

CS224W Homework 3

November 2, 2023

1 GraphRNN [20 points]

In class, we covered GraphRNN, a generative model for graph structures. Here we assume that the graph has no node types/features, and no edge types/features.

1.1 Edge-level RNN [12 points]

Remember that GraphRNN uses random BFS ordering to generate graphs by iteratively adding a node and predicting its connectivity to the nodes already in the graph. Suppose that the GraphRNN model is generating a grid graph:




Figure 1: GraphRNN grid graph.

If we wanted a GraphRNN to generate this graph, what predictions would each cell in the edge-level RNN need to make? Recall that a GraphRNN needs to predict, for each new node, which existing nodes it needs to wire an edge with. It outputs 1 when there should be an edge, and 0 when there should not. Nodes are added in BFS ordering starting from Node A. Assume that the neighbors of a node are explored in alphabetical order (i.e. ties are broken using alphabetical ordering).

Sample answer format (at a particular step of the node-level RNN):

Decodes node < node_Id > (edge-level RNN predicts < edge_value(s) > for connectivity to < node_Id(s) >)

*fill-in < node_Id >, < edge_value(s) >, and < node_Id(s) >

★ Solution ★

1.2 Advantages of BFS ordering [8 points]

Explain 2 advantages of graph generation with random BFS ordering of nodes in the graph, as opposed to generating with a random ordering of nodes in the graph.

★ Solution ★

2 LightGCN [25 points]

We learned in class about **LightGCN**, a GNN model for recommender systems. Given a bipartite user-item graph $G = (V, E)$, let $\mathbf{A} \in \mathbb{R}^{|V| \times |V|}$ be its unnormalized adjacency matrix, $\mathbf{D} \in \mathbb{R}^{|V| \times |V|}$ be its degree matrix and $\mathbf{E}^{(k)} \in \mathbb{R}^{|V| \times d}$ be its node embedding matrix at layer k where d is the embedding dimension. Let $\tilde{\mathbf{A}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ be the normalized adjacency matrix.

The original GCN updates node embeddings across layers according to $\mathbf{E}^{(k+1)} = \text{ReLU}(\tilde{\mathbf{A}} \mathbf{E}^{(k)} \mathbf{W}^{(k)})$, while LightGCN removes the non-linearity and uses the equation for each layer $k \in \{0, 1, \dots, K-1\}$:

$$\mathbf{E}^{(k+1)} = \tilde{\mathbf{A}} \mathbf{E}^{(k)} \quad (1)$$

Moreover, LightGCN adopts multi-scale diffusion to compute the final node embeddings for link prediction, averaging across layers:

$$\mathbf{E} = \sum_{i=0}^K \alpha_i \mathbf{E}^{(i)}, \quad (2)$$

where we have uniform coefficients $\alpha_i = \frac{1}{K+1}$.

2.1 Advantages of Average Embeddings [4 points]

Why does LightGCN average over layer embeddings? What benefits does it bring, in a recommendation systems setting?

What to submit? 1-3 sentences of explanation on the reasons and benefits of averaging across layers.

★ Solution ★

2.2 Self-connection [4 points]

We denote the embedding of an item i at layer- k $\mathbf{e}_i^{(k)}$ and that of a user u $\mathbf{e}_u^{(k)}$. The graph convolution operation (a.k.a., propagation rule) in LightGCN is defined as:

$$\mathbf{e}_u^{(k+1)} = \sum_{i \in \mathcal{N}_u} \frac{1}{\sqrt{|\mathcal{N}_u|} \sqrt{|\mathcal{N}_i|}} \mathbf{e}_i^{(k)}$$
$$\mathbf{e}_i^{(k+1)} = \sum_{u \in \mathcal{N}_i} \frac{1}{\sqrt{|\mathcal{N}_i|} \sqrt{|\mathcal{N}_u|}} \mathbf{e}_u^{(k)}$$

The symmetric normalization term $\frac{1}{\sqrt{|\mathcal{N}_u|} \sqrt{|\mathcal{N}_i|}}$ follows the design of standard GCN, which can avoid the scale of embeddings increasing with graph convolution operations.

However, from the equations above, we can find that in LGCN, we only aggregate the connected neighbors and do not integrate the target node itself (i.e., there is no **self-connection**). This is different from most existing graph convolution operations that typically aggregate extended neighbors and also specifically handle self-connection.

Does LightGCN contain implicit self-connection? If your answer is yes, which operation captures the same effect as self-connection? If no, what do you think is the reason why LightGCN doesn't need self-connection or similar effects?

What to submit? Yes or no and 1-2 sentences of justification.

★ Solution ★

2.3 Relation with APPNP [5 points]

There is a work that connects GCN with Personalized PageRank, where the authors propose a GCN variant named APPNP that can propagate long range without the risk of oversmoothing. Inspired by the teleport design in Personalized PageRank, APPNP complements each propagation layer with the starting features (i.e., the 0-th layer embeddings), which can balance the need of preserving locality (i.e., staying close to the root node to alleviate oversmoothing) and leveraging the information from a large neighborhood. The propagation layer in APPNP is defined as:

$$\mathbf{E}^{(k+1)} = \beta \mathbf{E}^{(0)} + (1 - \beta) \tilde{\mathbf{A}} \mathbf{E}^{(k)}$$

where β is called the “teleport probability” to control the retention of starting features in the propagation, and $\tilde{\mathbf{A}}$ denotes the normalized adjacency matrix.

Aligning with Equation (2), we can see that by setting α_k accordingly, LightGCN can fully recover the prediction embedding used by APPNP. As such, LightGCN shares the strength of APPNP in combating oversmoothing — by setting the α properly, LightGCN allows using a large K for long-range modeling with controllable oversmoothing.


Express the layer- K embeddings $\mathbf{E}^{(K)}$ of APPNP as a function of the initial embeddings $\mathbf{E}^{(0)}$ and the normalized adjacency matrix $\tilde{\mathbf{A}}$. Show all work.

What to submit? Multi-line mathematical derivation of the relationship between $\mathbf{E}^{(K)}$ and $\mathbf{E}^{(0)}$

★ Solution ★

2.4 Recommendation Task [12 points]

We are given the following bipartite graph. Solid edges represent already-observed user-item interactions, while dotted edges denote the set of positive interactions *in the future*.



One-dimensional embeddings for the nodes are listed in the following table.

a	b	c	d	e	f	g	h
0.3	0.96	0.7	0.6	0.4	0.8	0.7	0.64

We are given a recommendation model that uses L2 distance (lower distance means higher recommendation) between user and item embeddings to calculate the scores.

Compute the Recall@2 score for users a, b, c . For each user, explicitly write out the positive items P_u and recommended items R_u , as sets of the form $\{X, Y, Z\}$. Please exclude already-interacted items in R_u , as we are only interested in recommending not-yet-interacted items.

Fill in your answers in the following table. Where the first 5 columns represent user-item distances. What is the final Recall@2?

	d	e	f	g	h	P_u	R_u	Recall@2
a								
b								
c								

What to submit? The P_u , R_u , and the Recall@2 values for each of the users, as well as the overall final Recall@2.

★ **Solution** ★

3 Honor Code [0 points]

(X) I have read and understood Stanford Honor Code before I submitted my work.

Collaboration: Write down the names & SUNetIDs of students you collaborated with on Homework 2 (None if you didn't).

Note: Read our website on our policy about collaboration!