

Stanford CS520: Introduction to Graph Neural Networks

Jiaxuan You, Stanford University

(Slides adapted from CS224W: Machine Learning with Graphs)



Many Types of Data are Graphs




Image credit: [Maximilian Nickel et al](#)

Knowledge Graphs




Image credit: [ResearchGate](#)

Code Graphs




Image credit: ese.wustl.edu

Regulatory Networks




Image credit: [MDPI](#)

Molecules




Image credit: math.hws.edu

Scene Graphs




Image credit: [Wikipedia](#)

3D Shapes

Many Types of Data are Graphs




Image credit:





Image credit: [ResearchGate](#)


Code Graphs



Main question:

How do we take advantage of relational structure for better prediction?

Image credit: MDPI



math.hws.edu

Graphs




Image credit: Wikipedia

Molecules


3D Shapes

Graphs: Machine Learning

Complex domains have a rich relational structure, which can be represented as a **relational graph**

By explicitly modeling relationships we achieve better performance!

Modern ML Toolbox




Modern deep learning toolbox is designed for simple sequences & grids



Why is it Hard?

Networks are complex.


- Arbitrary size and complex topological structure (*i.e.*, no spatial locality like grids)



Networks




Images




Text

Node Embeddings

- **Intuition:** Map nodes to d -dimensional embeddings such that similar nodes in the graph are embedded close together



Input graph




2D node embeddings

How to learn mapping function f ?

Node Embeddings

Goal: $\text{similarity}(u, v) \approx \mathbf{z}_v^T \mathbf{z}_u$

Need to define!



Input network

d-dimensional
embedding space

Two Key Components

- **Encoder:** maps each node to a low-dimensional vector

$\text{ENC}(v) = \mathbf{z}_v$

d-dimensional embedding

node in the input graph

- **Similarity function:** specifies how the relationships in vector space map to the relationships in the original network

$\text{similarity}(u, v) \approx \mathbf{z}_v^T \mathbf{z}_u$


Similarity of u and v in the original network

Decoder

dot product between node embeddings

“Shallow” Encoding

Simplest encoding approach: **encoder is just an embedding-lookup**



KG: shallow encoder + similarity decoder

- Edges in KG are represented as **triplets** (h, r, t)
 - head (h) has **relation** (r) with tail (t)
- **Knowledge Graph embeddings**
 - **Shallow encoder**: Model entities and relations in the embedding/vector space \mathbb{R}^d via embedding lookup
 - Associate entities and relations with **shallow embeddings**
 - **Similarity-based decoder**: Given a true triple (h, r, t) , the goal is that the **embedding of (h, r) should be close** to the **embedding of t** .



Example: TransE

■ Translation Intuition:

For a triple (h, r, t) , $\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{R}^d$,

$\mathbf{h} + \mathbf{r} \approx \mathbf{t}$ if the given fact is true
else $\mathbf{h} + \mathbf{r} \neq \mathbf{t}$

Scoring function: $f_r(h, t) = -||\mathbf{h} + \mathbf{r} - \mathbf{t}||$



Shallow Encoders: Limitations

- Limitations of shallow embedding methods:
 - **$O(|V|)$ parameters are needed:**
 - No sharing of parameters between nodes
 - Every node has its own unique embedding
 - **Inherently “transductive”:**
 - Cannot generate embeddings for nodes that are not seen during training
 - **Do not incorporate node features:**
 - Many graphs have features that we can and should leverage


Today: Deep Graph Encoders

- **Today:** We will now discuss deep methods based on **graph neural networks (GNNs)**:

$\text{ENC}(\nu) =$ **multiple layers of
non-linear transformations
based on graph structure**

- **Note:** All these deep encoders can be **combined with node similarity functions** defined in Knowledge Graph models

Deep Graph Encoders



Stanford CS520: Graph Neural Networks

Jiaxuan You, Stanford University

(Slides adapted from CS224W: Machine Learning with Graphs)




Setup


- Assume we have a graph G :
 - V is the **vertex set**
 - A is the **adjacency matrix** (assume binary)
 - $X \in \mathbb{R}^{m \times |V|}$ is a matrix of **node features**
 - v : a node in V ; $N(v)$: the set of neighbors of v .
- **Node features:**
 - Social networks: User profile, User image
 - Biological networks: Gene expression profiles, gene functional information
 - When there is no node feature in the graph dataset:
 - Indicator vectors (one-hot encoding of a node)
 - Vector of constant 1: $[1, 1, \dots, 1]$

A Naïve Approach

- Join adjacency matrix and features
- Feed them into a deep neural net:




	A	B	C	D	E	Feat
A	0	1	1	1	0	1 0
B	1	0	0	1	1	0 0
C	1	0	0	1	0	0 1
D	1	1	1	0	1	1 1
E	0	1	0	1	0	1 0



- Issues with this idea:
 - $O(|V|)$ parameters
 - Not applicable to graphs of different sizes
 - Sensitive to node ordering

Idea: Convolutional Networks


CNN on an image:




Goal is to generalize convolutions beyond simple lattices
Leverage node features/attributes (e.g., text, images)

Real-World Graphs

But our graphs look like this:




or this:




- There is no fixed notion of locality or sliding window on the graph
- Graph is permutation invariant

From Images to Graphs

Single Convolutional neural network (CNN) layer with 3x3 filter:



Image




Graph


Idea: transform information at the neighbors and combine it:

- Transform “messages” h_i from neighbors: $W_i h_i$
- Add them up: $\sum_i W_i h_i$

Graph Convolutional Networks





Determine node computation graph



Propagate and transform information



Idea: Aggregate Neighbors

- **Key idea:** Generate node embeddings based on **local network neighborhoods**



Idea: Aggregate Neighbors


- **Intuition:** Nodes aggregate information from their neighbors using neural networks




Idea: Aggregate Neighbors

- **Intuition:** Network neighborhood defines a computation graph

Every node defines a computation graph based on its neighborhood!





INPUT GRAPH




Deep Model: Many Layers

- Model can be **of arbitrary depth**:
 - Nodes have embeddings at each layer
 - Layer-0 embedding of node u is its input feature, x_u
 - Layer- k embedding gets information from nodes that are K hops away




Neighborhood Aggregation

- **Basic approach:** Average information from neighbors and apply a neural network



The Math: Deep Encoder

- **Basic approach:** Average neighbor messages and apply a neural network




How to train a GNN

- GNN provides us node embedding \mathbf{z}_v
- **Supervised setting:**
- we want to minimize the loss \mathcal{L} :


$$\min_{\Theta} \mathcal{L}(\mathbf{y}, f(\mathbf{z}_v))$$

- \mathbf{y} : node/edge/graph label (from external sources)
- \mathcal{L} could be L2 if \mathbf{y} is real number, or cross entropy if \mathbf{y} is categorical
- **Unsupervised setting:**
 - Use graph structure itself as supervision
 - E.g.: random walks, KG objective (TransE, RotatE, ...)

Model Design: Overview




(1) Define a neighborhood aggregation function




(2) Define a loss function on the embeddings


Model Design: Overview



(3) Train on a set of nodes, i.e.,
a batch of computational graphs




Model Design: Overview



(4) Test time: Generate embeddings for nodes as needed


Even for nodes we never trained on!



Stanford CS520: Applications of GNNs

Jiaxuan You, Stanford University

(Slides adapted from CS224W: Machine Learning with Graphs)



Tasks on Networks

Tasks we will be able to solve:

- Node classification
 - Predict a type of a given node
- Link prediction
 - Predict whether two nodes are linked
- Community detection
 - Identify densely linked clusters of nodes
- Graph classification
 - Classify different graphs

Example of Node-level ML Tasks

Example (1): Protein Folding

A protein chain acquires its native 3D structure

Every protein is made up of a sequence of amino acids bonded together

These amino acids interact locally to form shapes like helices and sheets

These shapes fold up on larger scales to form the full three-dimensional protein structure

Proteins can interact with other proteins, performing functions such as signalling and transcribing DNA






Image credit: [DeepMind](#)

The Protein Folding Problem

Computationally predict a protein's 3D structure
based solely on its amino acid sequence



T1037 / 6vr4
90.7 GDT
(RNA polymerase domain)




T1049 / 6y4f
93.3 GDT
(adhesin tip)

- Experimental result
- Computational prediction

Image credit: [DeepMind](#)

AlphaFold: Solving Protein Folding

- **Key idea:** “Spatial graph”
 - **Nodes:** Amino acids in a protein sequence
 - **Edges:** Proximity between amino acids in 3D




Examples of Edge-level ML Tasks

Example (2): Recommender Systems

■ **Users interacts with items**


- Watch movies, buy merchandise, listen to music
- **Nodes:** Users and items
- **Edges:** User-item interactions

■ Goal: Recommend items users might like



PinSage: Graph-based Recommender



Task: Recommend related pins to users



Task: Learn node embeddings z_i such that


$$d(z_{cake1}, z_{cake2}) < d(z_{cake1}, z_{sweater})$$

Predict whether two nodes in a graph are related



Examples of Subgraph-level ML Tasks

Example (3): Traffic Prediction



Road Network as a Graph

- **Nodes:** Road segments
- **Edges:** Connectivity between road segments





Image credit: [DeepMind](#)

Traffic Prediction via GNN

Predict via Graph Neural Networks



- Used in Google Maps

THE MODEL ARCHITECTURE FOR DETERMINING OPTIMAL ROUTES AND THEIR TRAVEL TIME.






Image credit: [DeepMind](#)

Examples of Graph-level ML Tasks

Example (4): Drug Discovery

■ Antibiotics are small molecular graphs

- **Nodes:** Atoms
- **Edges:** Chemical bonds



Konaklieva, Monika I. "Molecular targets of β -lactam-based antimicrobials: beyond the usual suspects." *Antibiotics* 3.2 (2014): 128-142.





Image credit: [CNN](#)

Deep Learning for Antibiotic Discovery


- A graph classification task
- Predict promising molecules from a pool of existing candidates




Stokes, Jonathan M., et al. "A deep learning approach to antibiotic discovery." Cell 180.4 (2020): 688-702.

Molecule Generation / Optimization


Graph generation: Generating novel molecules




Use case 1: Generate novel molecules with high drug likeness




0.948



0.945




0.944




0.941


Use case 2: Optimize existing molecules to have desirable properties




-8.32



-0.71



-5.55



-1.78

Summary

- **Motivations for GNNs**
 - Expressive and scalable
- **What is a GNN**
 - **Key:** A node neighborhood aggregation function
 - Define losses and training procedure
- **Applications of GNNs**
 - **Different levels:** Node, edge, subgraph, graph
- **More materials:**
 - Stanford CS224W
 - Course website: <http://web.stanford.edu/class/cs224w/>