Imperial College London

Task-Agnostic Graph Neural Network Evaluation via Adversarial Collaboration

Xiangyu Zhao, Hannes Stärk, Dominique Beaini, Yiren Zhao, Pietro Liò

We propose Graph Adversarial Collaboration (GraphAC) – a conceptually novel, principled, task-agnostic, and stable framework for evaluating GNNs through contrastive self-supervision, without using any augmentations.

Overview

- It has been increasingly demanding to develop reliable methods to evaluate the progress of GNN research for molecular representation learning;
- Current GNN benchmarking approaches (evaluating GNNs on some classification/regression tasks) can be limited:
  1) classification/regression tasks are not challenging enough;
  2) GNNs are particularly vulnerable to noisy labels;
- Most of the existing graph SSL methods either:
  1) rely on applying handcrafted augmentations to the graphs, which has several severe difficulties, or
  2) exploit task-specific properties of the graphs, which cannot be transferred to other domains;
- GraphAC addresses both of the above issues, by having GNNs directly compete against each other on the same unlabelled graphs, in a contrastive self-supervised manner.
  It then ensures that more expressive GNNs can always win by producing more complex and informative graph embeddings.

Method

1. Intuition

If a GNN can predict another GNN's graph embeddings from its own graph embeddings better than the other way round, then its embeddings can be deemed more complex and informative than the other GNN's embeddings, and that GNN is more expressive.

The two GNNs collaborate by predicting each other's output graph embeddings, and compete adversarially to prevent the other GNNs from predicting their own graph embeddings.

2. Competitive Barlow Twins

Replace the off-diagonal sum in Barlow Twins (Zbontar et al., 2021) with the difference between the upper-triangle and the lower-triangle of the cross-correlation matrix:
(see bottom-left figure)

- Sum of each triangle measures how much a model's output embeddings correlate to the other model's output embeddings
- Smaller-indexed features of the models' output embeddings also become the more important features ⇒ further regularize training by ordering the features by importance

3. Overall Framework

Combine the Competitive Barlow Twins losses with the VICReg (Bardes et al., 2022) covariance regularization loss:

$$\mathcal{L}_{\mathrm{GNN}_A} = \alpha \mathcal{L}_{\mathrm{CBT}_A} + \beta \mathcal{L}_{\mathrm{Cov}}$$
$$\mathcal{L}_{\mathrm{GNN}_B} = \alpha \mathcal{L}_{\mathrm{CBT}_B} + \beta \mathcal{L}_{\mathrm{Cov}}$$

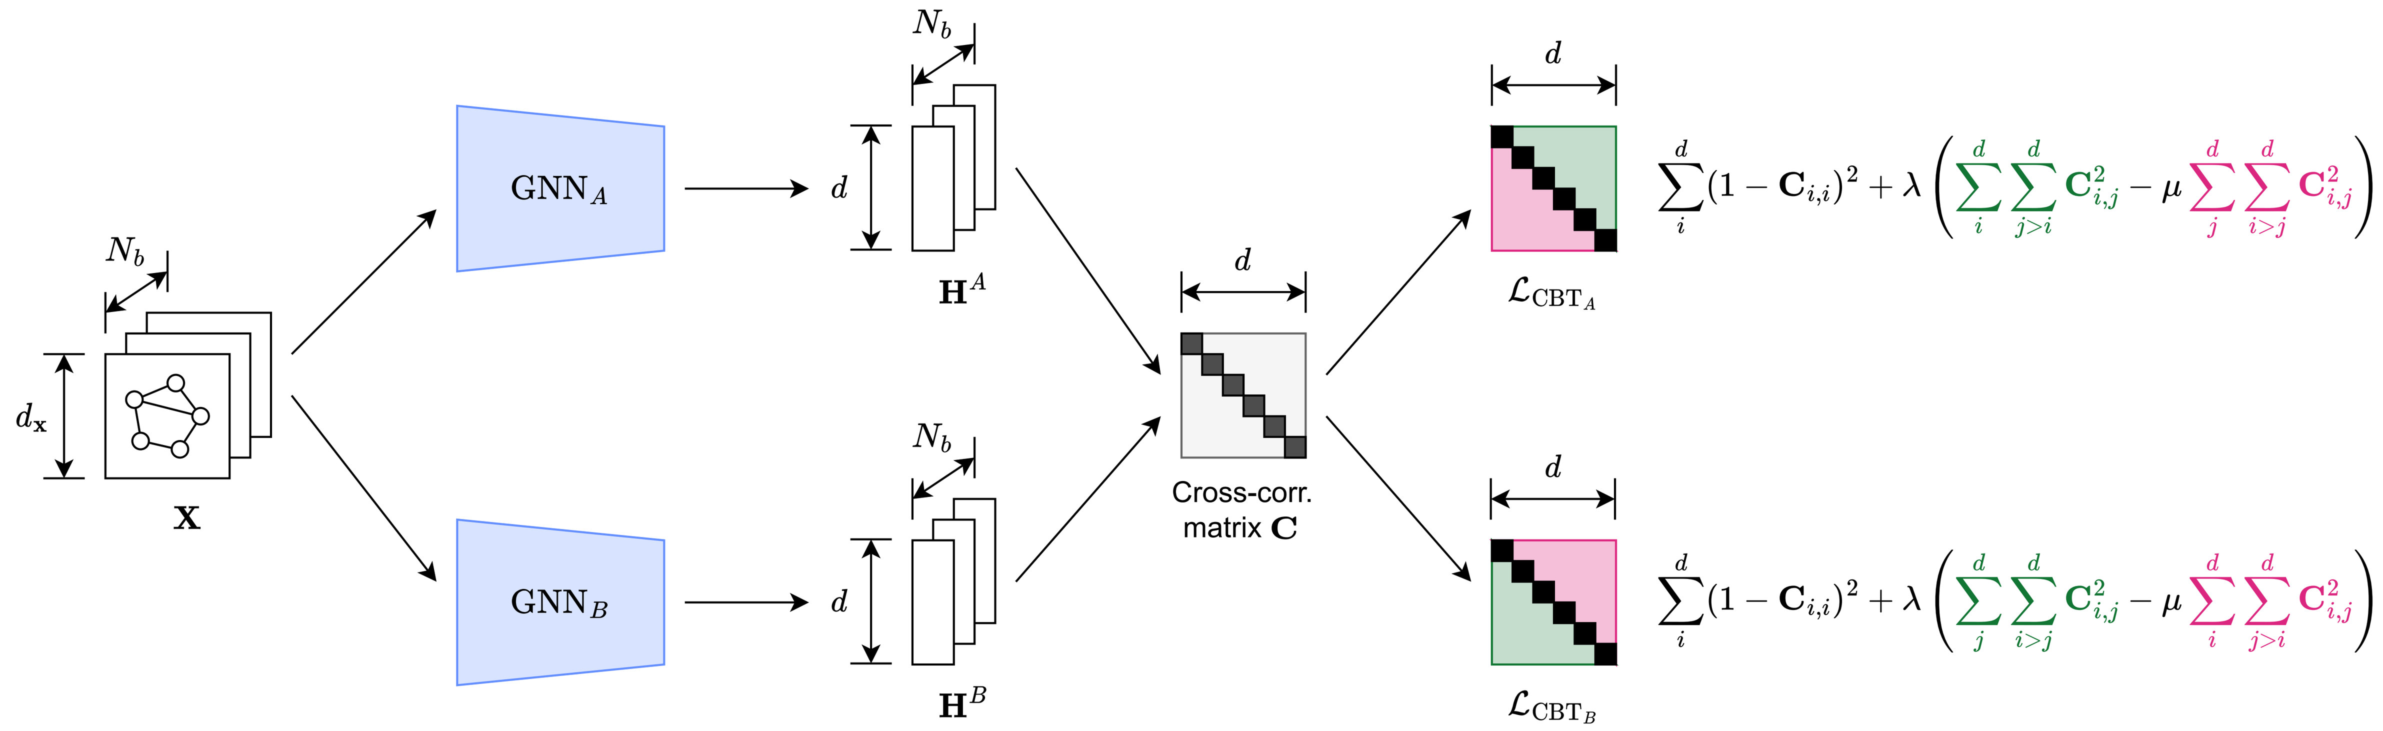

		#Layers in GNN_B				
		2	4	6	8	10
#Layers in GNN_A	2	-0.042	1.290	1.458	1.882	1.966
	4	-1.315	0.027	0.799	1.399	1.748
	6	-1.478	-0.845	-0.014	0.674	0.808
	8	-1.687	-1.276	-0.406	0.030	0.555
	10	-2.008	-1.751	-1.076	-0.668	-0.004

		Hidden dims in GNN_B				
		16	32	64	128	256
Hidden dims in GNN_A	16	0.022	1.390	1.883	2.356	2.554
	32	-1.240	0.001	0.932	1.614	2.093
	64	-2.290	-0.895	0.018	1.206	1.813
	128	-2.493	-1.609	-0.997	-0.010	1.490
	256	-2.541	-2.041	-1.704	-1.330	-0.006

		Aggregators in GNN_B			
		[max]	[mean]	[sum]	Combined
Aggregators in GNN_A	[max]	-0.026	0.182	0.305	0.342
	[mean]	-0.174	-0.007	0.241	0.328
	[sum]	-0.304	-0.248	0.047	0.228
	Combined	-0.329	-0.295	-0.239	0.018

		GNN_B architecture		
		GCN	GIN	PNA
GNN_A architecture	GCN	-0.069	0.338	0.467
	GIN	-0.337	-0.026	0.441
	PNA	-0.594	-0.419	0.018

Loss differences of the conducted experiments on the ogbg-molpcba dataset. Negative value means GNN_A wins the game, and positive value means GNN_B wins the game. Greater absolute value indicates larger gap in expressiveness determined by GraphAC. The consistent gradient from bottom left to upper right clearly indicates that GraphAC genuinely favors more expressive GNNs.

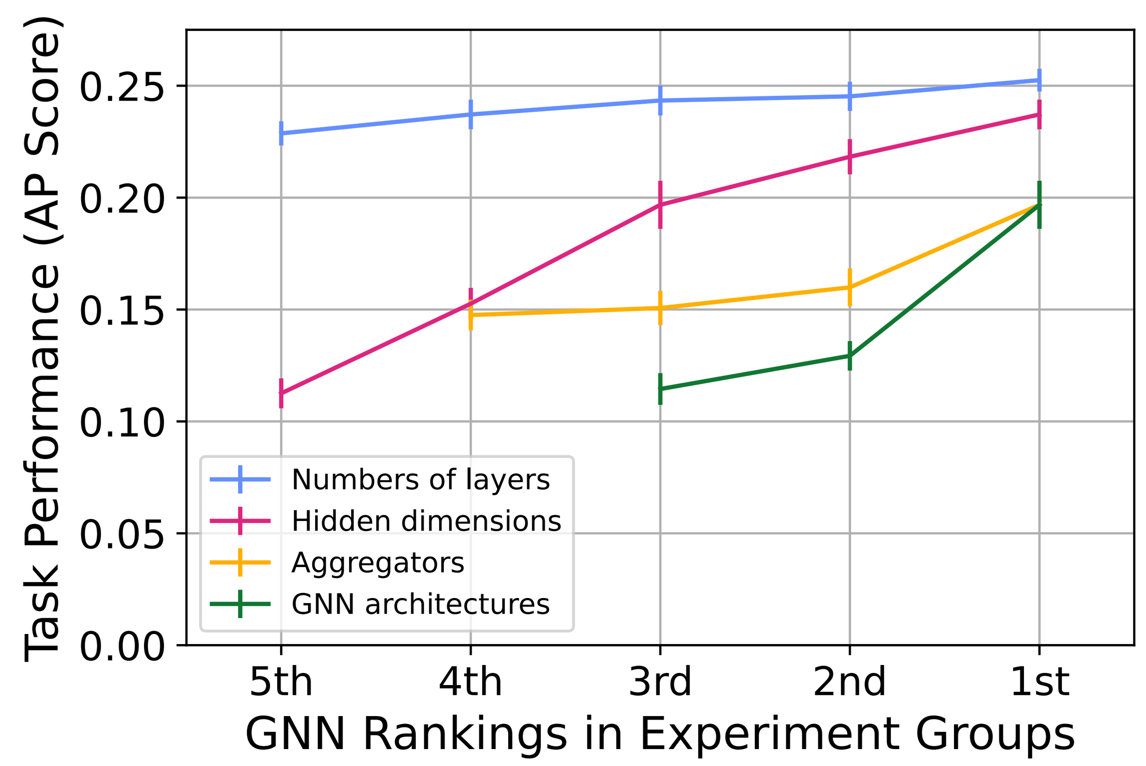

Results

GraphAC is evaluated on the OGB `molpcba` and `code2` datasets:

Expressiveness Rankings

GraphAC can successfully distinguish GNNs of different expressiveness across various aspects, and consistently favors more expressive GNNs

Total Ordering

If $\mathrm{GNN}_A > \mathrm{GNN}_B > \mathrm{GNN}_C$ in expressiveness, then GraphAC yields greater loss difference between (GNN_A, GNN_C) than (GNN_A, GNN_B), thereby producing a total ordering of all GNNs

No False Winners

GraphAC produces loss differences of GNNs with the same expressiveness close to zero, allowing them to tie

Order-Invariance

GraphAC can genuinely distinguish different GNNs regardless of their ordering in the framework

Task-Agnosticism

GraphAC can be successfully generalised to various distinct domains, without using labels or augmentations

Link with GNN Performance

GraphAC's expressiveness rankings on GNNs also correlate strongly with GNNs' actual task performances (see plots in the left figure)

References

Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised learning. In *10th International Conference on Learning Representations (ICLR 2022)*. OpenReview.net, 2022.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stephane Deny. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, volume 139, pp. 12310–12320. PMLR, 2021.

UNIVERSITY OF CAMBRIDGE

Valence