TUTORIALS

Graph Neural Networks for Recommendation: Reproducibility, Graph Topology, and Node Representation

Part 3. Node Representation

45 minutes

The 2nd Learning on Graphs Conference (LoG 2023)

TABLE OF CONTENTS

01
A FORMAL TAXONOMY OF GRAPH CF

02
STRATEGIES FOR NODE REPRESENTATION

03
MULTIMODAL FEATURES ON ITEMS' NODES

04
GNNs + MULTIMODALITY IN ELLIOT

HANDS-ON

05
TAKE-HOME MESSAGES

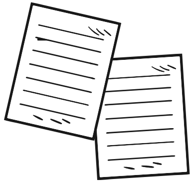

USEFUL RESOURCES

The content of the following slides is taken from:

- Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, Vincenzo Paparella, Claudio Pomo:
  Auditing Consumer- and Producer-Fairness in Graph Collaborative Filtering. ECIR (1) 2023: 33-48
- Daniele Malitesta, Giandomenico Cornacchia, Claudio Pomo, Felice Antonio Merra, Tommaso Di Noia, Eugenio Di Sciascio:
  Formalizing Multimedia Recommendation through Multimodal Deep Learning. CoRR abs/2309.05273 (2023)
- Daniele Malitesta, Giandomenico Cornacchia, Claudio Pomo, Tommaso Di Noia:
  On Popularity Bias of Multimodal-aware Recommender Systems: A Modalities-driven Analysis. MMIR@MM 2023: 59-68

Scan me!

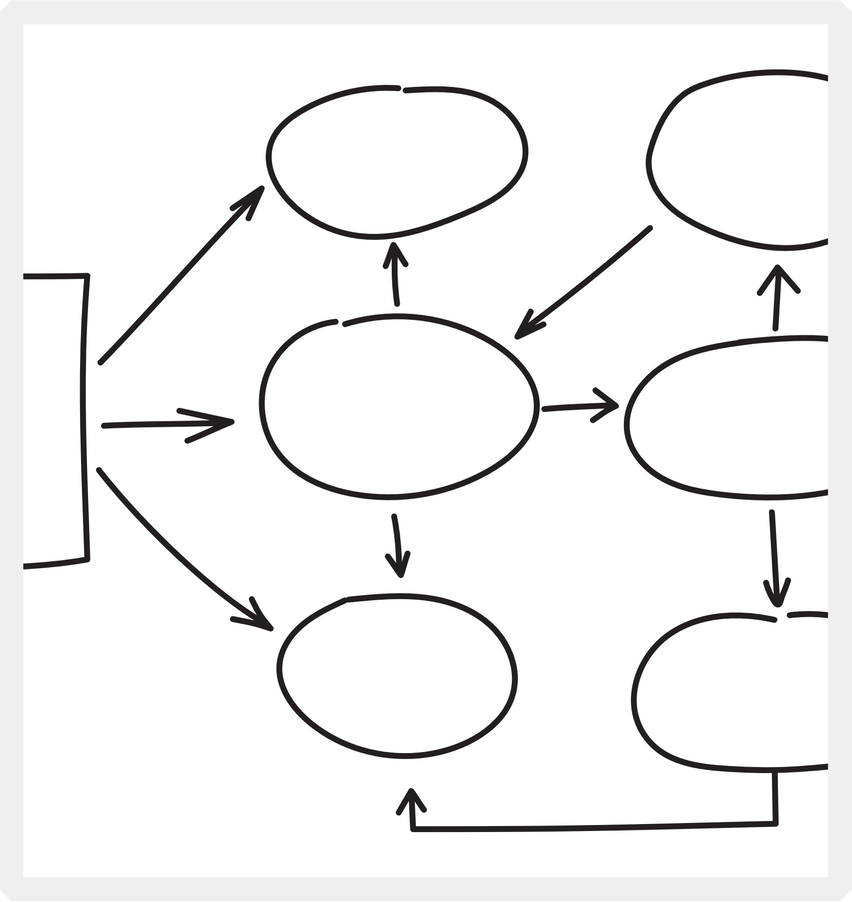

01

A FORMAL TAXONOMY OF GRAPH CF

THE GRAPH CF PIPELINE

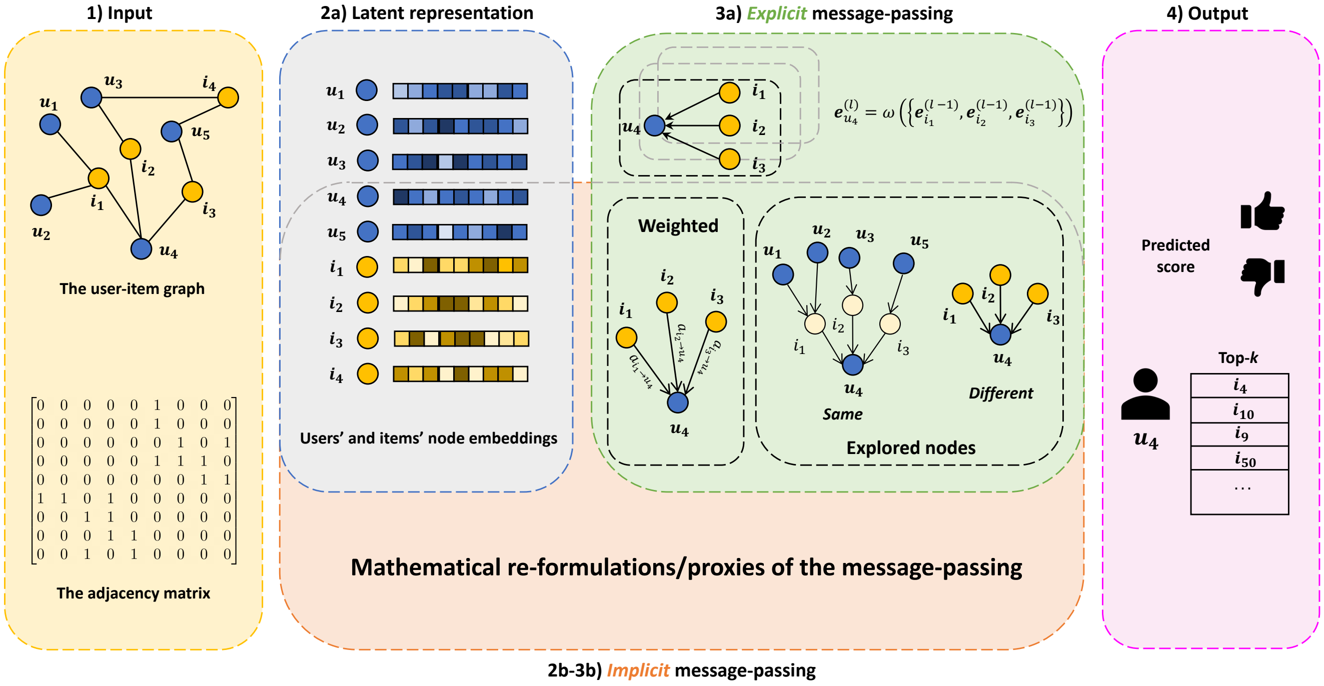

A NON-EXHAUSTIVE FORMAL TAXONOMY

Models	**Nodes Representation**				**Neighborhood Exploration**			
	Latent representation		**Weighting**		**Explored nodes**		**Message passing**	
	low	*high*	*weighted*	*unweighted*	*same*	*different*	*implicit*	*explicit*
NGCF	✓			✓		✓		✓
DGCF	✓		✓			✓		✓
LightGCN	✓			✓		✓		✓
SGL	✓			✓		✓		✓
UltraGCN	✓				✓	✓	✓	
GFCF						✓	✓	

Node representation indicates the representation **strategy** to model **users' and items' nodes**. It involves the **dimensionality** of node embeddings, and the possibility of **weighting the contributions** from neighbor nodes.

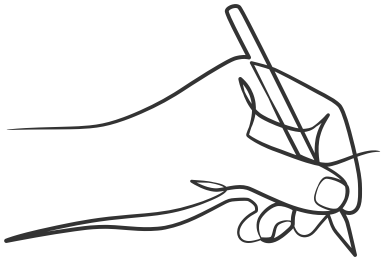

Neighborhood exploration refers to the procedures to **explore** the multi-hop **neighborhoods** of **each node** to update the node latent representation. It involves the **type of node-node connections** which are explored, and the **type of message-passing** schema.

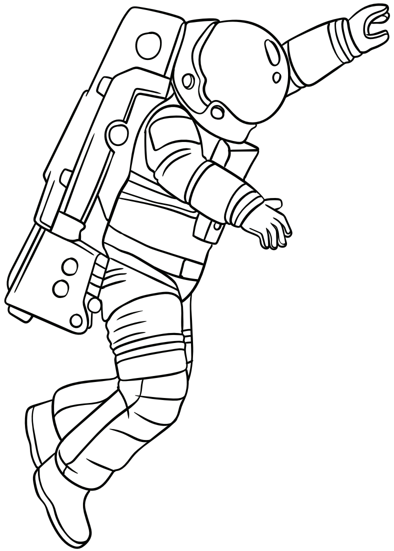

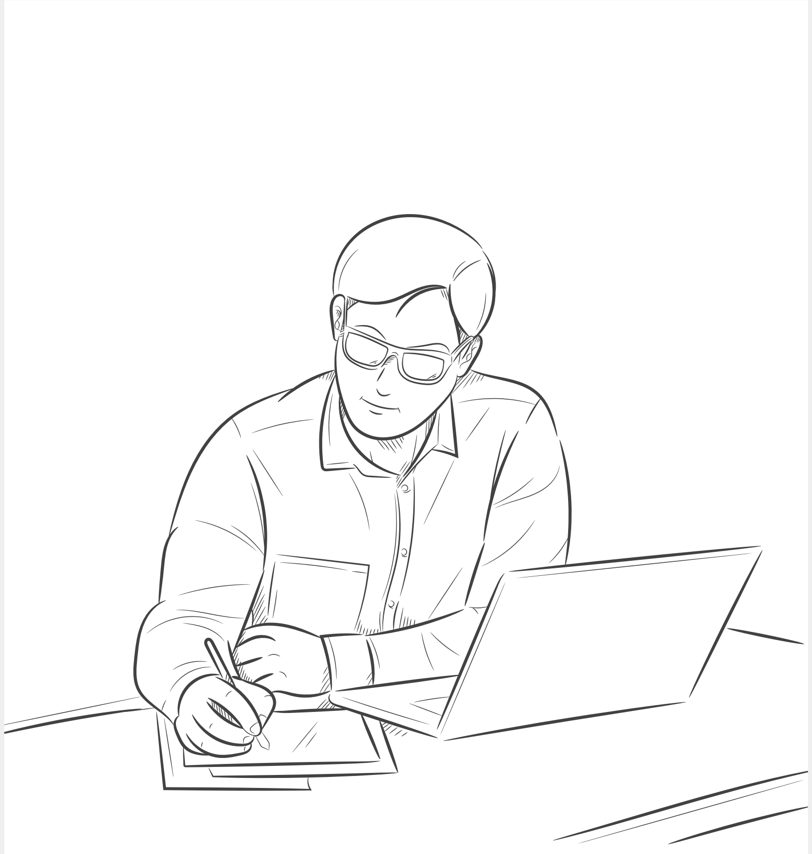

02

STRATEGIES FOR NODE REPRESENTATION

RULES OF THUMB FOR NODE REPRESENTATION

01

EMBEDDING DIMENSIONALITY

FROM THE ORIGINAL PAPERS

Commonly explored dimensionalities for the node embeddings: $\mathrm{d} \in [64, 128, 256, \dots]$. The usually chosen embedding dimensionality is $d = 64$.

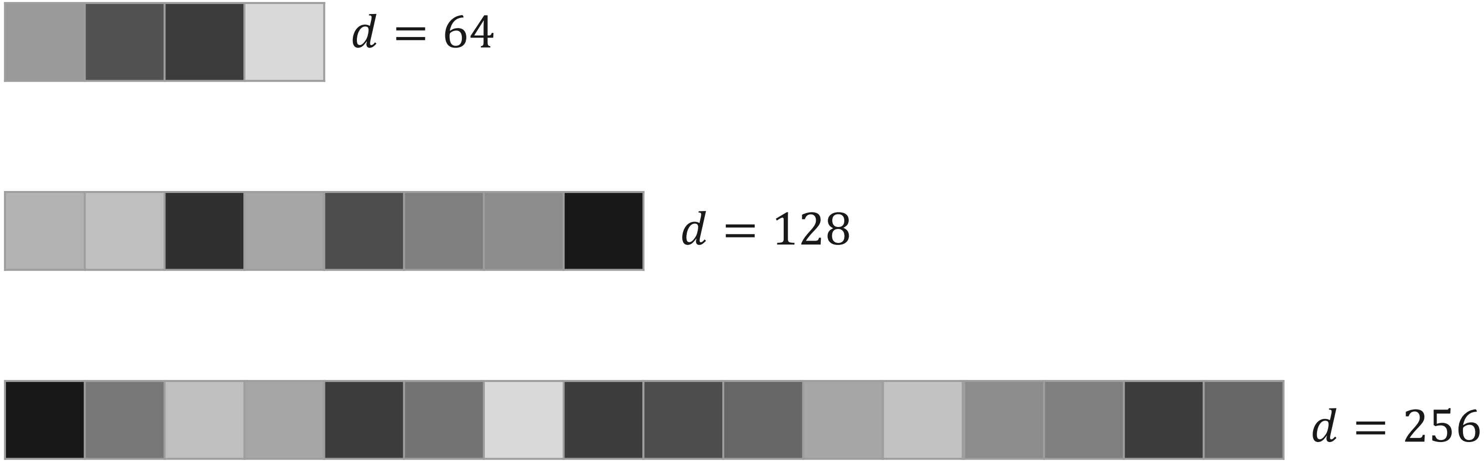

BENCHMARKS FROM THE LITERATURE (1/2)

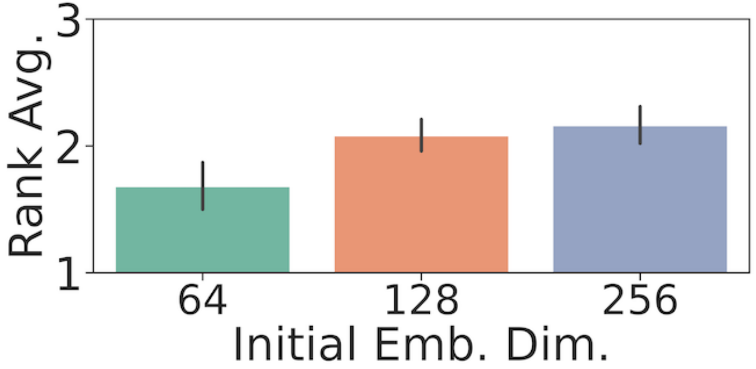

[Wang et al.]

When $d = 256$ the recommendation performance seems to be higher…

BENCHMARKS FROM THE LITERATURE (2/2)

Dataset	Gowalla			Yelp2018			Amazon-book		
Method	**recall**	**ndcg**	**training time**	**recall**	**ndcg**	**training time**	**recall**	**ndcg**	**training time**
LightGCN-64	0.1830	0.1554	2.77×10^4s	0.0649	0.0530	5.15×10^4s	0.0411	0.0315	1.27×10^5s
LightGCN-128	0.1878	0.1591	3.31×10^4s	0.0671	0.0550	5.66×10^4s	0.0459	0.0353	1.81×10^5s
LightGCN-256	0.1893	0.1606	4.54×10^4s	0.0689	0.0568	8.09×10^4s	0.0481	0.0371	2.98×10^5s
LightGCN-512	0.1892	0.1604	7.28×10^4s	0.0689	0.0569	1.33×10^5s	0.0485	0.0375	5.26×10^5s
GF-CF	0.1849	0.1518	30.5s	0.0697	0.0571	46.0s	0.0710	0.0584	65.8s

[Shen et al.]

… but this might not be always true and could come at the expense of training time.

RULES OF THUMB FOR NODE REPRESENTATION

01

EMBEDDING DIMENSIONALITY

02

INITIALIZATION

FROM THE ORIGINAL CODES

Models	Strategies	
	Normal	**Xavier**
NGCF		✓
DGCF		✓
LightGCN	✓	
SGL		✓
UltraGCN	✓	
GFCF		

Normal and Xavier are almost equally-preferrable.

RULES OF THUMB FOR NODE REPRESENTATION

01

EMBEDDING DIMENSIONALITY

02

INITIALIZATION

03

NODE WEIGHTING

NOT ALL NEIGHBORS HAVE THE SAME IMPORTANCE

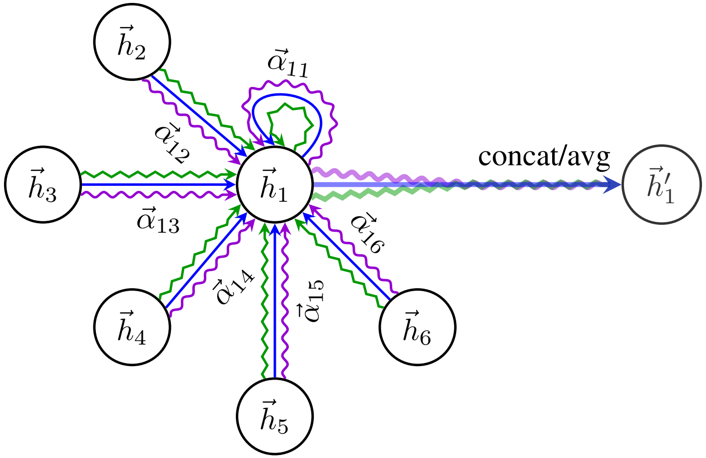

[Veličković et al.]

It might be useful to weight the contribution provided by each neighbour node before the aggregation.

DISENTANGLING THE NODE REPRESENTATION

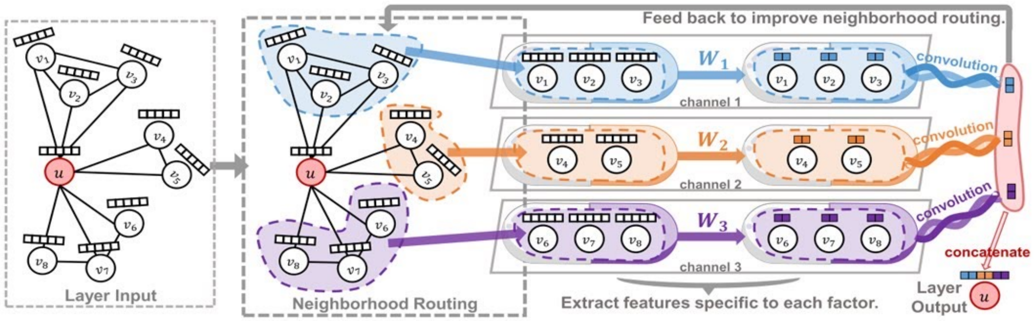

Neighbour nodes have specific features which might explain why they interacted with the ego node.

[Ma et al.]

DISENTANGLING IN GRAPH CF (1/2)

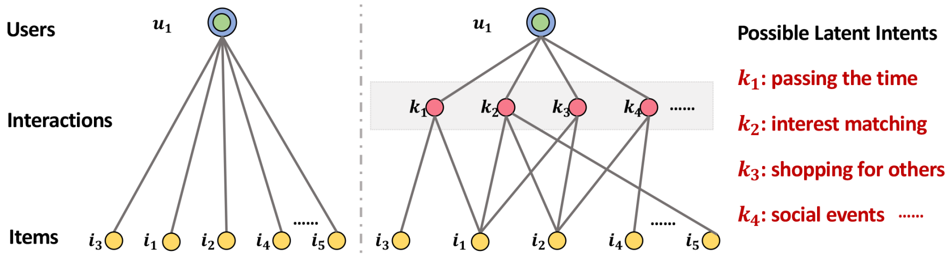

There might exist hidden intentions underlying each user-item interaction.

[Wang et al.]

DISENTANGLING IN GRAPH CF (2/2)

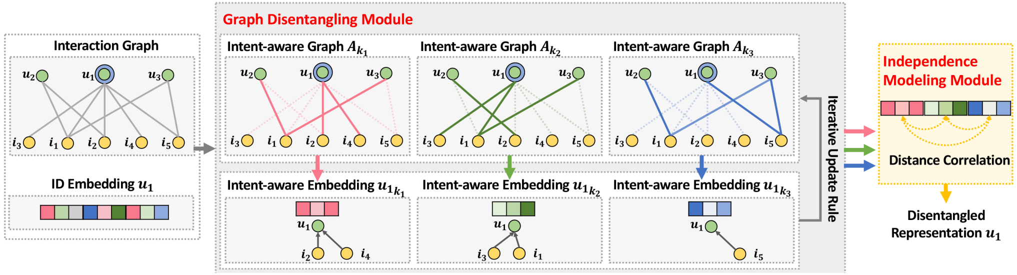

The node embeddings are split into intentions, and they are trained to be uncorrelated.

[Wang et al.]

RULES OF THUMB FOR NODE REPRESENTATION

01

EMBEDDING DIMENSIONALITY

02

INITIALIZATION

03

NODE WEIGHTING

04

EMBEDDING NORMALIZATION

FROM THE ORIGINAL CODES

Models	Normalize	When
NGCF	✓	`during message-passing`
DGCF	✓	`during message-passing`
LightGCN		
SGL	✓	`after message-passing`
UltraGCN		
GFCF		

The l2 normalization seems to stabilize the training, and it is performed during or after the message-passing.

RULES OF THUMB FOR NODE REPRESENTATION

01

EMBEDDING DIMENSIONALITY

02

INITIALIZATION

03

NODE WEIGHTING

04

EMBEDDING NORMALIZATION

05

LAYER COMBINATION

FROM THE ORIGINAL CODES

Models	How	Final dimension
NGCF	`concat`	$\sum_{l=1}^{L} d_l$
DGCF	`mean`	d
LightGCN	`mean`	d
SGL	`stack + mean`	d
UltraGCN		
GFCF		

Each layer combination comes with a final dimension which may increase the computational complexity.

BENCHMARKS FROM THE LITERATURE

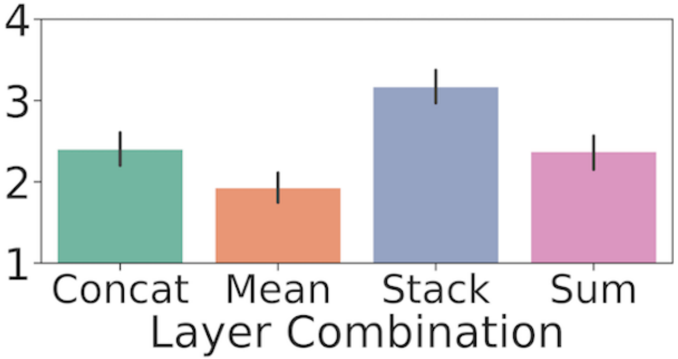

[Wang et al.]

Stack and concat seem to be the highest recommendation performance.

03

MULTIMODAL FEATURES ON ITEMS' NODES

RECSYS LEVERAGING MULTIMODAL DATA

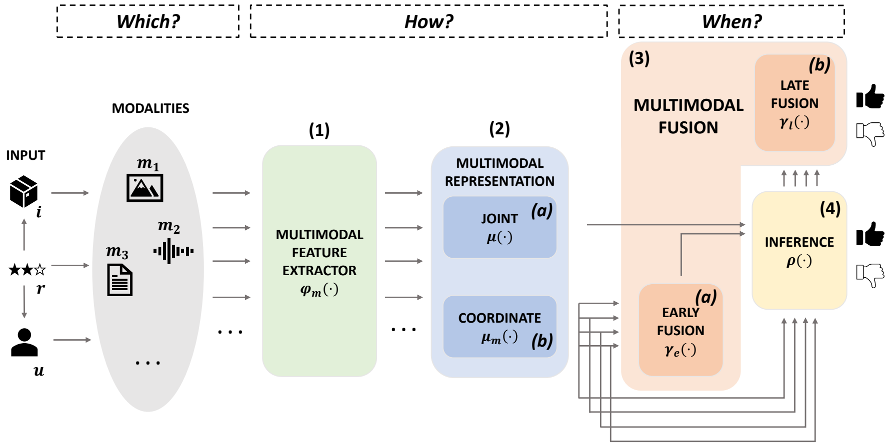

Multimodal-aware **recommender** systems exploit multimodal (i.e., audio, visual, textual) data to **augment** the representation of **items**, thus tackling **known issues** such as **dataset sparsity** and the **inexplicable nature** of users' actions (i.e., views, clicks) on online platforms.

MULTIMODAL EMBEDDINGS

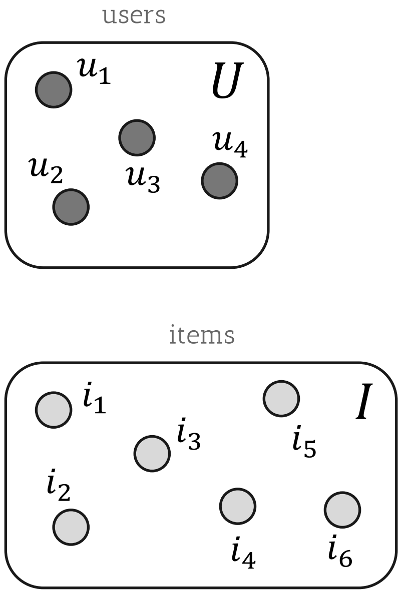

	i_1	i_2	i_3	i_4	i_5	i_6
u_1	1	0	1	1	1	0
u_2	0	0	0	0	1	0
u_3	1	1	1	0	0	0
u_4	0	0	1	1	0	1

user-item
interaction matrix

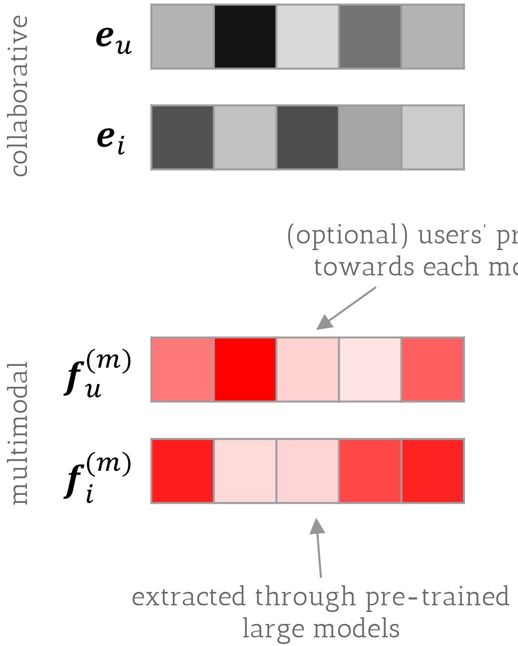

MULTIMODAL-AWARE RECSYS

Models	Venue	Multimodal embeddings		GNN	Multimodal graphs	
		Users	Items		U-I	I-I
VBPR	AAAI'16	✓	✓	✗		
MMGCN	MM'19	✓	✓	✓	✓	✗
GRCN	MM'20	✗	✓	✓	✓	✗
LATTICE	MM'21	✗	✓	✓	✗	✓
BM3	WWW'23	✗	✓	✓	✓	✗
FREEDOM	MM'23	✗	✓	✓	✗	✓

The recent tendency is to exploit GNN + multimodality, on the user-item and/or the item-item graphs.

LIGHTGCN AND MULTIMODALITY: HOW TO?

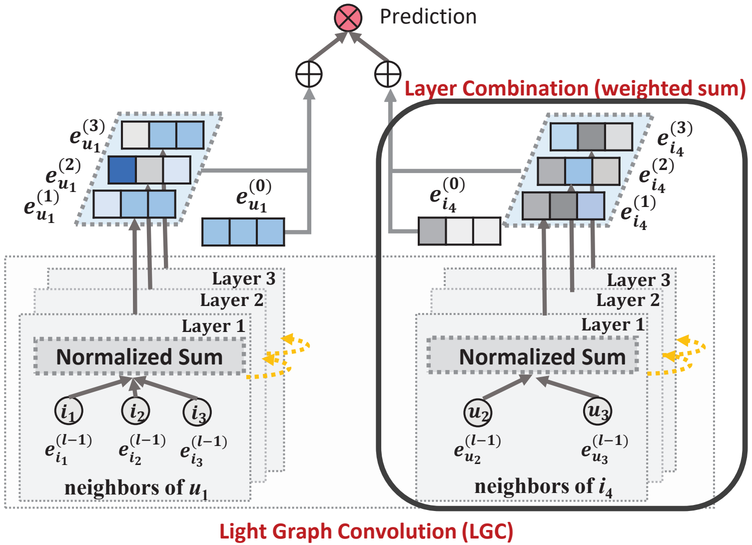

These are not pure collaborative, but extracted from multimodal content and trainable

[Wei et al.]

04

HANDS-ON #2

GNNs + MULTIMODALITY IN ELLIOT

SCAN ME AND GO TO GOOGLE COLAB!

or find me at:

https://sisinflab.github.io/tutorial-gnns-recsys-log2023/sections/node_representation/

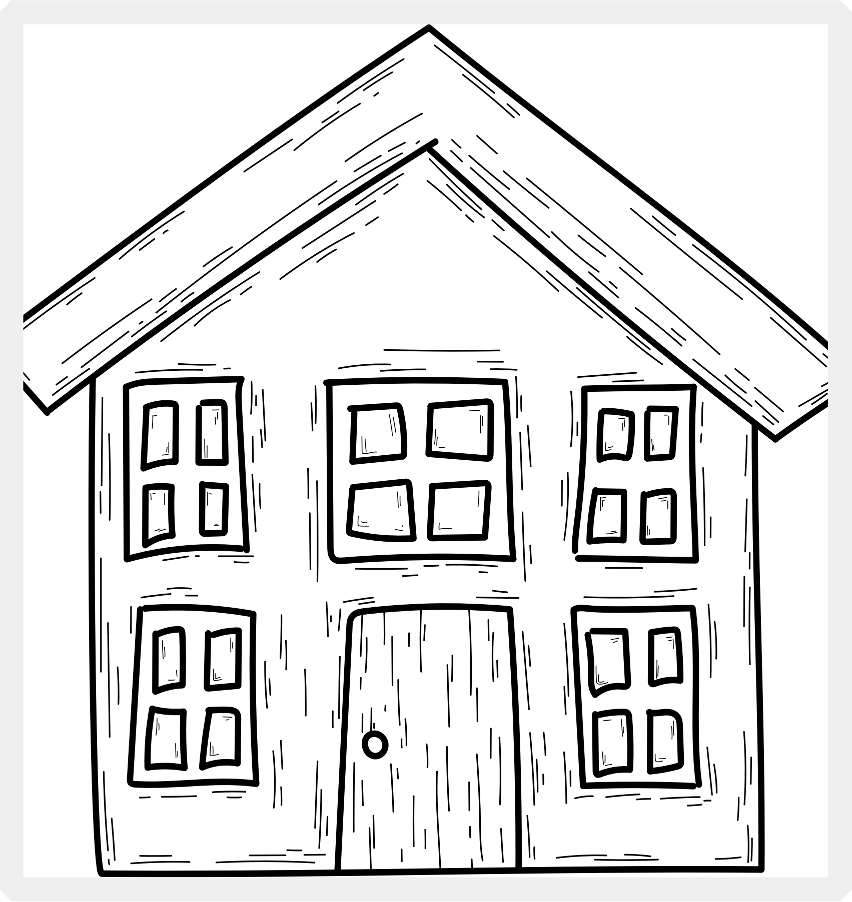

05

TAKE-HOME MESSAGES

WHAT WE HAVE LEARNED

01

NODE REPR. &
NEIGHBORHOOD EXPL.

WHAT WE HAVE LEARNED

01

NODE REPR. &
NEIGHBORHOOD EXPL.

02

NODE EMBEDDING SIZE

WHAT WE HAVE LEARNED

01

NODE REPR. &
NEIGHBORHOOD EXPL.

02

NODE EMBEDDING SIZE

03

INITIALIZATION

WHAT WE HAVE LEARNED

01

NODE REPR. & NEIGHBORHOOD EXPL.

02

NODE EMBEDDING SIZE

03

INITIALIZATION

04

WEIGHTING NEIGHBOUR NODES

WHAT WE HAVE LEARNED

01

NODE REPR. &
NEIGHBORHOOD EXPL.

02

NODE EMBEDDING SIZE

03

INITIALIZATION

04

WEIGHTING NEIGHBOUR
NODES

05

LAYER COMBINATION

WHAT WE HAVE LEARNED

01

NODE REPR. &
NEIGHBORHOOD EXPL.

02

NODE EMBEDDING SIZE

03

INITIALIZATION

04

WEIGHTING NEIGHBOUR NODES

05

LAYER COMBINATION

06

GNN + MULTIMODALITY FOR RECSYS

REFERENCES (1/2)

[Anelli et al.] Auditing Consumer- and Producer-Fairness in Graph Collaborative Filtering. ECIR (1) 2023: 33-48

[Malitesta et al.] Formalizing Multimedia Recommendation through Multimodal Deep Learning. CoRR abs/2309.05273 (2023)

[Malitesta et al.] On Popularity Bias of Multimodal-aware Recommender Systems: A Modalities-driven Analysis. MMIR@MM 2023: 59-68

[Wang et al.] Profiling the Design Space for Graph Neural Networks based Collaborative Filtering. WSDM 2022: 1109-1119

[Veličković et al.] Graph Attention Networks. ICLR (Poster) 2018

[Ma et al.] Disentangled Graph Convolutional Networks. ICML 2019: 4212-4221

[He and McAuley] VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback. AAAI 2016: 144-150

[Wei et al. (2019)] MMGCN: Multi-modal Graph Convolution Network for Personalized Recommendation of Micro-video. ACM Multimedia 2019: 1437-1445

[Wei et al. (2020)] Graph-Refined Convolutional Network for Multimedia Recommendation with Implicit Feedback. ACM Multimedia 2020: 3541-3549

[Zhang et al.] Mining Latent Structures for Multimedia Recommendation. ACM Multimedia 2021: 3872-3880

REFERENCES (2/2)

[Zhou et al.]Bootstrap Latent Representations for Multi-modal Recommendation. WWW 2023: 845-854
[Zhou and Shen] A Tale of Two Graphs: Freezing and Denoising Graph Structures for Multimodal Recommendation. ACM Multimedia 2023: 935-943

THANKS!

Scan me!