
LARGE-SCALE NODE CLASSIFICATION WITH BOOTSTRAPPING

OPEN GRAPH BENCHMARK LARGE-SCALE CHALLENGE ENTRY (MAG240M)

Ravichandra Addanki*
DeepMind
ravichandraa@deepmind.com

Peter W. Battaglia*
DeepMind
peterbattaglia@deepmind.com

David Budden*
DeepMind
budden@deepmind.com

Andrea Deac*
DeepMind, Mila, Université de Montréal
adeac@deepmind.com

Jonathan Godwin*
DeepMind
jonathangodwin@deepmind.com

Thomas Keck*
DeepMind
thomaskeck@deepmind.com

Alvaro Sanchez-Gonzalez*
DeepMind
alvarosg@deepmind.com

Jacklynn Stott*
DeepMind
jacklynnstott@deepmind.com

Shantanu Thakoor*
DeepMind
thakoor@deepmind.com

Petar Veličković*
DeepMind
petarv@deepmind.com

ABSTRACT

Effectively and efficiently deploying graph neural networks (GNNs) at scale remains one of the most challenging aspects of graph representation learning. Many powerful solutions have only ever been validated on comparatively small datasets, often with counter-intuitive outcomes—a barrier which has recently been broken by the Open Graph Benchmark Large-Scale Challenge (OGB-LSC). Here we describe our OGB-LSC entry for the MAG240M-LSC benchmark: a deep transductive node classifier powered by bootstrapping. Our model achieved an award-level (top-3) performance. In doing so, we demonstrate strong evidence of scalable self-supervised graph representation learning, and utility of deep GNNs—both very important open issues. Our code is publicly available at: https://github.com/deepmind/deepmind-research/tree/master/ogb_lsc/mag.

Keywords OGB-LSC · MPNNs · BGRL

1 Introduction

Effective high-dimensional representation learning necessitates properly exploiting the geometry of data [Bronstein et al., 2021]—otherwise, it is a cursed estimation problem. Indeed, early success stories of deep learning relied on imposing strong geometric assumptions, primarily that the data lives on a grid domain; either spatial or temporal. In these two respective settings, convolutional neural networks (CNNs) [LeCun et al., 1998] and recurrent neural networks (RNNs) [Hochreiter and Schmidhuber, 1997] have traditionally dominated.

While both CNNs and RNNs are demonstrably powerful models, with many applications of high interest, it can be recognised that most data coming from nature cannot be natively represented on a grid. Recent years are marked with a gradual shift of attention towards models that admit a more generic class of geometric structures [Masci et al., 2015, Veličković et al., 2017, Cohen et al., 2018, Battaglia et al., 2018, de Haan et al., 2020, Satorras et al., 2021].

In many ways, the most generic and versatile of these models are graph neural networks (GNNs). This is due to the fact that most discrete-domain inputs can be observed as instances of a graph structure. The corresponding area of graph

*All authors contributed equally.

representaton learning [Hamilton, 2020] has already seen immense success across industrial and scientific disciplines. GNNs have successfully been applied for drug screening [Stokes et al., 2020], modelling the dynamics of glass [Bapst et al., 2020], web-scale social network recommendations [Ying et al., 2018] and chip design [Mirhoseini et al., 2020].

While the above results are certainly impressive, they likely only scratch the surface of what is possible with well-tuned GNN models. Many problems of real-world interest require graph representation learning *at scale*: either in terms of the amount of graphs to process, or their sizes (in terms of numbers of nodes and edges). Perhaps the clearest motivation for this comes from the Transformer family of models [Vaswani et al., 2017]. Transformers operate a self-attention mechanism over a complete graph, and can hence be observed as a specific instance of GNNs [Joshi, 2020]. At very large scales of natural language data, Transformers have demonstrated significant returns with the increase in capacity, as exemplified by models such as GPT-3 [Brown et al., 2020]. Transformers enjoy favourable scalability properties at the expense of their functional complexity: each node’s features are updated with *weighted sums* of neighbouring node features. In contrast, GNNs that rely on message passing [Gilmer et al., 2017]—passing vector signals across edges that are conditioned on *both* the sender and receiver nodes—are an empirically stronger class of models, especially on tasks requiring complex reasoning [Veličković et al., 2019] or simulations [Sanchez-Gonzalez et al., 2020, Pfaff et al., 2020].

One reason why generic message-passing GNNs have not been scaled up as widely as Transformers is the lack of appropriate datasets. Only recently has the field advanced from simple transductive benchmarks of only few thousands of nodes [Sen et al., 2008, Shchur et al., 2018, Morris et al., 2020] towards larger-scale real-world and synthetic benchmarks [Dwivedi et al., 2020, Hu et al., 2020], but important issues still remain. For example, on many of these tasks, randomly-initialised GNNs [Veličković et al., 2018], shallow GNNs [Wu et al., 2019] or simple label propagation-inspired GNNs [Huang et al., 2020] can perform near the state-of-the-art level at only a fraction of the parameters. When most bleeding-edge expressive methods are unable to improve on the above, this can often lead to controversial discussion in the community. One common example is: *do we even need deep, expressive GNNs?*

Breakthroughs in deep learning research have typically been spearheaded by impactful large-scale competitions. For image recognition, the most famous example is the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [Russakovsky et al., 2015]. In fact, the very “deep learning revolution” has partly been kickstarted by the success of the AlexNet CNN model of Krizhevsky et al. [2012] at the ILSVRC 2012, firmly establishing deep CNNs as the workhorse of image recognition for the forthcoming decade.

Accordingly, we have entered the recently proposed Open Graph Benchmark Large-Scale Challenge (OGB-LSC) [Hu et al., 2021]. OGB-LSC provides graph representation learning tasks at a previously unprecedented scale—millions of nodes, billions of edges, and/or millions of graphs. Further, the tasks have been designed with immediate practical relevance in mind, and it has been verified that expressive GNNs are likely to be necessary for strong performance. Here we detail our submitted model for the MAG240M task, and our empirical observations while developing it. Namely, we find that the dataset’s immense scale provides a great platform for demonstrating clear outperformance of self-supervised GNN setups such as bootstrapping [Thakoor et al., 2021]. In doing so, we have provided meaningful evidence towards a positive resolution to the above discussion: deep and expressive GNNs are, indeed, necessary at the right level of task scale and/or complexity. Our final model has achieved award-level (top-3) ranking on MAG240M.

2 Dataset description

The **MAG240M-LSC** dataset is a transductive node classification dataset, based on the Microsoft Academic Graph (MAG) [Wang et al., 2020]. It is a heterogeneous graph containing paper, author and institution nodes, with edges representing the relations between them: paper-cites-paper, author-writes-paper, author-affiliated-with-institution. All paper nodes are endowed with 768-dimensional input features, corresponding to the RoBERTa sentence embedding [Liu et al., 2019, Reimers and Gurevych, 2019] of their title and abstract. MAG240M is currently the largest-scale publicly available node classification dataset by a wide margin, at ~ 240 million nodes and ~ 1.8 billion edges. The aim is to classify the ~ 1.4 million arXiv papers into their corresponding topics, according to a temporal split: papers published up to 2018 are used for training, with validation and test sets including papers from 2019 and 2020, respectively.

3 GNN Architectures

We rely on a common *encode-process-decode* blueprint [Hamrick et al., 2018]. This implies that our input features are encoded into a latent space using node-, edge- and graph-wise *encoder* functions, and latent features are decoded to node-, edge- and graph- level predictions using appropriate *decoder* functions. The bulk of the computational processing is powered by a *processor* network, which performs multiple graph neural network layers over the encoded latents.

To formalise this, assume that our input graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, has node features $\mathbf{x}_u \in \mathbb{R}^n$, edge features $\mathbf{x}_{uv} \in \mathbb{R}^m$ and graph-level features $\mathbf{x}_{\mathcal{G}} \in \mathbb{R}^l$, for nodes $u, v \in \mathcal{V}$ and edges $(u, v) \in \mathcal{E}$. Our encoder functions $f_n : \mathbb{R}^n \rightarrow \mathbb{R}^k$, $f_e : \mathbb{R}^m \rightarrow \mathbb{R}^k$ and $f_g : \mathbb{R}^l \rightarrow \mathbb{R}^k$ then transform these inputs into the latent space:

$$\mathbf{h}_u^{(0)} = f_n(\mathbf{x}_u) \quad \mathbf{h}_{uv}^{(0)} = f_e(\mathbf{x}_{uv}) \quad \mathbf{h}_{\mathcal{G}}^{(0)} = f_g(\mathbf{x}_{\mathcal{G}}) \quad (1)$$

Our processor network then transforms these latents over several rounds of message passing:

$$\mathbf{H}^{(t+1)} = P_{t+1}(\mathbf{H}^{(t)}) \quad (2)$$

where $\mathbf{H}^{(t)} = \left(\left\{ \mathbf{h}_u^{(t)} \right\}_{u \in \mathcal{V}}, \left\{ \mathbf{h}_{uv}^{(t)} \right\}_{(u,v) \in \mathcal{E}}, \mathbf{h}_{\mathcal{G}}^{(t)} \right)$ contains all of the latents at a particular processing step $t \geq 0$.

The processor network P is iterated for T steps, recovering final latents $\mathbf{H}^{(T)}$. These can then be decoded into node-, edge-, and graph-level predictions (as required), using analogous decoder functions g_n, g_e and g_g :

$$\mathbf{y}_u = g_n(\mathbf{h}_u^{(T)}) \quad \mathbf{y}_{uv} = g_e(\mathbf{h}_{uv}^{(T)}) \quad \mathbf{y}_{\mathcal{G}} = g_g(\mathbf{h}_{\mathcal{G}}^{(T)}) \quad (3)$$

We will detail the specific design of f, P and g in the following sections. Generally, f and g are simple MLPs, whereas we use highly expressive GNNs for P in order to maximise the advantage of the large-scale datasets. Specifically, we use message passing neural networks (MPNNs) [Gilmer et al., 2017]. All of our models have been implemented using the jraph library [Godwin et al., 2020].

4 MAG240M-LSC setup

Subsampling Running graph neural networks over datasets that are even a fraction of MAG240M’s scale is already prone to multiple scalability issues, which necessitated either aggressive subsampling [Hamilton et al., 2017, Chen et al., 2018, Zeng et al., 2019, Zou et al., 2019], graph partitioning [Liao et al., 2018, Chiang et al., 2019] or less expressive GNN architectures [Rossi et al., 2020, Bojchevski et al., 2020, Yu et al., 2020].

As we would like to leverage expressive GNNs, and be able to pass messages across any partitions, we opted for the subsampling approach. Accordingly, we subsample moderately-sized patches around the nodes we wish to compute latents for, execute our GNN model over them, and use the latents in the central nodes to train or evaluate the model.

We adapt the standard GraphSAGE subsampling algorithm of Hamilton et al. [2017], but make several modifications to it in order to optimise it for the specific features of MAG240M. Namely:

- We perform separate subsampling procedures across edge types. For example, an author node will separately sample a pre-specified number of papers written by that author and a pre-specified number of institutions that author is affiliated with.
- GraphSAGE mandates sampling an exact number of neighbours for every node, and uses sampling with replacement to achieve this even when the neighbourhood size is variable. We find this to be wasteful for smaller neighbourhoods, and hence use our pre-specified neighbour counts only as an upper bound. Denoting this upper bound as K , and node u ’s original neighbourhood as \mathcal{N}_u , we proceed² as follows:
 - For nodes that have fewer neighbours of a particular type than the upper bound ($|\mathcal{N}_u| \leq K$), we simply take the entire neighbourhood, without any subsampling;
 - For nodes that have a moderate amount of neighbours ($K < |\mathcal{N}_u| \leq 5K$) we subsample K neighbours *without* replacement, hence we do not wastefully duplicate nodes when the memory costs are reasonable.
 - For all other nodes ($|\mathcal{N}_u| > 5K$), we resort to the usual GraphSAGE strategy, and sample K neighbours with replacement, which doesn’t require an additional row-copy of the adjacency matrix.
- GraphSAGE directed the edges in the patch from the subsampled neighbours to the node which sampled them, and run their GNN for the exact same number of steps as the sampling depth. We instead modify the message passing update rule to scalably make the edges *bidirectional*, which naturally allows us to run deeper GNNs over the patch. The exact way in which we performed this will be detailed in the model architecture.

Taking all of the above into account, our model’s subsampling strategy proceeds, starting from paper nodes as central nodes, up to a depth of two (sufficient for institution nodes to become included). We did not observe significant benefits from sampling deeper patches. Instead, we sample significantly larger patches than the original GraphSAGE paper, to exploit the wide context available for many nodes:

²Note that, according to the previous bullet point, K and \mathcal{N}_u are defined on a per-edge-type basis.

Depth-0 Contains the chosen central paper node.

Depth-1 We sample up to $K = 40$ citing papers, $K = 40$ cited papers, and up to $K = 20$ authors for this paper.

Depth-2 We sample according to the following strategy, for all paper and author nodes sampled at depth-1:

Papers Identical strategy as for depth-1 papers: up to $K = 40$ cited, $K = 40$ citing, $K = 20$ authors.

Authors We sample up to $K = 40$ written papers, and up to $K = 10$ affiliations for this author.

Overall, this inflates our maximal patch size to nearly 10,000 nodes, which makes our *patches* of a comparable size to traditional *full-graph* datasets [Sen et al., 2008, Shchur et al., 2018]. Coupled with the fact that MAG240M has hundreds of millions of papers to sample these patches from, our setting enables transductive node classification at previously unexplored scale. We have found that such large patches were indeed necessary for our model’s performance.

One final important remark for MAG240M subsampling concerns the existence of *duplicated paper nodes*—i.e. nodes with exactly the same RoBERTa embeddings. This likely corresponds to identical papers submitted to different venues (e.g. conference, journal, arXiv). For the purposes of enriching our subsampled patches, we have combined the adjacency matrix rows and columns to “fuse” all versions of duplicated papers together.

Input preprocessing As just described, we seek to support execution of expressive GNNs on large quantities of large-scale subsampled patches. This places further stress on the model from a computational and storage perspective. Accordingly, we found it very useful to further compress the input nodes’ RoBERTa features. Our qualitative analysis demonstrated that their 129-dimensional PCA projections already account for 90% of their variance. Hence we leverage these PCA vectors as the actual input paper node features.

Further, only the paper nodes are actually provided with any features. We adopt the identical strategy from the baseline LSC scripts provided by Hu et al. [2021] to featurise the authors and institutions. Namely, for authors, we use the average PCA features across all papers they wrote. For institutions, we use the average features across all the authors affiliated with them. We found this to be a simple and effective strategy that performed empirically better than using structural features. This is contrary to the findings of Yu et al. [2020], probably because we use a more expressive GNN.

Besides the PCA-based features, our input node features \mathbf{x}_u also contain the one-hot representation of the node’s type (paper/author/institution), the node’s depth in the sampled patch (0/1/2), and a bitwise representation of the papers’ publication year (zeroed out for other nodes). Lastly, and according to an increasing body of research that argues for the utility of labels in transductive node classification tasks [Zhu and Ghahramani, 2002, Stretcu et al., 2019, Huang et al., 2020], we use the arXiv paper labels as features [Wang et al., 2021] (zeroed out for other nodes). We make sure that the validation labels are not observed at training time, and that the central node’s own label is not provided. It is possible to sample the central node at depth 2, and we make sure to mask out its label if this happens.

We also endow the patches’ edges with a simple edge type feature, \mathbf{x}_{uv} . It is a 7-bit binary feature, where the first three bits indicate the one-hot type of the sampling node (paper, author or institution) and the next four bits indicate the one-hot type of the sampled neighbour (cited paper, citing paper, author or institution). We found running a standard GNN over these edge-type features more performant than running a heterogeneous GNN—once again contrary to existing baseline results [Hu et al., 2021], and likely because of the expressivity of our processor GNN.

Model architecture For the GNN architecture we have used on MAG240M, our encoders and decoders are both two-layer MLPs, with a hidden size of 512 features. The node and edge encoders’ output layers compute 256 features, and we retain this dimensionality for $\mathbf{h}_u^{(t)}$ and $\mathbf{h}_{uv}^{(t)}$ across all steps t .

Our processor network is a deep message-passing neural network (MPNN) [Gilmer et al., 2017]. It computes message vectors, $\mathbf{m}_{uv}^{(t)}$, to be sent across the edge (u, v) , and then aggregates them in the receiver nodes as follows:

$$\mathbf{m}_{uv}^{(t+1)} = \psi_{t+1} \left(\mathbf{h}_u^{(t)}, \mathbf{h}_v^{(t)}, \mathbf{h}_{uv}^{(0)} \right) \tag{4}$$

$$\mathbf{h}_u^{(t+1)} = \phi_{t+1} \left(\mathbf{h}_u^{(t)}, \sum_{v \in \mathcal{N}_v} \mathbf{m}_{vu}^{(t+1)}, \sum_{v \in \mathcal{N}_u} \mathbf{m}_{uv}^{(t+1)} \right) \tag{5}$$

Taken together, Equations 4–5 fully specify the operations of the P_{t+1} network in Equation 2. The message function ψ_{t+1} and the update function ϕ_{t+1} are both two-layer MLPs, with identical hidden and output sizes to the encoder network. We note two specific aspects of the chosen MPNN:

- We did not find it useful to use global latents or update edge latents (Equation 4 uses $\mathbf{h}_{uv}^{(0)}$ at all times and does not include \mathbf{h}_G). This is likely due to the fact that the prediction is strongly centred at the central node, and that the edge features and types do not encode additional information.

- Note the third input in Equation 5, which is not usually included in MPNN formulations. In addition to pooling all incoming messages, we also pool all *outgoing* messages a node sends, and concatenate that onto the input to the sender node’s update function. This allowed us to simulate bidirectional edges without introducing additional scalability issues, allowing us to prototype MPNNs whose depth exceeded the subsampling depth.

The process is repeated for $T = 4$ message passing layers, after which $\mathbf{h}_u^{(4)}$ for the central node is sent to the decoder network for predictions.

Bootstrapping objective The non-arXiv papers within MAG240M are unlabelled and hence, under a standard node classification training regime, would contribute only implicitly to the learning algorithm (as neighbours of labelled papers). Early work on self-supervised graph representation learning [Veličković et al., 2018] had already shown this could be a wasteful approach, even on small-scale transductive benchmarks. Appropriately using the unlabelled nodes can provide the model with a wealth of information about the feature and network structure, which cannot be easily recovered from supervision alone. On a dataset like MAG240M—which contains $\sim 120\times$ more unlabelled papers than labelled ones—we have been able to observe significant gains from deploying such methods.

Especially, we leverage bootstrapped graph latents (BGRL) [Thakoor et al., 2021], a recently-proposed scalable method for self-supervised learning on graphs. Rather than contrasting several node representations across multiple views, BGRL *bootstraps* the GNN to make a node’s embeddings be predictive of its embeddings from another view, under a target GNN. The target network’s parameters are always set to an exponential moving average (EMA) of the GNN parameters. Formally, let f^- and P^- be the target versions of the encoder and processor networks (periodically updated to the EMA of f and P ’s parameters), and \mathbf{X} and $\tilde{\mathbf{X}}$ be two views of an input patch (in terms of features, adjacency structure or both). Then, BGRL performs the following computations:

$$\mathbf{H}^{(T)} = P(f(\mathbf{X})) \quad \tilde{\mathbf{H}}^{(T)} = P^-(f^-(\tilde{\mathbf{X}})) \tag{6}$$

where $P(f(\mathbf{X}))$ is short-hand for applying Equation 1, followed by repeatedly applying Equations 4–5 for T steps. The BGRL loss is then optimised to make the central node embedding $\mathbf{h}_u^{(T)}$ predictive of its counterpart, $\tilde{\mathbf{h}}_u^{(T)}$. This is done by projecting $\mathbf{h}_u^{(T)}$ to another representation using a *projector network*, p , as follows:

$$\mathbf{z}_u = p(\mathbf{h}_u^{(T)}) \tag{7}$$

where p is a two-layer MLP with identical hidden and output size as our encoder MLPs. We then optimise the cosine similarity between the projector output and $\tilde{\mathbf{h}}_u^{(T)}$:

$$\mathcal{L}_{\text{BGRL}} = -2 \frac{\mathbf{z}_u^\top \tilde{\mathbf{h}}_u^{(T)}}{\|\mathbf{z}_u\| \|\tilde{\mathbf{h}}_u^{(T)}\|} \tag{8}$$

using stochastic gradient ascent. Once training is completed, the projector network p is discarded.

This approach, inspired by BYOL [Grill et al., 2020], eliminates the need for crafting negative samples, reduces the storage requirements of the model, and its pointwise loss aligns nicely with our patch-wise learning setting, as we can focus on performing the bootstrapping objective on each central node separately. All of this made BGRL a natural choice in our setting, and we have found that we can easily apply it at scale.

Previously, BGRL has been applied on moderately-sized graphs with less expressive GNNs, showing modest returns. Conversely, we find the benefits of BGRL were truly demonstrated with stronger GNNs on the large-scale setting of MAG240M. Not only does BGRL monotonically improve when increasing proportions of unlabelled-to-labelled nodes during training, it consistently outperformed relevant self-supervised GNNs such as GRACE [Zhu et al., 2020].

Ultimately, our submitted model is trained with an auxiliary BGRL objective, with each batch containing a 10 : 1 ratio of unlabelled to labelled node patches. Just as in the BGRL paper, we obtain the two input patch views by applying dropout [Srivastava et al., 2014] on the input features (with $p = 0.4$) and DropEdge [Rong et al., 2019] (with $p = 0.2$), independently on each view. The target network (f^-, P^-) parameters are updated with EMA decay rate $\epsilon = 0.999$.

Training and regularisation We train our GNN to minimise the cross-entropy for predicting the correct topic over the labelled central nodes in the training patches, added together with the BGRL objective for the unlabelled central nodes. We use the AdamW SGD optimiser [Loshchilov and Hutter, 2017] with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$ and weight decay rate of $\lambda = 10^{-5}$. We use a cosine learning rate schedule with base learning rate $\eta = 0.01$ and 50,000 warm-up steps, decayed over 500,000 training iterations. Optimisation is performed over dynamically-batched data: we fill up each training minibatch with sampled patches until any of the following limits are exceeded: $\sim 84,000$ nodes, $\sim 185,000$ edges, or 256 patches.

To regularise our model, we perform early stopping on the accuracy over the validation dataset, and apply feature dropout (with $p = 0.3$) and DropEdge [Rong et al., 2019] (with $p = 0.25$) at every message passing layer of the GNN. We further apply layer normalisation [Ba et al., 2016] to intermediate outputs of all of our MLP modules.

Evaluation At evaluation time, we make advantage of the transductive and subsampled learning setup to enhance our predictions even further: first, we make sure that the model has access to **all** validation labels as inputs at test time, as this knowledge may be highly indicative. Further, we make sure that any “fused” copies of duplicated nodes also provide that same label as input. As our predictions are potentially conditioned on the specific topology of the subsampled patch, for each test node we average our predictions over 50 subsampled patches—an ensembling trick which consistently improved our validation performance. Lastly, given that we already use EMA as part of BGRL’s target network, for our evaluation predictions we use the EMA parameters, as they are typically slightly more stable.

5 Ensembling and training on validation

Once we established the top single-model architecture for our MAG240M entry, we found it very important to perform two post-processing steps: (a) re-train on the validation set, (b) ensemble various models together.

Re-training on validation data offers a great additional wealth of learning signal, even just by the sheer volume of data available in the OGB-LSC. But aside from this, the way in which the data was split offers even further motivation. On MAG240M, for example, the temporal split implies that validation papers (from 2019) are most relevant to classifying test papers (from 2020)—simply put, because they both correspond to the latest trends in scholarship.

However, training on the full validation set comes with a potentially harmful drawback: no held-out dataset would remain to early-stop on. In a setting where overfitting can easily occur, we found the risk to vastly outweigh the rewards. Instead, we decided to randomly partition the validation data into $k = 10$ equally-sized folds, and perform a cross-validation-style setup: we train k different models, each one observing the training set *and* $k - 1$ validation folds as its training data, validating and early stopping on the held-out fold. Each model holds out a different fold, allowing us to get an overall validation estimate over the entire dataset by combining their respective predictions.

While this approach may not correspond to the intended dataset splits, we have verified that the scores on individual held-out folds match the patterns observed on models that did not observe any validation data. This gave us further reassurance that no unintended strong overfitting had happened as a result.

Another useful outcome of our k -fold approach is that it allowed us a very natural way to perform ensembling as well: simply aggregating all of the k models’ predictions would give us a certain mixture of experts, as each of the k models had been trained on a slightly modified training set. Our final ensembled models employ exactly this strategy, with the inclusion of two seeds per fold. This brings our overall number of ensembled models to 20, and these ensembles correspond to our final submission on MAG240M.

6 Results and Discussion

Our final ensembled model achieved a *validation accuracy of 77.10%* on MAG240M. Translated on the LSC test sets, we recover **75.19% test accuracy** on MAG240M. We incur a minimal amount of distribution shift, which is a testament to our principled ensembling and post-processing strategies, in spite of using labels as inputs or training on validation.

Our entry has been designated as an *awardee* (ranked **in the top-3**) on MAG240M, solidifying the impact that very deep expressive graph neural networks can have on large scale datasets of industrial relevance. Further, we demonstrate how several recently proposed auxiliary objectives for GNN training, such as BGRL [Thakoor et al., 2021] can be highly impactful at the right dataset scales. We hope that our work serves towards resolving several open disputes in the community, such as the utility of deep GNNs, and the influence of self-supervision in this setting.

We would like to note that this is a *draft technical report* submitted to the OGB-LSC organisers shortly after the competition ended. We intend to release a substantially more detailed write-up in terms of experimental results (especially ablation studies demonstrating the influence of our various design choices) in the near-term.

In many ways, the OGB has been to graph representation learning what ImageNet has been to computer vision. We hope that OGB-LSC is only the first in a series of events designed to drive research on GNN architectures forward, and sincerely thank the OGB team for all their hard work and effort in making a contest of this scale possible and accessible.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Victor Bapst, Thomas Keck, A Grabska-Barwińska, Craig Donner, Ekin Dogus Cubuk, Samuel S Schoenholz, Annette Obika, Alexander WR Nelson, Trevor Back, Demis Hassabis, et al. Unveiling the predictive power of static structure in glassy systems. *Nature Physics*, 16(4):448–454, 2020.
- Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Aleksandar Bojchevski, Johannes Klicpera, Bryan Perozzi, Amol Kapoor, Martin Blais, Benedek Rózemerczki, Michal Lukasik, and Stephan Günnemann. Scaling graph neural networks with approximate pagerank. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2464–2473, 2020.
- Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Jie Chen, Tengfei Ma, and Cao Xiao. Fastgcn: fast learning with graph convolutional networks via importance sampling. *arXiv preprint arXiv:1801.10247*, 2018.
- Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. Cluster-gcn: An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 257–266, 2019.
- Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018.
- Pim de Haan, Maurice Weiler, Taco Cohen, and Max Welling. Gauge equivariant mesh cnns: Anisotropic convolutions on geometric graphs. *arXiv preprint arXiv:2003.05425*, 2020.
- Vijay Prakash Dwivedi, Chaitanya K Joshi, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *arXiv preprint arXiv:2003.00982*, 2020.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272. PMLR, 2017.
- Jonathan Godwin, Thomas Keck, Peter Battaglia, Victor Bapst, Thomas Kipf, Yujia Li, Kimberly Stachenfeld, Petar Veličković, and Alvaro Sanchez-Gonzalez. Jraph: A library for graph neural networks in jax., 2020. URL <http://github.com/deepmind/jraph>.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- William L Hamilton. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. *arXiv preprint arXiv:1706.02216*, 2017.
- Jessica B Hamrick, Kelsey R Allen, Victor Bapst, Tina Zhu, Kevin R McKee, Joshua B Tenenbaum, and Peter W Battaglia. Relational inductive bias for physical construction in humans and machines. *arXiv preprint arXiv:1806.01203*, 2018.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.
- Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-lsc: A large-scale challenge for machine learning on graphs. *arXiv preprint arXiv:2103.09430*, 2021.
- Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R Benson. Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint arXiv:2010.13993*, 2020.
- Chaitanya Joshi. Transformers are graph neural networks. *The Gradient*, 2020.

- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Renjie Liao, Marc Brockschmidt, Daniel Tarlow, Alexander L Gaunt, Raquel Urtasun, and Richard Zemel. Graph partition neural networks for semi-supervised classification. *arXiv preprint arXiv:1803.06272*, 2018.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 37–45, 2015.
- Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Jiang, Ebrahim Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Sungmin Bae, et al. Chip placement with deep reinforcement learning. *arXiv preprint arXiv:2004.10746*, 2020.
- Christopher Morris, Nils M Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. *arXiv preprint arXiv:2007.08663*, 2020.
- Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning mesh-based simulation with graph networks. *arXiv preprint arXiv:2010.03409*, 2020.
- Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*, 2019.
- Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. Dropedge: Towards deep graph convolutional networks on node classification. *arXiv preprint arXiv:1907.10903*, 2019.
- Emanuele Rossi, Fabrizio Frasca, Ben Chamberlain, Davide Eynard, Michael Bronstein, and Federico Monti. Sign: Scalable inception graph neural networks. *arXiv preprint arXiv:2004.11198*, 2020.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning*, pages 8459–8468. PMLR, 2020.
- Victor Garcia Satorras, Emiel Hooeboom, and Max Welling. E (n) equivariant graph neural networks. *arXiv preprint arXiv:2102.09844*, 2021.
- Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Jonathan M Stokes, Kevin Yang, Kyle Swanson, Wengong Jin, Andres Cubillos-Ruiz, Nina M Donghia, Craig R MacNair, Shawn French, Lindsey A Carfrae, Zohar Bloom-Ackermann, et al. A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702, 2020.
- Otilia Stretcu, Krishnamurthy Viswanathan, Dana Movshovitz-Attias, Anthony Platanios, Sujith Ravi, and Andrew Tomkins. Graph agreement models for semi-supervised learning. 2019.
- Shantanu Thakoor, Corentin Tallec, Mohammad Gheshlaghi Azar, Rémi Munos, Petar Veličković, and Michal Valko. Bootstrapped representation learning on graphs. *arXiv preprint arXiv:2102.06514*, 2021.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

- Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. *arXiv preprint arXiv:1809.10341*, 2018.
- Petar Veličković, Rex Ying, Matilde Padovano, Raia Hadsell, and Charles Blundell. Neural execution of graph algorithms. *arXiv preprint arXiv:1910.10593*, 2019.
- Kuansan Wang, Zhihong Shen, Chiyuan Huang, Chieh-Han Wu, Yuxiao Dong, and Anshul Kanakia. Microsoft academic graph: When experts are not enough. *Quantitative Science Studies*, 1(1):396–413, 2020.
- Yangkun Wang, Jiarui Jin, Weinan Zhang, Yong Yu, Zheng Zhang, and David Wipf. Bag of tricks for node classification with graph neural networks. *arXiv preprint arXiv:2103.13355*, 2021.
- Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR, 2019.
- Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 974–983, 2018.
- Lingfan Yu, Jiajun Shen, Jinyang Li, and Adam Lerer. Scalable graph neural networks for heterogeneous graphs. *arXiv preprint arXiv:2011.09679*, 2020.
- Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931*, 2019.
- Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. 2002.
- Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. Deep graph contrastive representation learning. *arXiv preprint arXiv:2006.04131*, 2020.
- Difan Zou, Ziniu Hu, Yewen Wang, Song Jiang, Yizhou Sun, and Quanquan Gu. Layer-dependent importance sampling for training deep and large graph convolutional networks. *arXiv preprint arXiv:1911.07323*, 2019.