

GNNs don't need Backprop

Benoit Gouil, Fabian Jögl, Pascal Welke

Ecole Polytechnique

TU Wien



Lancaster University Leipzig

What is the goal
of GNN training?


What happens
during GNN training?

Does GNN
work how
we think
it does?

Supervised Learning on Graphs



Neural Networks on Graphs



Message Passing Graph Neural Network Layer 1

Message Passing Graph Neural Network Layer k

Graph level pooling

Feed forward neural network



To train, we usually use

♥
♥
Stochastic Gradient Descent
♥
♥
♥
♥
♥

To train, we usually use

♥
♥
Stochastic Gradient Descent
♥
♥
♥
♥
♥


But do we need to?

Usually, we use a loss function that is some sort of distance (well, more or less) to a target value.

$$\text{Loss}(t) = |f_{\theta}(t) - y(t)|^2$$

$$y(\text{triangle}) = 1$$

$$y(\text{circle}) = 0$$




Usually, we use a loss function that is some sort of distance (well, more or less) to a target value.

$$\text{Loss}(t) = |f_{\theta}(t) - y(t)|^2$$

$$y(\triangle) = 1$$

$$y(\circ) = 0$$




Usually, we use a loss function that is some sort of distance (well, more or less) to a target value.

$$\text{Loss}(t) = |f_{\theta}(t) - y(t)|^2$$

$$y(\Delta) = 1$$

$$y(\circ) = 0$$




Usually, we use a loss function that is some sort of distance (well, more or less) to a target value.

$$\text{Loss}(t) = |f_{\theta}(t) - y(t)|^2_2$$

$$y(\Delta) = 1$$

$$y(\circ) = 0$$




Usually, we use a loss function that is some sort of distance (well, more or less) to a target value.

$$\text{Loss}(t) = |f_{\theta}(t) - y(t)|^2_2$$

$$y(\Delta) = 1$$

$$y(\circ) = 0$$





Often repeated assumption in learning -

"Similar graphs behave similarly"



— Often repeated assumption in learning —

“Similar graphs behave similarly”

What about:

“Two graphs are similar
if they behave similarly
under γ ”













"Two graphs are similar if they behave similarly under y "


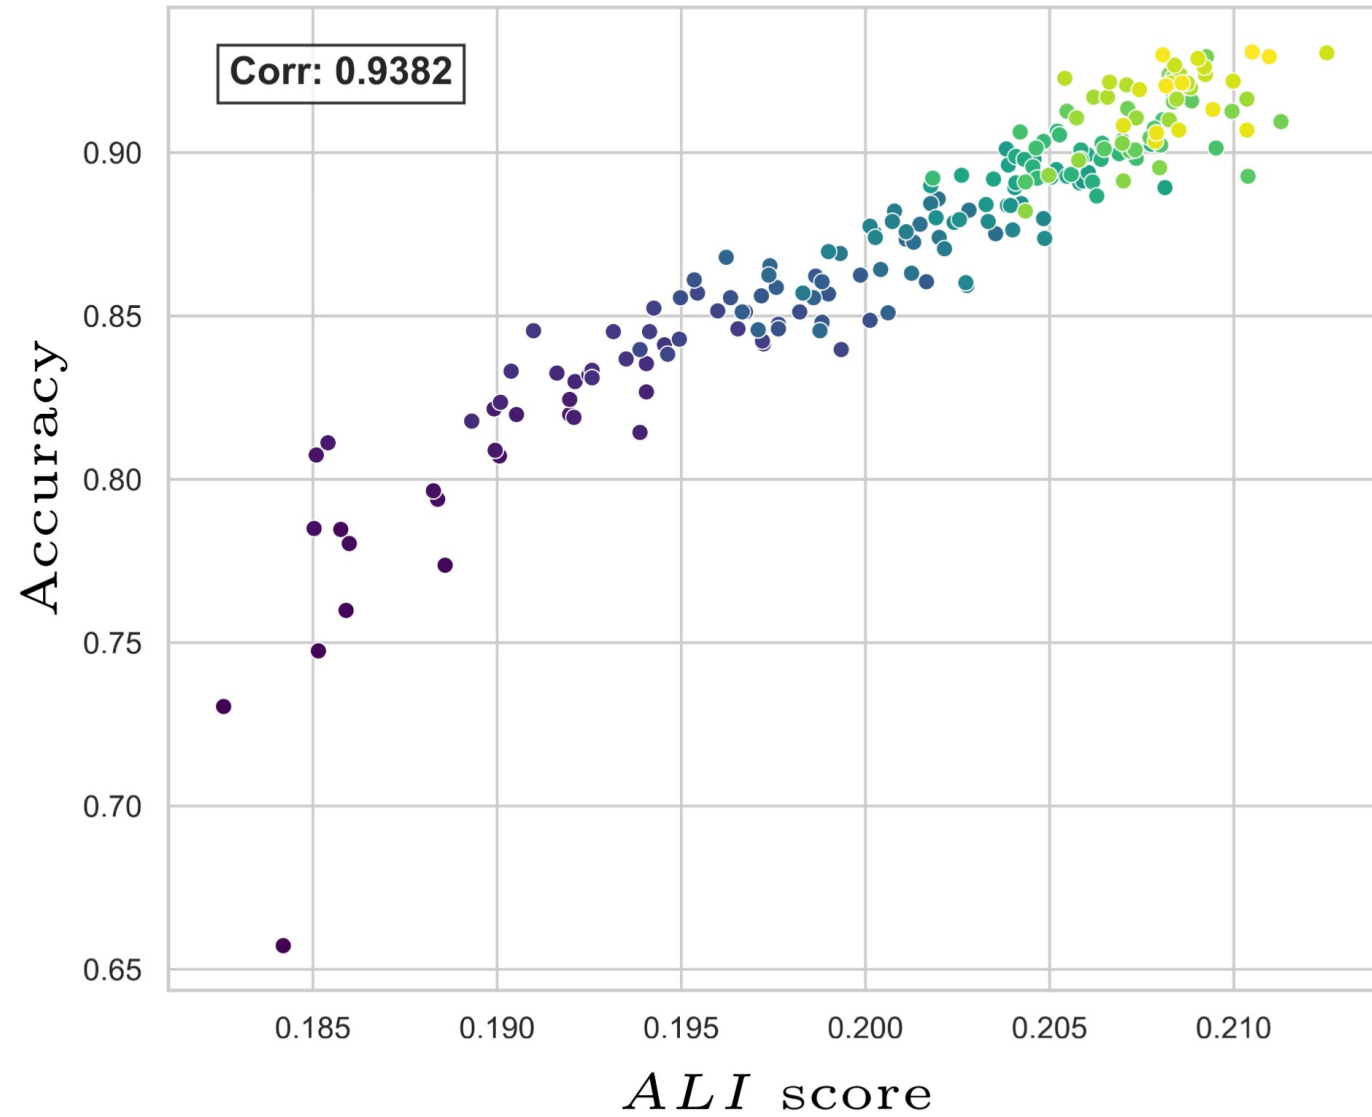
The best distance function we can define for a given learning problem is

$$d_y(G, H) = d'(y(G), y(H))$$

where d' is some sensible distance function on the target labels







Now to the technical part

If we want a distance function

$$d(G, H) = d'(y(G), y(H))$$

Now to the technical part

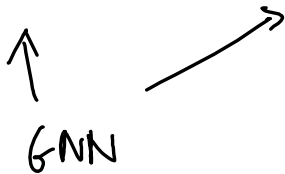
If we want a distance function

$$d(G, H) \approx d'(y(G), y(H))$$

ss

$$|f_{\theta}(G) - f_{\theta}(H)|_2$$

↑
GNN



Now to the technical part

If we want a distance function

$$d(G, H) \approx d'(y(G), y(H))$$

ss

$$|f_{\theta}(G) - f_{\theta}(H)|_2$$

↑
GNN

Do we really need to
use SGD?

Navarin et al. An Untrained Graph Neural Network for fast and accurate Graph Classification ICANN '23

"Take a randomly initialized GNN and just train a linear model on top. Done."

Böker et al. Fine-grained Expressivity of GNNs. Neurips '23

"Untrained GNNs behave as well as trained ones"

Folklore:

Randomly initialized GNNs are basically Weisfeiler Leman

(and that works well for learning)

Bui et al. Random Propagations in GNNs. Neurips '24

Huong et al:

You can have better GNNs by not training at all LOG '22

Zambon et al.

Graph Random Neural Features for Distance Preserving Graph Representations ICLR '20

Outside of GNNs: (a few examples)

Rahimi & Recht:

Weighted Sums of Random Kitchen Sinks

NeurIPS '08

Jaeger & Haas:

Harvesting Nonlinearity: Predicting Chaotic Systems and saving
Energy in Wireless Communication

Science '04

My Data Mining work:

Sampling a few graphs (trees, low treewidth graphs, ...)
and using these as features in learning works well

Wahl et al '14-'23

Can we avoid SGD?

... using pretrained
GNNs?

While exploiting
Graph Similarity
under the target
function?

Can alignment between

the distance of representations

①

obtained from randomly sampled GNNs

Can alignment between

the distance of representations

①

obtained from randomly sampled GNNs

and

distance between target values y

②

be used for learning?

The Algorithm TM

① Sample a gazillion randomly initialized GNNs

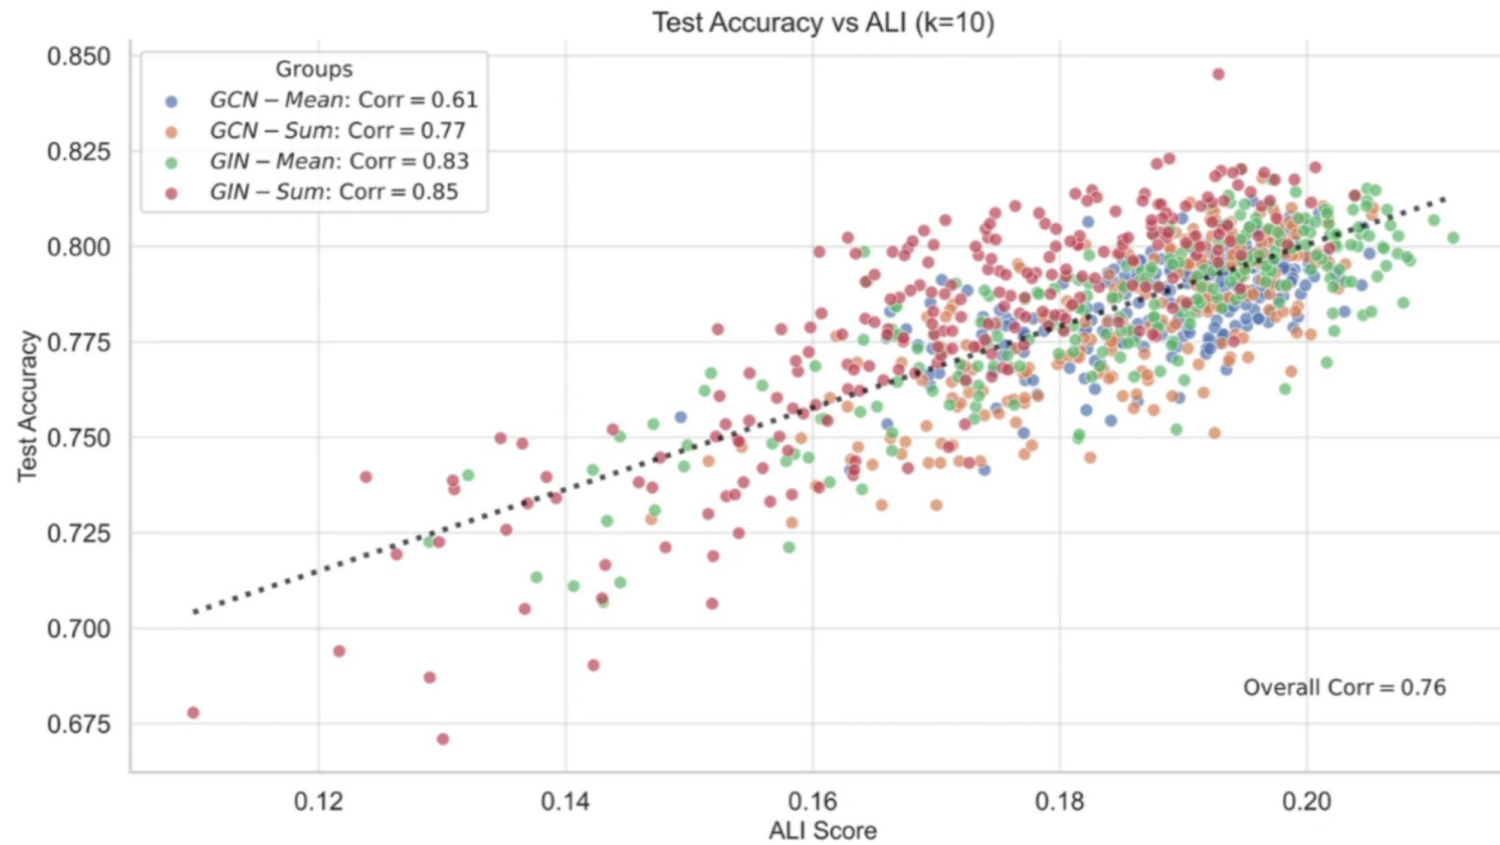
The Algorithm ^(Tm)

- ① Sample a gazillion randomly initialized GNNs
- ② Return the GNN f_{θ} that maximizes the alignment $ALI_k(f_{\theta}, d')$ on the training dataset \mathcal{G}

The Algorithm ^(Tm)

- ① Sample a gazillion randomly initialized GNNs
- ② Return the GNN f_θ that maximizes the alignment $ALI_k(f_\theta, d')$ on the training dataset \mathcal{G}
- ③ Train a classifier on top of the representations given by f_θ

ALI-Accuracy Correlation



$$A(G, f_{\theta}, y) = \frac{1}{k} \sum d^1(y(G), y(H))$$

$H \in S_k(f_{\theta}, G)$

\uparrow
k Nearest Neighbors under f_{θ}

distance between targets
 \downarrow

$$A(G, f_{\theta}, y) = \frac{1}{k} \sum_{H \in S_k(f_{\theta}, G)} d'(y^{(G)}, y^{(H)})$$

distance between targets

$H \in S_k(f_{\theta}, G)$

\uparrow
 k Nearest Neighbors under f_{θ}

$$B(G, f_{\theta}, y) = \frac{1}{|y| - k - 1} \sum_{H \in T S_k(f_{\theta}, G)} d'(y^{(G)}, y^{(H)})$$

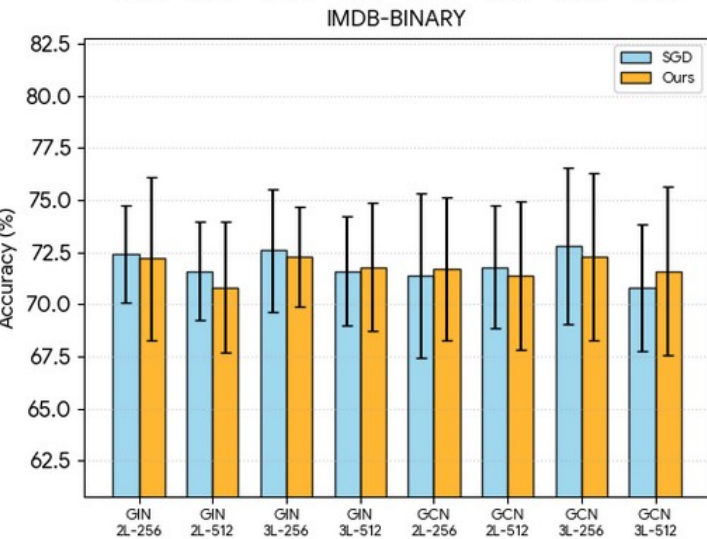

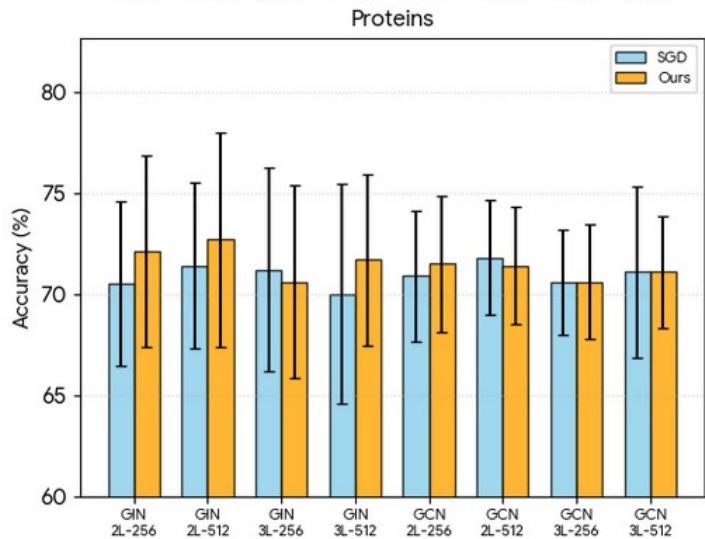
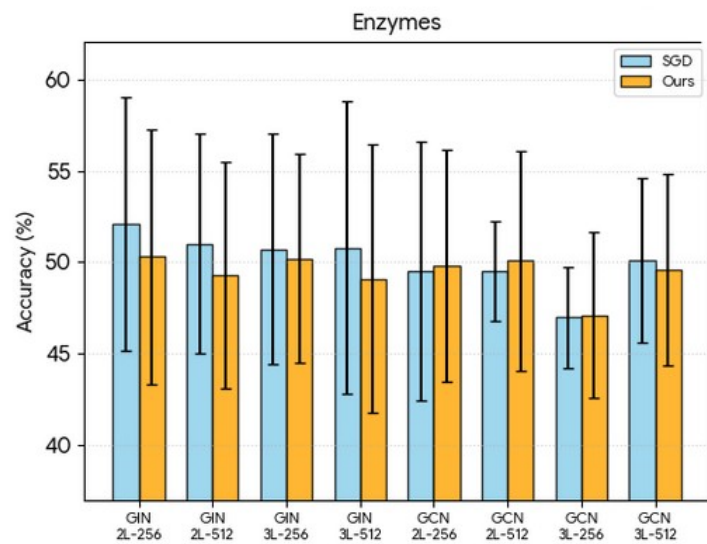

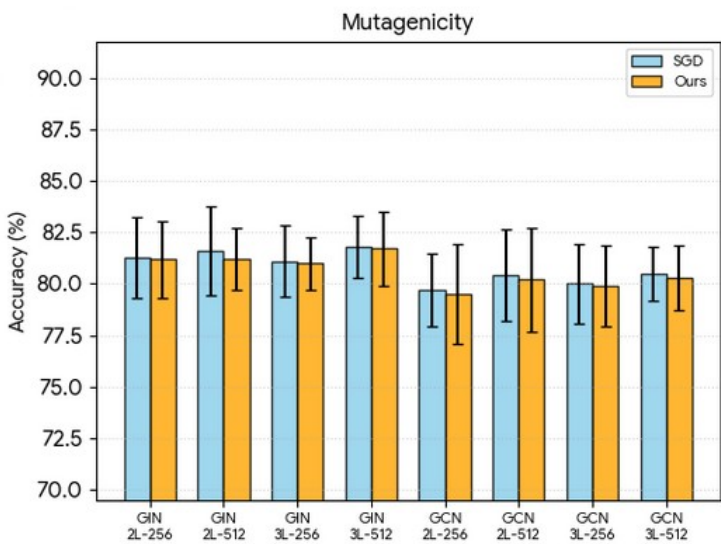
$$A(G, f_{\theta}, y) = \frac{1}{k} \sum_{H \in S_k(f_{\theta}, G)} d'(y^{(G)}, y^{(H)})$$

distance between targets

\uparrow
 k Nearest Neighbors under f_{θ}

$$B(G, f_{\theta}, y) = \frac{1}{|y| - k - 1} \sum_{H \in T S_k(f_{\theta}, G)} d'(y^{(G)}, y^{(H)})$$

$$ALI_k(f_{\theta}, d') = \sum_{G \in \mathcal{G}} -A(G, f_{\theta}, y) + B(G, f_{\theta}, y)$$



Conclusion

- Alignment is a good predictor of GNN performance
- Our sampling approach can achieve similar accuracy to SGD

Conclusion

- Alignment is a good predictor of GNN performance
- Our sampling approach can achieve similar accuracy to SGD

Limitations

- Struggles with deeper GNNs
- Runtime of alignment computation is quadratic in training set size

Conclusion

- Alignment is a good predictor of GNN performance
- Our sampling approach can achieve similar accuracy to SGD

Limitations

- Struggles with deeper GNNs
- Runtime of alignment computation is quadratic in training set size

Future work

- Ensembles of sampled GNNs
- Evaluate on regression tasks + node level tasks
- Evaluate on larger datasets
- Robustness to noisy data?

Last page as a bonus

in case

it vanishes