

# Scaling Up Graph Neural Network

PRESENTER: Liao Ningyi

SUPERVISOR: Asst Prof Luo Siqiang

27 Jun 2023

# Contents

**Introduction**

**Scalable GNN with Feature-Oriented Optimization**

**Scalable Heterophilous GNN with Decoupled Embedding**

**Conclusion and Future Works**

**Q&A**

# Contents

## **Introduction**

**Graph Data and Graph Representation**

**Graph Neural Network**

**Challenge and Motivation**

Scalable GNN with Feature-Oriented Optimization

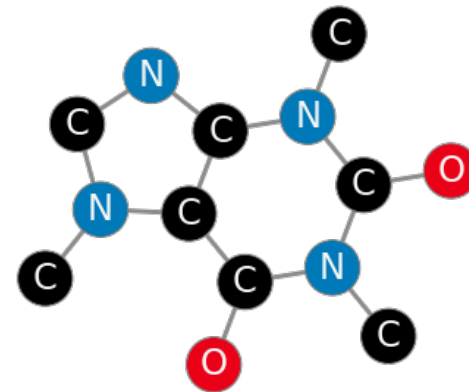
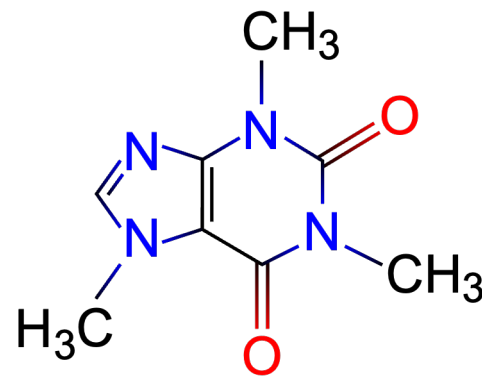
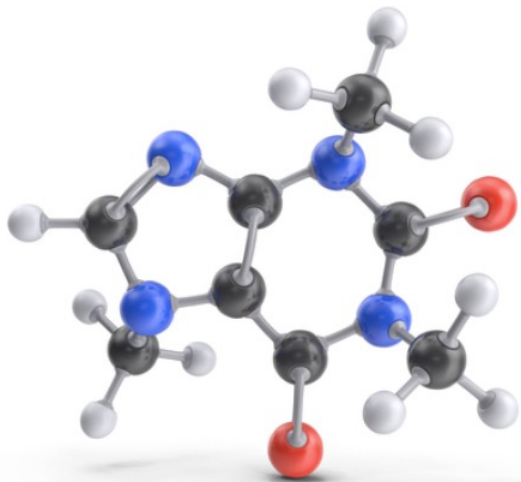
Scalable Heterophilous GNN with Decoupled Embedding

Conclusion and Future Works

Q&A

# Graph: A Ubiquitous Data Structure

- Graphs model entities (**nodes**) and relationship (**edges**)
- Graph data structures are **non-Euclidean**



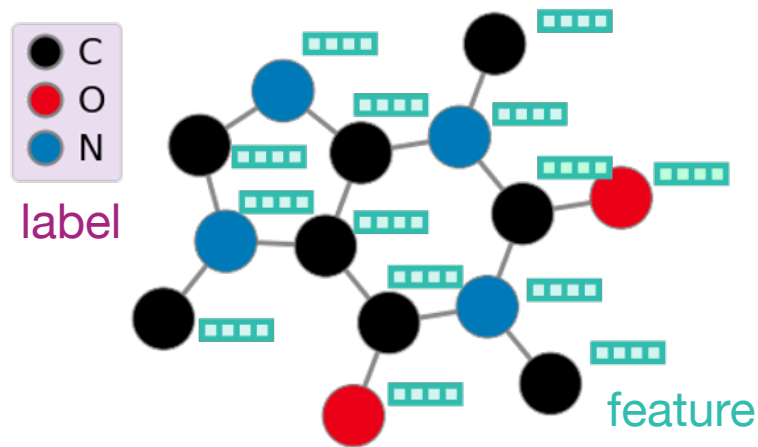
c	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0
c	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0
N	0	0	0	0	1	0	0	0	0	1	0	0	1	0	0
O	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
c	0	0	0	0	0	0	0	0	0	1	0	1	1	0	0
N	0	0	0	0	0	0	1	0	0	1	0	0	0	0	1
O	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
c	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0
c	0	1	0	0	1	0	1	0	0	0	0	0	0	0	0
c	0	0	0	1	0	1	0	0	0	0	0	0	1	0	0
N	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0
c	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0
N	1	0	1	0	0	1	0	0	0	0	0	0	0	0	0
c	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
	c	N	C	N	C	C	C	O	N	C	O	N	C	C	

Different representations of the caffeine molecule



# Graph Representation

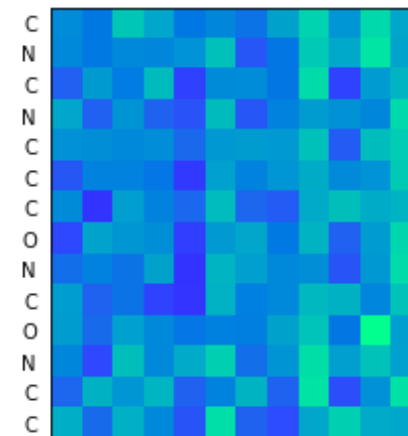
- Nodes  $V$ , Edges  $E$
- Each node: Feature  $x$ , Label  $y$
- Adjacency matrix  $A$ , Feature matrix  $X$



c	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0
c	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0
N	0	0	0	0	1	0	0	0	0	1	0	0	1	0	0
o	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0
c	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0
N	0	0	0	0	0	0	1	0	0	1	0	0	0	0	1
o	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0
c	0	0	0	0	0	1	0	1	1	0	0	0	0	0	0
c	0	1	0	0	1	0	1	0	0	0	0	0	0	0	0
c	0	0	0	1	0	1	0	0	0	0	0	0	1	0	0
N	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0
c	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0
N	1	0	1	0	0	1	0	0	0	0	0	0	0	0	0
c	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0

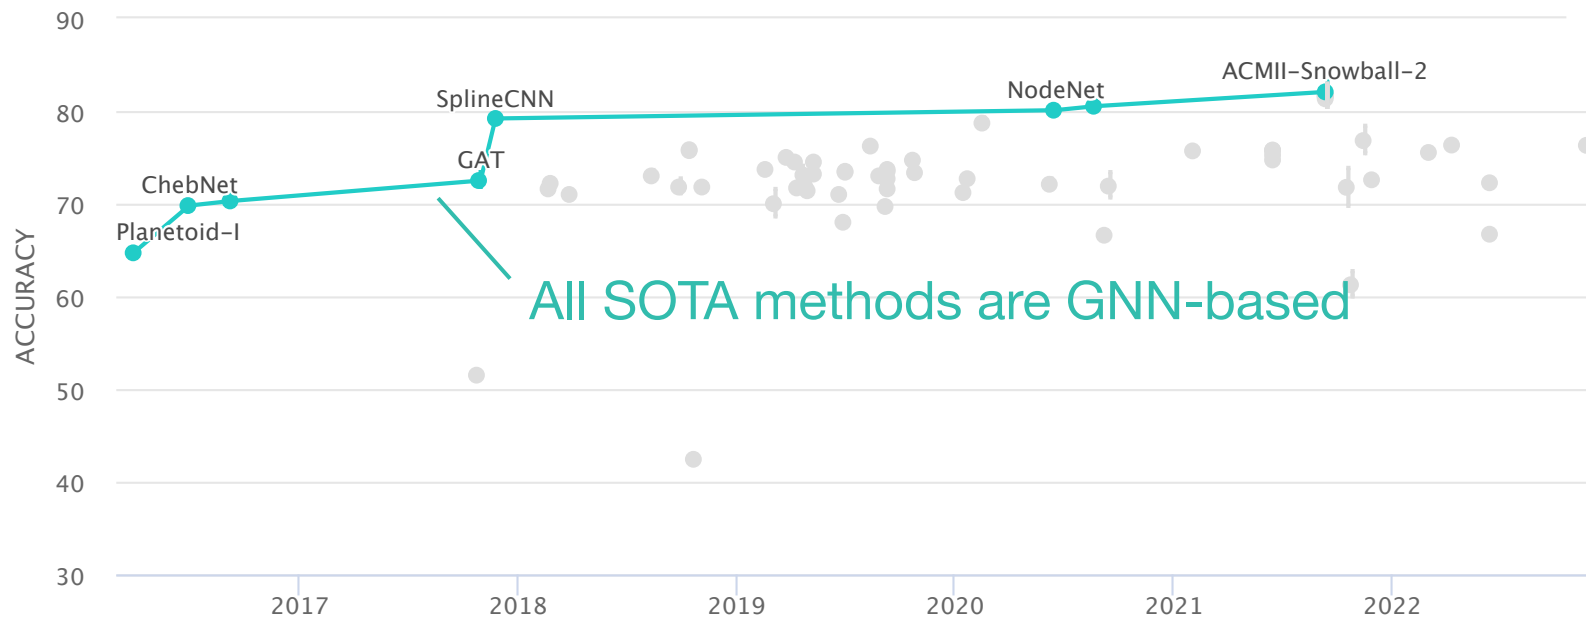
C N C N C C C O N C O N C C

Adjacency matrix



# GNN: Graph Neural Network

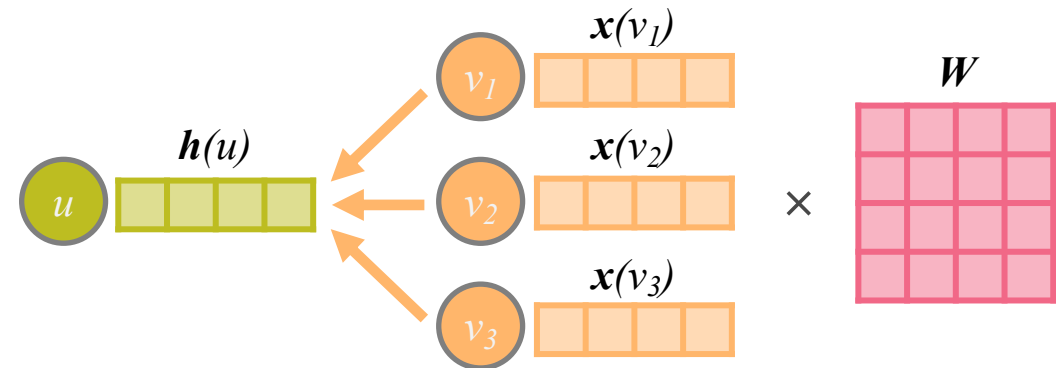
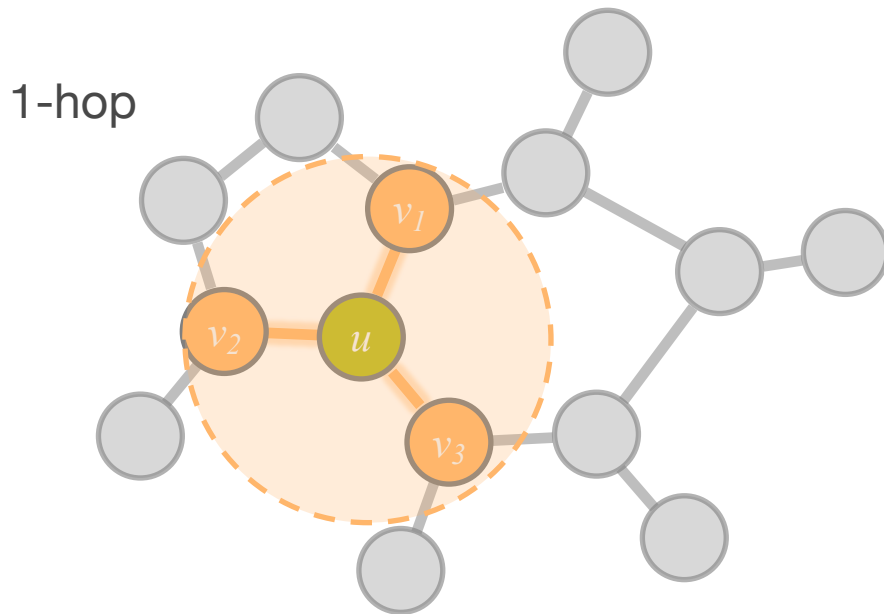
- Apply deep learning architectures to graph data
- Achieve strong performance on graph learning tasks



Leaderboard of Node Classification on Citeseer

# GNN: Graph Neural Network

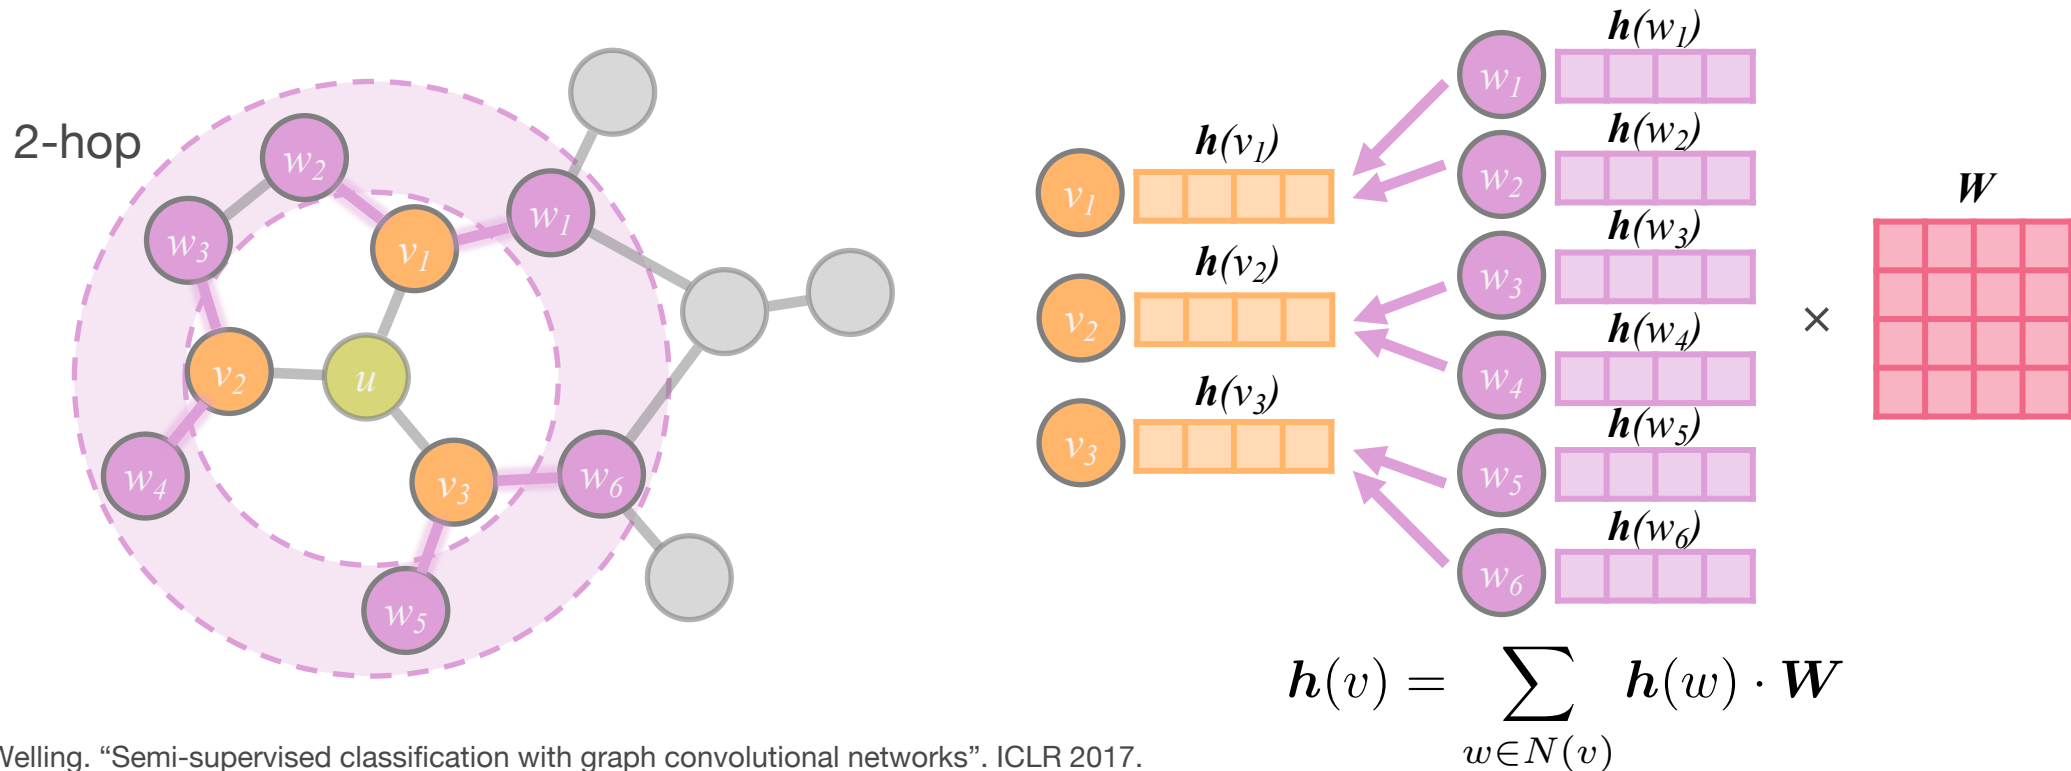
- Graph convolution: aggregate neighbor information  $N(u)$  by learnable weights  $W$  to each node as node representation  $h$



$$h(u) = \sum_{v \in N(u)} x(v) \cdot W$$

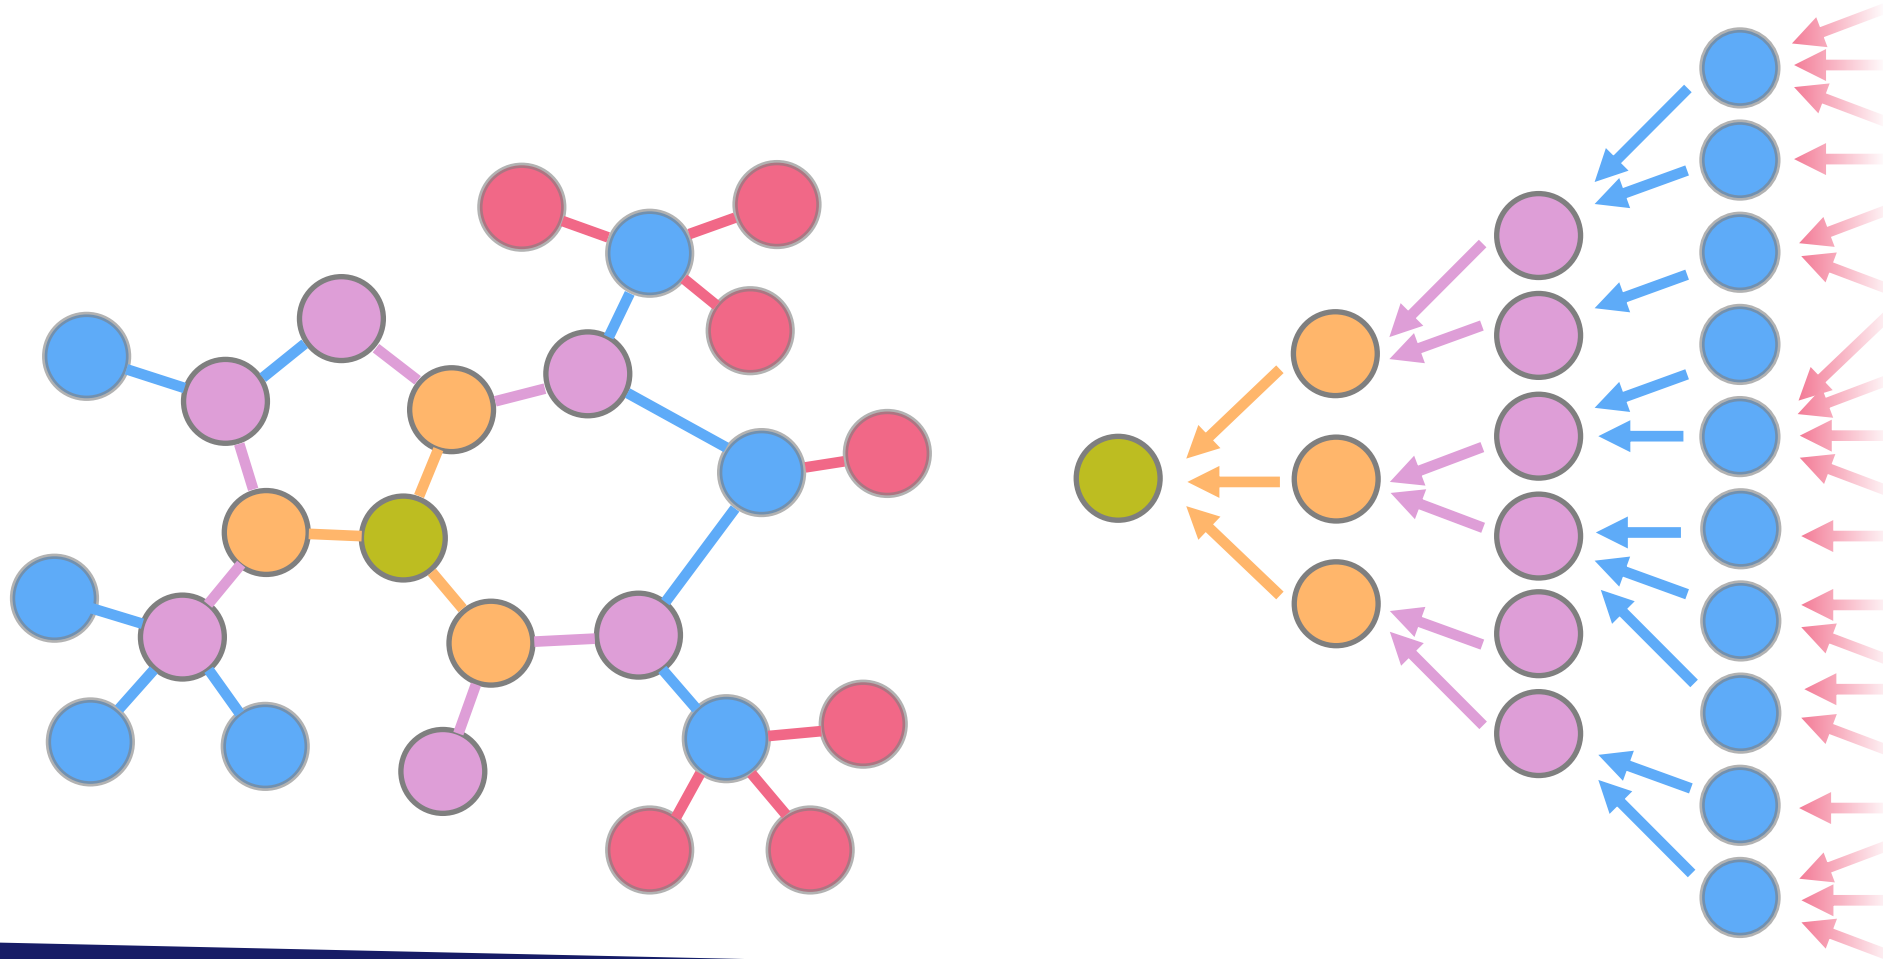
# GNN: Graph Neural Network

- Stacking multi-hop graph convolutions as layers



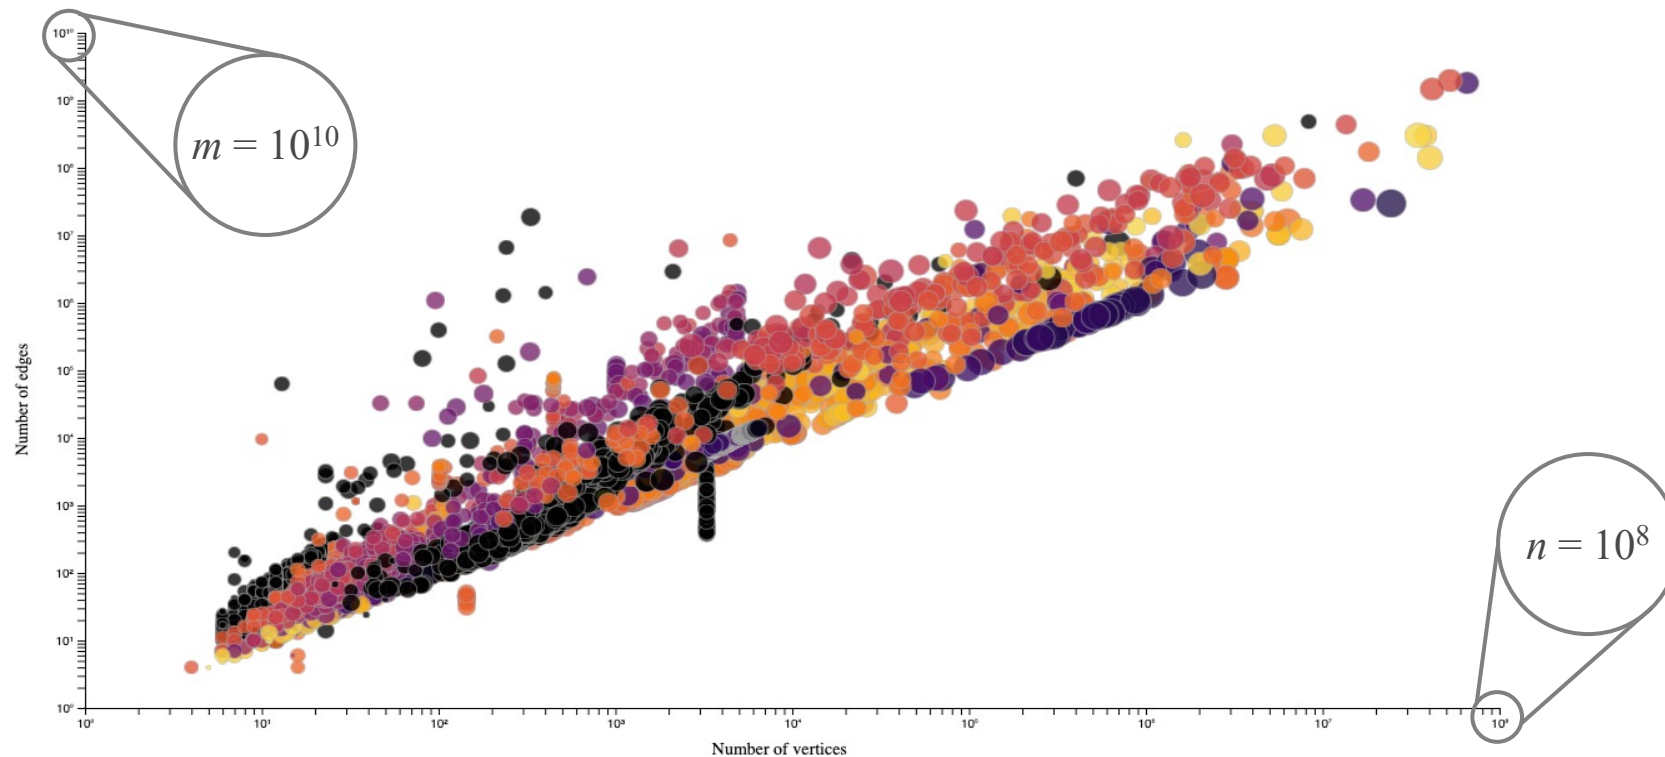
# Challenge: GNN Scalability

- Curse of scale: computation overhead greatly increase!



# Challenge: GNN Scalability

- Modern real-world graphs are on the scale of millions or billions



Scales of recent graph datasets

# Motivations and Objectives

- How to scale up GNN to learn million- or billion-scale graphs?
  - What are the **bottlenecks of GNN scalability** with respect to computation time and memory complexity?
  - How to simplify and optimize **GNN graph propagation** using advanced graph management techniques?
  - How to address the issue of GNN scalability in different **variants of graph data**?

# SCARA: Scalable GNN with Feature-Oriented Optimization

**Ningyi Liao\***, Dingheng Mo\*, Siqiang Luo, Xiang Li, Pengcheng Yin. “SCARA: Scalable Graph Neural Networks with Feature-Oriented Optimization”. *Proceedings of the VLDB Endowment*, Vol. 15, No. 11, pp. 3240-3248, 2022.



# Contents

Introduction

## **Scalable GNN with Feature-Oriented Optimization**

**Motivation**

**Method**

**Experimental Evaluation**

Scalable Heterophilous GNN with Decoupled Embedding

Conclusion and Future Works

Q&A

# Scalability Issue of GNNs

- GNN in matrix form:

Feature Matrix

$$H^{(0)} = X$$

$$H^{(l+1)} = \sigma \left( \tilde{A} H^{(l)} W^{(l)} \right), \quad l = 0, 1, \dots, L - 1$$

Adjacency Matrix

Layer Representation

Learnable Weights

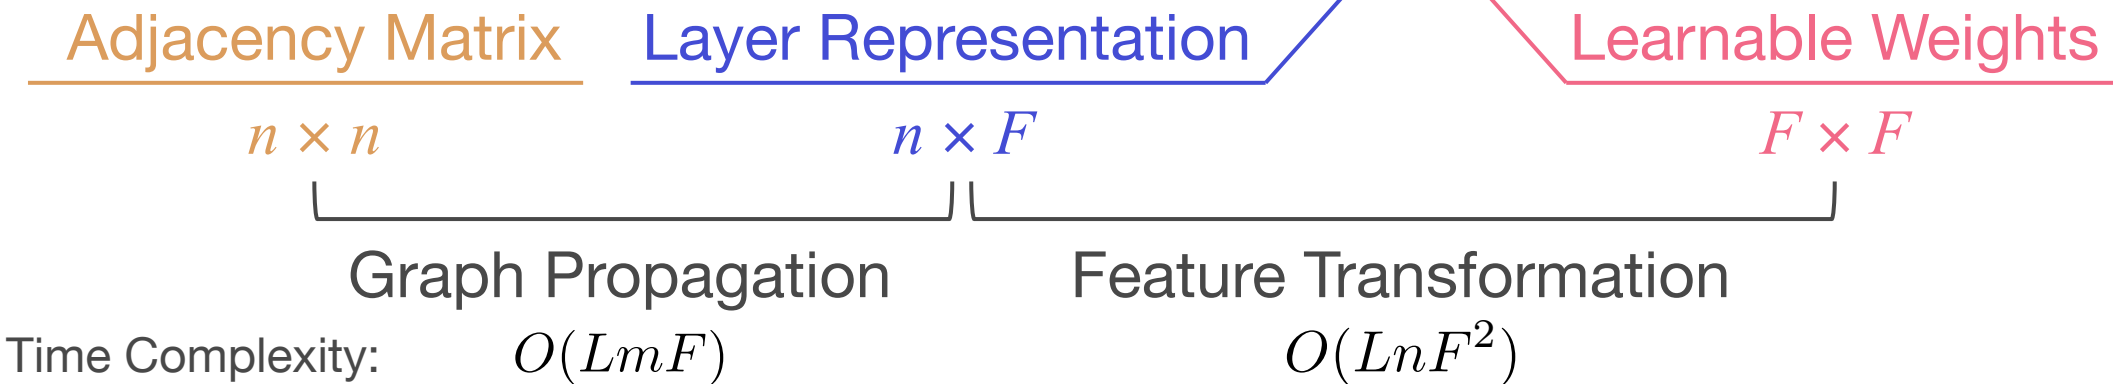
Graph Propagation

Feature Transformation

# Scalability Issue of GNNs

- Bottleneck of GNN computation is graph propagation

$$H^{(l+1)} = \sigma \left( \tilde{A} H^{(l)} W^{(l)} \right), \quad l = 0, 1, \dots, L - 1$$



- |  |                          |
|--|--------------------------|
| • Sparse-Dense Matrix Mul ☹️ <i>Not optimized for minibatch</i>                                    | • Dense-Dense Matrix Mul |
| • $m$ at a larger scale than $n$ ☹️ <i>Not scalable to large <math>m</math> and <math>n</math></i> | • Similar to common NN   |

[1] T Kipf & M Welling. "Semi-supervised classification with graph convolutional networks". ICLR 2017.

[2] W Chiang et al. "Cluster-GCN: An efficient algorithm for training deep and large graph convolutional networks". KDD 2019.

# Existing Work: Decoupling

- GCN: iterative propagation and transformation

$$\mathbf{H}^{(L)} = \underbrace{\sigma(\tilde{\mathbf{A}} \sigma(\tilde{\mathbf{A}} \cdots \tilde{\mathbf{A}} \mathbf{X} \mathbf{W}^{(1)} \cdots))}_{L \text{ Graph Propagation}} \underbrace{\mathbf{W}^{(L-1)}}_{L \text{ Feature Transformation}}$$

- SGC: decoupled propagation as precomputation

Graph Embedding

Precomputed Propagation

$$\mathbf{H}^{(0)} = \mathbf{P} = \tilde{\mathbf{A}}^L \cdot \mathbf{X}$$

$$\mathbf{H}^{(l+1)} = \sigma(\mathbf{H}^{(l)} \mathbf{W}^{(l)}), \quad l = 0, 1, \cdots, L - 1$$

[1] T Kipf & M Welling. "Semi-supervised classification with graph convolutional networks". ICLR 2017.

[2] F Wu et al. "Simplifying graph convolutional networks". ICML 2019.

# Existing Work: Decoupling

- Time complexity of SGC is still not scalable

$$P = \tilde{A}^L \cdot X = \tilde{A} \left( \tilde{A} \cdots \left( \tilde{A} X \right) \right)$$

Adjacency Matrix Feature Matrix

$n \times n$   $n \times F$

- $L$  Sparse-Dense Matrix Mul  $\rightarrow$  Time Complexity:  $O(LmF)$
- Result stored in  $P$   $\rightarrow$  Memory Complexity:  $O(nF)$

# Our Model: SCARA

- Propagation as precomputation

$$\mathbf{H}^{(0)} = \mathbf{P} = \sum_{l=0}^{\infty} \alpha(1 - \alpha)^l \tilde{\mathbf{A}}^l \cdot \mathbf{X}$$

Time Complexity:  $O(F\sqrt{m \log n}/\lambda)$

😊 *Only sublinear to  $m$*

Memory Complexity:  $O(nF)$

😊 *Efficient for CPU precomputation*

- Feature transformation

$$\mathbf{H}^{(l+1)} = \sigma(\mathbf{H}^{(l)} \mathbf{W}^{(l)}), \quad l = 0, 1, \dots, L - 1$$

Time Complexity:  $O(LnF^2)$

😊 *Efficient GPU training*

Memory Complexity:  $O(Ln_b F + LF^2)$

# Our Model: SCARA

- **Time Complexity:**
  - Sub-linear precomputation time complexity
  - Efficient decoupled training and inference

TABLE 2.1: Precomputation, training, and inference time complexity of common GNN models.

Model	Precomp. Time	Training Time	Inference Time
GCN [3]	–	$O(ILmF + ILnF^2)$	$O(LmF + LnF^2)$
Cluster-GCN [22]	$O(m)$	$O(ILmF + ILnF^2)$	$O(LmF + LnF^2)$
GraphSAINT [23]	–	$O(IL_P LnF^2)$	$O(LmF + LnF^2)$
GAS [24]	$O(m + LnF)$	$O(ILmF + ILnF^2)$	$O(nF)$
APPNP [25]	$O(m)$	$O(IL_P mF + ILnF^2)$	$O(L_P mF + LnF^2)$
PPRGo [26]	$O(m/\delta)$	$O(IKnF + ILnF^2)$	$O(KnF + LnF^2)$
SGC [27]	$O(L_P mF)$	$O(ILnF^2)$	$O(LnF^2)$
GBP [11]	$O(L_P F \sqrt{L_P m \log(L_P n)} / \epsilon)$	$O(ILnF^2)$	$O(LnF^2)$
<b>SCARA (ours)</b>	$O(F \sqrt{m \log n} / \lambda)$	$O(ILnF^2)$	$O(LnF^2)$

*Not scalable to large  $m$  and  $n$*

# Our Model: SCARA

- **Memory Complexity:**
  - Efficient precomputation memory usage
  - Minibatch training and inference

TABLE 2.2: Precomputation, training, and inference memory complexity of common GNN models.

Model	Precomp. Mem.	Training Mem.	Inference Mem.
GCN [3]	–	$O(LnF + LF^2)$	$O(LnF + LF^2)$
Cluster-GCN [22]	$O(n)$	$O(Ln_bF + LF^2)$	$O(LnF + LF^2)$
GraphSAINT [23]	–	$O(L_P Ln_bF + LF^2)$	$O(LnF + LF^2)$
GAS [24]	$O(LnF)$	$O(Ldn_bF + LF^2)$	$O(Ldn_bF + LF^2)$
APPNP [25]	$O(m)$	$O(Ln_bF + LF^2 + nn_b)$	$O(Ln_bF + LF^2 + nn_b)$
PPRGo [26]	$O(n/\delta)$	$O(Ln_bF + LF^2 + Kn_b)$	$O(Ln_bF + LF^2 + Kn_b)$
SGC [27]	$O(m)$	$O(Ln_bF + LF^2)$	$O(Ln_bF + LF^2)$
GBP [11]	$O(nF)$	$O(Ln_bF + LF^2)$	$O(Ln_bF + LF^2)$
<b>SCARA (ours)</b>	$O(nF)$	$O(Ln_bF + LF^2)$	$O(Ln_bF + LF^2)$

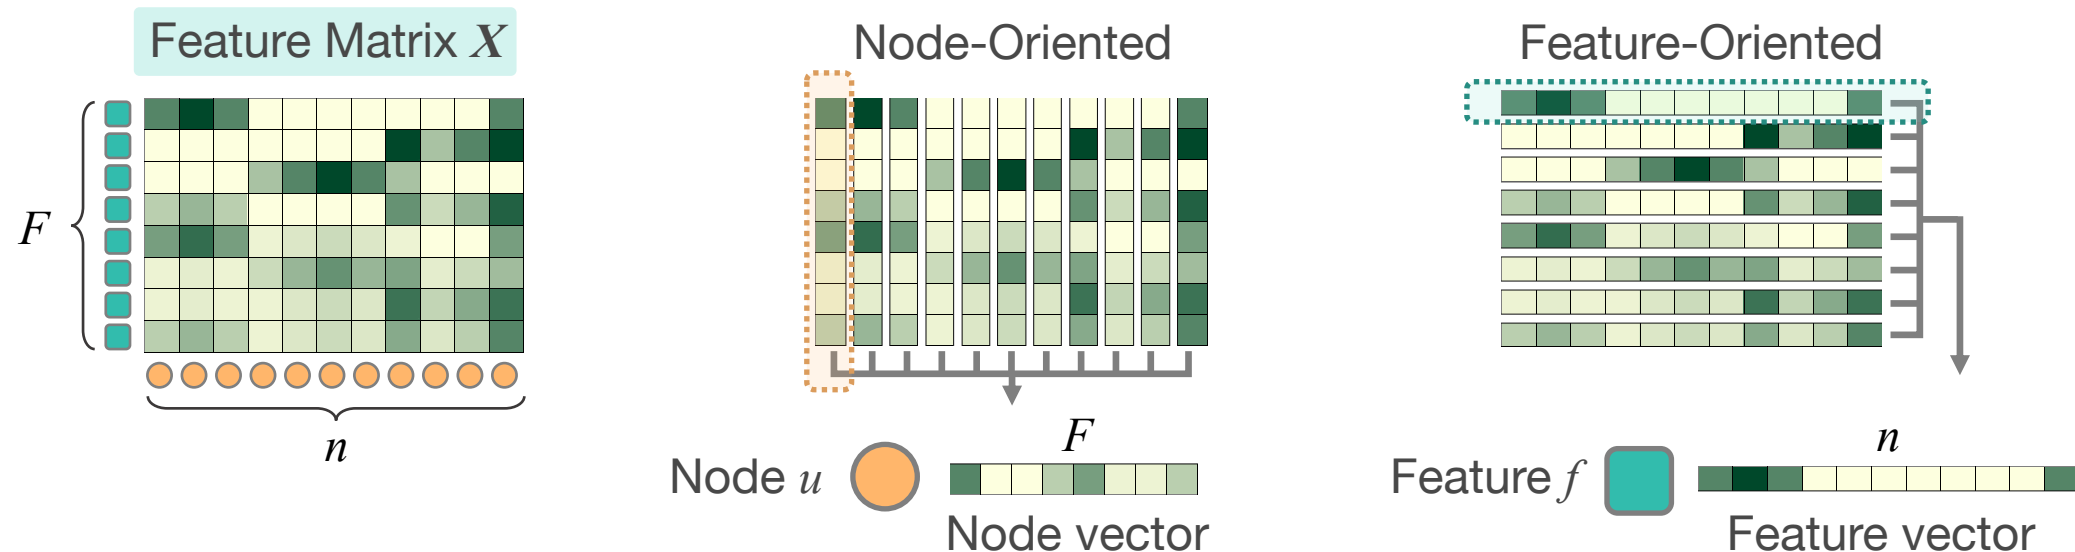
*Not suitable for minibatch*

*Not scalable to large  $m$  and  $n$*



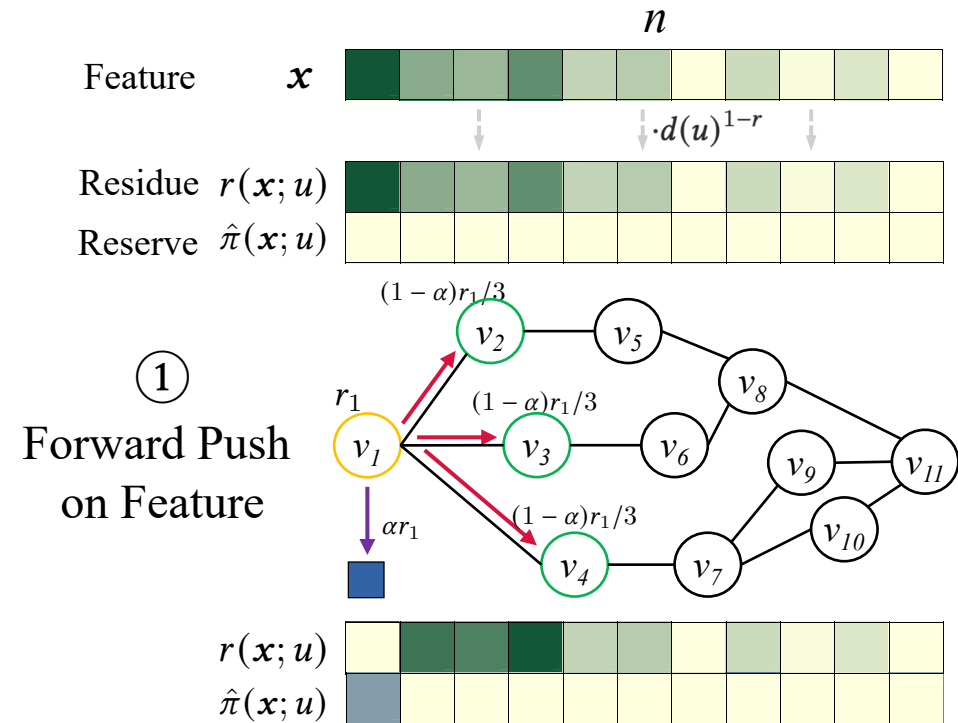
# Method: SCARA Framework

- Efficient propagation with feature-oriented algorithms
  - **FEATURE-PUSH**: *single* feature vector-based propagation
  - **FEATURE-REUSE**: reuse among *multiple* features



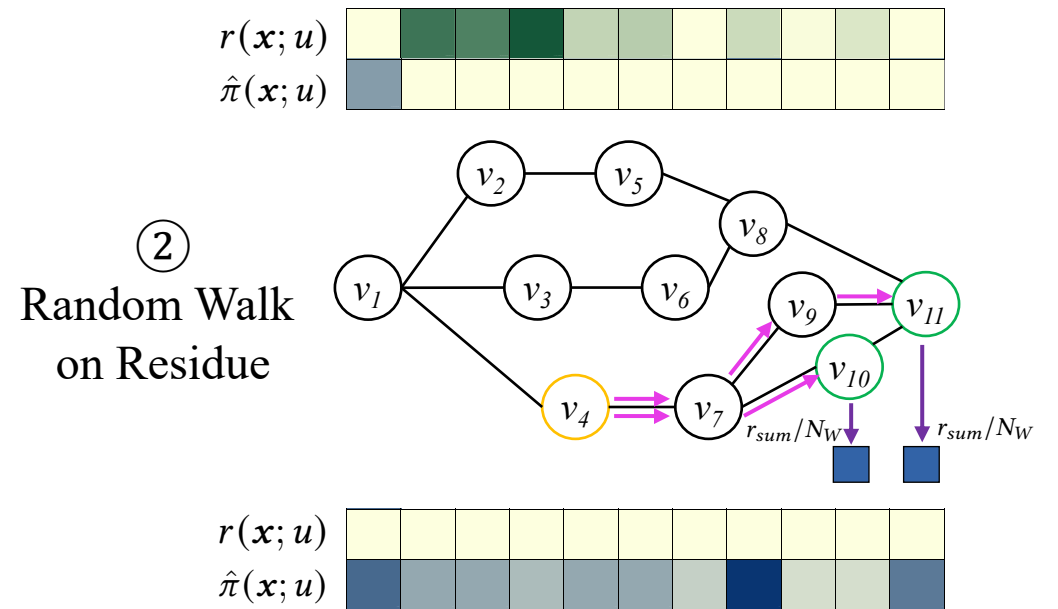
# Method: FEATURE-PUSH

- INITIALIZATION: feature vector as residue variable
- STEP ①: Forward Push on Feature Value
  - Residue  $r$ : values pending push
  - Reserve  $\pi$ : underestimation of embedding value
  - Complexity:  $O(\|\mathbf{x}\|_1/r_{max})$



# Method: FEATURE-PUSH

- STEP ②: Random Walk on Feature Residue
  - Increase reserves of end nodes based on walks
  - Precision guarantee
  - Complexity:  $O(m \cdot r_{max}/\beta)$
- COMBINATION: Push Parameter  $\beta$ 
  - Overall complexity:  $O(\sqrt{\frac{m\|\mathbf{x}\|_1}{\beta}})$   
 $\implies O(\sqrt{m \log n}/\lambda)$



# Method: FEATURE-REUSE

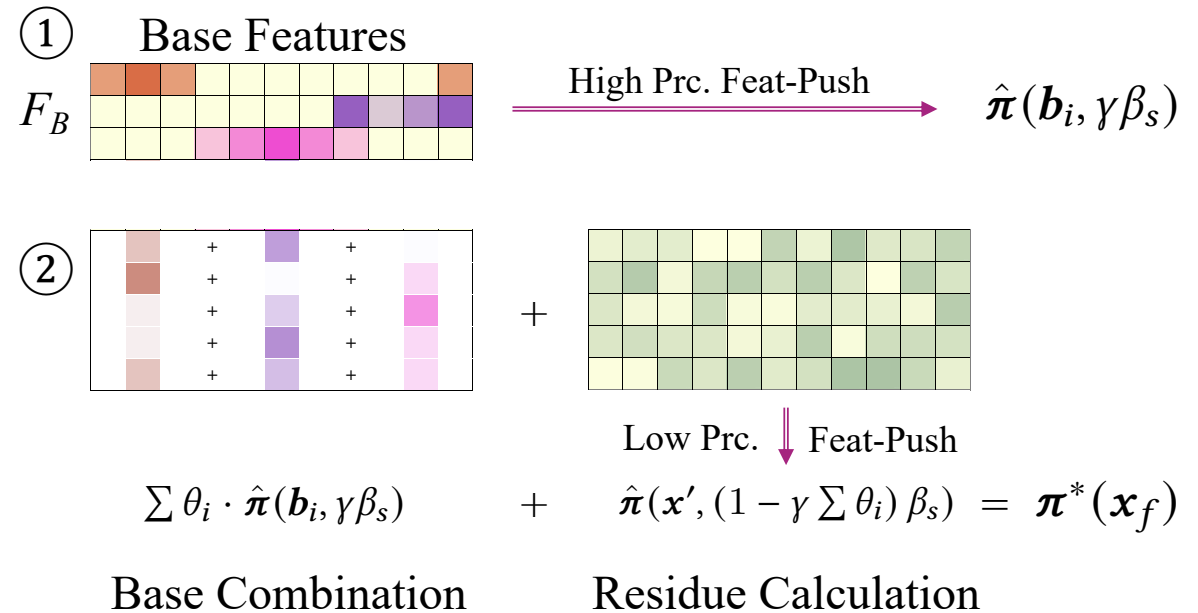
- Base Features

- Select  $F_B \ll F$  base features
- High Prc. FEAT-PUSH with  $\beta = \gamma\beta_s$

- Non-Base Features

- Linear combination of base calculation results
- Low Prc. FEAT-PUSH: sparse vectors thus faster

- Complexity:  $O\left(\sqrt{\frac{m(1 - \theta_{sum})}{\beta_s(1 - \gamma\theta_{sum})}}\right)$



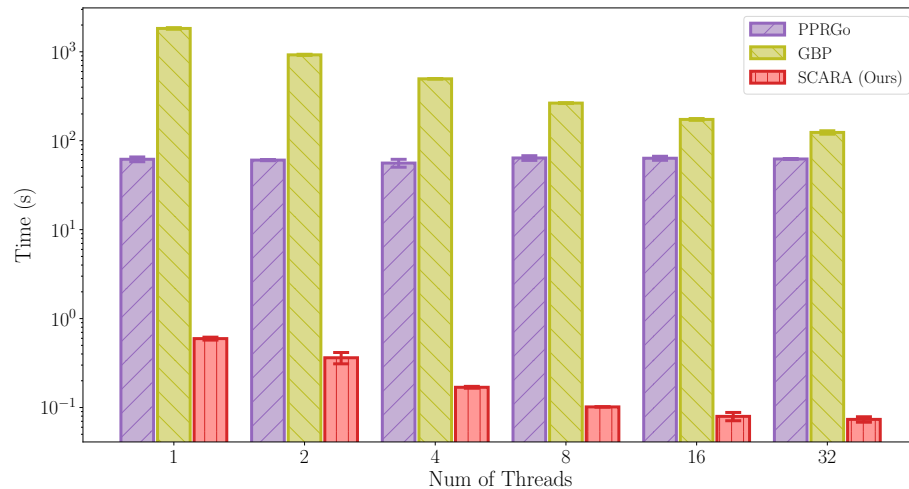
# Experimental Evaluation

- **Time Efficiency:** 30-800× faster parallel precomputation, comparable or better training and inference clock time
- **Memory Efficiency:** Paper100M with 72GB without OOM
- **Effectiveness:** similar or better F1-score, fast convergence

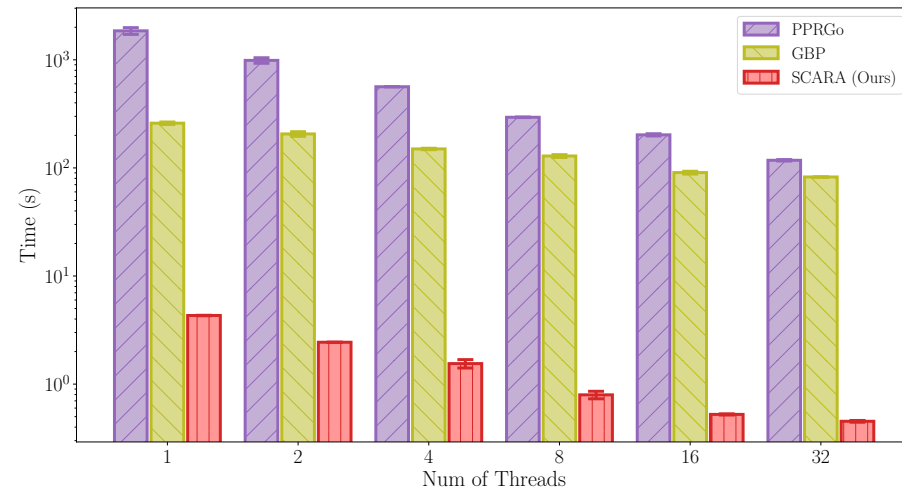
	Dataset	Nodes $n$	Edges $m$	Features $F$		Dataset	Nodes $n$	Edges $m$	Features $F$		
	Reddit [15]	232,965	114,615,892	602		Papers100M [16]	111,059,956	1,615,685,872	128		
Transductive	Reddit					Papers100M					
	Learn	( Pre. + Train)	Infer	Mem.	F1	Learn	( Pre. + Train)	Infer	Mem.	F1	
GraphSAINT	14.4	( - 14.4)	166.2	13.7	$41.6 \pm 4.8$	-	-	-	-	OOM	-
GAS	1151	( - 1151)	<b>2.2</b>	14.0	$38.2 \pm 0.3$	-	-	-	-	OOM	-
PPRGo	79.4	(62.3 + 17.1)	29.1	9.4	$41.5 \pm 2.3$	-	-	-	-	OOM	-
GBP	138	( 124 + 13.7)	13.5	<u>7.9</u>	$38.8 \pm 0.3$	-	-	-	-	OOM	-
<b>SCARA (ours)</b>	<b>13.9</b>	(0.07 + 13.8)	<u>10.7</u>	<b>5.6</b>	<b><math>44.1 \pm 0.4</math></b>	<b>1346</b>	(12.7 + 1333)	<b>4.7</b>	<b>71.4</b>	<b><math>35.7 \pm 0.9</math></b>	

# Experimental Evaluation

- **Time Efficiency:** 30-800× faster parallel precomputation, comparable or better training and inference clock time
- **Memory Efficiency:** Paper100M with 72GB without OOM
- **Effectiveness:** similar or better F1-score, fast convergence



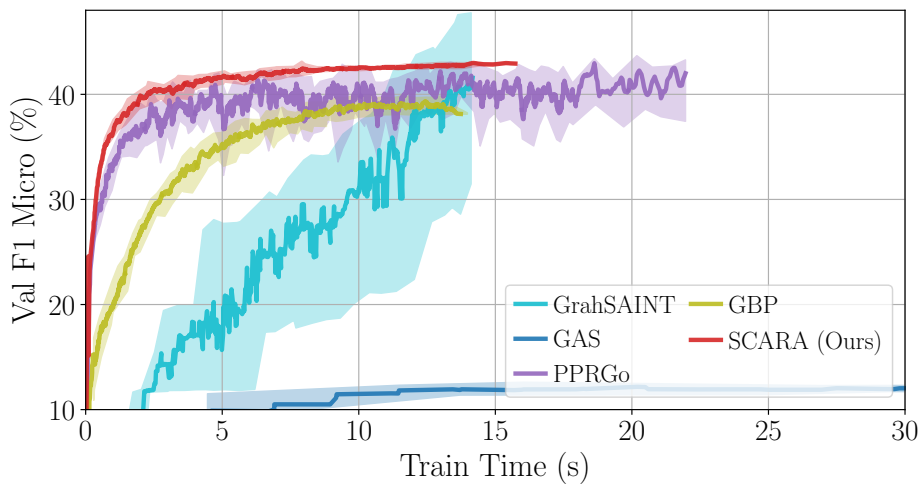
(A) Reddit



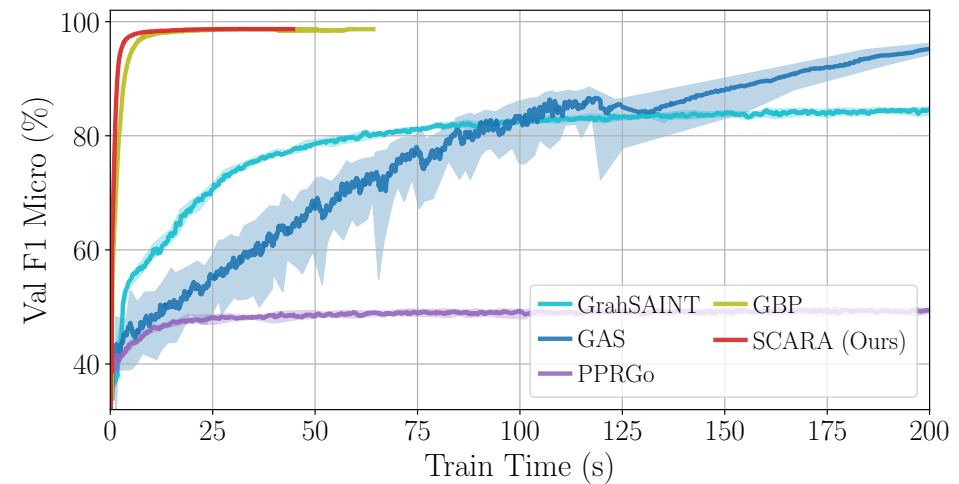
(B) Amazon

# Experimental Evaluation

- **Time Efficiency:** 30-800× faster parallel precomputation, comparable or better training and inference clock time
- **Memory Efficiency:** Paper100M with 72GB without OOM
- **Effectiveness:** similar or better F1-score, fast convergence



(a) Reddit



(b) PPI

# Summary

- **SCARA Framework:** Pre-Propagation Decoupled Model, optimized propagation with fast GPU batch training and inference
- **FEATURE-PUSH:** feature-oriented fast vector-based propagation, sub-linear precomputation complexity
- **FEATURE-REUSE:** efficient reuse among multiple features, further saves computation with guaranteed precision
- **Performance Evaluation:** up to 800× faster precomputation, able to process billion-scale Papers100M in 13 seconds



# LD<sup>2</sup>: Scalable Heterophilous GNN with Decoupled Embedding

**Ningyi Liao**, Siqiang Luo, Xiang Li, Jieming Shi. “LD<sup>2</sup>: Scalable Heterophilous Graph Neural Network with Decoupled Embedding”. Under submission of *the 37<sup>th</sup> Conference on Neural Information Processing Systems, 2023*.

# Contents

Introduction

Scalable GNN with Feature-Oriented Optimization

## **Scalable Heterophilous GNN with Decoupled Embedding**

**Motivation**

**Method**

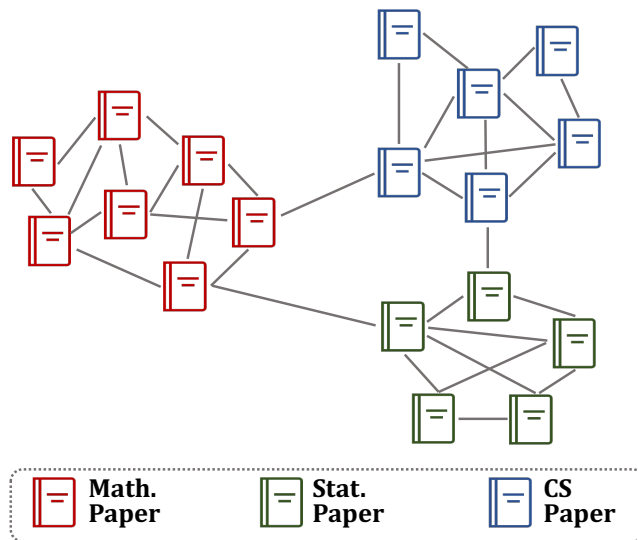
**Experimental Evaluation**

Conclusion and Future Works

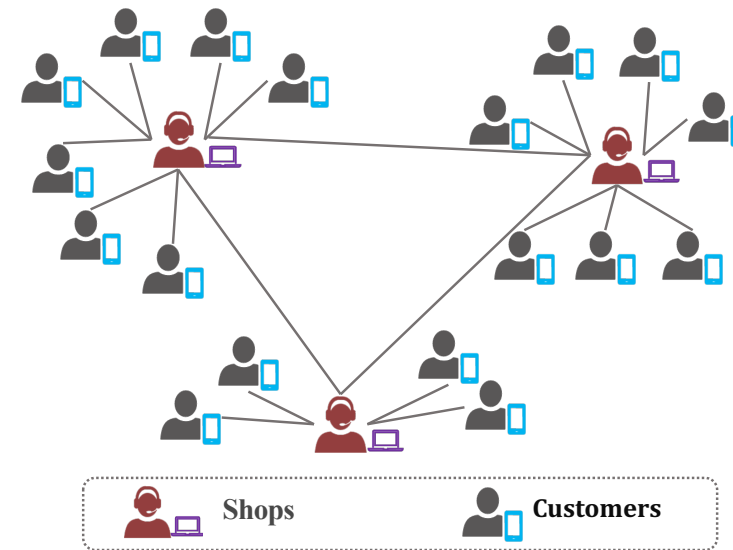
Q&A

# Graph Heterophily

- Homophily: connected nodes tend to be of *similar* classes
- Heterophily: connected nodes tend to be of *dissimilar* classes



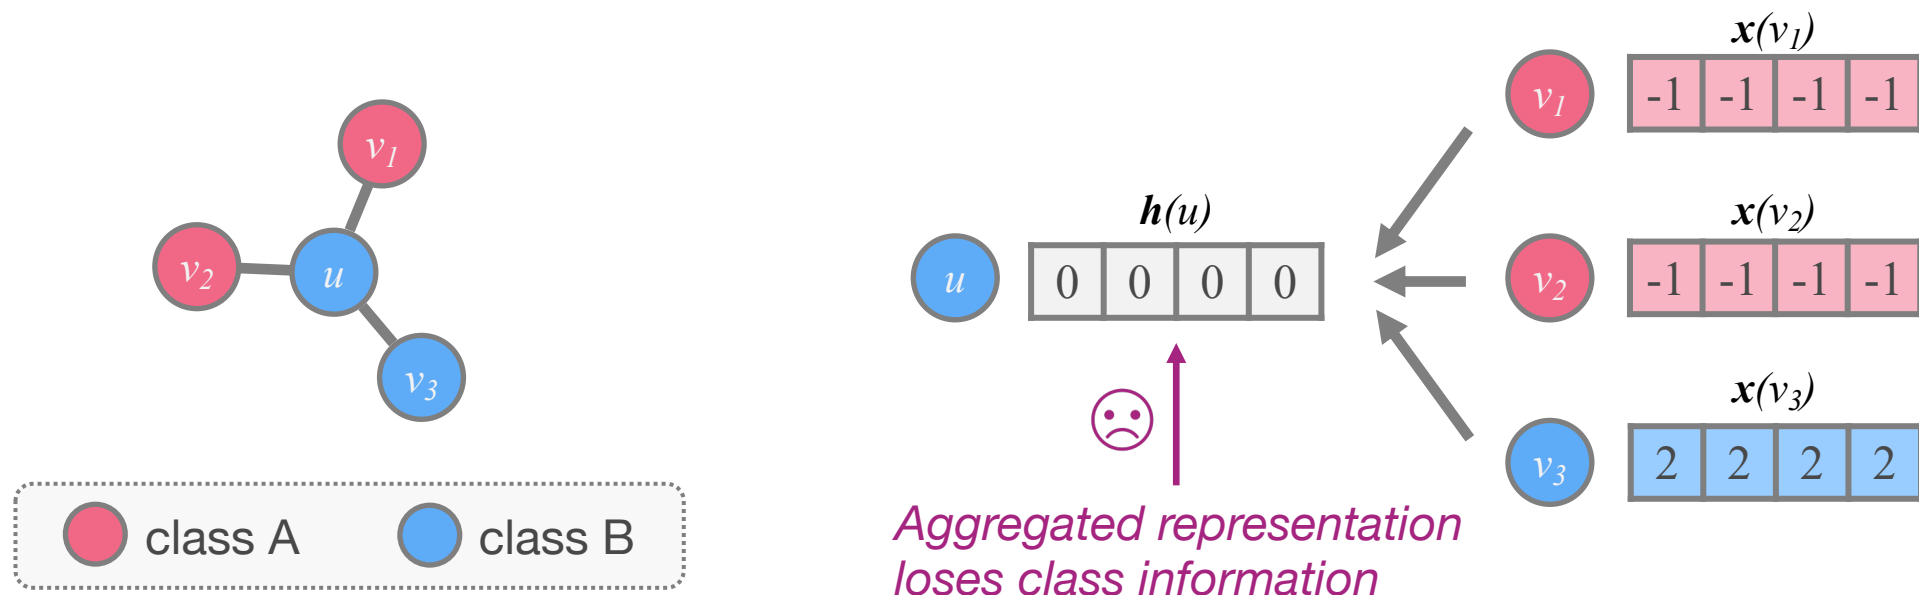
Homophilous Graph: citation network



Heterophilous Graph: transaction network

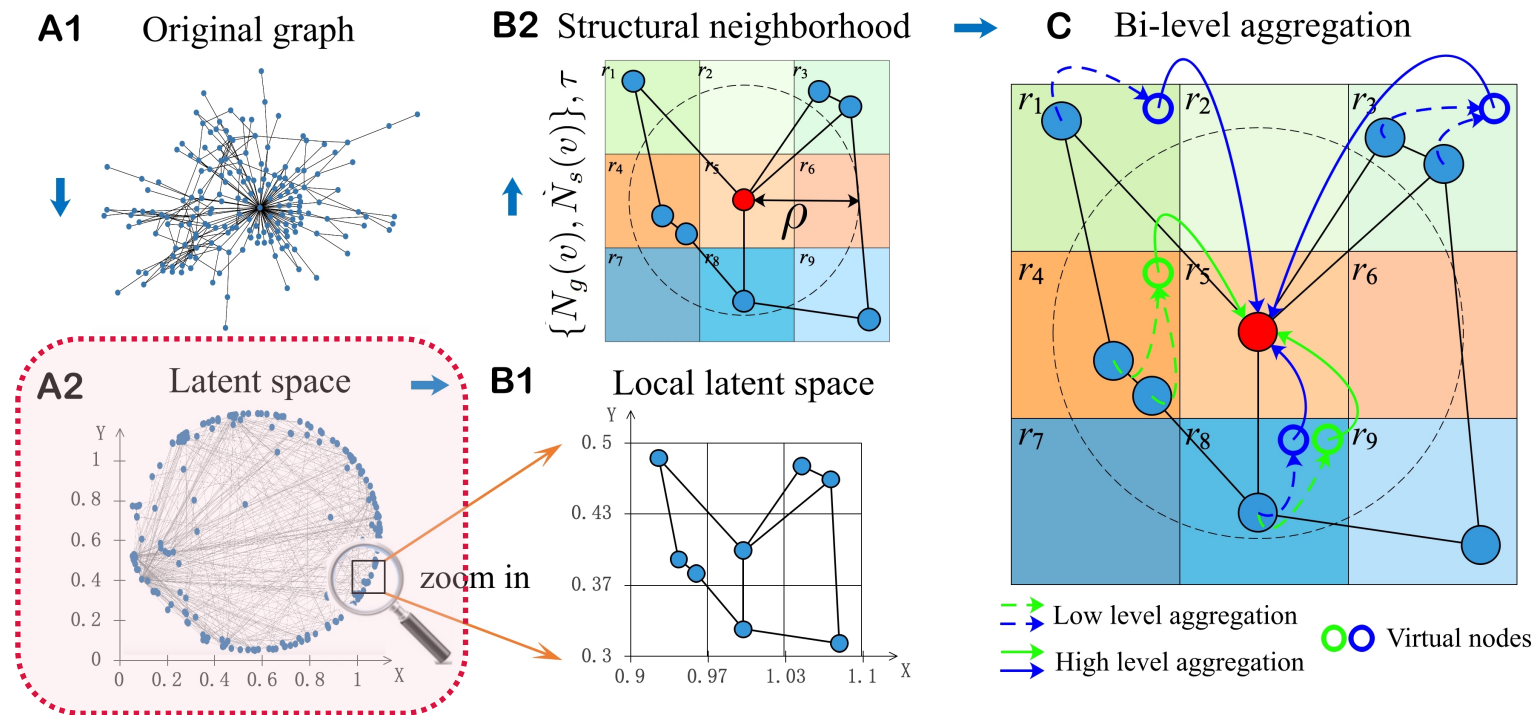
# Graph Heterophily

- Locality-based GNNs not suitable for heterophilous graphs



# Existing Work: Full-Graph Information

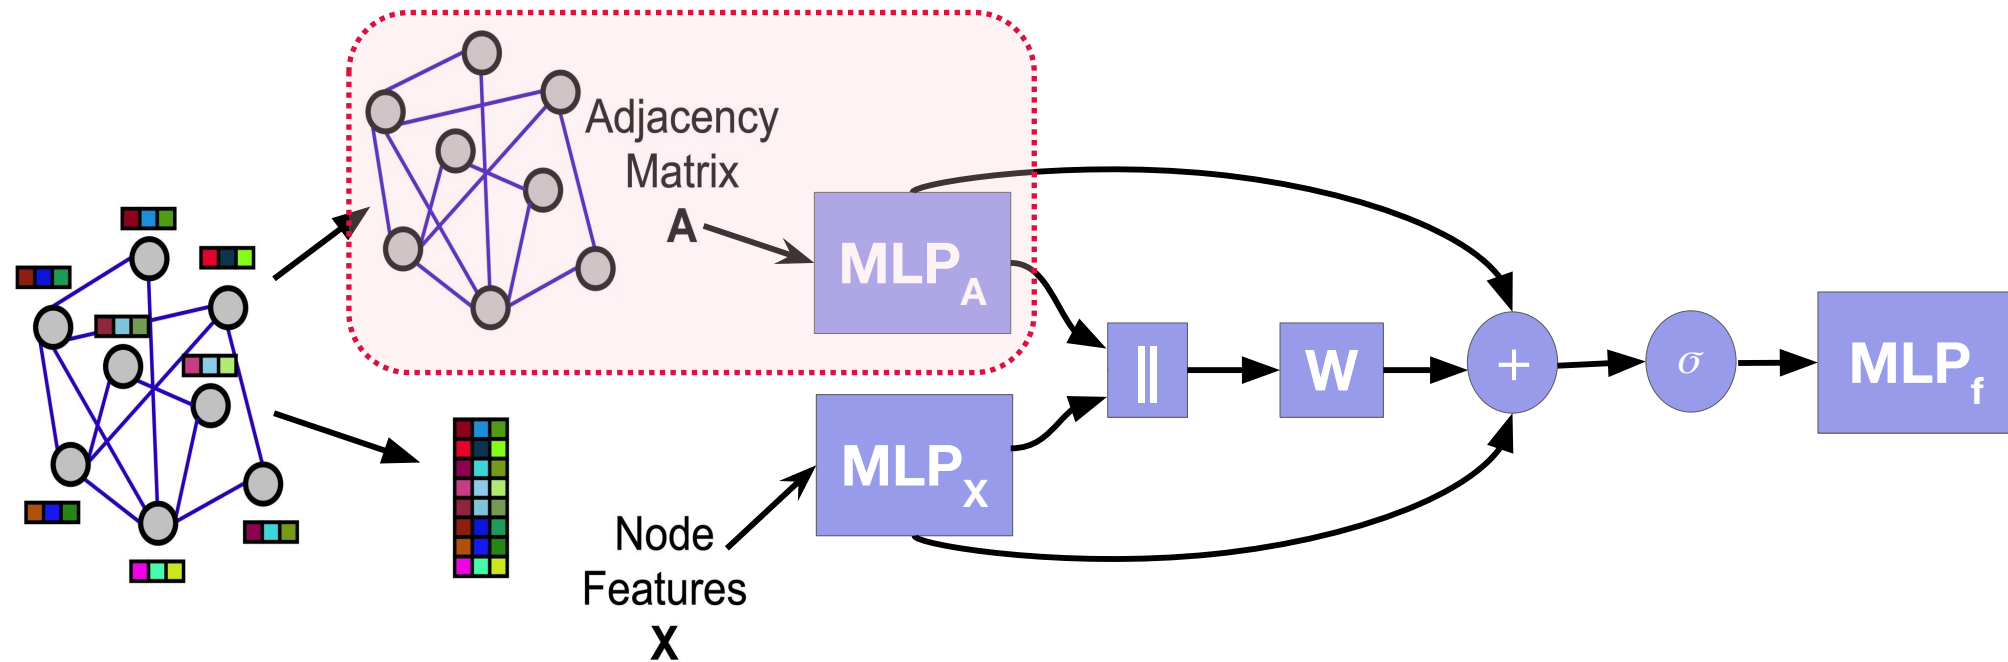
- Heterophilous GNNs usually rely on whole-graph information



Geom-GCN: full graph embedding

# Existing Work: Full-Graph Information

- Heterophilous GNNs usually rely on whole-graph information



LINKX: entire graph adjacency matrix

# Scalability Issue of Heterophilous GNNs

- Natural conflict: whole-graph processing not suitable for large-scale and minibatch tasks

TABLE 2.3: Precomputation, training, and inference time complexity of heterophilous GNNs.

Model	Precomp. Time	Training Time	Inference Time
GPRGNN [39]	$O(m)$	$O(IL_P m F + ILnF^2)$	$O(L_P m F + LnF^2)$
GCNJK [40]	–	$O(ILmF + ILnF^2)$	$O(LmF + LnF^2)$
MixHop [41]	–	$O(IL_P LmF + ILnF^2)$	$O(L_P LmF + LnF^2)$
LINKX [42]	–	$O(ImF + ILnF^2)$	$O(mF + LnF^2)$
<b>LD<sup>2</sup> (ours)</b>	$O(L_P m F)$	$O(ILnF^2)$	$O(LnF^2)$

☹ Terms that not scalable to large  $m$  and  $n$

TABLE 2.4: Precomputation, training, and inference memory complexity of heterophilous GNNs.

Model	Precomp. Mem.	Training Mem.	Inference Mem.
GPRGNN [39]	$O(m)$	$O(LnF + LF^2 + m)$	$O(LnF + LF^2 + m)$
GCNJK [40]	–	$O(L_C n F + L_C F^2)$	$O(L_C n F + L_C F^2)$
MixHop [41]	–	$O(CLnF + CLF^2)$	$O(CLnF + CLF^2)$
LINKX [42]	–	$O(L_C n_b F + L_C F^2 + nF)$	$O(L_C n_b F + L_C F^2 + nF)$
<b>LD<sup>2</sup> (ours)</b>	$O(CnF)$	$O(L_C n_b F + L_C F^2)$	$O(L_C n_b F + L_C F^2)$

☹ Terms that not suitable for minibatch

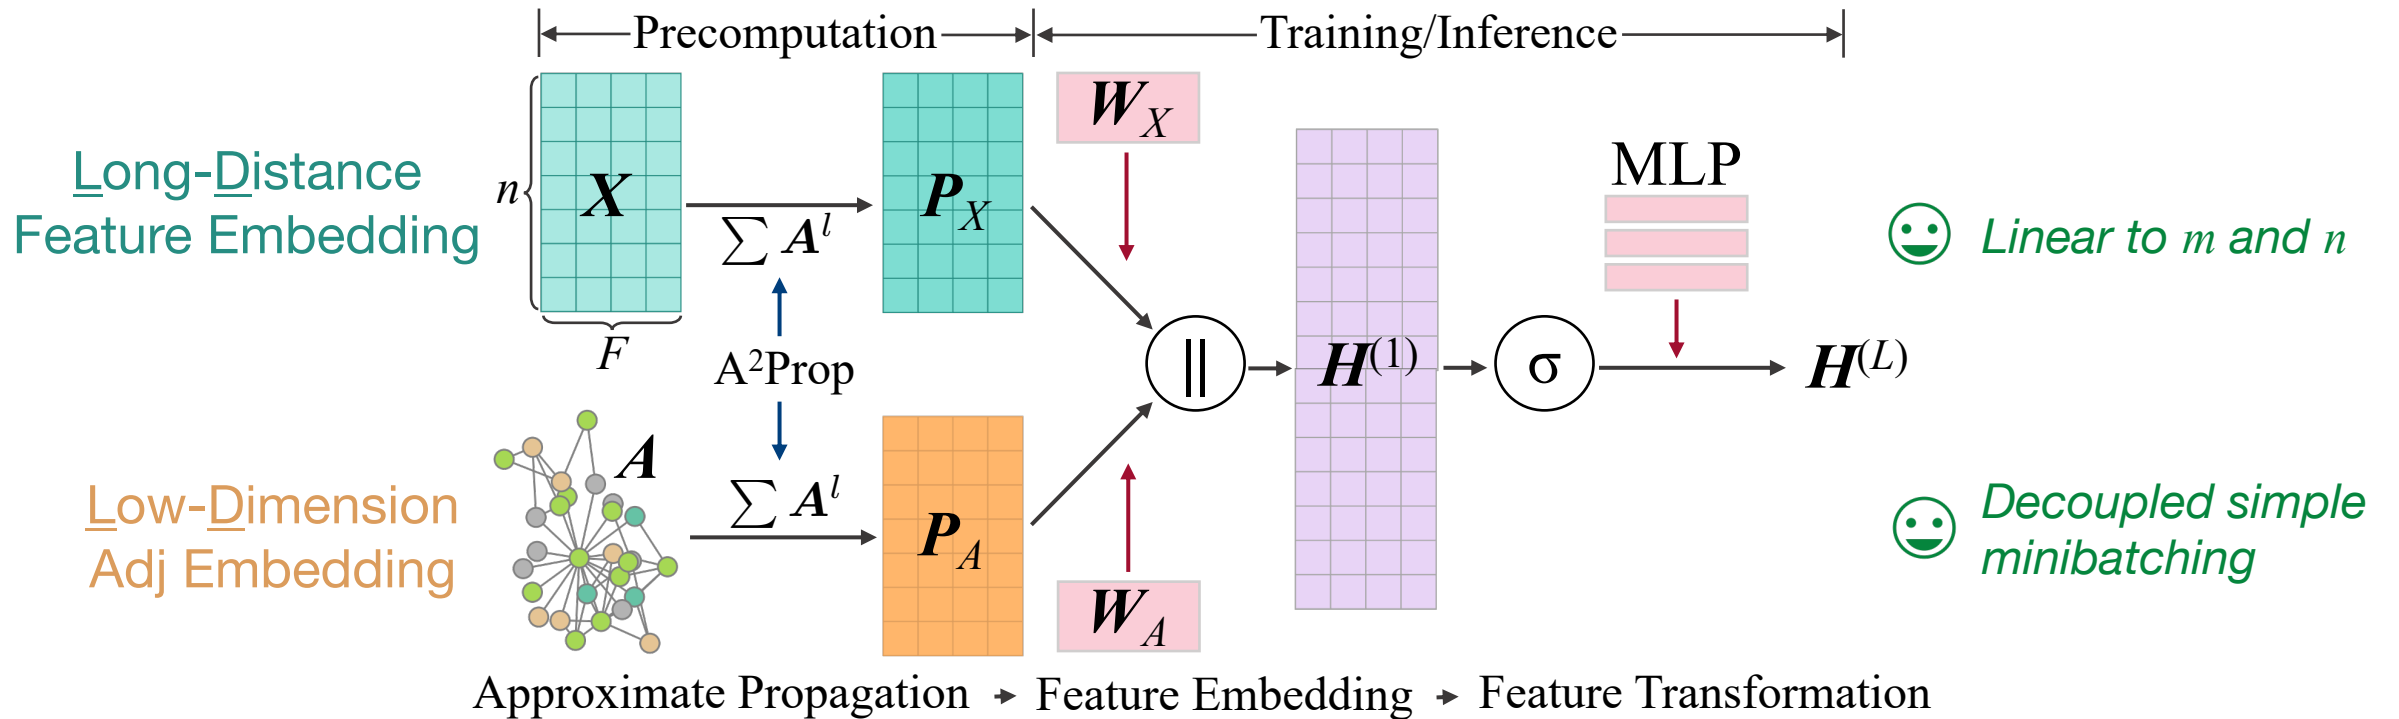
# Our Model: LD<sup>2</sup>

- Precomputation:

$$P_A, P_X = A^2\text{Prop}(A, X)$$

- Feature Transformation:

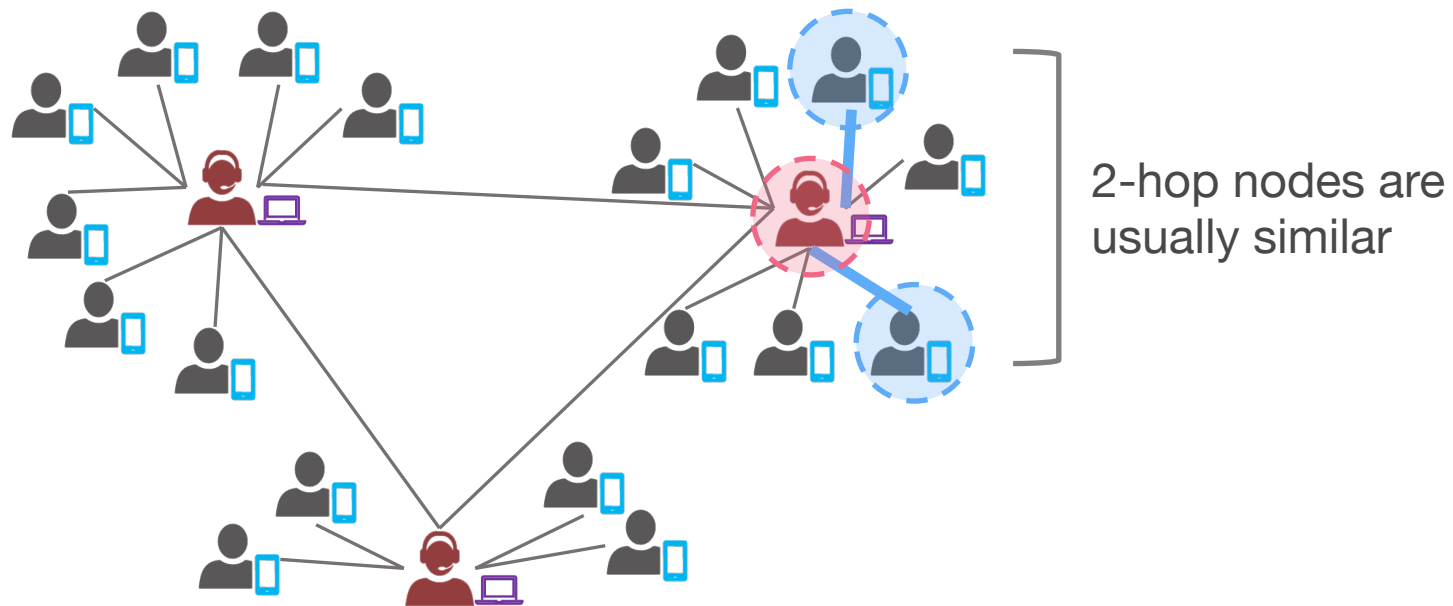
$$H^{(L)} = \text{MLP}(P_A W_A \| P_X W_X)$$





# Method: Adjacency Embedding

- 2-hop neighbors are proved to be homophily-dominant regardless of 1-hop neighbor distribution



# Method: Adjacency Embedding

- Low-Dimension 2-hop adj decomposition as embedding:

$$\underbrace{P_A}_{n \times F} \cdot \underbrace{P_A^\top}_{n \times F} \approx \underbrace{A^2}_{n \times n}$$

Adjacency Embedding 2-hop Adjacency Matrix

- Calculation:

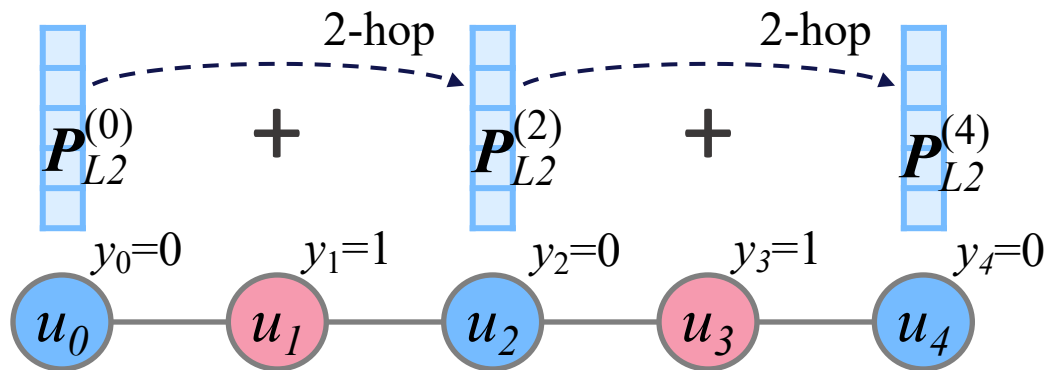
$$\begin{aligned}
 P_A^{(0)} &= \mathcal{N} \quad \text{Gaussian Matrix} \\
 P_A^{(i+1)} &= \phi(A^2 \cdot P_A^{(i)}) \quad \text{Power Iteration}
 \end{aligned}
 \Leftrightarrow
 \begin{aligned}
 P_A &= \phi(A^2 \phi(A^2 \dots \phi(A^2 \mathcal{N}))) \quad \text{Iterative Adj Multiplication} \\
 &\quad \text{Orthogonalize + Normalize}
 \end{aligned}$$

# Method: Feature Embedding

- Long-Distance generalized graph propagation

CHANNEL ①: *Constant* 2-hop  
Adjacency Propagation

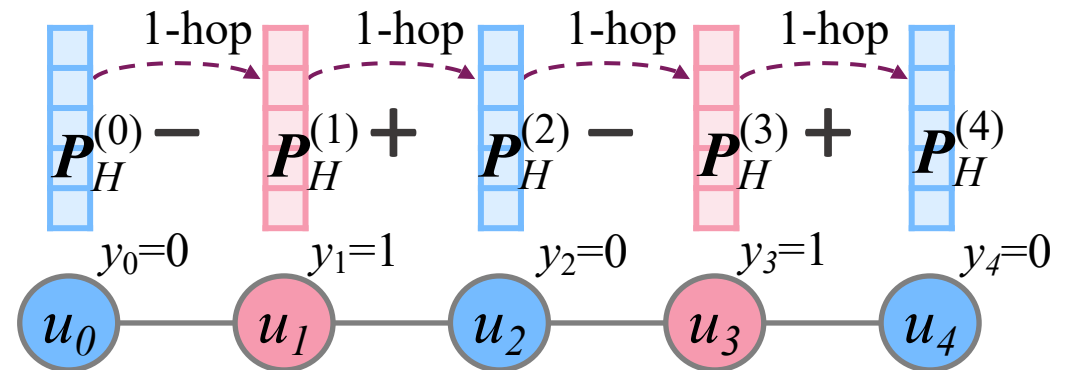
$$P_{X,L2} = \sum_{l=1}^L \bar{A}^{2l} \cdot X$$



CHANNEL ②: *Inverse* 1-hop  
Laplacian Propagation

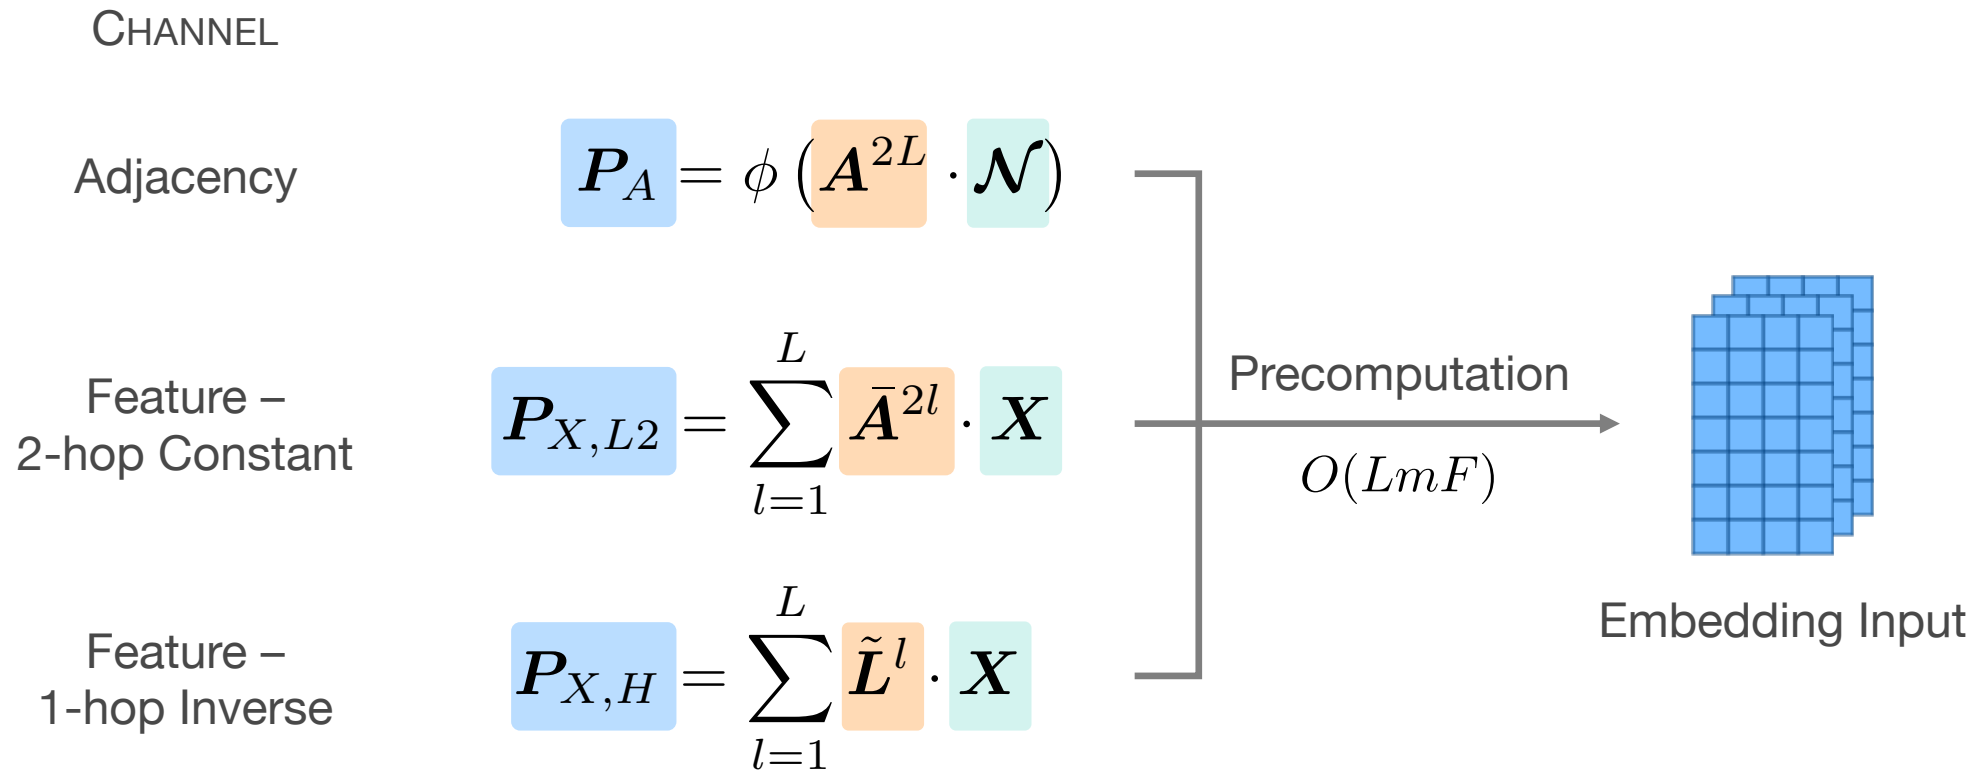
$$P_{X,H} = \sum_{l=1}^L \tilde{L}^l \cdot X$$

Laplacian Matrix  $L = I - A$



# Method: Embedding Calculation

- Multi-channel embedding, one-time computation:



# Experimental Evaluation

- **Effectiveness:**
  - Top 1 accuracy on 6/8 graphs
  - No accuracy drop for minibatch

Dataset	squirrel	genius	penn94	arxiv-year	twitch-gamers	pokec	snap-patents	wiki
Nodes $n$	5,201	421,858	41,536	169,343	168,114	1,632,803	<b>2,738,035</b>	1,770,981
Edges $m$	401,907	1,344,722	1,403,756	1,327,142	6,965,671	23,934,767	16,705,984	<b>244,278,050</b>
$F / N_c$	2,089 / 5	12 / 2	4,814 / 2	128 / 5	7 / 2	65 / 2	269 / 5	600 / 5
GCNJK-GS	27.63 $\pm$ 4.72	80.65 $\pm$ 0.07	65.91 $\pm$ 0.16	48.26 $\pm$ 0.64	59.91 $\pm$ 0.42	59.38 $\pm$ 0.21	33.64 $\pm$ 0.05	42.95 $\pm$ 0.39
MixHop-GS	33.24 $\pm$ 2.44	80.63 $\pm$ 0.04	75.00 $\pm$ 0.37	49.26 $\pm$ 0.16	61.80 $\pm$ 0.00	64.02 $\pm$ 0.02	34.73 $\pm$ 0.15	45.52 $\pm$ 0.11
LINKX	60.14 $\pm$ 0.92	82.51 $\pm$ 0.10	<b>78.63</b> $\pm$ 0.25	<b>50.44</b> $\pm$ 0.30	64.15 $\pm$ 0.18	68.64 $\pm$ 0.65	52.69 $\pm$ 0.05	50.59 $\pm$ 0.12
<b>LD<sup>2</sup> (ours)</b>	<b>66.87</b> $\pm$ 0.02	<b>85.31</b> $\pm$ 0.06	<u>75.52</u> $\pm$ 0.10	<u>50.29</u> $\pm$ 0.11	<b>64.33</b> $\pm$ 0.19	<b>74.93</b> $\pm$ 0.10	<b>58.58</b> $\pm$ 0.34	<b>52.91</b> $\pm$ 0.16

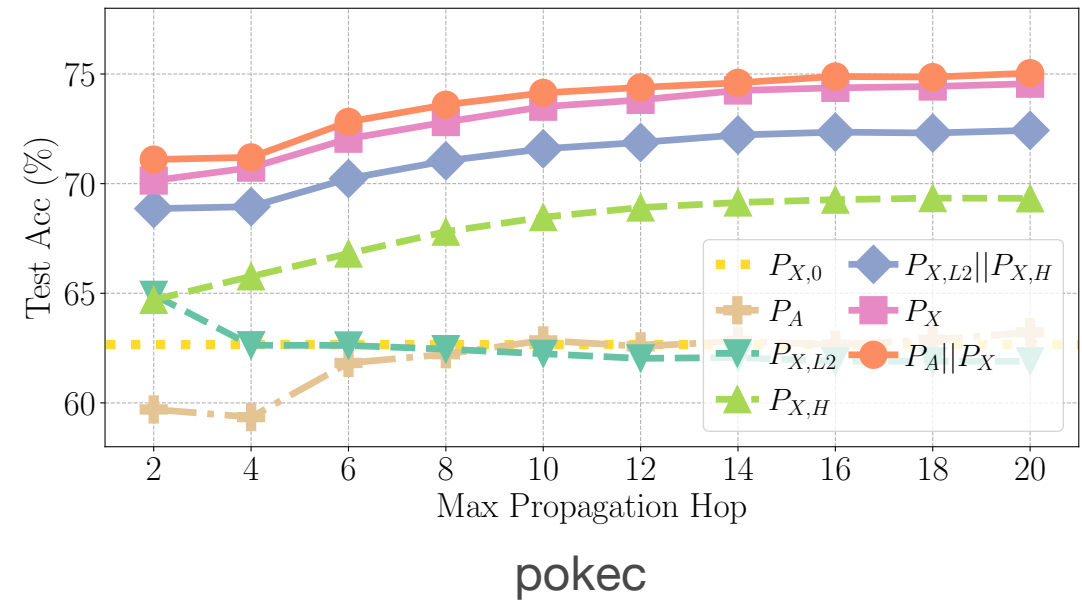
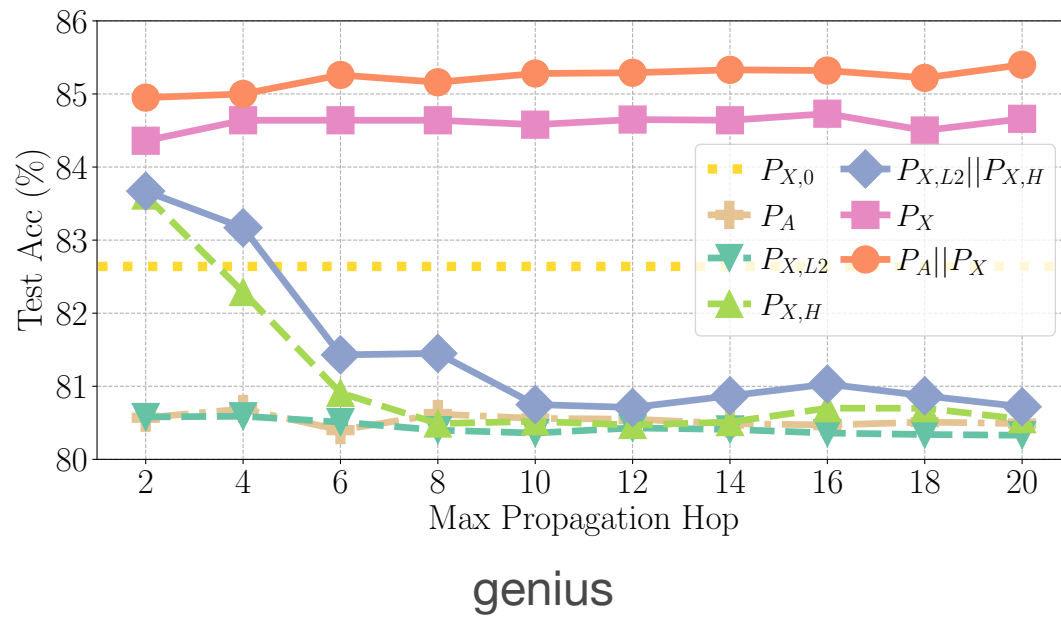
# Experimental Evaluation

- **Efficiency:**
  - 3-15× faster minibatch training, significantly fast inference
  - Up to 5× lower memory for large graphs

Dataset	twitch-gamers			pokec			snap-patents			wiki		
	Learn	Infer	Mem.	Learn	Infer	Mem.	Learn	Infer	Mem.	Learn	Infer	Mem.
MLP	6.36	0.02	0.61	47.86	0.11	13.77	27.39	0.28	9.33	133.55	0.62	18.15
PPRGo	10.46+15.88	0.41	9.64	121.95+56.11	2.69	3.82	(>12h)			(>12h)		
SGC	0.09+0.74	0.01	0.28	1.05+8.08	0.01	0.28	4.94+23.54	0.01	0.42	12.66+7.98	0.01	0.52
GCNJK-GS	71.48	0.02*	7.33	27.33	0.09*	9.03	19.02	0.23*	9.21	95.52	0.69*	16.36
MixHop-GS	52.12	0.01*	1.49	71.35	0.03*	12.91	45.24	0.16*	19.58	84.22	0.23*	16.28
LINKX	10.99	0.19	2.35	28.77	0.33	9.03	39.80	0.22	21.53	180.71	1.14	14.53
<b>LD<sup>2</sup> (ours)</b>	0.85+ <b>1.96</b>	<b>0.01</b>	<b>1.44</b>	17.95+ <b>6.18</b>	<b>0.01</b>	<b>3.82</b>	31.32+ <b>6.96</b>	<b>0.02</b>	<b>3.96</b>	28.12+ <b>6.50</b>	<b>0.01</b>	<b>4.47</b>

# Experimental Evaluation

- **Ablation Study:**
  - Combination of channels is effective for different graphs



# Summary

- **LD<sup>2</sup> Framework:** Pre-Propagation Decoupled Heterophilous GNN, optimized minibatch training
- **Low Dimension Adjacency Embedding:** embed full graph topology in low dimensional matrix representation
- **Long Distance Feature Embeddings:** multi-channel node features by different graph propagation schemes
- **Performance Evaluation:** 3–15× faster minibatch training and inference, up to 5× smaller memory footprint



# Contents

Introduction

Scalable GNN with Feature-Oriented Optimization

Scalable Heterophilous GNN with Decoupled Embedding

**Conclusion and Future Works**

Q&A

# Conclusion

- CONTRIBUTION ①: **SCARA**
  - Scalable GNN with decoupled efficient graph propagation and feature-oriented optimizations
  - Sub-linear complexity, fast precomputation, 1.6B graph in 13 sec
- CONTRIBUTION ②: **LD<sup>2</sup>**
  - Scalable heterophilous GNN with decoupled multi-channel topology and feature embeddings
  - Minibatch ability, simplified learning, 240M graph in 40 sec

# Future Works

- **Broader Range of Models:**
  - Extend propagation schemes: iterative-, post-propagation
  - Apply on more architectures: GAT, Graph Transformer
- **Variants of Graph Data:**
  - Heterogeneous Graph: dissimilar edges and propagations
  - Dynamic Graph: nodes and edges change with time
- **Benchmarking GNN Scalability:**
  - Evaluate performance of different scalable GNN designs
  - Explore efficacy-efficiency tradeoff

# Contents

**Introduction**

**Scalable GNN with Feature-Oriented Optimization**

**Scalable Heterophilous GNN with Decoupled Embedding**

**Conclusion and Future Works**

**Q&A**

THANK YOU