NEURAL INFORMATION PROCESSING SYSTEMS

Learning Pseudo-Riemannian Manifolds for Heterophilic Graphs via Graph Signature

Yun Young Choi[1*] **Asung Kil**[1*] **Sun Woo Park**[2*] **Minho Lee**[1] **Seokhwan Kim**[1]

[1]SolverX, Seoul [2]Max Planck Institute for Mathematics, Bonn [*]Equal contribution

NeurIPS 2026

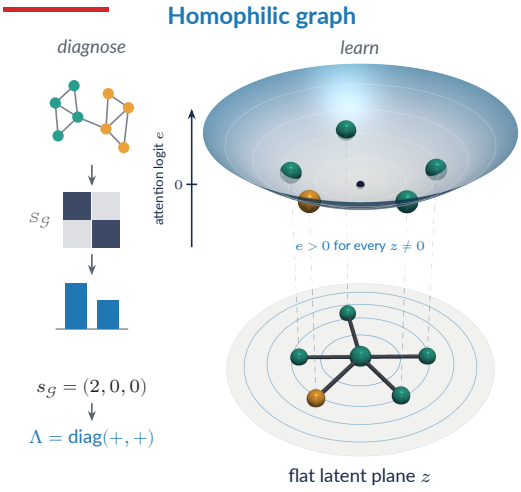
SolverX

Why: many GNNs carry a positive-definite prior

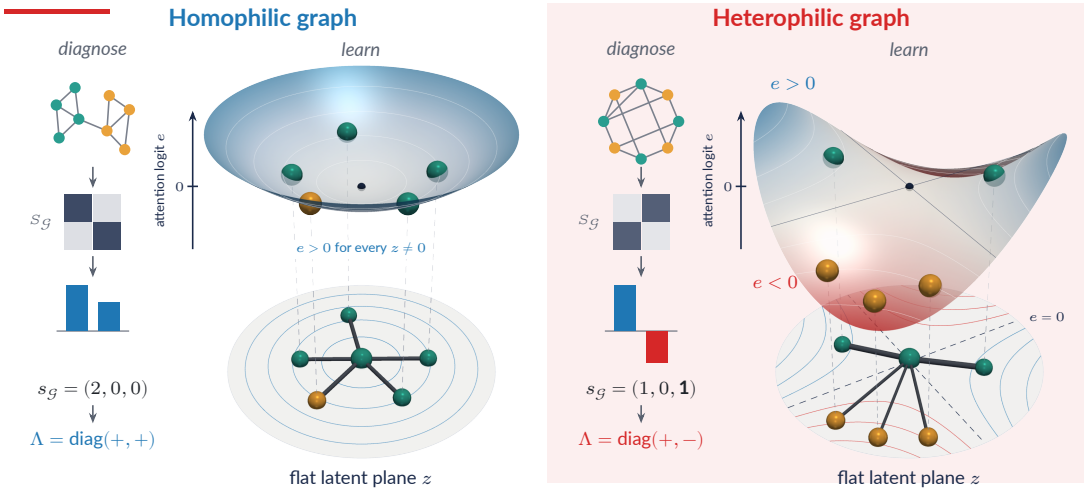
- Heat-diffusion convolution: **positive semi-definite** operator
- Hyperbolic GNNs: **positive-definite** manifolds
- Dot-product and GAT attention: **positive-definite** inner products

Heterophilic edges often join different classes, but **no positive-definite form gives a negative squared distance** to a direction the labels separate.

PRAT at a glance

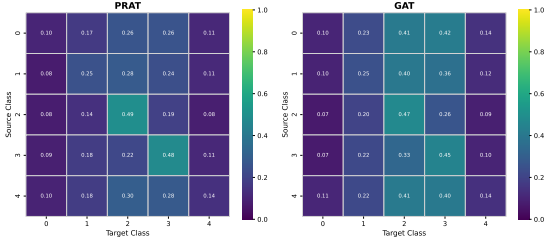


PRAT at a glance



Measured attention: homophilic, then heterophilic

Homophilic, $h_E = 0.9$

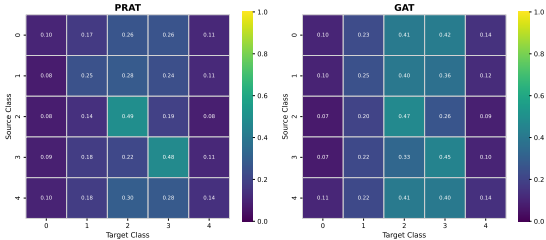


accuracy **PRAT 90.6%** GAT 90.2%

Synthetic graphs (Section 4.4). Mean attention from source class (row) to target class (column).

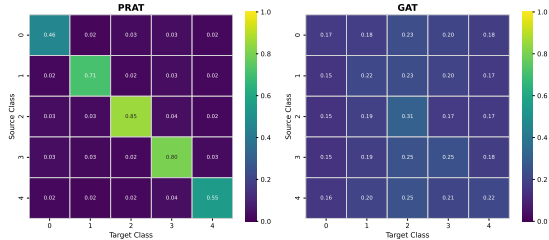
Measured attention: homophilic, then heterophilic

Homophilic, $h_E = 0.9$



accuracy **PRAT 90.6%** GAT 90.2%

Heterophilic, $h_E = 0.1$



accuracy **PRAT 80.3%** GAT 28.6%

PRAT moves its attention onto the diagonal (same class); GAT's map goes flat.

Synthetic graphs (Section 4.4). Mean attention from source class (row) to target class (column).

Diagnose: the Graph Signature Index

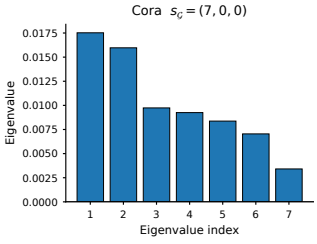
Class-pair edge density:

$$s_{i,j} = \frac{|\{(v,w) \in \mathcal{E} : l(v) = i,\ l(w) = j\}|}{n_i\, n_j}$$

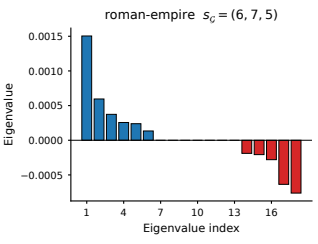
$s_{\mathcal{G}}$ = (#positive, #zero, #negative) eigenvalues of $S_{\mathcal{G}}$

Negative eigenvalue ⇒ **indefinite**

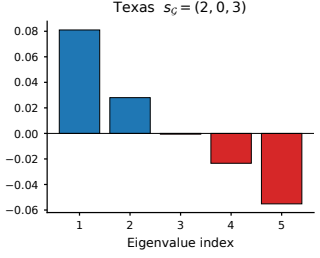
Positive definite



Indefinite, with zero eigenvalues



Indefinite, without zero eigenvalues



$|\lambda| < 10^{-4}$ counted as zero.

Diagnose: the Graph Signature Index

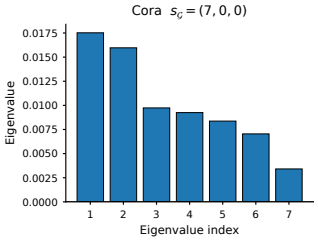
Class-pair edge density:

$$s_{i,j} = \frac{|\{(v,w) \in \mathcal{E} : l(v) = i,\ l(w) = j\}|}{n_i\, n_j}$$

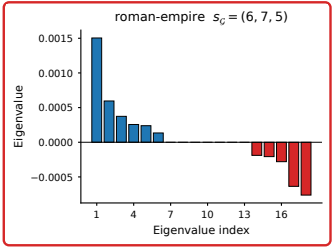
$s_{\mathcal{G}}$ = (#positive, #zero, #negative) eigenvalues of $S_{\mathcal{G}}$

Negative eigenvalue ⇒ **indefinite**

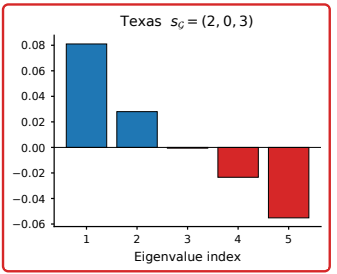
Positive definite



Indefinite, with zero eigenvalues



Indefinite, without zero eigenvalues



$|\lambda| < 10^{-4}$ counted as zero.

Two theorems

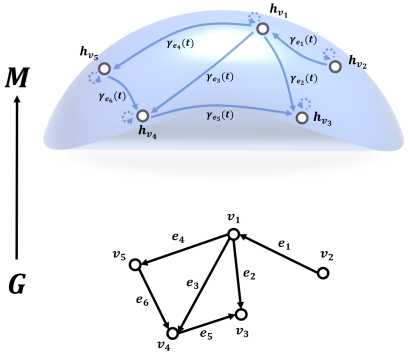
Theorem 2.4 (Limitation of homophily indices)
For every rational $h \in (0, 1)$, some graph has $h_E = h_V = h$ and a negative eigenvalue of $S_{\mathcal{G}}$.

$\Rightarrow$ Homophily cannot rule out an indefinite $S_{\mathcal{G}}$.

Theorem 2.5 (Geometric mismatch)
If $S_{\mathcal{G}}$ is indefinite and $T_p f$ is surjective for all p, no diffeomorphism ϕ gives $(\phi \circ f)^* I = f^* S_{\mathcal{G}}$.

$\Rightarrow$ The mismatch with a Riemannian prior is intrinsic.

Learn: Pseudo-Riemannian Attention Network (PRAT)



Nodes become points of a pseudo-Riemannian manifold $(\mathcal{M}, g)$; edges become curves.

$$e(h_v, h_u) = \begin{cases} \Delta(h_{v,u})^\top \mathbf{D}\, \Delta(h_{v,u}) & u \in \mathcal{N}_v \\ a\, h_v^\top \mathbf{D}\, h_v & u = v \end{cases}$$

$$\Delta(h_{v,u}) = h_v - h_u, \qquad \mathbf{D} = \mathbf{W}'^\top \Lambda\, \mathbf{W}'$$

$$\Lambda_{\text{init}} = \frac{\Lambda_{\mathcal{G},\text{train}}}{\text{std}(\Lambda_{\mathcal{G},\text{train}})} \cdot \frac{1}{\sqrt{|L|}}$$

- **D**: pullback of a **diagonal** Λ on a $|L|$-dimensional latent space
- Λ: initialized from the eigenvalues of $S_{\mathcal{G},\text{train}}$, then **learned**
- Cost: **GAT's asymptotic cost**

Fifteen node-classification benchmarks

Best mean rank on both suites: **2.63** vs. 2.88 (BernNet), **1.86** vs. 4.57 (GT-sep)

Standard benchmarks

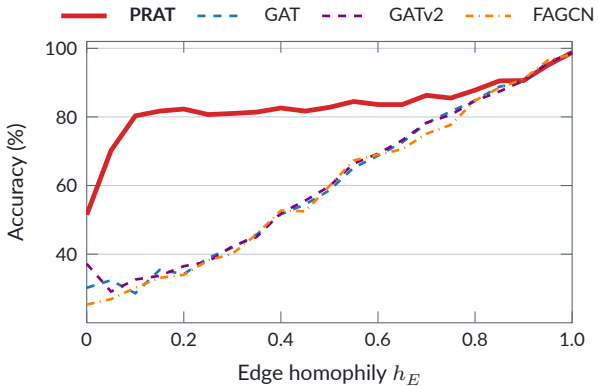
	$s_{\mathcal{G}}$	PRAT	Rank
Cora	(7,0,0)	88.42	6th
PubMed	(3,0,0)	90.03	3rd
Citeseer	(6,0,0)	78.06	**1st**
Chameleon	(3,0,**2**)	68.73	3rd
Squirrel	(2,2,**1**)	56.96	2nd
Actor	(1,4,0)	37.26	**1st**
Cornell	(3,0,**2**)	71.89	4th
Texas	(2,0,**3**)	84.32	**1st**

Platonov benchmarks

	$s_{\mathcal{G}}$	PRAT	Rank
Cham.-filt.	(3,0,**2**)	42.69	**1st**
Sq.-filt.	(2,1,**2**)	39.60	2nd
roman-empire	(6,7,**5**)	90.08	**1st**
amazon-ratings	(3,2,0)	53.78	**1st**
minesweeper*	(1,1,0)	92.24	4th
tolokers*	(2,0,0)	85.04	**1st**
questions*	(1,1,0)	78.04	3rd

Rank among all models of each table. *ROC AUC.

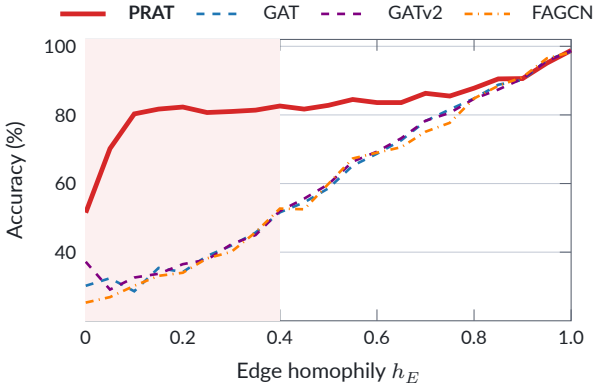
Synthetic graphs: vary only the homophily



PRAT: above 80% down to $h_E = 0.1$

Others: fall away below $h_E = 0.5$

Synthetic graphs: vary only the homophily

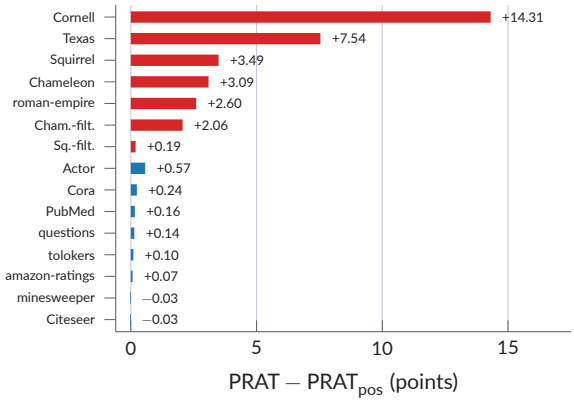


PRAT: above 80% down to $h_E = 0.1$

Others: fall away below $h_E = 0.5$

Gap to FAGCN at $h_E \leq 0.4$: **40.6 points** on average

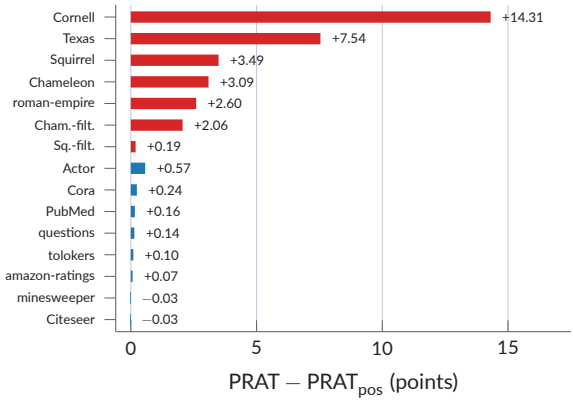
Ablation: the sign structure carries the gain



PRAT$_{pos}$: Λ kept positive definite (scaled-identity start)

- graphs with negative eigenvalues ($n_- > 0$)
- graphs without ($n_- = 0$)

Ablation: the sign structure carries the gain



PRAT$_{pos}$: Λ kept positive definite (scaled-identity start)

- graphs with negative eigenvalues ($n_- > 0$)
- graphs without ($n_- = 0$)

Paired Platonov ablation: every significant loss is on a graph with $n_- > 0$.

A negative signature is **necessary for a detectable loss**.

Takeaways

- Geometry is a **modeling choice**, not a default
- **Diagnose** with $s_{\mathcal{G}}$, then **learn** the metric with PRAT
- $s_{\mathcal{G}}$ is a **diagnostic**, not a predictor of accuracy

Limitation: the signature uses training labels; k-means pseudo-labels recovered the accuracy on both indefinite benchmarks tested.

Paper arxiv.org/abs/2611.06265
Code github.com/voltwin-dev/PRAT_NeurIPS_2026