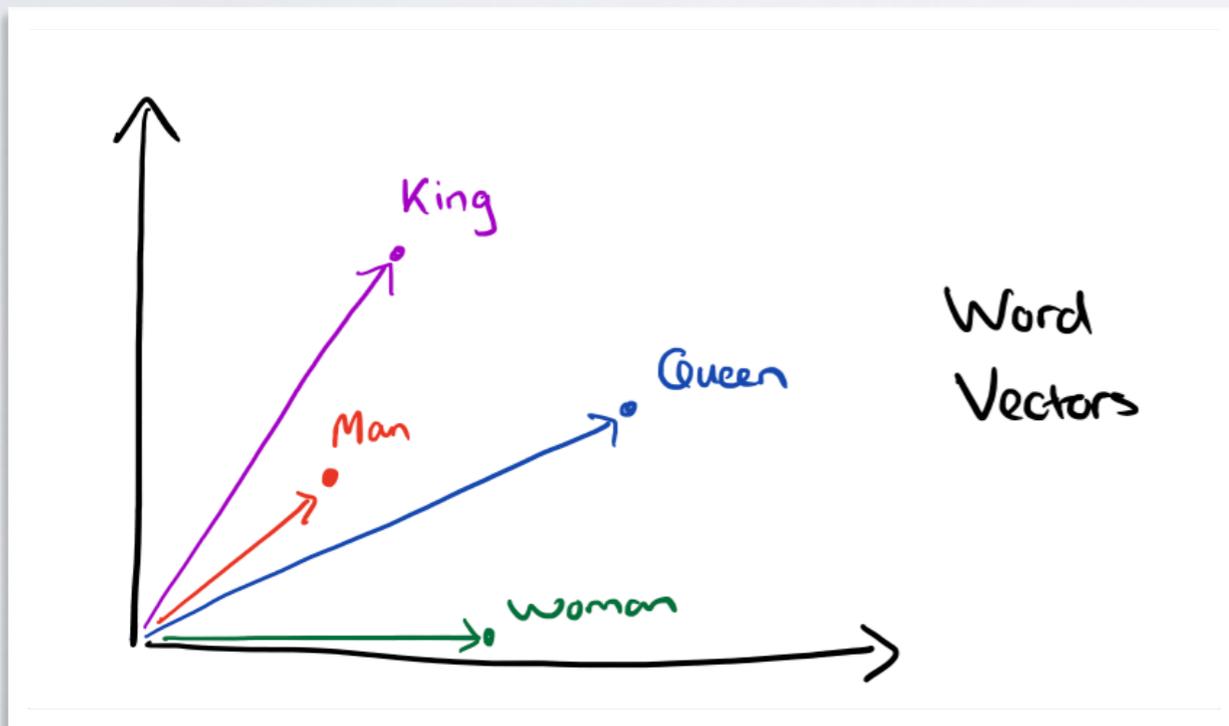


SKIPGRAM



SKIPGRAM

	King	Queen	Woman	Princess	...
Royalty	0.99	0.99	0.02	0.98	
Masculinity	0.99	0.05	0.01	0.02	
Femininity	0.05	0.93	0.999	0.94	
Age	0.7	0.6	0.5	0.1	
...	⋮				



[<https://blog.aolyer.org/2016/04/21/the-amazing-power-of-word-vectors/>]

SKIPGRAM

Table 8: *Examples of the word pair relationships, using the best word vectors from Table 4 (Skip-gram model trained on 783M words with 300 dimensionality).*

Relationship	Example 1	Example 2	Example 3
France - Paris	Italy: Rome	Japan: Tokyo	Florida: Tallahassee
big - bigger	small: larger	cold: colder	quick: quicker
Miami - Florida	Baltimore: Maryland	Dallas: Texas	Kona: Hawaii
Einstein - scientist	Messi: midfielder	Mozart: violinist	Picasso: painter
Sarkozy - France	Berlusconi: Italy	Merkel: Germany	Koizumi: Japan
copper - Cu	zinc: Zn	gold: Au	uranium: plutonium
Berlusconi - Silvio	Sarkozy: Nicolas	Putin: Medvedev	Obama: Barack
Microsoft - Windows	Google: Android	IBM: Linux	Apple: iPhone
Microsoft - Ballmer	Google: Yahoo	IBM: McNealy	Apple: Jobs
Japan - sushi	Germany: bratwurst	France: tapas	USA: pizza

[<https://blog.acolyer.org/2016/04/21/the-amazing-power-of-word-vectors/>]

GENERIC “SKIPGRAM”

- Algorithm that takes an input:
 - The element to embed
 - A list of “context” elements
- Provide as output:
 - An embedding with interesting properties
 - Works well for machine learning
 - Similar elements are close in the embedding
 - Somewhat preserves the overall structure

DEEPWALK

- Skipgram for graphs:
 - ▶ 1) Generate “sentences” using random walks
 - ▶ 2) Apply Skipgram
- Parameters:
 - ▶ Embedding dimensions d
 - ▶ Context size
 - ▶ More technical parameters: length of random walks, number of walks starting from each node, etc.

NODE2VEC

- Use biased random walk to tune the context to capture *what we want*
 - ▶ “Breadth first” like RW => local neighborhood (edge probability ?)
 - ▶ “Depth-first” like RW => global structure ? (Communities ?)
 - ▶ 2 parameters to tune:
 - **p**: bias towards revisiting the previous node
 - **q**: bias towards exploring undiscovered parts of the network

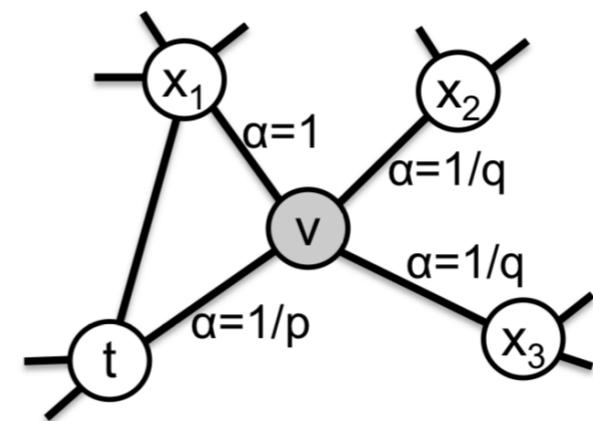


Figure 2: Illustration of the random walk procedure in *node2vec*. The walk just transitioned from t to v and is now evaluating its next step out of node v . Edge labels indicate search biases α .

RANDOM WALK METHODS

- What is the objective function ?
- How to interpret the distance between nodes in the embedding ?

RANDOM WALK METHODS

Approximately

$$y = \min \sum_{(i,j)} p(n_j | n_i) - \sigma(y_i y_j^T)$$

with $p(w_j | w_i)$ the probability to encounter node n_j in a random walk of a chosen length starting from node n_i . Its objective is therefore to make the distance in the embedding proportional to a random walk based distance in the graph.

with σ the softmax function defined as $\frac{e^x}{\sum e^x}$, a function commonly used in neural networks to add non-linearity and to ensure that the solution is a probability.

RANDOM WALK METHODS

- Scalability:
 - Skipgram uses techniques from machine learning developed for very large datasets: highly **scalable** (not necessarily *fast* or *cost efficient*)
- Matrix factorization methods require the similarity matrix S as input
 - Computing all random walk distance: $\mathcal{O}(n^2)$
 - k random walks of length ℓ from each node: $\mathcal{O}(n)$

SOME REMARKS ON WHAT
ARE EMBEDDINGS

ADJACENCY MATRIX

- An adjacency matrix is an embedding... in high dimension
- That represents the structural equivalence
 - 2 nodes have similar “embeddings” if they have similar neighborhoods
 - Distance \Rightarrow # of different neighbors (Manhattan Distance)
- Standard dimensionality reduction (T-SNE, PCA) of this matrix?
 - Small dimensions
 - But still unintuitive notion of distance

GRAPH LAYOUT

- Graph layouts are also embeddings.
 - Force layout, kamada-kawai
- They try to put connected nodes close to each other and non-connected ones “not close”
- Problem: they try to avoid overlaps
- Usually not scalable

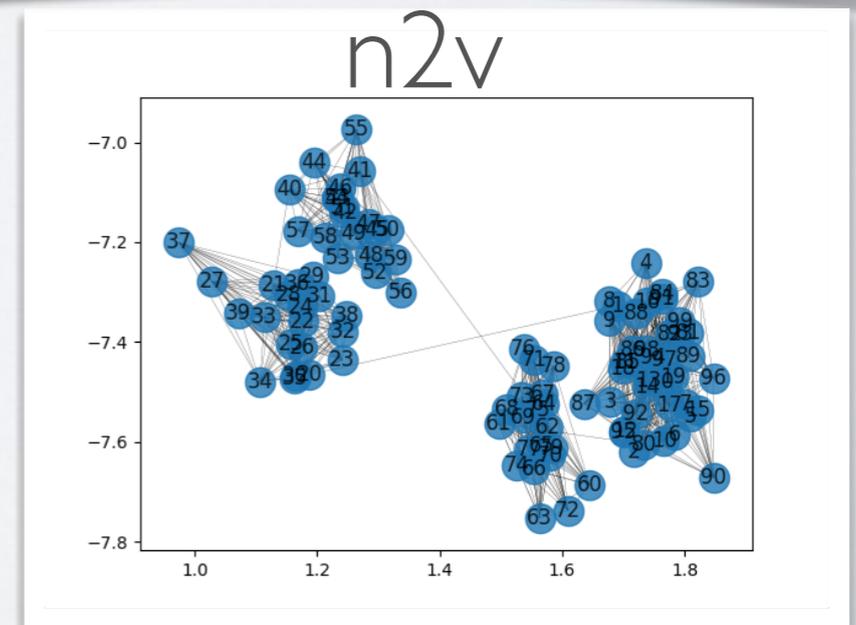
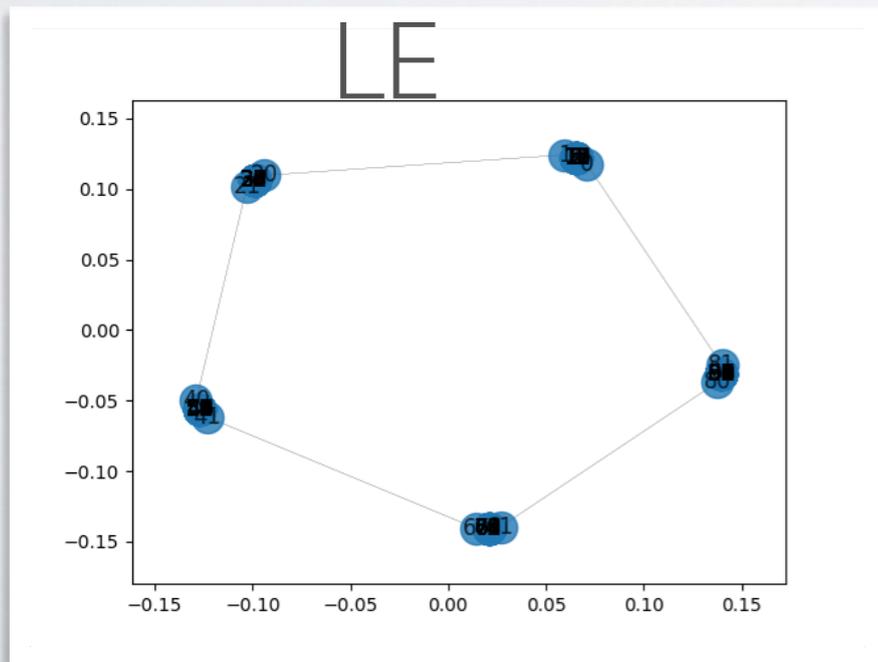
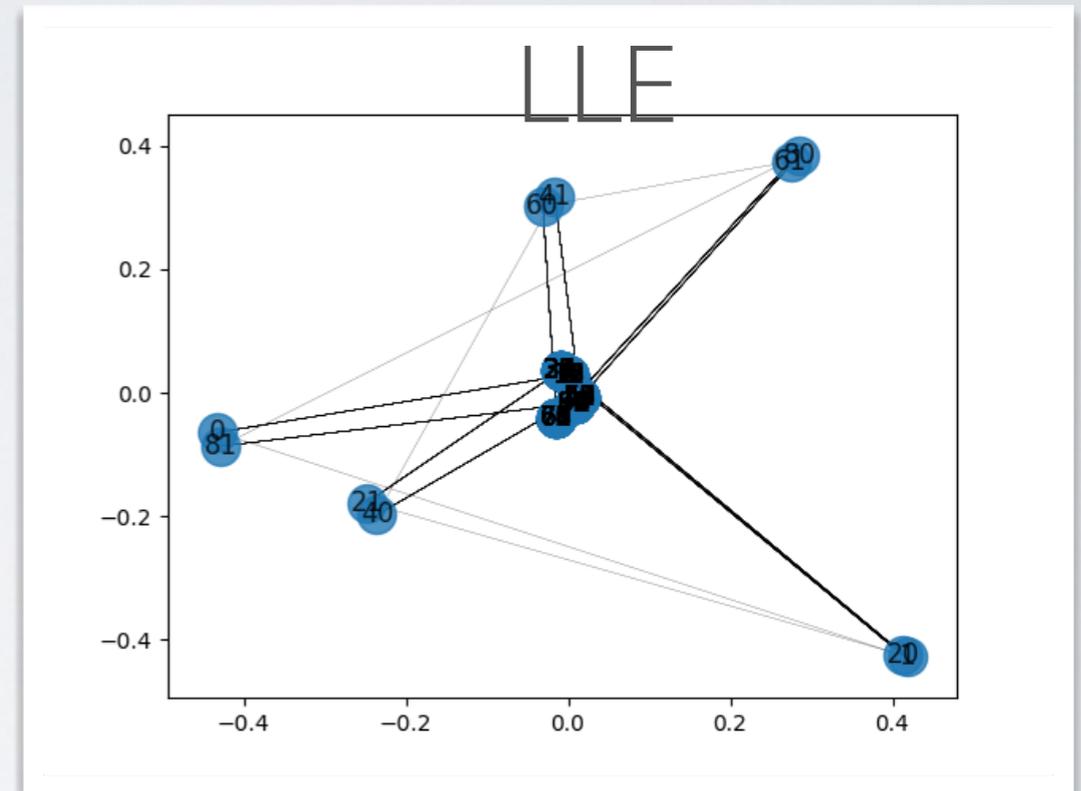
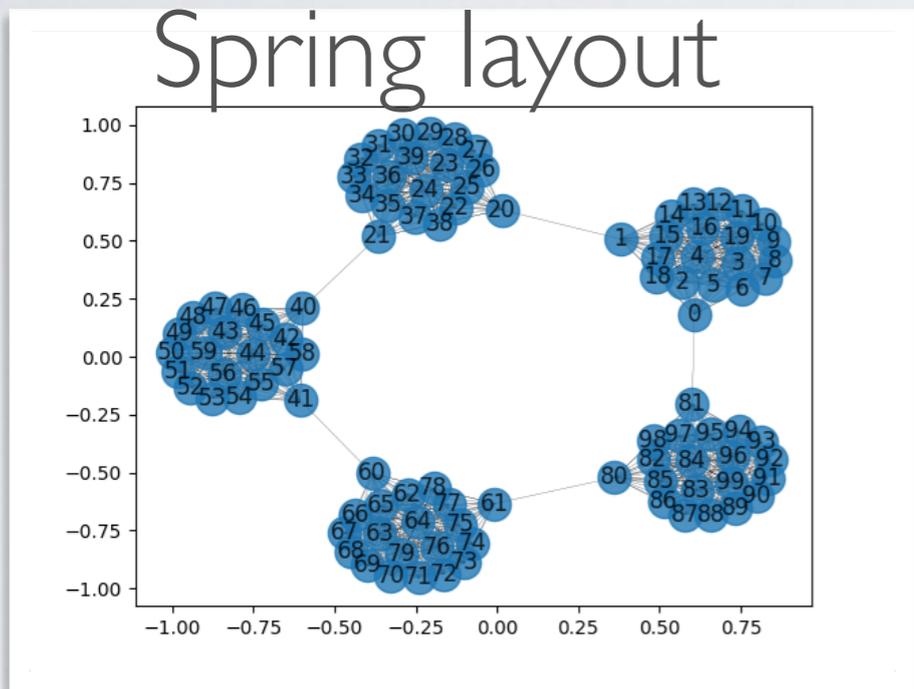
NODE EMBEDDING: VISUALIZATION

FROM D TO 2

- Graph embedding can be used to visualize graphs
- Requires to reduce the embedding from d to 2
 - TSNE
 - PCA
 - ...
- Interpretable positions of nodes
- But not necessarily optimized for human reading

CLIQUE RING

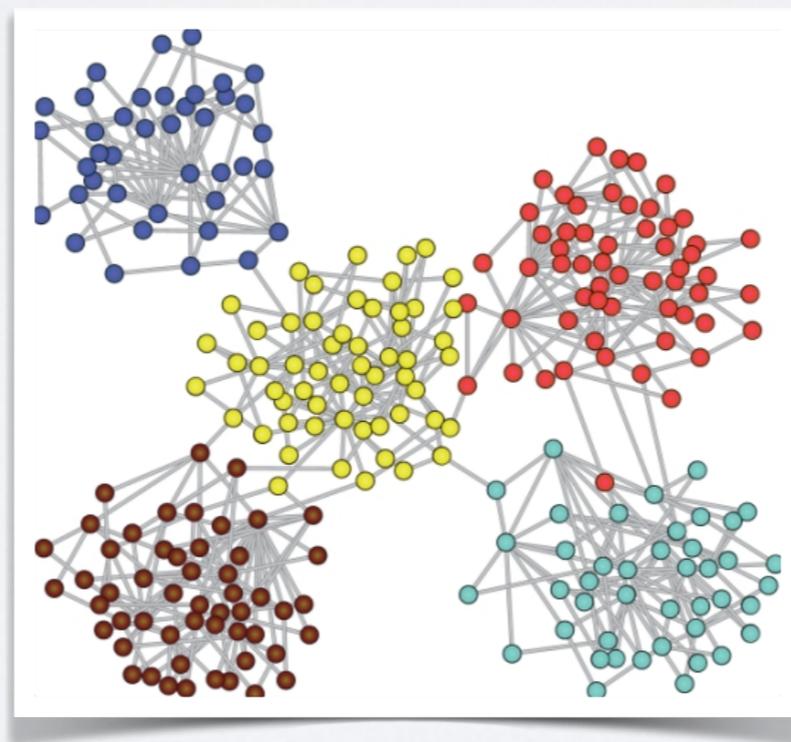
5 cliques of size 20 with 1 edge between them



NODE EMBEDDING: COMMUNITY DETECTION

CLUSTERING EMBEDDINGS

- Many algorithm exists for **clustering** non-network data
 - K-means, DBscan, etc.
- Clustering: group nodes that are close in the feature space.

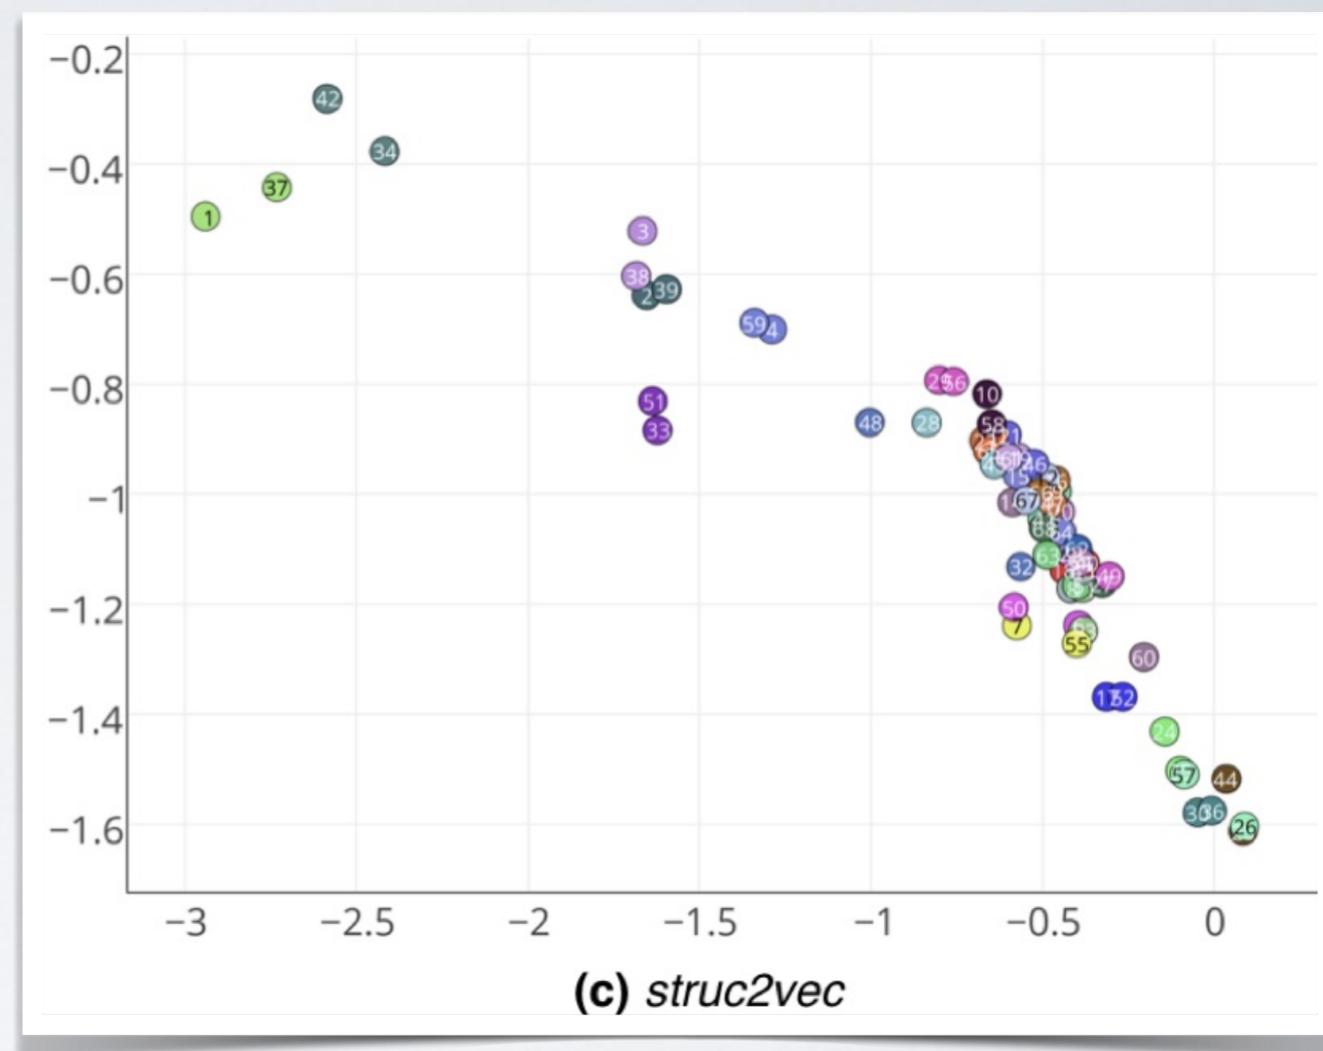
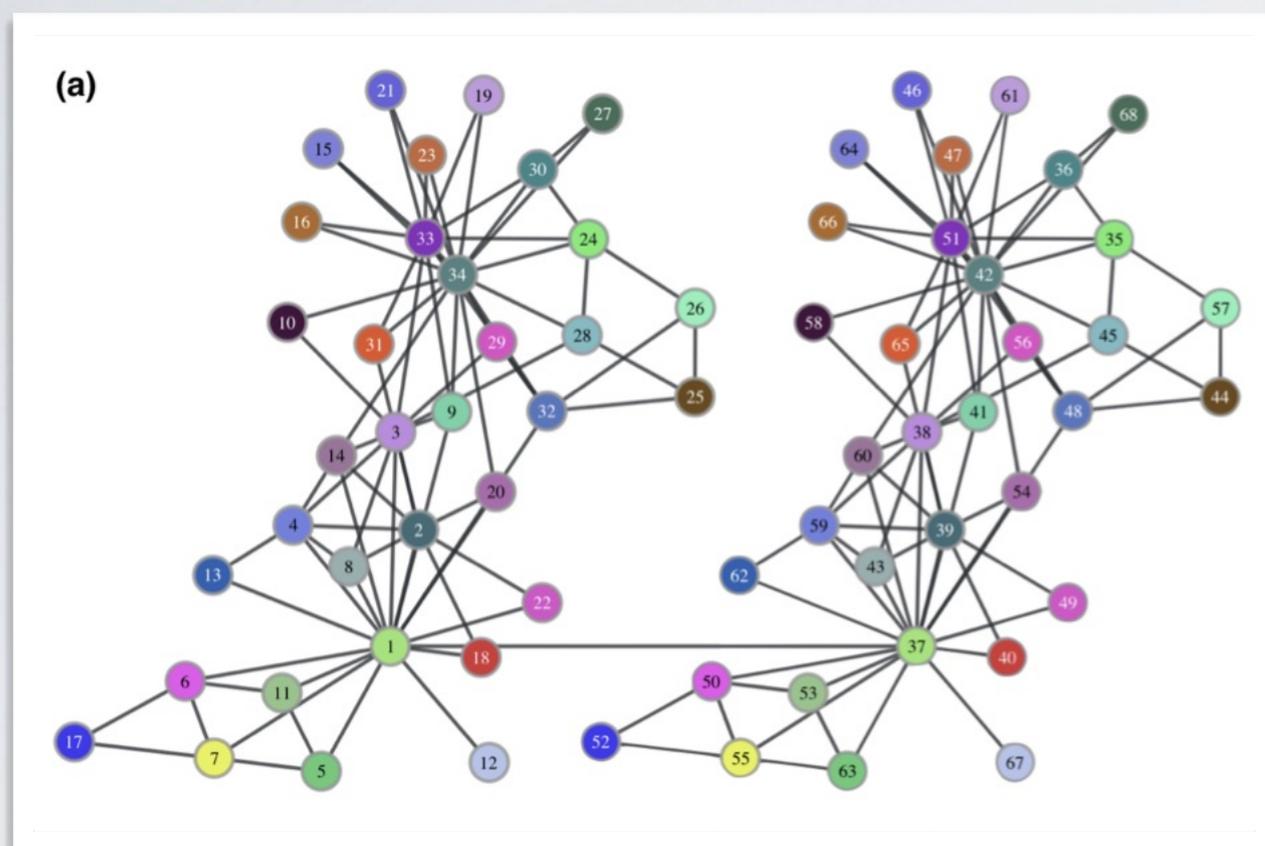


EMBEDDING ROLES

STRUC2VEC/ROLE2VEC

- In node2vec/Deepwalk, the context collected by RW contain the **labels** of encountered nodes
- Instead, we could memorize the **properties** of the nodes: attributes if available, or computed attributes (degrees, CC, ...)
- => Nodes with a same context will be nodes in a same “position” in the graph
- => Capture the role of nodes instead of proximity

STRUCT2VEC : DOUBLE ZKC



NODE CLASSIFICATION WITH EMBEDDINGS

NODE CLASSIFICATION

- To each node is associated a vector in the embedding
 - This vector corresponds to topological features of the node, used instead of, for instance, centralities
 - Both types of features can be combined
- As usual, a classifier can be trained using those features

NODE CLASSIFICATION

Algorithm	Dataset		
	BlogCatalog	PPI	Wikipedia
Spectral Clustering	0.0405	0.0681	0.0395
DeepWalk	0.2110	0.1768	0.1274
LINE	0.0784	0.1447	0.1164
<i>node2vec</i>	0.2581	0.1791	0.1552
<i>node2vec</i> settings (p,q)	0.25, 0.25	4, 1	4, 0.5
Gain of <i>node2vec</i> [%]	22.3	1.3	21.8

Controversies...

LINK PREDICTION WITH EMBEDDINGS

UNSUPERVISED LINK PREDICTION

- Unsupervised link prediction **from embeddings**
- => Compute the distance between nodes in the embedding
- => Use it as a similarity score

SUPERVISED LINK PREDICTION

- Supervised link prediction **from embeddings**
- => embeddings provide features for nodes (nb features: dimensions)
 - Combine nodes features to obtain edge features
- => Train a classifier to predict edges based on features from the embedding

SUPERVISED LINK PREDICTION

Operator	Result
Average	$(\mathbf{a} + \mathbf{b})/2$
Concat	$[\mathbf{a}_1, \dots, \mathbf{a}_d, \mathbf{b}_1, \dots, \mathbf{b}_d]$
Hadamard	$[\mathbf{a}_1 * \mathbf{b}_1, \dots, \mathbf{a}_d * \mathbf{b}_d]$
Weighted L1	$[\mathbf{a}_1 - \mathbf{b}_1 , \dots, \mathbf{a}_d - \mathbf{b}_d]$
Weighted L2	$[(\mathbf{a}_1 - \mathbf{b}_1)^2, \dots, (\mathbf{a}_d - \mathbf{b}_d)^2]$

Combining nodes vectors into edge vectors

SUPERVISED LINK PREDICTION

- How well does it work ?
- According to recent articles
 - Node2vec (2016)
 - VERSE (2018)
- => These methods are better than the state of the art

Op	Algorithm	Dataset		
		Facebook	PPI	arXiv
	Common Neighbors	0.8100	0.7142	0.8153
	Jaccard's Coefficient	0.8880	0.7018	0.8067
	Adamic-Adar	0.8289	0.7126	0.8315
	Pref. Attachment	0.7137	0.6670	0.6996
(a)	Spectral Clustering	0.5960	0.6588	0.5812
	DeepWalk	0.7238	0.6923	0.7066
	LINE	0.7029	0.6330	0.6516
	<i>node2vec</i>	0.7266	0.7543	0.7221
(b)	Spectral Clustering	0.6192	0.4920	0.5740
	DeepWalk	0.9680	0.7441	0.9340
	LINE	0.9490	0.7249	0.8902
	<i>node2vec</i>	0.9680	0.7719	0.9366
(c)	Spectral Clustering	0.7200	0.6356	0.7099
	DeepWalk	0.9574	0.6026	0.8282
	LINE	0.9483	0.7024	0.8809
	<i>node2vec</i>	0.9602	0.6292	0.8468
(d)	Spectral Clustering	0.7107	0.6026	0.6765
	DeepWalk	0.9584	0.6118	0.8305
	LINE	0.9460	0.7106	0.8862
	<i>node2vec</i>	0.9606	0.6236	0.8477

(a) Average, (b) Hadamard, (c) Weighted-L1, and (d) Weighted-L2

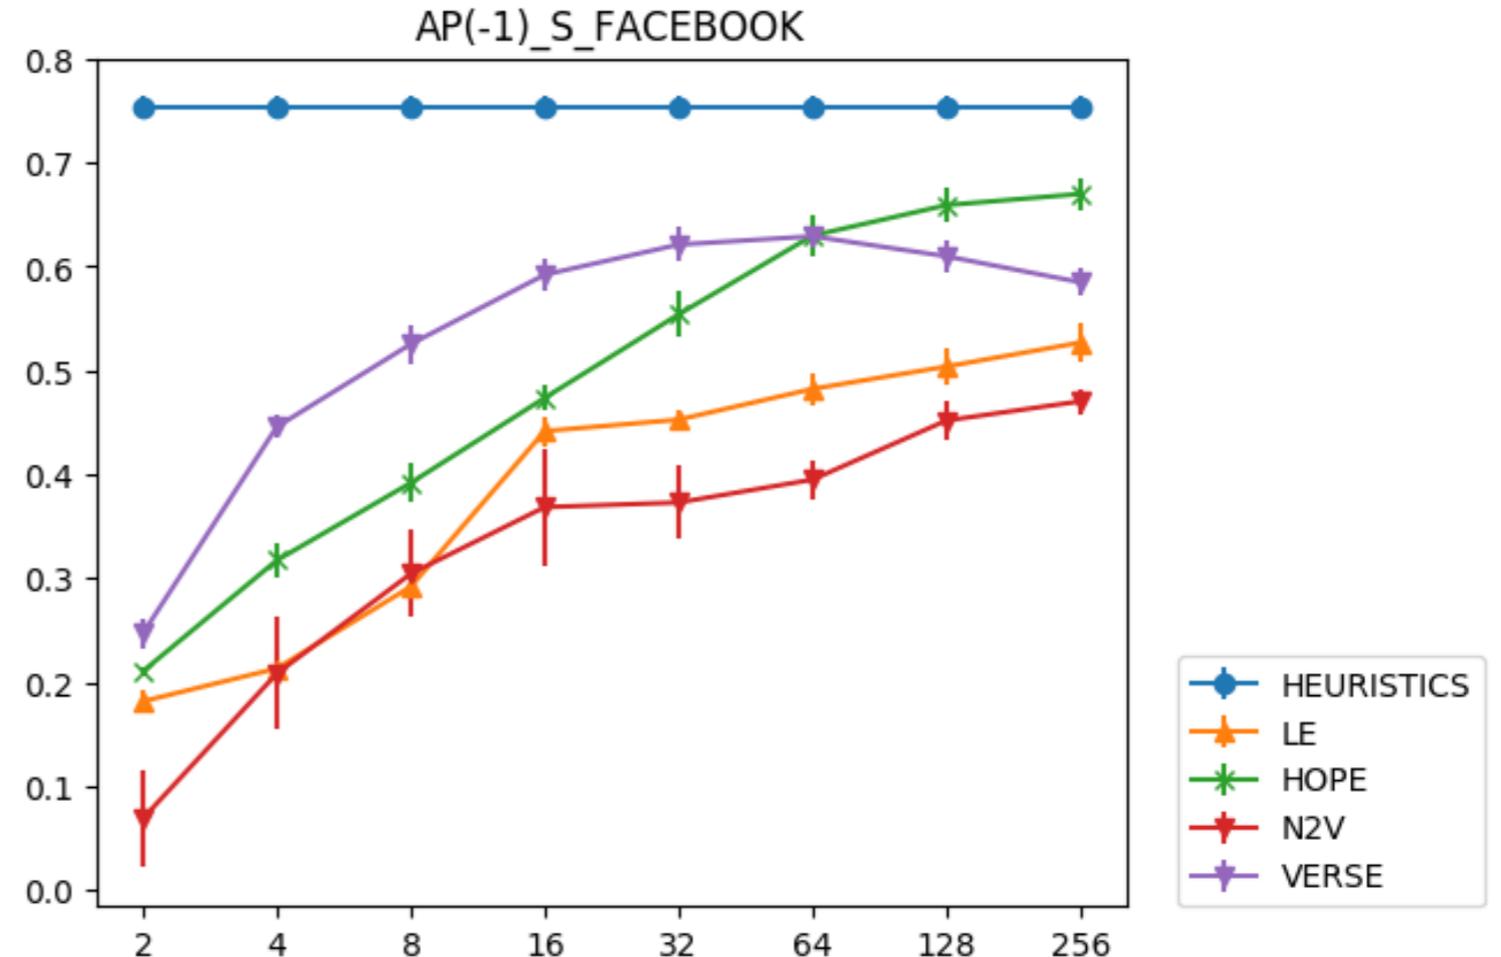
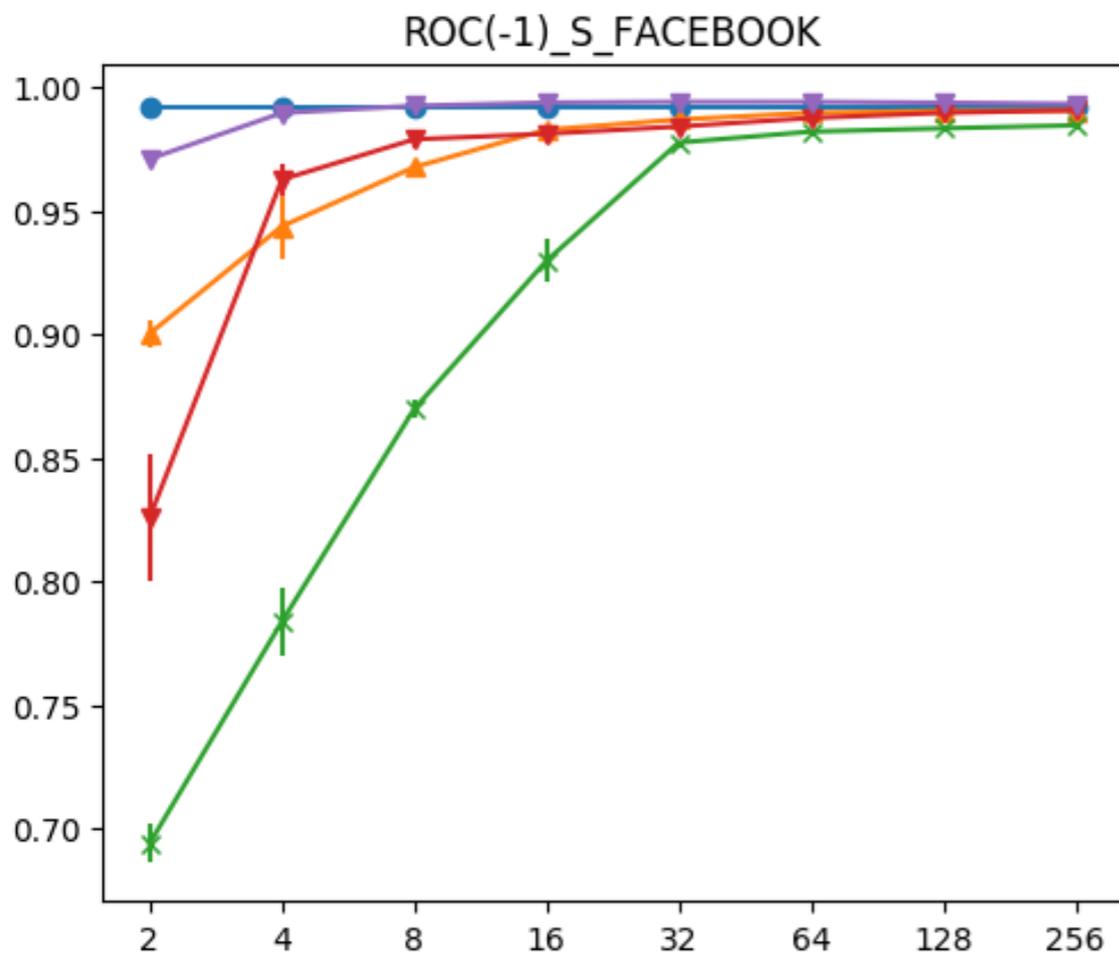
(AUC)

LINK PREDICTION

- Personal opinion: not that simple

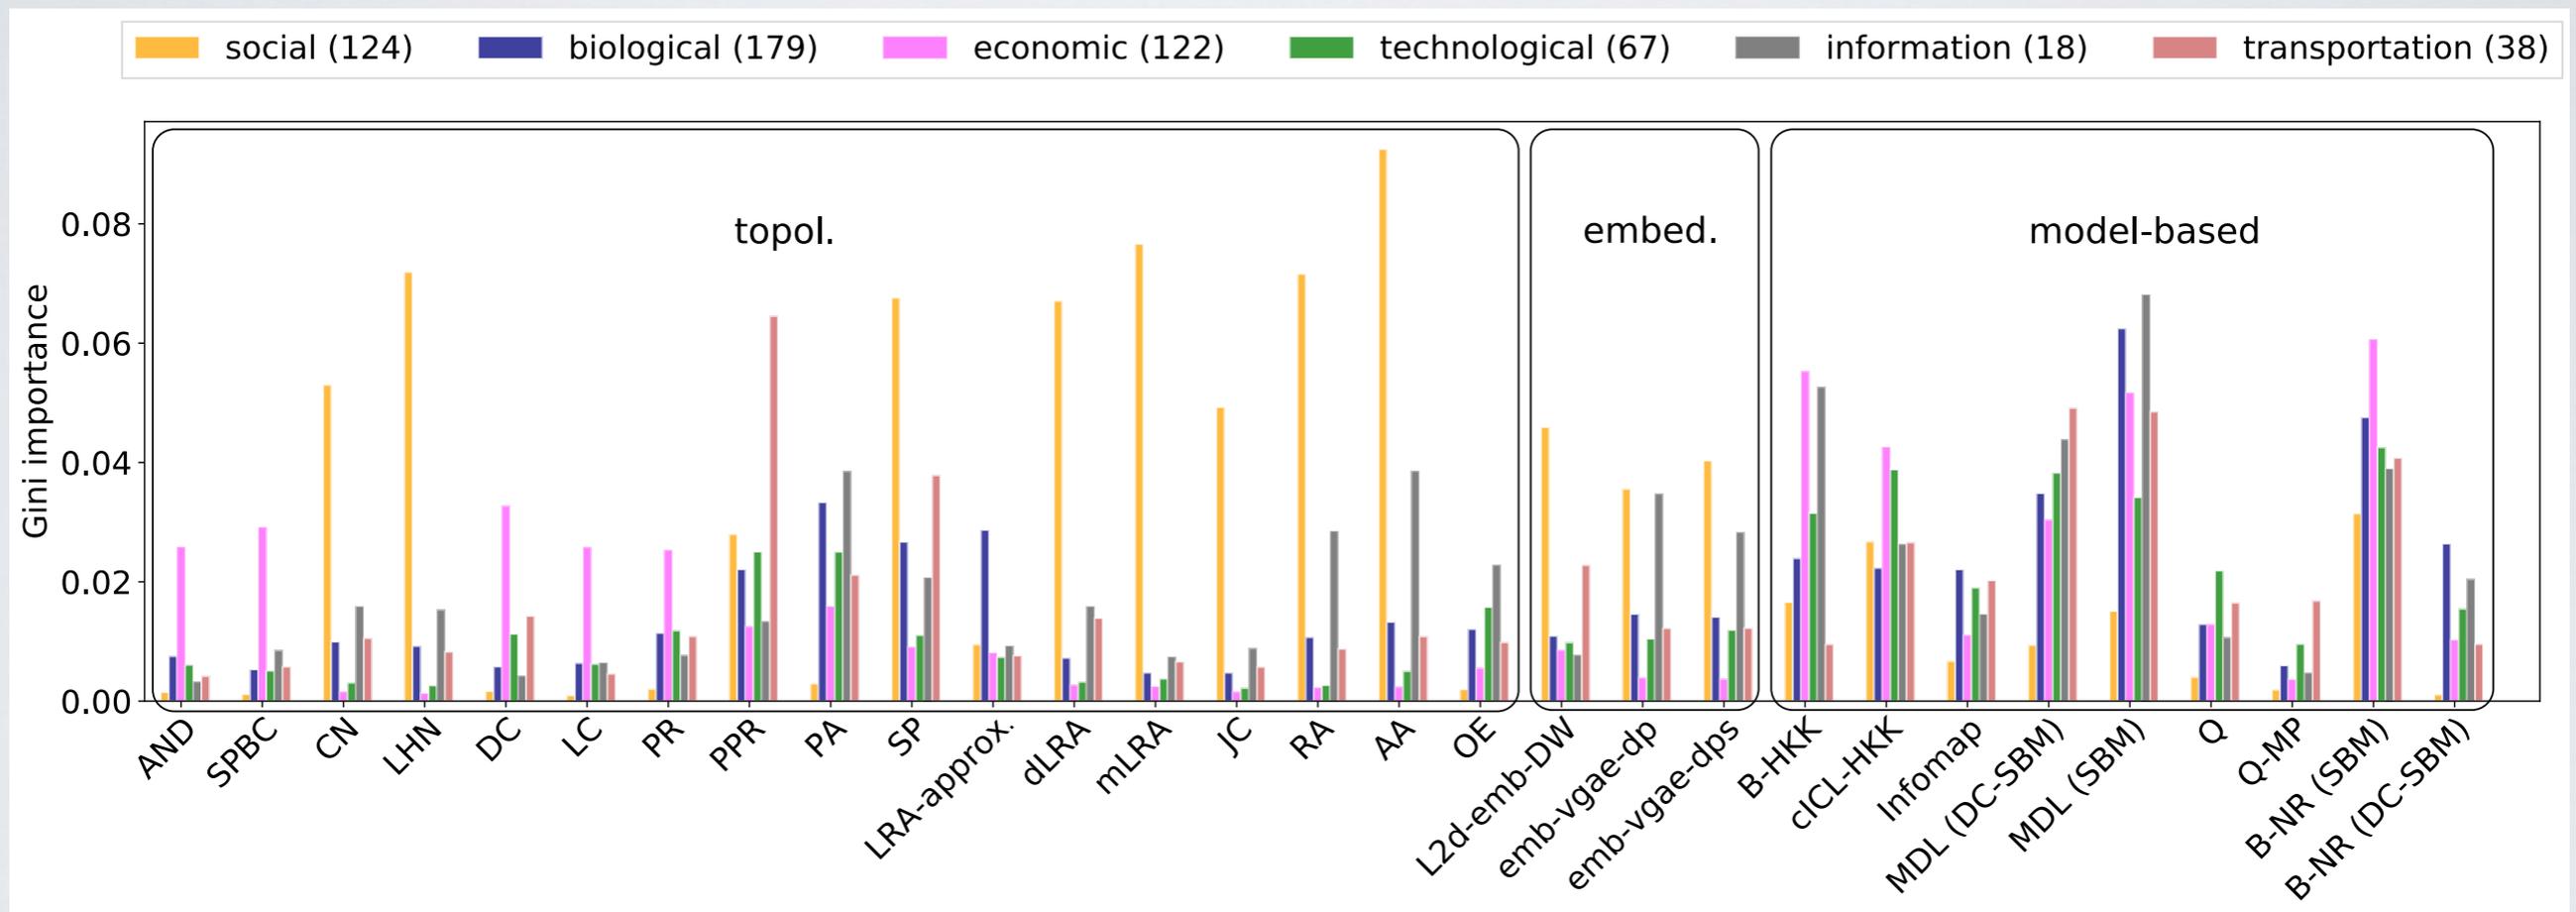
LINK PREDICTION

- Personal opinion: not that simple



MODEL STACKING

Ghasemian, A., Hosseinmardi, H., Galstyan, A., Airoidi, E. M., & Clauset, A. (2020). Stacking models for nearly optimal link prediction in complex networks. *Proceedings of the National Academy of Sciences*, 117(38), 23393-23400.



MODEL STACKING

Table S12. Average AUC, precision, and recall performances of the link prediction algorithms over 124 social networks as a subset of CommunityFitNet corpus. A random forest is used for supervised stacking of methods. Here, the predictors are adjusted for maximum F measure using a model selection through a cross validation on training set. The results are reported on 20% holdout test set.

Algorithm	AUC	Precision	Recall
Q	0.89 ± 0.07	0.42 ± 0.13	0.85 ± 0.08
Q-MR	0.87 ± 0.07	0.38 ± 0.16	0.78 ± 0.07
Q-MP	0.86 ± 0.08	0.25 ± 0.07	0.83 ± 0.09
B-NR (SBM)	0.93 ± 0.06	0.3 ± 0.08	0.85 ± 0.12
B-NR (DC-SBM)	0.93 ± 0.07	0.28 ± 0.08	0.88 ± 0.08
cICL-HKK	0.93 ± 0.08	0.34 ± 0.1	0.85 ± 0.14
B-HKK	0.88 ± 0.07	0.17 ± 0.05	0.79 ± 0.17
Infomap	0.91 ± 0.04	0.29 ± 0.08	0.83 ± 0.05
MDL (SBM)	0.94 ± 0.07	0.31 ± 0.09	0.87 ± 0.16
MDL (DC-SBM)	0.93 ± 0.09	0.26 ± 0.09	0.89 ± 0.11
S-NB	0.94 ± 0.07	0.3 ± 0.1	0.87 ± 0.08
mean model-based	0.91 ± 0.08	0.3 ± 0.12	0.84 ± 0.12
mean indiv. topol.	0.64 ± 0.19	0.2 ± 0.27	0.56 ± 0.33
mean indiv. topol. & model	0.7 ± 0.21	0.22 ± 0.25	0.62 ± 0.32
emd-DW	0.95 ± 0.1	0.45 ± 0.16	0.92 ± 0.13
emb-vgae	0.95 ± 0.08	0.09 ± 0.02	0.96 ± 0.09
all topol.	0.97 ± 0.08	0.89 ± 0.21	0.88 ± 0.2
all model-based	0.95 ± 0.07	0.76 ± 0.2	0.68 ± 0.17
all embed.	0.95 ± 0.11	0.75 ± 0.23	0.74 ± 0.23
all topol. & model	0.98 ± 0.06	0.89 ± 0.22	0.88 ± 0.19
all topol. & embed.	0.96 ± 0.1	0.86 ± 0.22	0.83 ± 0.25
all model & embed.	0.96 ± 0.09	0.78 ± 0.21	0.74 ± 0.22
all topol., model & embed.	0.97 ± 0.09	0.86 ± 0.23	0.84 ± 0.23

Table 1. Link prediction performance (mean ± std. err.), measured by AUC, precision, and recall, for link prediction algorithms applied to the 548 structurally diverse networks in our corpus.

algorithm	AUC	precision	recall
Q	0.7 ± 0.14	0.14 ± 0.17	0.67 ± 0.15
Q-MR	0.67 ± 0.15	0.12 ± 0.17	0.63 ± 0.13
Q-MP	0.64 ± 0.15	0.09 ± 0.11	0.59 ± 0.17
B-NR (SBM)	0.81 ± 0.13	0.13 ± 0.12	0.65 ± 0.22
B-NR (DC-SBM)	0.7 ± 0.2	0.12 ± 0.12	0.61 ± 0.24
cICL-HKK	0.79 ± 0.13	0.14 ± 0.14	0.58 ± 0.25
B-HKK	0.77 ± 0.13	0.11 ± 0.1	0.51 ± 0.26
Infomap	0.73 ± 0.14	0.12 ± 0.12	0.68 ± 0.13
MDL (SBM)	0.79 ± 0.15	0.14 ± 0.13	0.57 ± 0.3
MDL (DC-SBM)	0.84 ± 0.1	0.13 ± 0.11	0.78 ± 0.12
S-NB	0.71 ± 0.19	0.12 ± 0.13	0.66 ± 0.17
mean model-based	0.74 ± 0.16	0.12 ± 0.13	0.63 ± 0.21
mean indiv. topol.	0.6 ± 0.13	0.09 ± 0.16	0.53 ± 0.35
mean indiv. topol. & model	0.63 ± 0.15	0.09 ± 0.16	0.55 ± 0.33
emb-DW	0.63 ± 0.23	0.17 ± 0.19	0.42 ± 0.35
emb-vgae	0.69 ± 0.19	0.05 ± 0.05	0.69 ± 0.21
all topol.	0.86 ± 0.11	0.42 ± 0.33	0.44 ± 0.32
all model-based	0.83 ± 0.12	0.39 ± 0.34	0.3 ± 0.29
all embed.	0.77 ± 0.16	0.32 ± 0.32	0.32 ± 0.31
all topol. & model	0.87 ± 0.1	0.48 ± 0.36	0.35 ± 0.35
all topol. & embed.	0.84 ± 0.13	0.4 ± 0.34	0.39 ± 0.33
all model & embed.	0.84 ± 0.13	0.36 ± 0.32	0.36 ± 0.31
all topol., model & embed.	0.85 ± 0.14	0.42 ± 0.34	0.39 ± 0.33

GRAPH CONVOLUTIONAL NETWORKS

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2019). A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*.

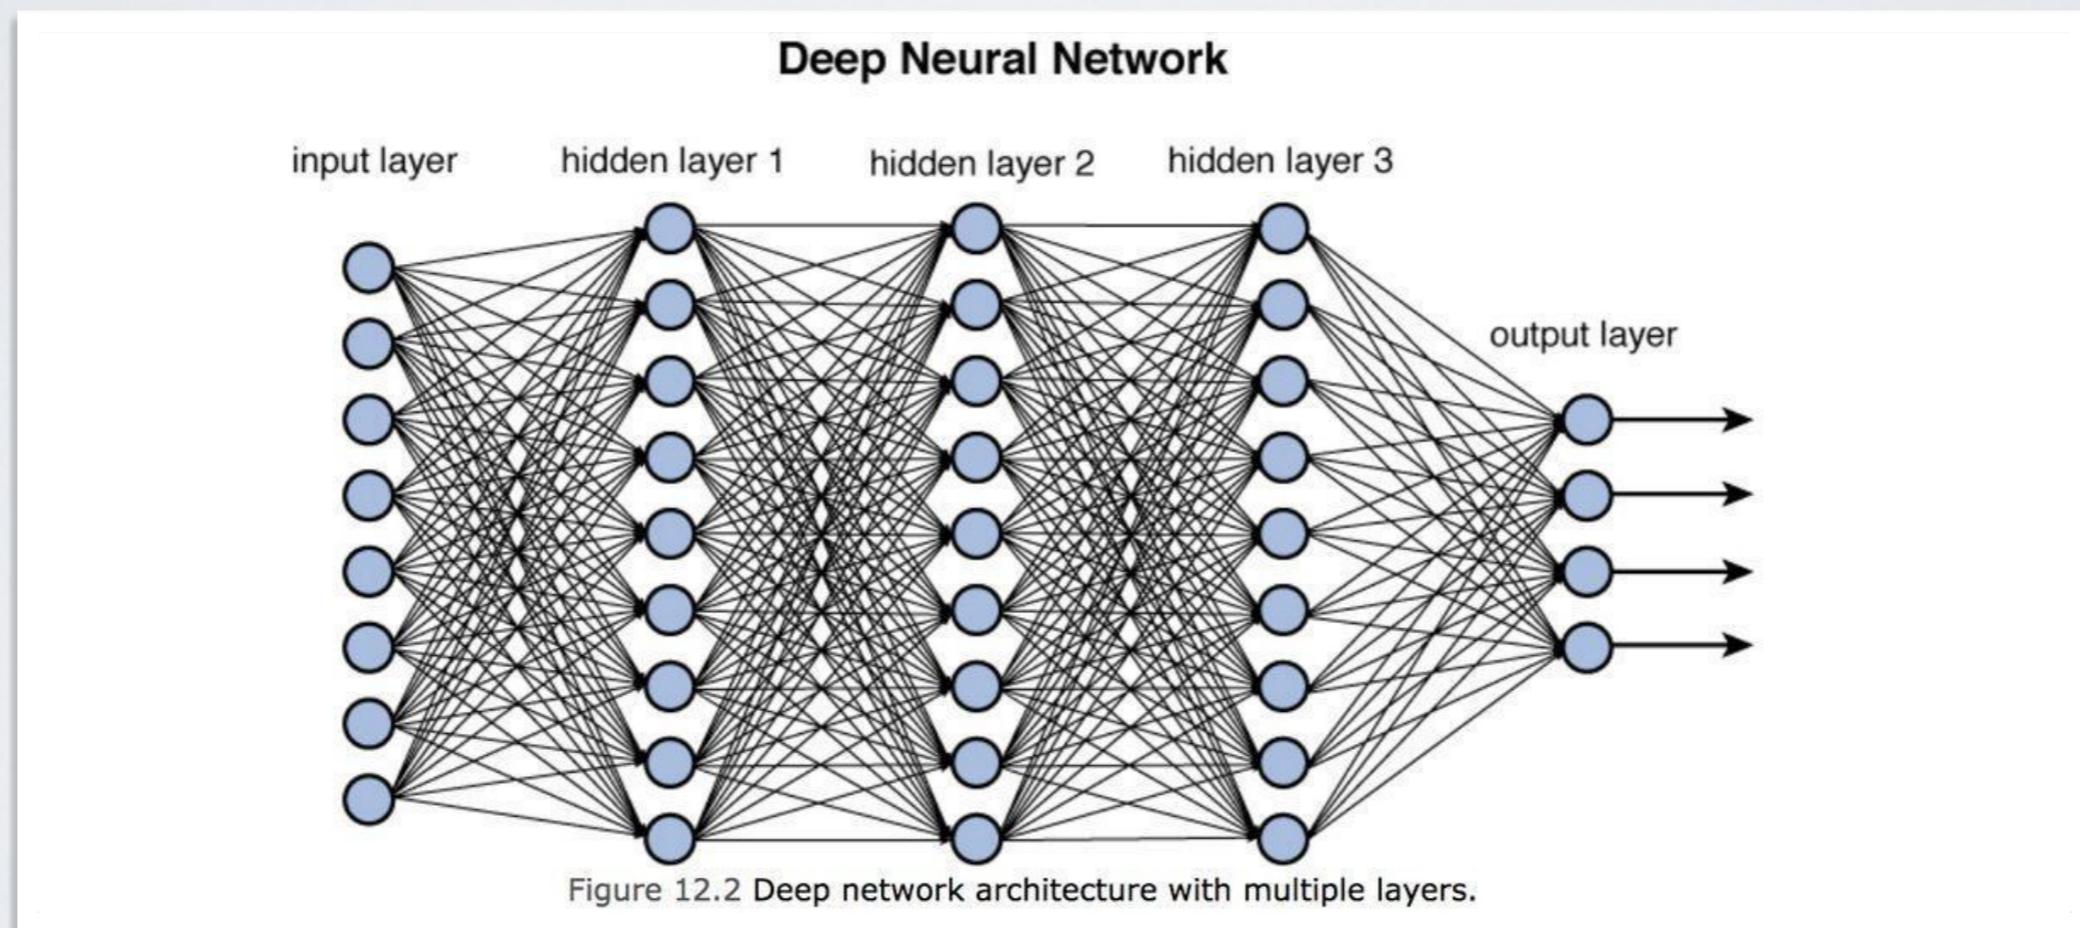
Zhang, Z., Cui, P., & Zhu, W. (2018). Deep learning on graphs: A survey. *arXiv preprint arXiv:1812.04202*.

Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

(DEEP) NEURAL NETWORKS

A deep neural networks can be seen as the chaining of multiple simple machine learning models (e.g., logistic classifier).

The output of a model is the input of the other, all weights optimized simultaneously (backpropagation)

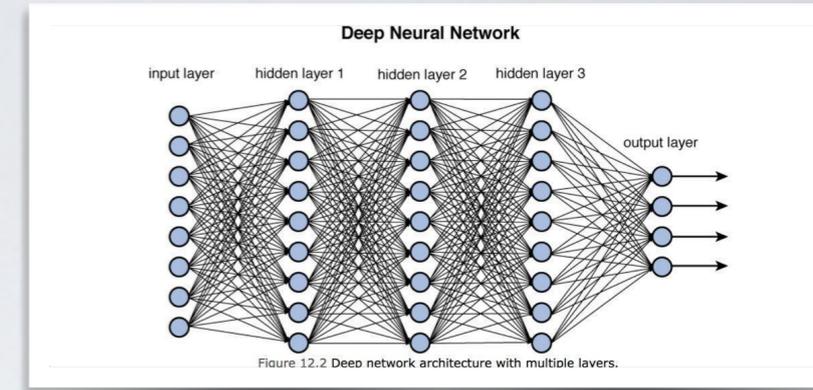


<https://medium.com/tebs-lab/introduction-to-deep-learning-a46e92cb0022>

<https://en.wikipedia.org/wiki/Backpropagation>

CONVOLUTIONAL NEURAL NETWORK

NETWORK



- All outputs of a layer connected to all inputs of the next is called **fully connected** layer
 - Learned weights will “cut” some edges (zero weights)
- In input data is structured, one can already use this structure
- **Convolutions** were introduced to work with pictures
 - Adjacency in pixels is meaningful

CONVOLUTION

1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image

4		

Convolved
Feature

▶ Extract “features” of “higher level”

- Pixels => lines, curves, dots => circles, long lines, curvy shapes => eye, hand, leaves => Animal, Car, sky ...

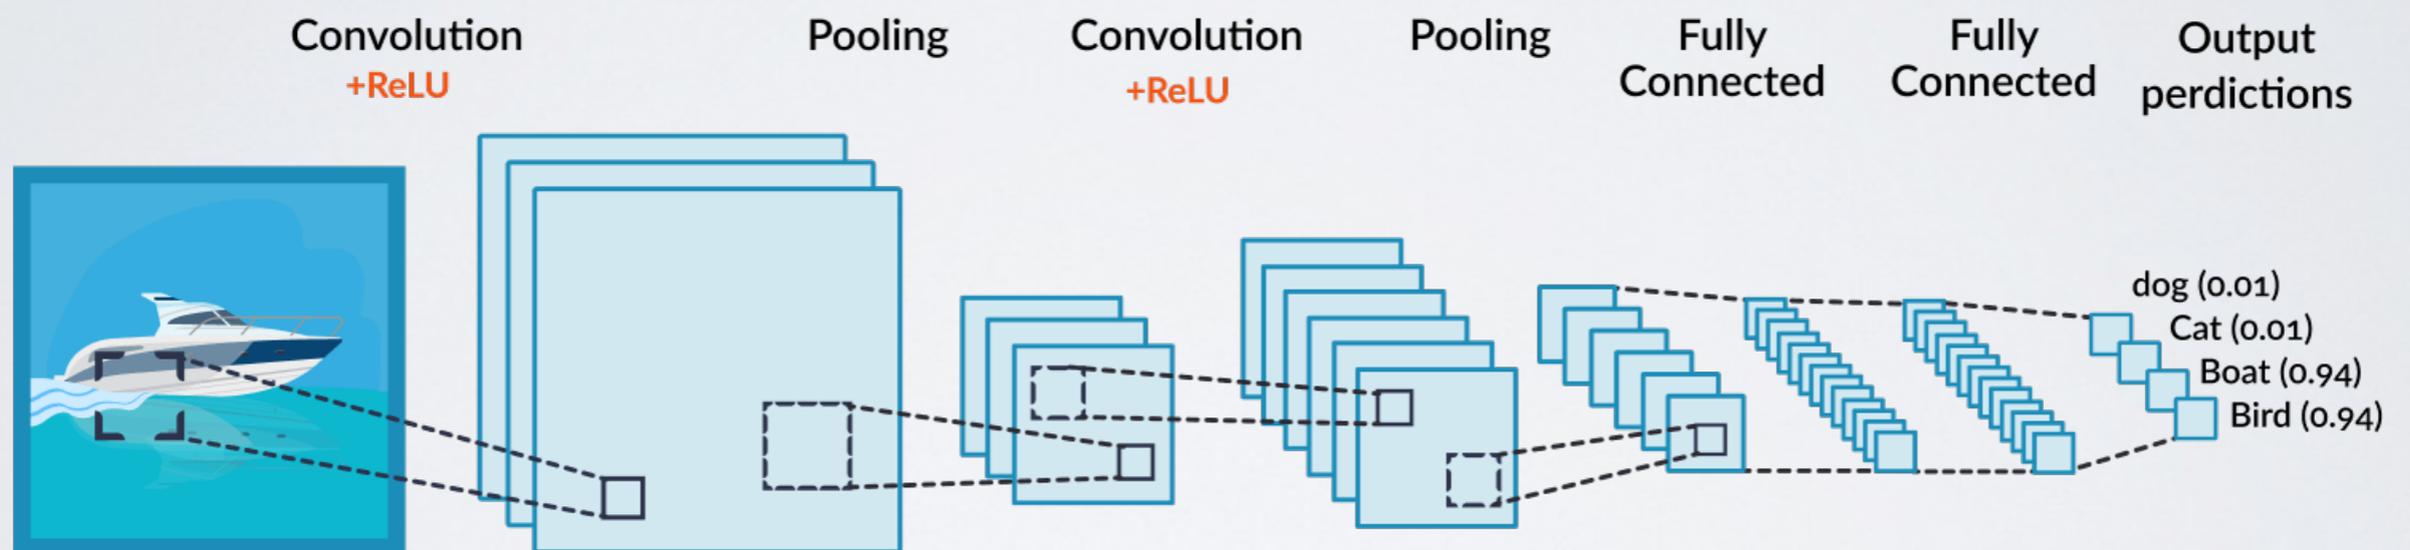
CONVOLUTION

- A convolution is defined by the weights of its kernel
- Which kernel(s) should we use?
- Weights of the kernel can be learnt, too

Identity	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	
Edge detection	$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$	
	$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$	
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	
Sharpen	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
Box blur (normalized)	$\frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$	
Gaussian blur 3 x 3 (approximation)	$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$	

[https://en.wikipedia.org/wiki/Kernel_\(image_processing\)](https://en.wikipedia.org/wiki/Kernel_(image_processing))

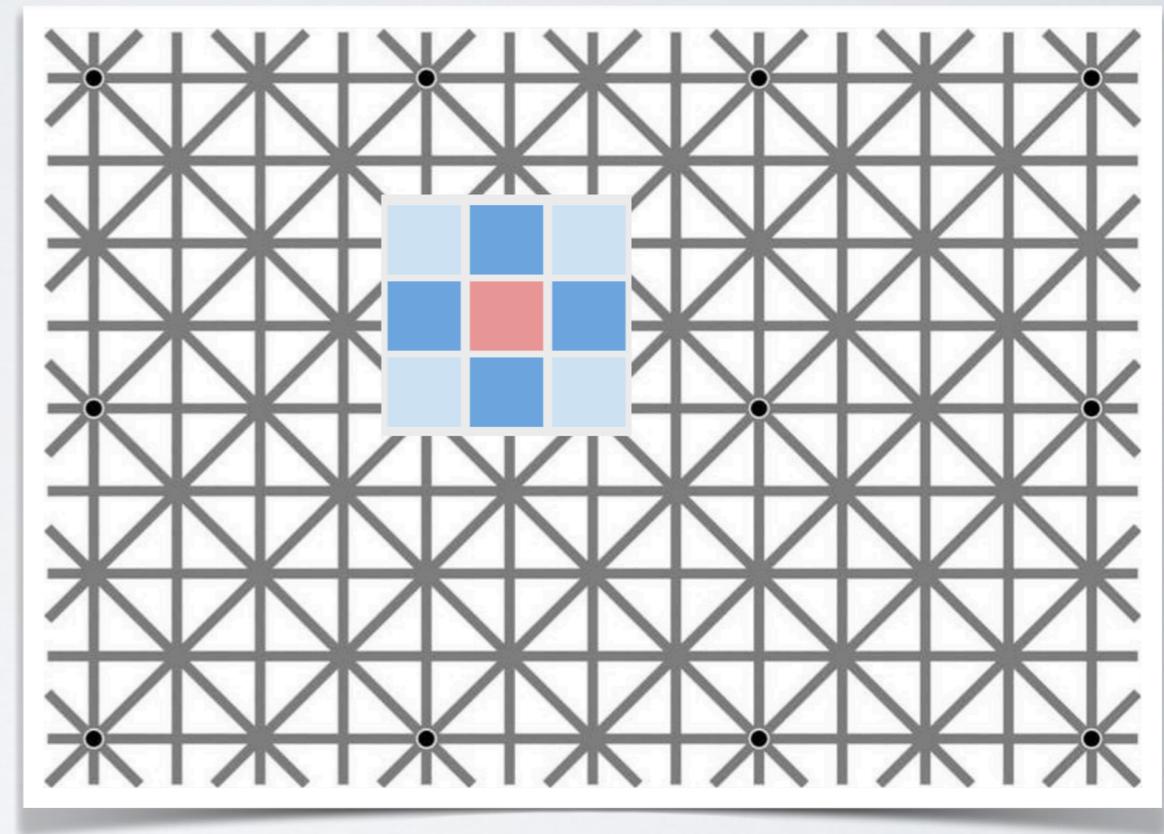
CONVOLUTIONAL NEURAL NETWORK



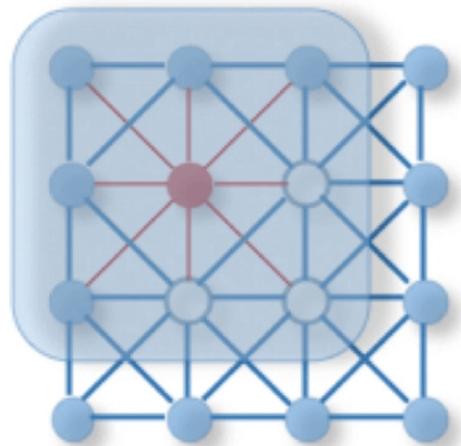
CONVOLUTIONAL NEURAL NETWORK

NETWORK

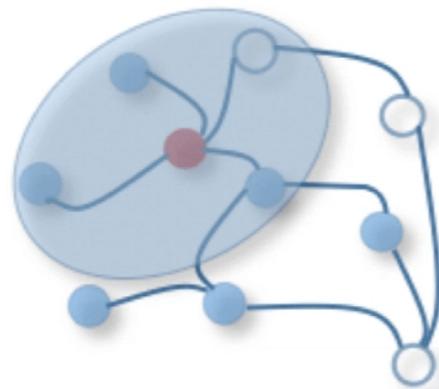
- Convolution on a picture can be seen as a special case of a graph operation:
 - Combine weights of neighbors
 - With an image represented as a regular grid
- Define convolutions on networks



GRAPH CONVOLUTION



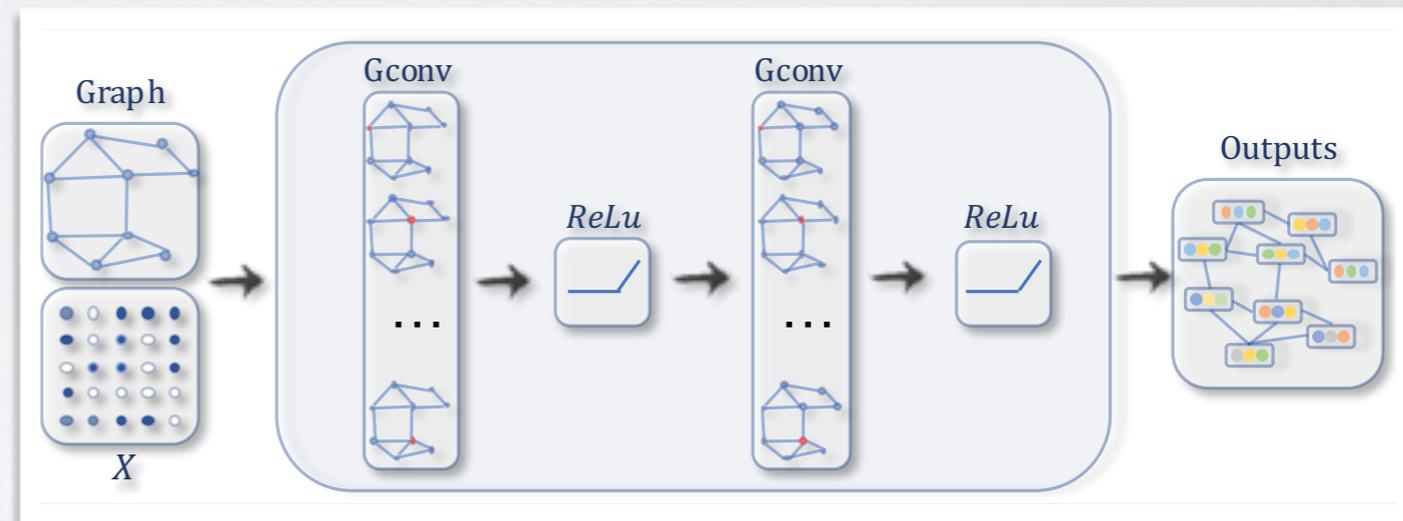
(a) 2D Convolution. Analogous to a graph, each pixel in an image is taken as a node where neighbors are determined by the filter size. The 2D convolution takes a weighted average of pixel values of the red node along with its neighbors. The neighbors of a node are ordered and have a fixed size.



(b) Graph Convolution. To get a hidden representation of the red node, one simple solution of graph convolution operation takes the average value of node features of the red node along with its neighbors. Different from image data, the neighbors of a node are unordered and variable in size.

Fig. 1: 2D Convolution vs. Graph Convolution.

Stacking convolution layers



GRAPH CONVOLUTION

$$H^{(l+1)} = f(H^{(l)}, A)$$

$$f(H^{(l)}, A) = \sigma \left(\hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right)$$

H : node features

A : adjacency matrix ($\hat{A} = A + I$)

l : layer index

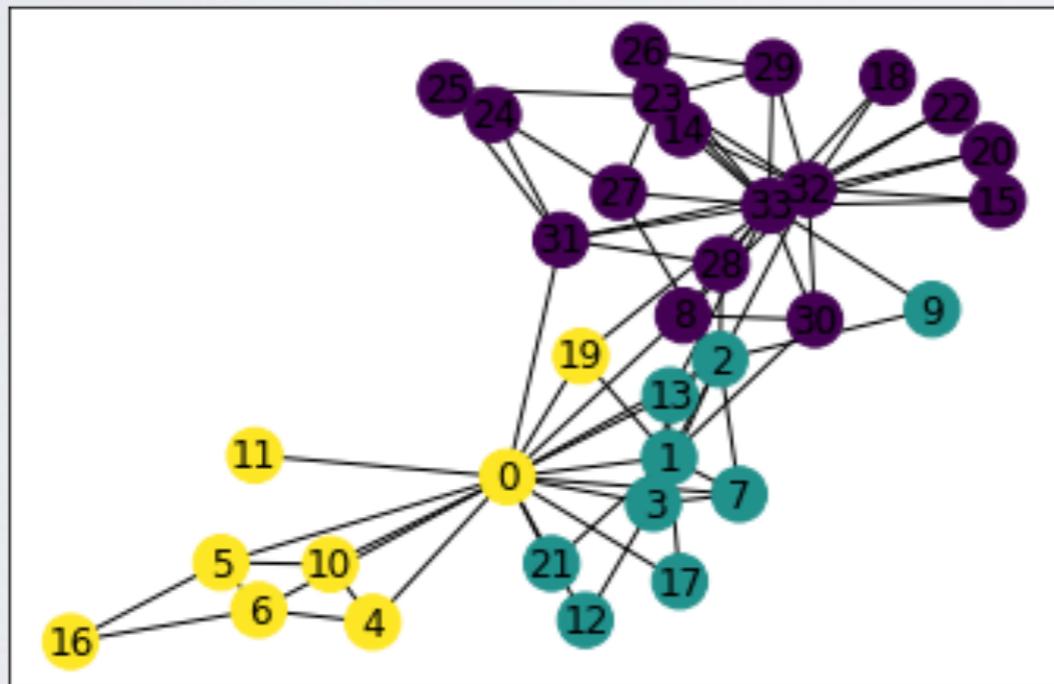
D : Degree matrix (degrees on the diagonal)

W : learnable weights

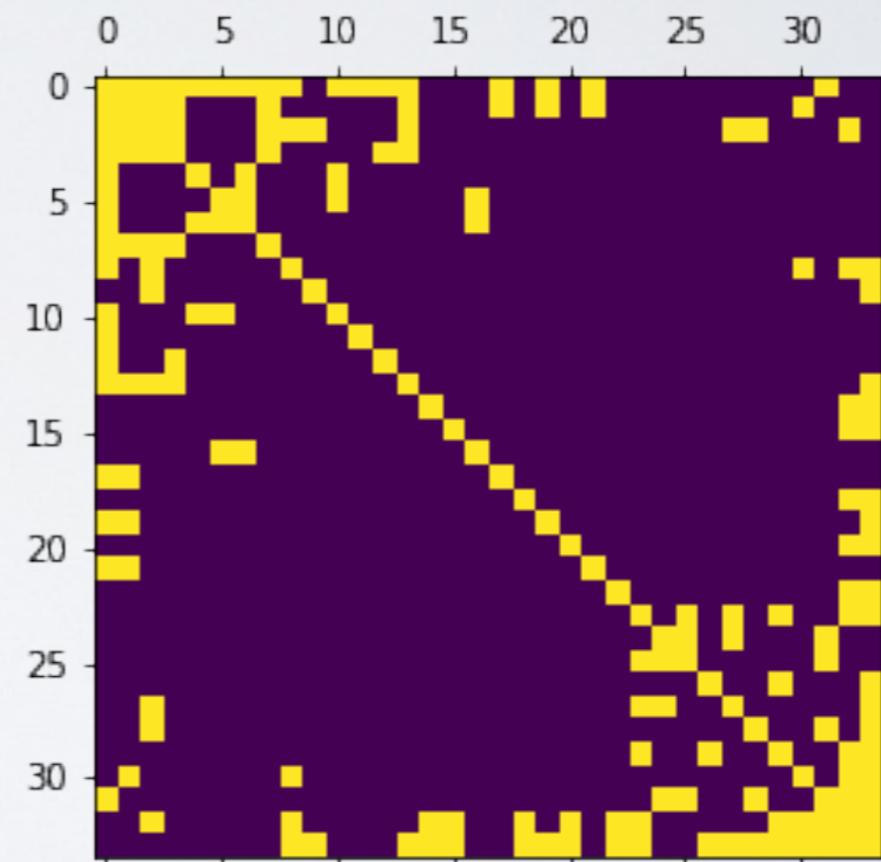
σ : activation function (often ReLU)

GRAPH CONVOLUTION

- Going through an example of the typical GCN

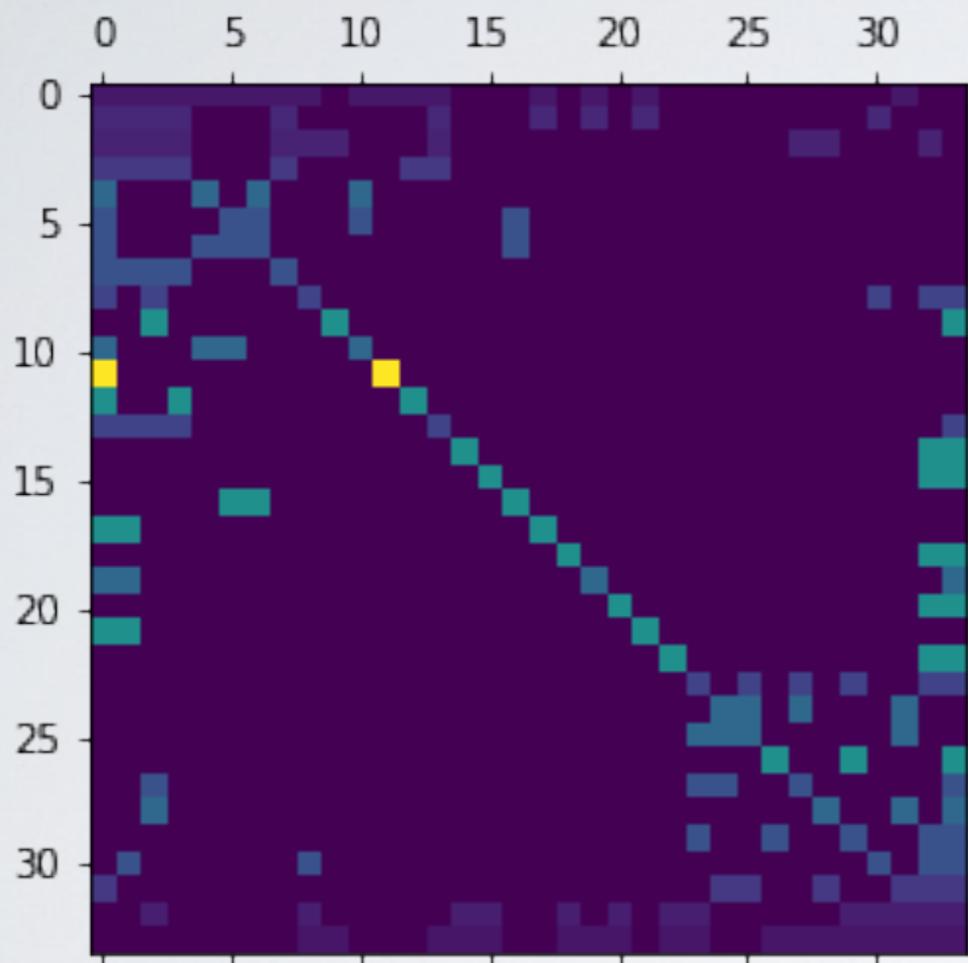


Zackary Karate club
(with communities for reference)

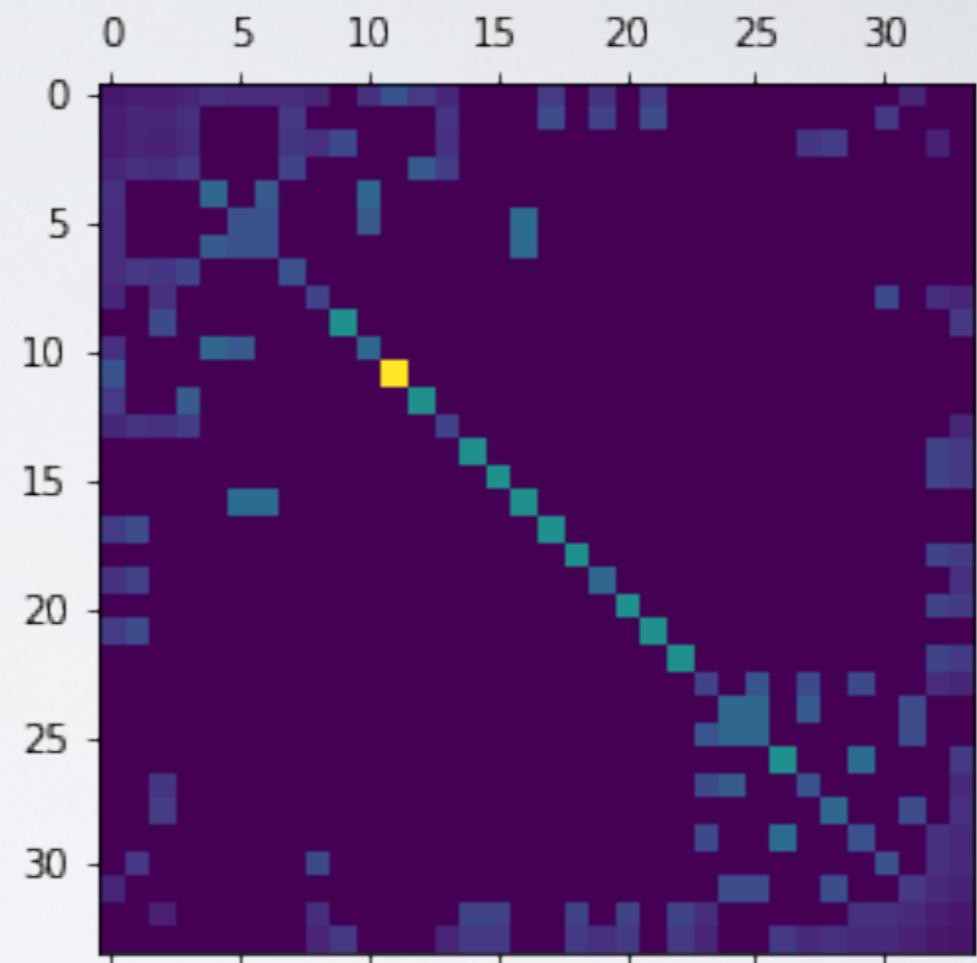


\hat{A}

GRAPH CONVOLUTION



$D^{-1}\hat{A}$
Simple average



$D^{-\frac{1}{2}}\hat{A}D^{-\frac{1}{2}}$
Weighted average

Normalisation of the adjacency matrix

GRAPH CONVOLUTION

$$f(H^{(l)}, A) = \sigma \left(\hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right)$$

$$D^{-\frac{1}{2}} \hat{A} D^{-\frac{1}{2}} H$$

Features of the nodes become the (weighted) average of the features of the neighbors

W has shape $(X \times Y)$, with X the number of features in input and Y the **desired** number of features in output

GRAPH CONVOLUTION

$$f(H^{(l)}, A) = \sigma \left(\hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right)$$

Size of the weight matrices by layer

$$W_0 : d_0 \times d_1$$

$$W_1 : d_1 \times d_2$$

...

$$W_n : d_n \times d_{n+1}$$

d_0 is the number of features per node in the original network data,
 d_{n+1} is the number of desired features (usually followed by a normal classifier, e.g., logistic)

GRAPH CONVOLUTION

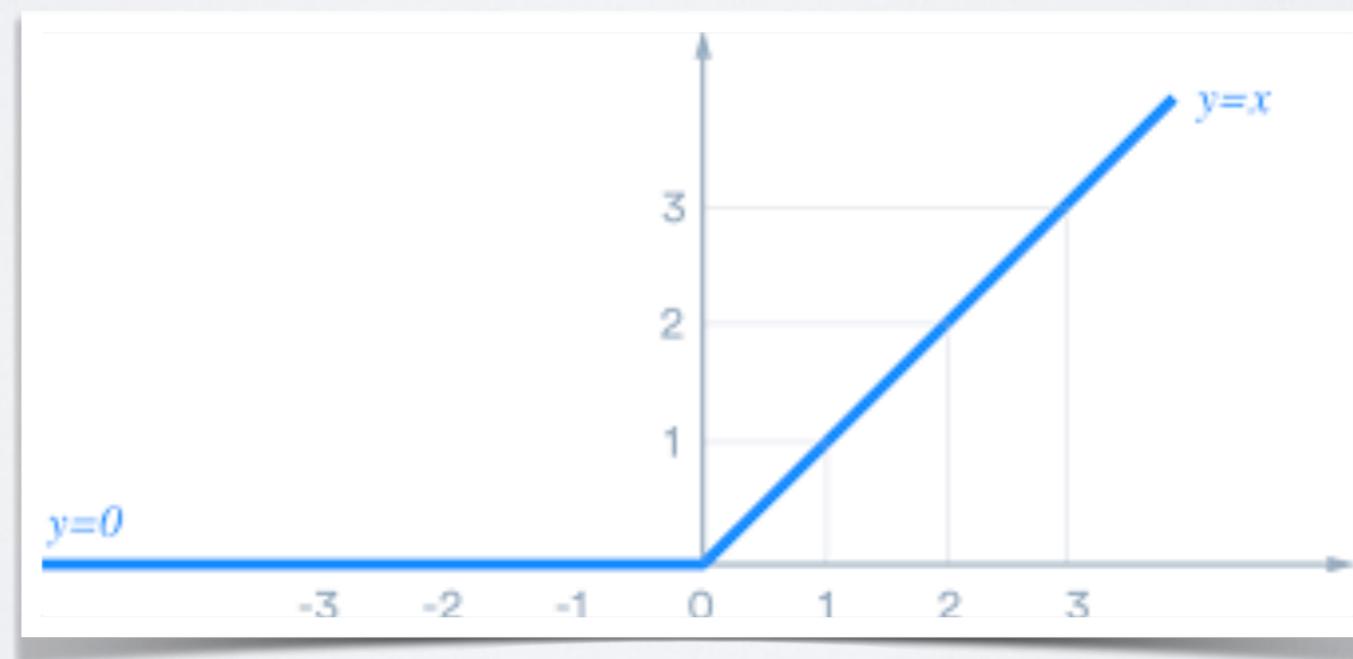
$$f(H^{(l)}, A) = \sigma \left(\hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right)$$

σ is called an activation function.

It is used to introduce non-linearity.

As of 2019, the most common choice is to use the **ReLU**,
(Rectified Linear Unit)

=> Simple to differentiate and to compute



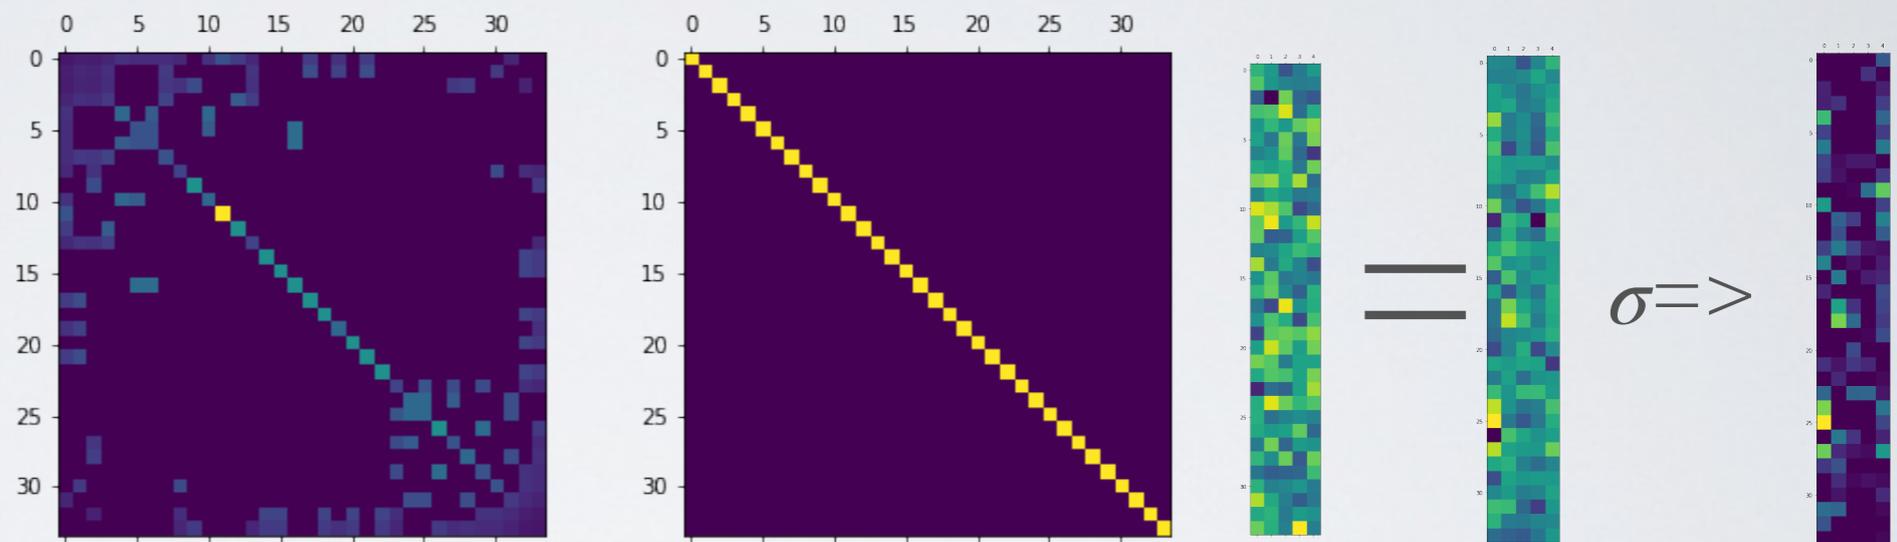
FORWARD STEP

- We can first look at what happens **without weight learning**, i.e., doing only the forward step.
- We set the original features to the identity matrix, $H_0 = I$. Each node's features is a *one hot vector* of itself (1 at its position, 0 otherwise)
- Weights are random (normal distribution centered on 0)
- Two layers, with W sizes $n \times 5, 5 \times 2$

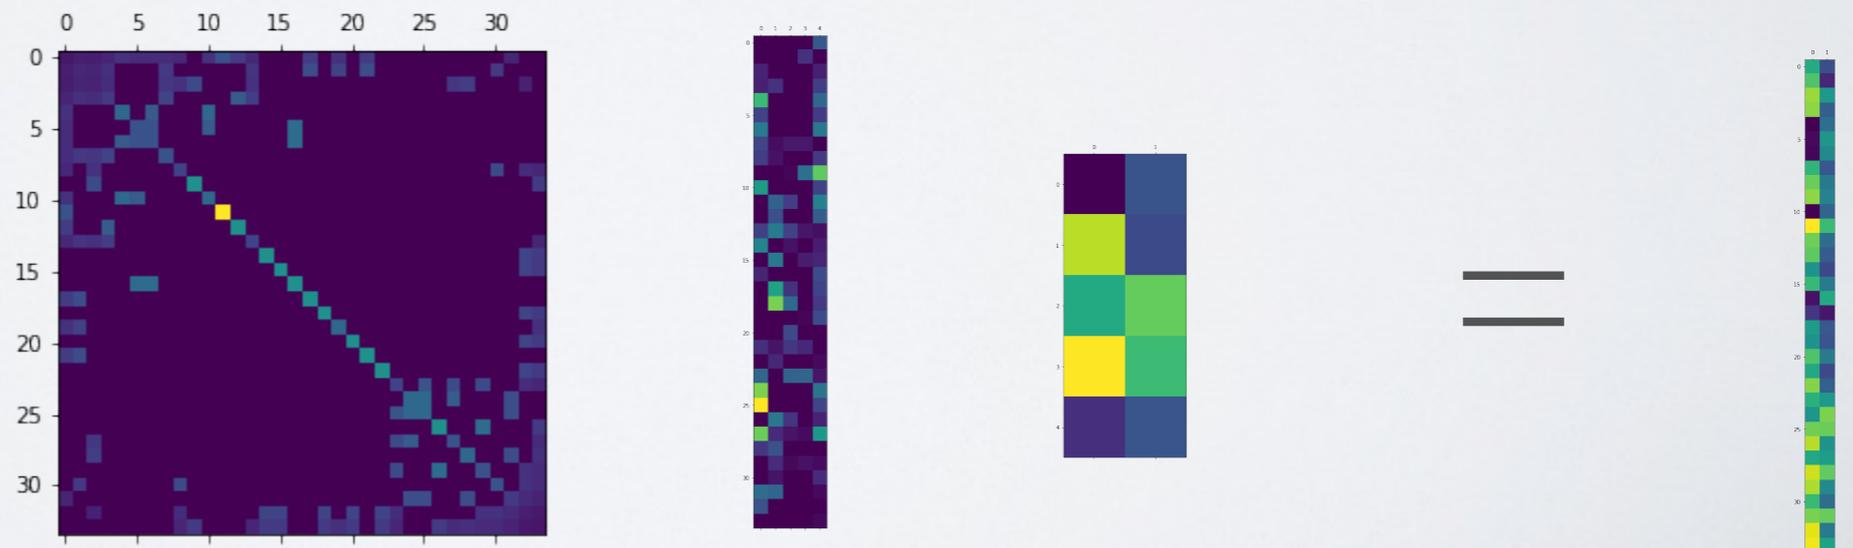
FORWARD STEP

$$f(H^{(l)}, A) = \sigma \left(\hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right)$$

LI = n to 5 features

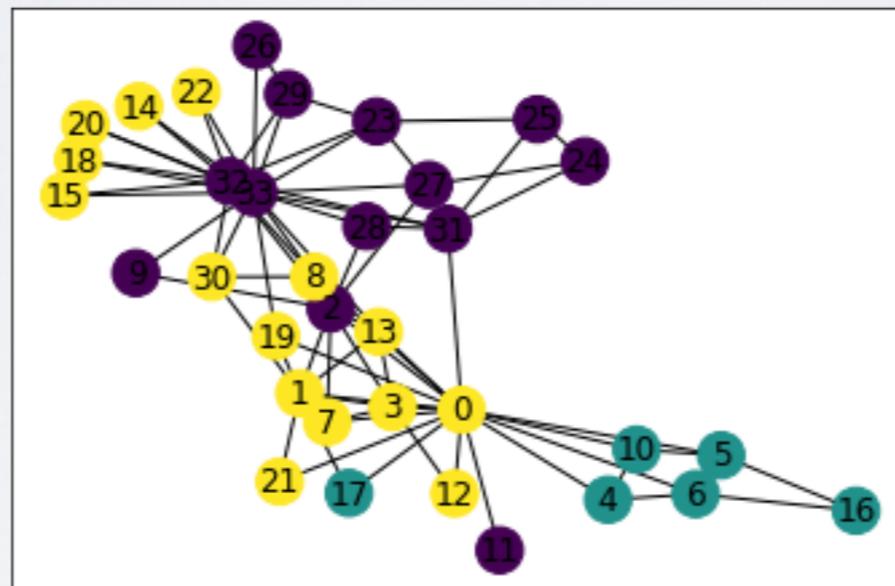


LI = 5 to 2 features



FORWARD STEP

K-means on the 2D “embedding”
(parameter $k=3$ clusters)



(Node positions based on spring layout)

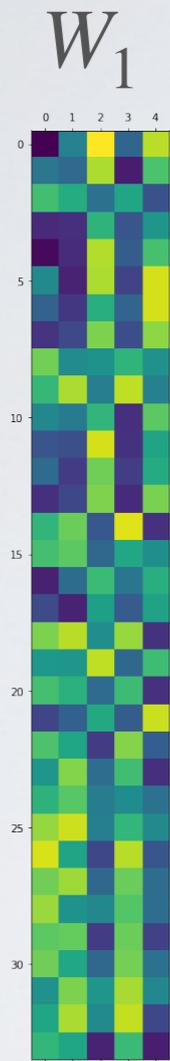
BACKWARD STEP

- To learn the weights, we use a mechanism called **back-propagation**
- Short summary
 - A **loss** function is defined to compare the “predicted values” with ground truth labels (at this point, we need some labels...)
 - Typically, log-likelihood
 - The **derivative** of the cost function relative to weights is computed
 - Weights are updated using **gradient descent** (i.e., weights are modified in the direction that will minimize the loss)

FITTING THE GCN

- We define the same GCN as before
- We define a “semi-supervised” process:
 - Labels are known only for a few nodes (the 2 instructors)
 - The loss is computed only for them
- We run e steps (“epoch”) of back-propagation, until convergence

FITTING THE GCN

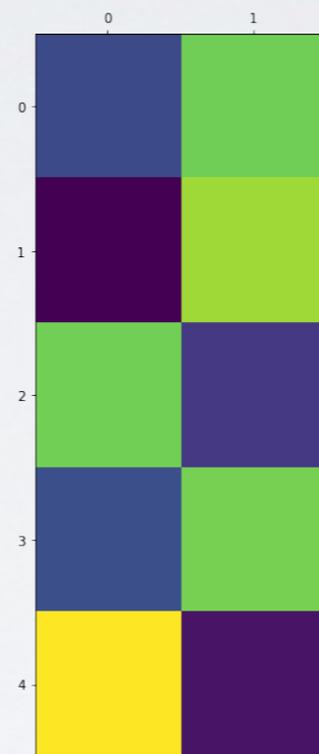


Step 1:

Each node takes the average features of its neighbors.

W_1 can be seen as “computed” features (this is because we used I as original features)

W_2

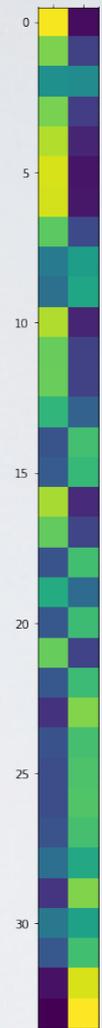


Step 2:

After averaging over results of step 1 (AH), each node combines its aggregated features according to this matrix

68

H



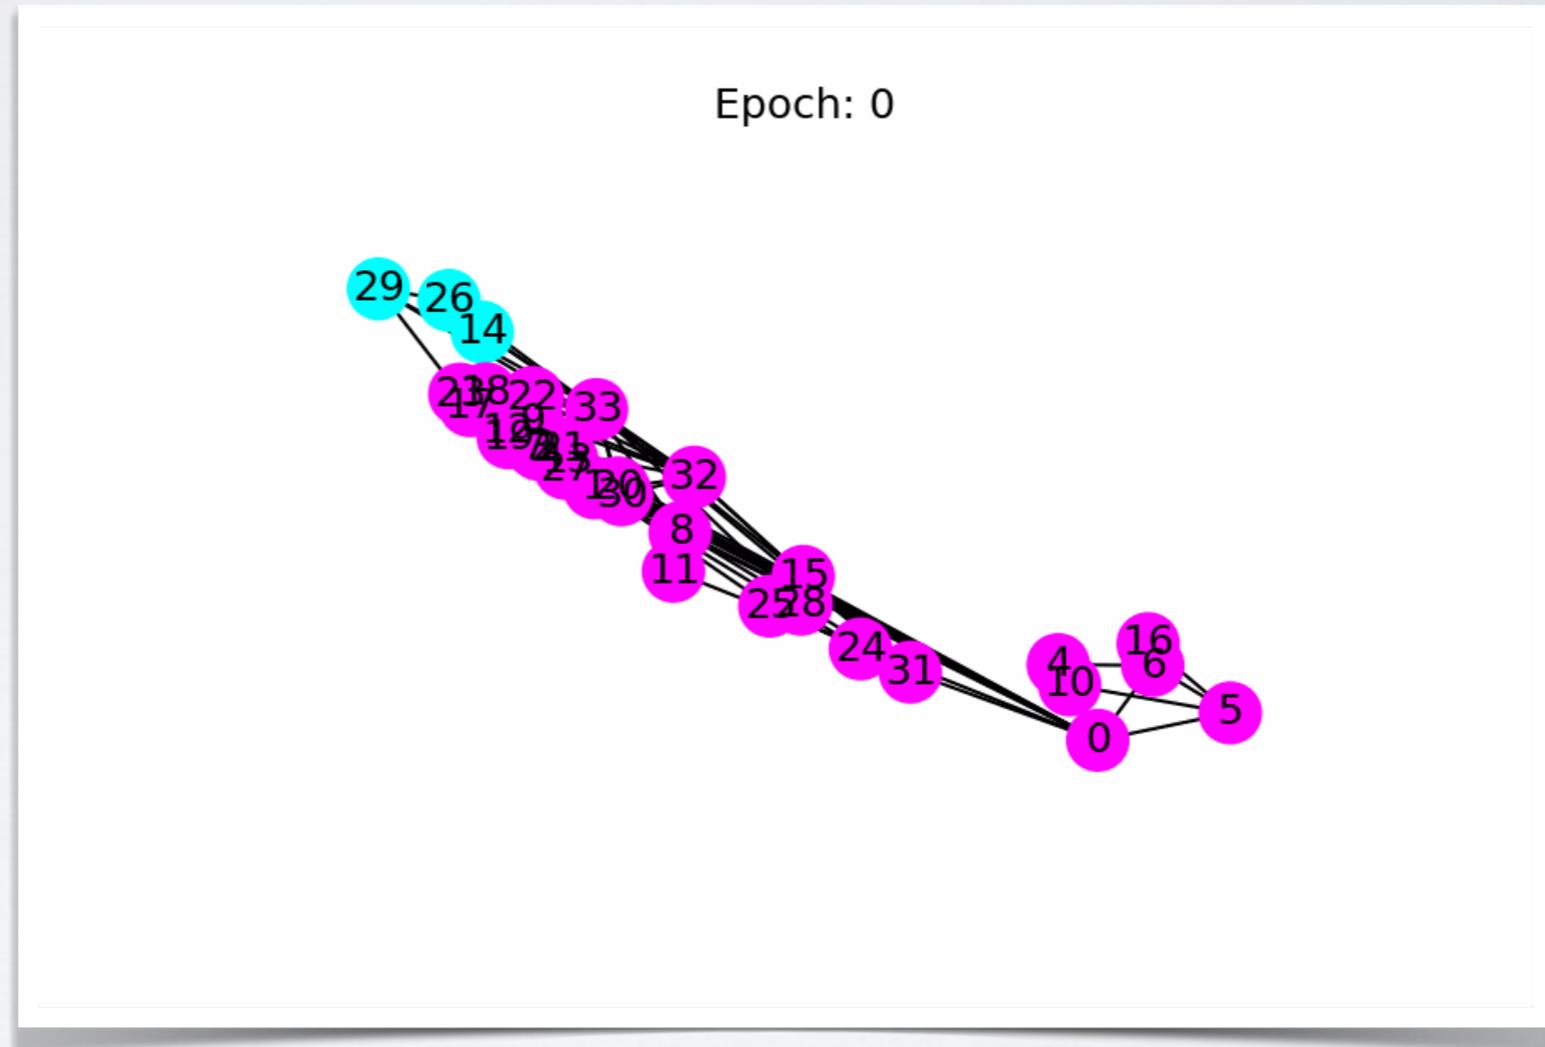
Result:

This is the computed feature vector.

As expected, values for nodes 0 and 33 are opposed

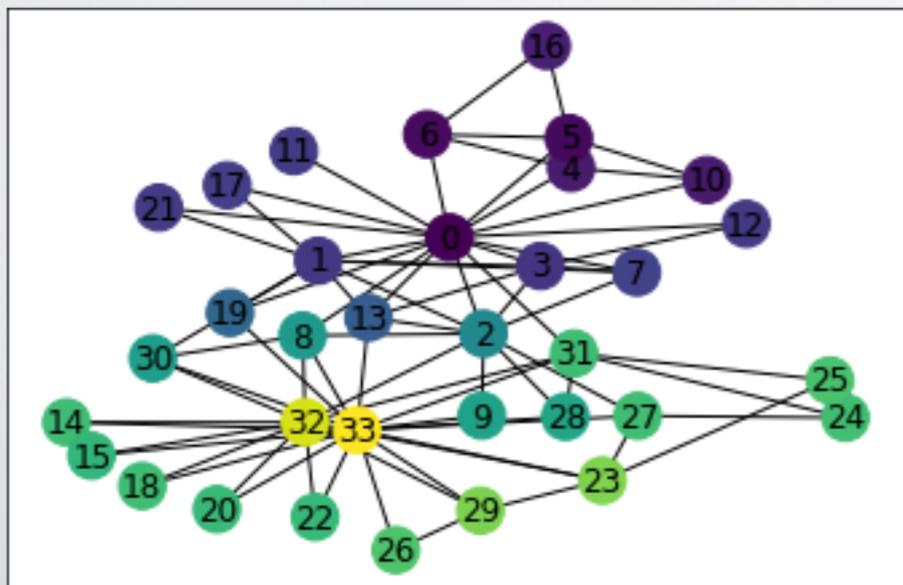
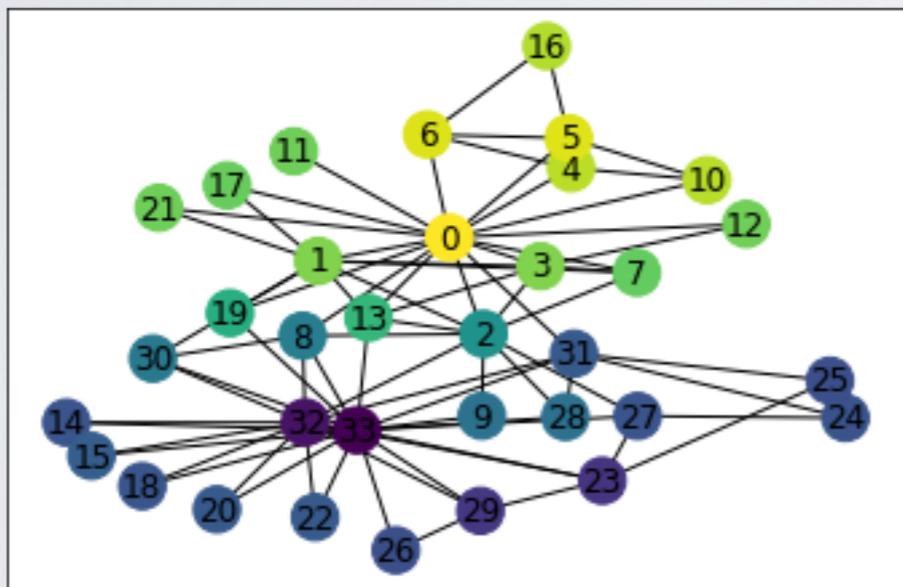
FITTING THE GCN

Epoch 0	Loss: 0.6987
Epoch 1	Loss: 0.6804
Epoch 2	Loss: 0.6634
Epoch 3	Loss: 0.6476
Epoch 4	Loss: 0.6326
Epoch 5	Loss: 0.6174
Epoch 6	Loss: 0.6017
Epoch 7	Loss: 0.5852
Epoch 8	Loss: 0.5684
Epoch 9	Loss: 0.5513
Epoch 10	Loss: 0.5338
Epoch 11	Loss: 0.5158
Epoch 12	Loss: 0.4976
Epoch 13	Loss: 0.4792
Epoch 14	Loss: 0.4605
Epoch 15	Loss: 0.4416
Epoch 16	Loss: 0.4225
Epoch 17	Loss: 0.4033
Epoch 18	Loss: 0.3842
Epoch 19	Loss: 0.3652
Epoch 20	Loss: 0.3464
Epoch 21	Loss: 0.3279
Epoch 22	Loss: 0.3096
Epoch 23	Loss: 0.2916
Epoch 24	Loss: 0.2741
Epoch 25	Loss: 0.2571
Epoch 26	Loss: 0.2407
Epoch 27	Loss: 0.2248
Epoch 28	Loss: 0.2095
Epoch 29	Loss: 0.1946
Epoch 30	Loss: 0.1803
Epoch 31	Loss: 0.1668
Epoch 32	Loss: 0.1541
Epoch 33	Loss: 0.1422
Epoch 34	Loss: 0.1312
Epoch 35	Loss: 0.1209
Epoch 36	Loss: 0.1113
Epoch 37	Loss: 0.1024
Epoch 38	Loss: 0.0940
Epoch 39	Loss: 0.0863
Epoch 40	Loss: 0.0793
Epoch 41	Loss: 0.0727
Epoch 42	Loss: 0.0667
Epoch 43	Loss: 0.0611
Epoch 44	Loss: 0.0560
Epoch 45	Loss: 0.0513
Epoch 46	Loss: 0.0470
Epoch 47	Loss: 0.0432
Epoch 48	Loss: 0.0396
Epoch 49	Loss: 0.0363
Epoch 50	Loss: 0.0333

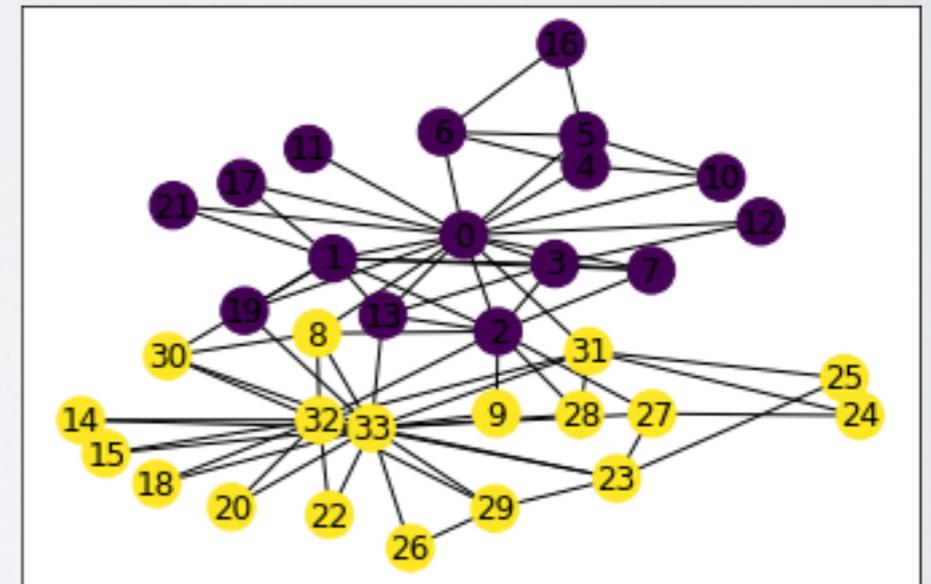


RESULTS

Features values



Highest feature as label



We retrieve the expected "communities"

GCN LITERATURE

- Results are claimed to be above the state of the art
 - Controversies, which is normal for such recent methods

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)
GCN (this paper)	70.3 (7s)	81.5 (4s)	79.0 (38s)	66.0 (48s)
GCN (rand. splits)	67.9 ± 0.5	80.1 ± 0.5	78.9 ± 0.7	58.4 ± 1.7

TO CONCLUDE

Many variations proposed already

Very active since 2017

Spawned renewed interest in networks in the ML literature

Hard to predict the future of these techniques.

Approach	Category	Inputs	Pooling	Readout	Time Complexity
GNN* (2009) [15]	RecGNN	A, X, X^e	-	a dummy super node	-
GraphESN (2010) [16]	RecGNN	A, X	-	mean	-
GGNN (2015) [17]	RecGNN	A, X	-	attention sum	-
SSE (2018) [18]	RecGNN	A, X	-	-	-
Spectral CNN (2014) [19]	Spectral-based ConvGNN	A, X	spectral clustering+max pooling	max	$O(n^3)$
Henaff et al. (2015) [20]	Spectral-based ConvGNN	A, X	spectral clustering+max pooling	-	$O(n^3)$
ChebNet (2016) [21]	Spectral-based ConvGNN	A, X	efficient pooling	sum	$O(m)$
GCN (2017) [22]	Spectral-based ConvGNN	A, X	-	-	$O(m)$
CayleyNet (2017) [23]	Spectral-based ConvGNN	A, X	mean/gracius pooling	-	$O(m)$
AGCN (2018) [40]	Spectral-based ConvGNN	A, X	max pooling	sum	$O(n^2)$
DualGCN (2018) [41]	Spectral-based ConvGNN	A, X	-	-	$O(m)$
NN4G (2009) [24]	Spatial-based ConvGNN	A, X	-	sum/mean	$O(m)$
DCNN (2016) [25]	Spatial-based ConvGNN	A, X	-	mean	$O(n^2)$
PATCHY-SAN (2016) [26]	Spatial-based ConvGNN	A, X, X^e	-	concat	-
MPNN (2017) [27]	Spatial-based ConvGNN	A, X, X^e	-	attention sum/ set2set	$O(m)$
GraphSage (2017) [42]	Spatial-based ConvGNN	A, X	-	-	-
GAT (2017) [43]	Spatial-based ConvGNN	A, X	-	-	$O(m)$
MoNet (2017) [44]	Spatial-based ConvGNN	A, X	-	-	$O(m)$
PGC-DGCNN (2018) [46]	Spatial-based ConvGNN	A, X	sort pooling	attention sum	$O(n^3)$
CGMM (2018) [47]	Spatial-based ConvGNN	A, X	-	concat	-
LGCN (2018) [45]	Spatial-based ConvGNN	A, X	-	-	-
GAAN (2018) [48]	Spatial-based ConvGNN	A, X	-	-	$O(m)$
FastGCN (2018) [49]	Spatial-based ConvGNN	A, X	-	-	-
StoGCN (2018) [50]	Spatial-based ConvGNN	A, X	-	-	-
Huang et al. (2018) [51]	Spatial-based ConvGNN	A, X	-	-	-
DGCNN (2018) [52]	Spatial-based ConvGNN	A, X	sort pooling	-	$O(m)$
DiffPool (2018) [54]	Spatial-based ConvGNN	A, X	differential pooling	mean	$O(n^2)$
GeniePath (2019) [55]	Spatial-based ConvGNN	A, X	-	-	$O(m)$
DGI (2019) [56]	Spatial-based ConvGNN	A, X	-	-	$O(m)$
GIN (2019) [57]	Spatial-based ConvGNN	A, X	-	concat+sum	$O(m)$
ClusterGCN (2019) [58]	Spatial-based ConvGNN	A, X	-	-	-