Archetypal Graph Generative Models

Explainable & Identifiable Communities via Anchor-Dominant Convex Hulls

Nikolaos Nakis[1] · Chrysoula Kosma[2] · Panagiotis Promponas[1]
Michail Chatzianastasis[3] · Giannis Nikolentzos[4]

[1] Yale Institute for Network Science · [2] Université Paris-Saclay / Centre Borelli
[3] École Polytechnique · [4] University of Peloponnese

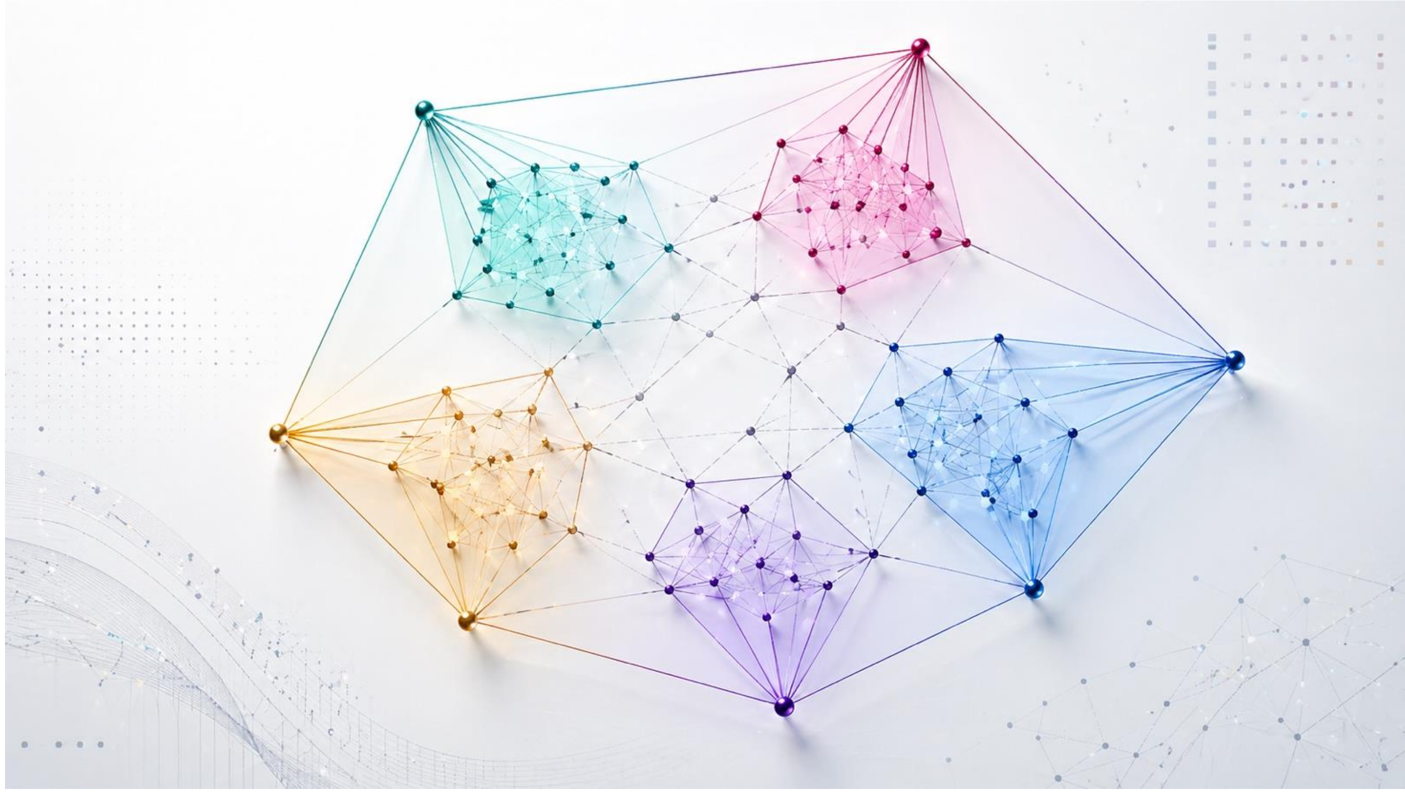

Motivation: predictions need geometric explanations

A generative framework for explainable graph representation learning via archetypes.

i

Hierarchical design

Global archetypes define pure communities; local hulls refine them with prototypes.

ii

Identifiability

Anchor dominance makes local hulls disjoint, yielding unique node reconstructions.

Truncated Dirichlet with $\varepsilon < ½$ ⇒ disjoint local hulls ⇒ unique node reconstruction.

iii

Stable & scalable

Boxed-SVD stabilizes optimization; subsampling reduces likelihood cost to linear in edges.

Boxed-SVD ⇒ Lipschitz gradient. Non-edge subsampling: $O(N^2) \Rightarrow O(|E|)$.

iv

Principled diversity

DPP priors spread archetypes and prevent degenerate hulls.

v

Empirical validation

Competitive link prediction and community detection across real networks.

Geometry is the explanation: global archetypes explain communities, **local prototypes** explain variation, **anchor dominance** gives identifiability, and **generative logits** make the geometry predictive.

Model: two levels of convex hulls

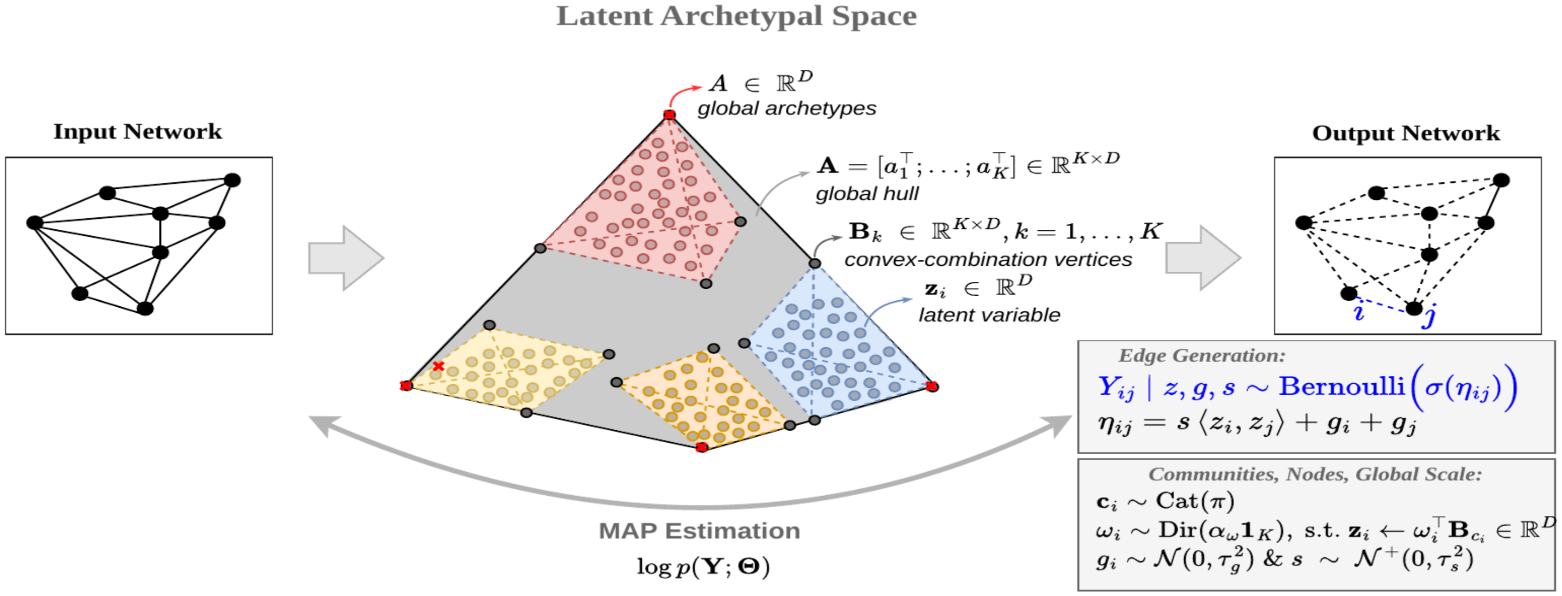

Global hull

$$\mathbf{A} = [\boldsymbol{a}_1^\top; \ldots; \boldsymbol{a}_K^\top] \in \mathbb{R}^{K \times D}$$

Stable basis

$$\mathbf{A} = \mathbf{U}\,\text{diag}(\boldsymbol{\sigma})\,\mathbf{V}^\top,$$

$$\sigma_j \in [\sigma_{\min}, \sigma_{\max}]$$

Local hull k

$$\mathbf{B}_k = \widetilde{\boldsymbol{W}}_k\,\boldsymbol{A} = \begin{bmatrix} \mathbf{W}_k \\ \boldsymbol{e}_k^\top \end{bmatrix} \mathbf{A}$$

Node i

$$c_i \sim \text{Cat}(\boldsymbol{\pi}), \quad \boldsymbol{\omega}_i \sim \text{Dir}(\alpha_\omega \mathbf{1}_K)$$

$$\boldsymbol{z}_i = \boldsymbol{\omega}_i^\top \mathbf{B}_{c_i}$$

Identifiability: anchor dominance separates the hulls

Anchor-dominant rows

$$\boldsymbol{w} \sim \mathrm{Dir}(\alpha\,\mathbf{1}_K)\,\Big|\, w_k \geq 1-\varepsilon, \quad 0<\varepsilon<\tfrac{1}{2}$$

Implementation used in the model

$$(\mathbf{W}_k)_{r,:} = (1-s_{k,r})\,\boldsymbol{e}_k^\top + s_{k,r}\,\widetilde{\boldsymbol{q}}_{k,r}, \qquad s_{k,r} = \varepsilon\, t_{k,r}$$

Non-overlap lemma

$$\varepsilon < \frac{1}{2} \quad \Longrightarrow \quad \mathrm{conv}(\mathbf{B}_k) \cap \mathrm{conv}(\mathbf{B}_\ell) = \varnothing \quad (k \neq \ell)$$

ε is the anchor mass: smaller ε → tighter, fully separated hulls.

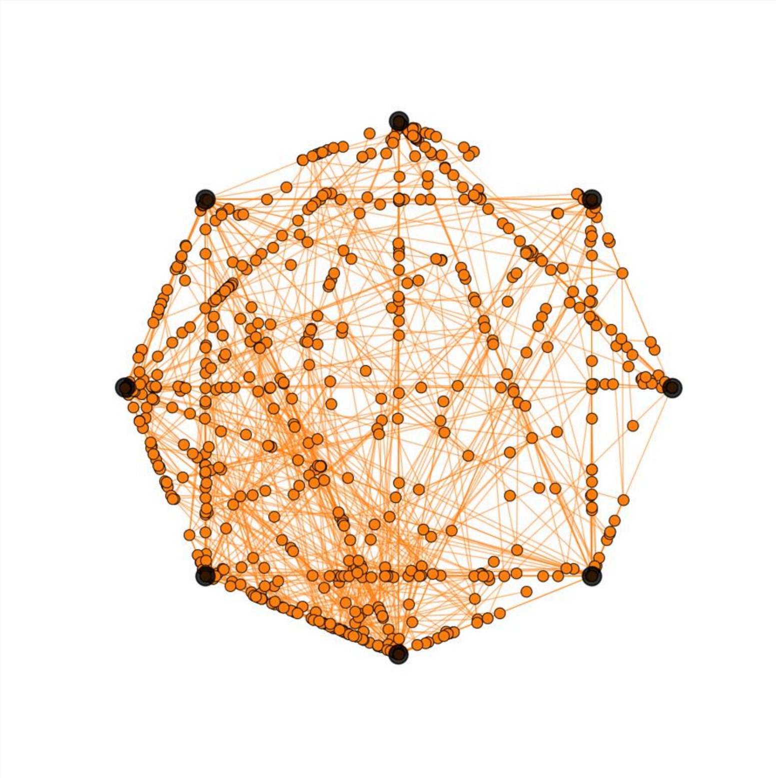

One local hull anchored at a_k: K−1 prototypes + the anchor itself.

Boxed SVD: bounded archetypes & stable optimization

Boxed-SVD parameterization

$$\mathbf{A} = \mathbf{U}\,\mathrm{diag}(\boldsymbol{\sigma})\,\mathbf{V}^\top$$

$$\sigma_j \in [\sigma_{\min}, \sigma_{\max}], \quad j = 1, \ldots, K$$

$$\|\mathbf{A}\|_2 \le \sigma_\star \implies \|\boldsymbol{z}_i\| \le \kappa_\star$$

Spectral box ⇒ embeddings stay in a bounded ball.

Lipschitz gradient (Theorem)

$$L_\omega \;\le\; \tfrac{1}{4}\Big[\, s\,\kappa_\star^2\,\deg_{\max} \;+\; \tfrac{s^2}{2}\,\kappa_\star^4\,\deg_{\max} \Big]$$

Projected gradient with $\eta \le 1/L\omega$ is provably safe — monotone descent.

Generative process

Algorithm — six steps

1. Sample $A \in \mathbb{R}^{K\times D}$ via boxed-SVD parameterization.
2. For each k: build $B_k = \tilde{W}_k A$ with anchor-dominant rows (truncated Dirichlet, $\varepsilon < ½$).
3. Place DPP priors on rows of A and each B_k to encourage spread.
4. Draw $c_i \sim \mathrm{Cat}(\pi)$, $\omega_i \sim \mathrm{Dir}(\alpha\_\omega\cdot 1)$, set $z_i = \omega_i^T B\_\{c_i\}$.
5. Draw degree bias $g_i \sim \mathcal{N}(0, \tau\_g^2)$ and global scale $s \sim \mathrm{HalfNormal}(\tau\_s)$.
6. Sample edges $Y_{ij} \sim \mathrm{Bern}(\sigma(s\langle z_i, z_j\rangle + g_i + g_j))$.

$$\varepsilon < ½ \;\Rightarrow\; \mathrm{conv}(B_k) \cap \mathrm{conv}(B_l) = \emptyset$$

Generative edge likelihood

$$\eta_{ij} = s\left\langle \boldsymbol{z}_i, \boldsymbol{z}_j \right\rangle + g_i + g_j$$

$$Y_{ij} \mid \mathbf{Z}, \boldsymbol{g}, s \;\sim\; \mathrm{Bernoulli}\big(\sigma(\eta_{ij})\big)$$

$$\log p(\mathbf{Y} \mid \mathbf{Z}, \boldsymbol{g}, s) = \textstyle\sum_{i<j}\big[\, Y_{ij}\eta_{ij} - \log(1+\exp \eta_{ij})\,\big]$$

GraphHull uses $\sigma(\eta_{ij})$ directly for link prediction: no post-hoc classifier.

MAP with geometry priors

$$\begin{aligned}\log p(\mathbf{Y}, \boldsymbol{\Theta}) = &\sum_{(i,j)\in\mathcal{D}} \left[Y_{ij}\eta_{ij} - \mathrm{softplus}(\eta_{ij})\right] \\ &+ \sum_{k=1}^{K}\left[\log\det(\mathbf{L}_k) - \log\det(\mathbf{I}+\mathbf{L}_k)\right] \\ &+ \left[\log\det(\mathbf{L}) - \log\det(\mathbf{I}+\mathbf{L})\right] \\ &- \frac{1}{2\tau_g^2}\sum_{i=1}^{N} g_i^2 - \frac{s^2}{2\tau_s^2}\end{aligned}$$

DPP priors repel vertices; non-edge subsampling reduces likelihood cost from $O(N^2)$ to $O(|E|)$.

What the model recovers: geometry and adjacency

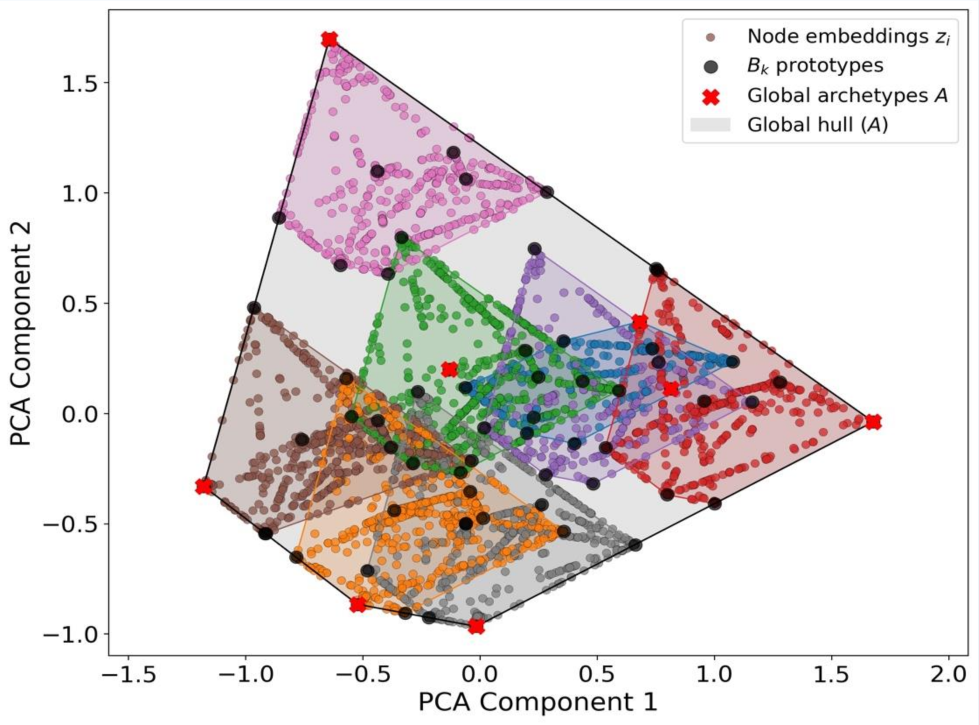

Latent hulls: global archetypes define the outer hull; each local hull captures one community.

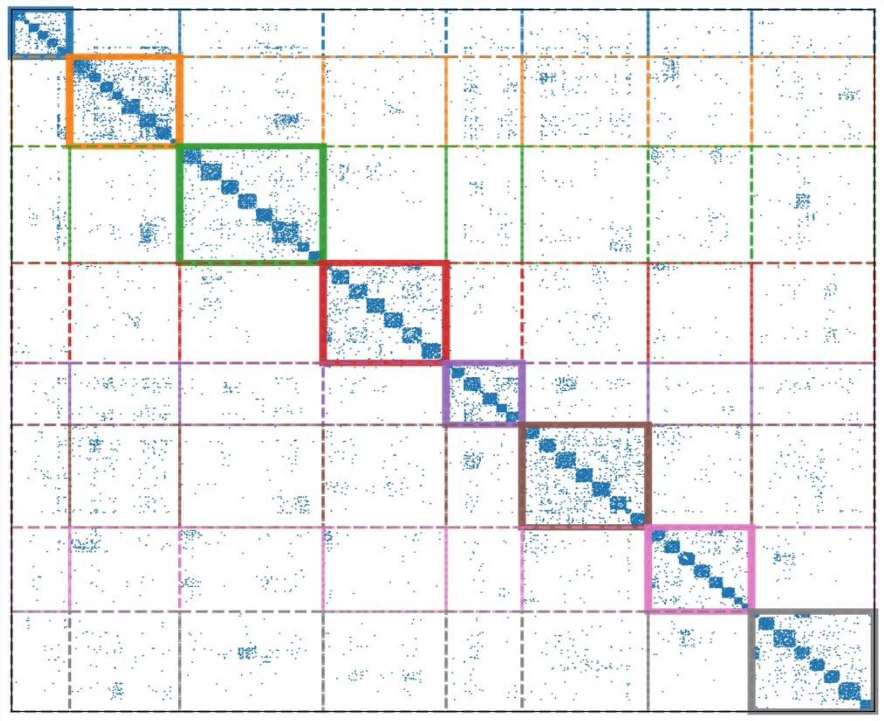

Reordered adjacency: first by global hull assignments, then within blocks by argmax ω_i.

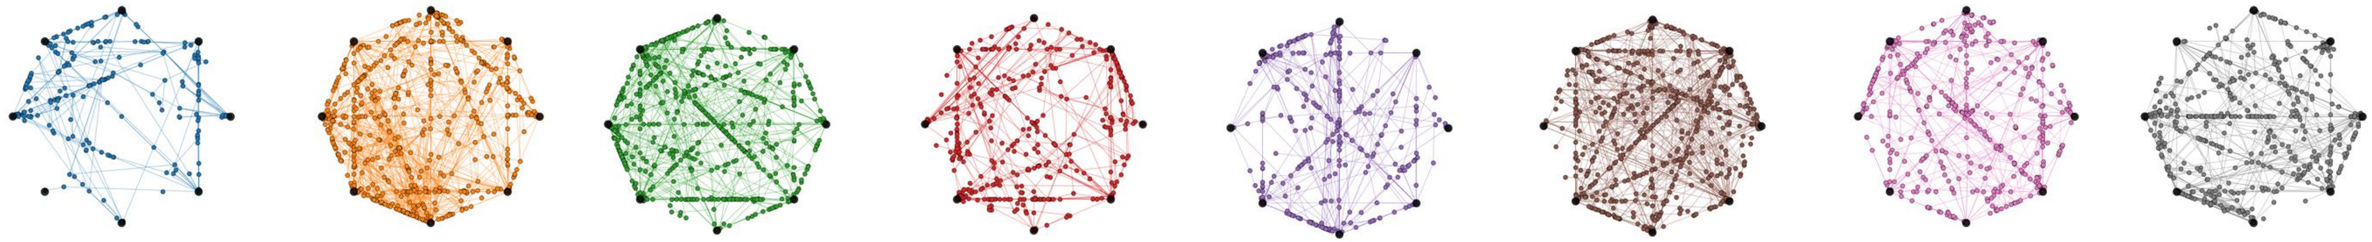

Circular plots for each community-specific polytope

Effect of ε: separation, volume, and performance

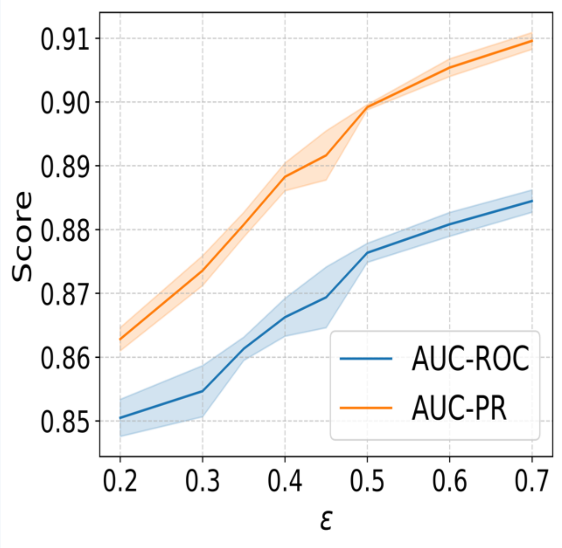

Predictive performance as ε varies.

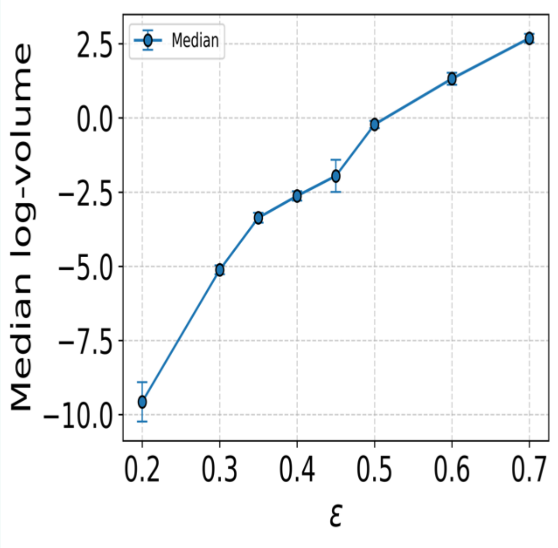

Effective log-volume increases as anchor dominance relaxes.

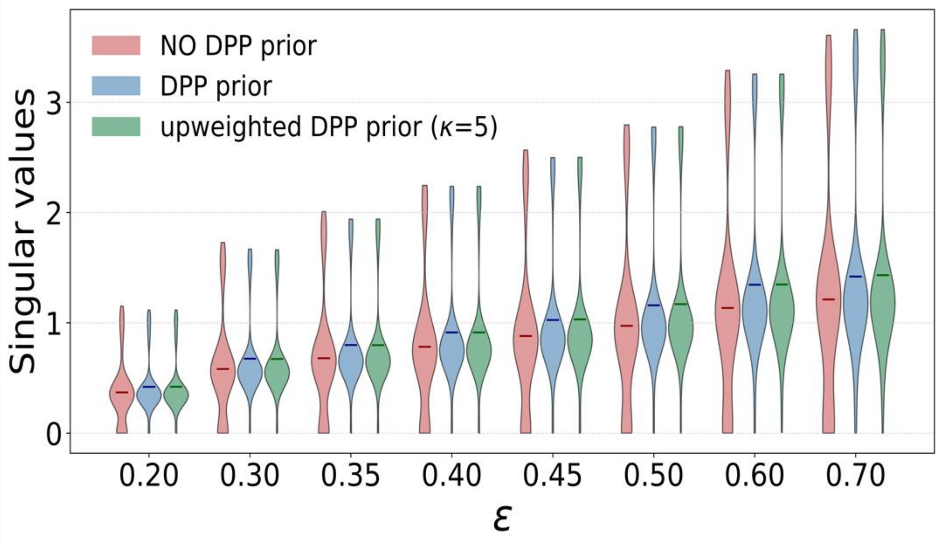

Local-hull singular values under no DPP, DPP, and upweighted DPP (κ=5).

Takeaway: ε < 1/2 keeps identifiability; increasing ε expands hulls and improves expressivity. On HepTh, high performance is already reached near the identifiable boundary.

Results: link prediction and node clustering

Link prediction: AUC-ROC at D=64

Method	AstroPh	GrQc	HepTh	Cora	DBLP
Node2Vec	.962	.936	.892	.777	.941
Role2Vec	.965	.934	.895	.752	.944
MNMF	.954	.918	.875	.680	.931
Dmon	.861	.866	.807	.788	.791
GraphHull	.965	.945	.904	.803	.948

5/5 wins or ties at D=64
with direct generative logits

Node clustering: NMI / ARI

Dataset	Cora	Citeseer	LastFM	Pol
Best NMI	.486	.320	.605	.727
GraphHull NMI	.422	.347	.616	.757
Best ARI	.417	.320	.522	.791
GraphHull ARI	.333	.372	.514	.820

5/8 best metrics
strongest on Citeseer, LastFM, Pol

Accuracy is paired with explanations: c_i gives the community, ω_i gives prototype membership, and η_{ij} explains each predicted edge.

Takeaway

GraphHull = generative embeddings with explanations built in.

Global archetypes define pure community profiles.

Local prototypes capture within-community variation.

Anchor dominance yields disjoint hulls and identifiable assignments.

Generative edge logits connect the geometry directly to link prediction.

Identifiable. Interpretable. Competitive.

GitHub

https://github.com/Nicknakis/GraphHull

arXiv

https://arxiv.org/abs/2602.21342

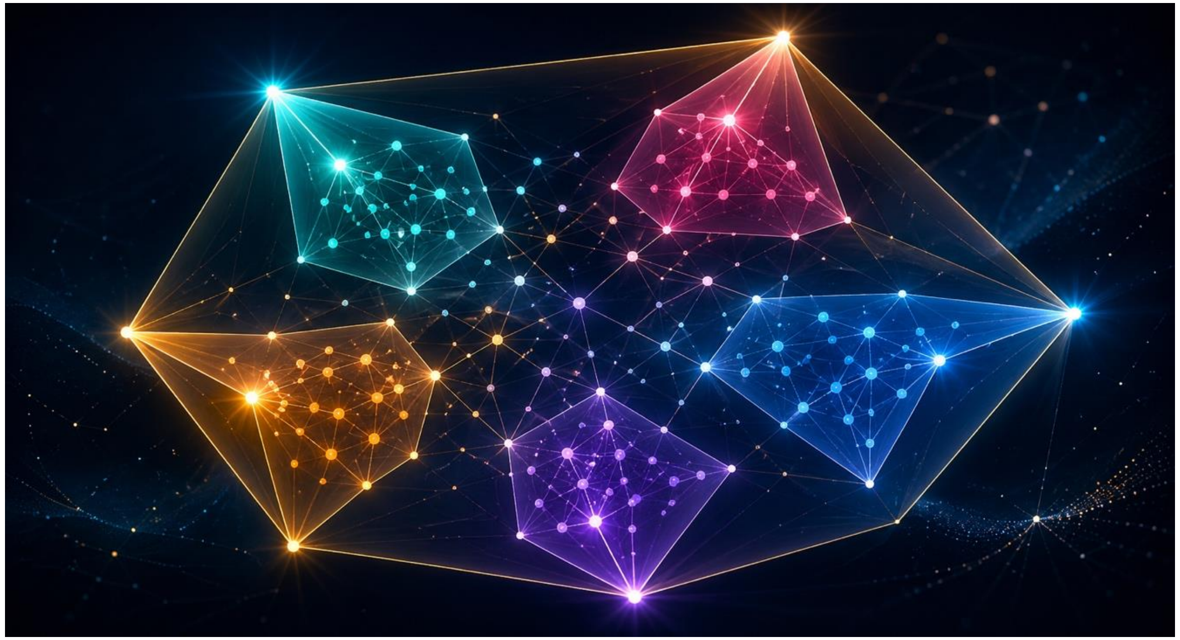

Thank you!

See you at the poster!