

Hierarchical Inter-Message Passing for Learning on Molecular Graphs

Matthias Fey^{*1} Jan-Gin Yuen^{*1} Frank Weichert¹

Abstract

We present a hierarchical neural message passing architecture for learning on molecular graphs. Our model takes in two complementary graph representations: the raw molecular graph representation and its associated junction tree, where nodes represent meaningful clusters in the original graph, *e.g.*, rings or bridged compounds. We then proceed to learn a molecule’s representation by passing messages inside each graph, and exchange messages between the two representations using a coarse-to-fine and fine-to-coarse information flow. Our method is able to overcome some of the restrictions known from classical GNNs, like detecting cycles, while still being very efficient to train. We validate its performance on the ZINC dataset and datasets stemming from the MOLECULENET benchmark collection.

1. Introduction

Machine learning algorithms offer great potentials in reducing the computation time required for predicting molecular properties from several hours to just a few milliseconds (Wu et al., 2018). In particular, graph neural networks (GNNs) have been proven to be very successful for this task, exceeding the previously predominated approach of manual feature engineering by a large margin (Gilmer et al., 2017; Schütt et al., 2017). In contrast to hand-crafted features, GNNs learn high-dimensional embeddings of atoms that are able to represent their complex interactions by exchanging and aggregating messages between them.

In this work, we present a *hierarchical* variant of message passing on molecular graphs. Here, we utilize two separate graph neural networks that operate on complementary representations of a molecule simultaneously: its raw molecular representation and its corresponding (coarsened) junction

^{*}Equal contribution ¹Department of Computer Science, TU Dortmund University, Dortmund, Germany. Correspondence to: Matthias Fey <matthias.fey@udo.edu>, Jan-Gin Yuen <jan-gin.yuen@udo.edu>.

tree representation. Each of the two GNN’s intra-message passing step is strengthened by an inter-message passing step that exchanges intermediate information between the two representations. This allows the network to reason about hierarchy, *e.g.*, rings, in molecules in a natural fashion, and enables the GNN to overcome some of its restrictions, *e.g.*, detecting cycles (Loukas, 2020), without relying on more sophisticated architectures to do so (Morris et al., 2019; Murphy et al., 2019; Maron et al., 2019). We show that this simple scheme can drastically increase the performance of a GNN, reaching state-of-the-art performance on a variety of different datasets. Despite its higher-order nature, our proposed network architecture is still very efficient to train and causes only marginal additional costs in terms of memory and execution time.

2. Learning on Molecular Graphs

Graph neural networks operate on graph representations of molecules $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where nodes $\mathcal{V} = \{1, \dots, n\}$ represent atoms and edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ are defined by a predefined structure or by connecting atoms that lie within a certain cutoff distance. Given atom features $\mathbf{X}^{(0)} \in \mathbb{R}^{|\mathcal{V}| \times f}$ and edge features $\mathbf{E}^{(0)} \in \mathbb{R}^{|\mathcal{E}| \times d}$, a GNN iteratively updates node embeddings $\mathbf{x}_i^{(l)}$ in layer $l+1$ by aggregating localized information via the parametrized functions

$$\begin{aligned} \mathbf{m}_v^{(l+1)} &= \text{AGGREGATE}_{\theta_1}^{(l+1)} \left(\left\{ \left\{ \left(\mathbf{x}_w^{(l)}, \mathbf{x}_v^{(l)}, \mathbf{e}_{w,v}^{(l)} \right) \right\}_{w \in \mathcal{N}(v)} \right\} \right) \\ \mathbf{x}_v^{(l+1)} &= \text{UPDATE}_{\theta_2}^{(l+1)} \left(\mathbf{x}_v^{(l)}, \mathbf{m}_v^{(l+1)} \right), \end{aligned}$$

where $\{\dots\}$ denotes a multiset and $\mathcal{N}(v) \subseteq \mathcal{V}$ defines the neighborhood set of node $v \in \mathcal{V}$ (Gilmer et al., 2017). After L layers, a graph representation is obtained via global aggregation of $\mathbf{X}^{(L)}$, *e.g.*, summation or averaging.

Many existing GNNs can be expressed using this neural message passing scheme (Kipf & Welling, 2017; Veličković et al., 2018). A GNN called GIN (GIN-E in case edge features are present) (Xu et al., 2019; Hu et al., 2020b)

$$\mathbf{x}_v^{(l+1)} = \text{MLP}_{\theta}^{(l+1)} \left((1 + \epsilon) \cdot \mathbf{x}_v^{(l)} + \sum_{w \in \mathcal{N}(v)} \mathbf{x}_w^{(l)} + \mathbf{e}_{w,v}^{(l)} \right),$$

defines its most expressive form, showing high similarities




Figure 1. Example of a molecular graph and its cluster assignment for obtaining junction trees. Cluster colors refer to ■ singletons, ■ bonds and ■ rings.

to the popular WL-test (Weisfeiler & Lehman, 1968) while being able to operate on continuous node and edge features.

3. Methodology

It has been shown that GNNs are unable to distinguish certain molecules when operating on the molecular graph or using limited cutoff distances, *e.g.*, Cyclohexane and two Cyclopropane molecules (Xu et al., 2019; Klicpera et al., 2020). These restrictions mostly stem from the fact that GNNs are not capable of detecting cycles (Loukas, 2020) since they are unable to maintain information about which vertex in its receptive field has contributed what to the aggregated information (Hy et al., 2018). In this section, we present a simple hierarchical scheme to overcome this restriction, which strengthens the GNN’s performance with minimal computational overhead in return.

Our method involves learning on two molecular graph representations simultaneously in an end-to-end fashion: the original graph representation and its associated junction tree. The junction tree representation encodes the tree structure of molecules and defines how clusters (singletons, bonds, rings, bridged compounds) are mutually connected, while the graph structure captures its more fine-grained connectivity (Jin et al., 2018). We briefly revisit how junction trees are obtained from molecular graphs before describing our method in detail.

Tree Decomposition. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a *tree decomposition* maps \mathcal{G} into a *junction tree* $\mathcal{T} = (\mathcal{C}, \mathcal{R})$ with node set $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_m\}$, $\mathcal{C}_i \subseteq \mathcal{V}$ for all $i \in \{1, \dots, m\}$, and edge set $\mathcal{R} \subseteq \mathcal{C} \times \mathcal{C}$ so that:

1. $\bigcup_i \mathcal{C}_i = \mathcal{V}$ and $\bigcup_i \mathcal{E}[\mathcal{C}_i] = \mathcal{E}$, where $\mathcal{E}[\mathcal{C}_i] \subseteq \mathcal{C}_i \times \mathcal{C}_i$ represents the edge set of the induced subgraph $\mathcal{G}[\mathcal{C}_i]$
2. $\mathcal{C}_i \cap \mathcal{C}_j \subseteq \mathcal{C}_k$ for all clusters $\mathcal{C}_i, \mathcal{C}_k, \mathcal{C}_j$ with connections $(\mathcal{C}_i, \mathcal{C}_k) \in \mathcal{R}$ and $(\mathcal{C}_k, \mathcal{C}_j) \in \mathcal{R}$.

The assignment of atoms $v \in \mathcal{V}$ to clusters $\mathcal{C}_i \in \mathcal{C}$ is given by the matrix $\mathbf{S} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{C}|}$ with $S_{v,i} = 1$ iff. $v \in \mathcal{C}_i$.

We closely follow the tree decomposition algorithm of related works (Rarey & Dixon, 1998; Jin et al., 2018). We first group all simple cycles and all edges that do not belong to any cycle into clusters in \mathcal{C} . Two rings are merged together if they share more than two overlapping atoms (bridged compounds). For atoms lying inside more than three clusters, we add the intersecting atom as a singleton cluster. A cluster graph is constructed by adding edges between all intersecting clusters, and the final junction tree \mathcal{T} is then given as one its spanning trees. Figure 1 visualizes how clusters are formed on an exemplary molecule. For each cluster, we additionally hold its respective category (singleton, bond, ring, bridged compound) as one-hot encodings $\mathbf{Z}^{(0)} \in \{0, 1\}^{|\mathcal{C}| \times 4}$.

Inter-Message Passing with Junction Trees. Our method is able to extend any GNN model for molecular property prediction by making use of intra-message passing *in* and inter-message passing *to* a complementary junction tree representation. Here, instead of using a single GNN operating on the molecular graph, we make use of *two* GNN models: one operating on the original graph \mathcal{G} and one operating on its associated junction tree \mathcal{T} , each passing intra-messages to their respective neighbors. We further enhance this scheme by making use of *inter-message passing*: Let $\mathbf{X}^{(l)} \in \mathbb{R}^{|\mathcal{V}| \times h}$ and $\mathbf{Z}^{(l)} \in \mathbb{R}^{|\mathcal{C}| \times h}$ denote the intermediate representations of \mathcal{G} and \mathcal{T} , respectively. Then, we enhance both representations $\mathbf{X}^{(l)}$ and $\mathbf{Z}^{(l)}$ by an additional coarse-to-fine information flow from \mathcal{T} to \mathcal{G}

$$\mathbf{X}^{(l)} \leftarrow \mathbf{X}^{(l)} + \sigma(\mathbf{S}\mathbf{Z}^{(l)}\mathbf{W}_1^{(l)})$$

and reverse fine-to-coarse information flow from \mathcal{G} to \mathcal{T}

$$\mathbf{Z}^{(l)} \leftarrow \mathbf{Z}^{(l)} + \sigma(\mathbf{S}^\top \mathbf{X}^{(l+1)} \mathbf{W}_2^{(l)}),$$

with $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{h \times h}$ denoting trainable weights and σ being a non-linearity. This leads to a hierarchical-variant of message passing for learning on molecular graphs, similar to the ones applied in computer vision (Ronneberger et al., 2015; Newell et al., 2016; Lin et al., 2017). Furthermore, each atom is able to know about its cluster assignment, and, more importantly, which other nodes are part of the same cluster. Specially, this leads to an increased expressivity of GNNs. For example, the popular example of a Cyclohexane molecule and two Cyclopropane molecules (a single ring and two disconnected rings) (Klicpera et al., 2020) are distinguishable by our scheme since the junction tree representations are distinguishable by the most expressive GNN.

The *readout* of the model is then given via

$$\sum_{v \in \mathcal{V}} \mathbf{x}_v^{(L)} \parallel \sum_{\mathcal{C}_i \in \mathcal{C}} \mathbf{z}_i^{(L)},$$




Figure 2. Overview of the proposed approach. Two GNNs are operating on the distinct graph representations \mathcal{G} and \mathcal{T} , and receive coarse-to-fine and fine-to-coarse information before another round of message passing starts.

with \parallel denoting the concatenation operator. A high-level overview of our method is visualized in Figure 2.

4. Related Work

We briefly review some of the related work and their relation to our proposed approach.

Learning on Molecular Graphs. Instead of using hand-crafted representations (Bartók et al., 2013), recent advancements in deep graph learning rely on an end-to-end learning of representations which has quickly led to major breakthroughs in machine learning on molecular graphs (Duvenaud et al., 2015; Gilmer et al., 2017; Schütt et al., 2017; Jørgensen et al., 2018; Unke & Meuwly, 2019; Chen et al., 2019). Most of these works are especially designed for learning on the molecular geometry. Here, earlier models (Schütt et al., 2017; Gilmer et al., 2017; Jørgensen et al., 2018; Unke & Meuwly, 2019; Chen et al., 2019) fulfill rotational invariance constraints by relying on interatomic distances, while recent models employ more expressive equivariant models. For example, DIMENET (Klicpera et al., 2020) deploys directional message passing between node triplets to also model angular potentials. Another line of work breaks symmetries by taking permutations of nodes into account (Murphy et al., 2019; Hy et al., 2018; Albooyeh et al., 2019). Recently, it has been shown that strategies for pre-training models on molecular graphs can effectively increase their performance for certain downstream tasks (Hu et al., 2020b). Our approach fits nicely into these lines of work since it also increases the expressiveness of GNNs while being orthogonal to further advancements in this field.

Junction Trees. So far, junction trees have solely been used for molecule generation based on a coarse-to-fine generation procedure (Jin et al., 2018; 2019). In contrast to the generation of SMILES strings (Gómez-Bombarelli et al., 2018), this allows the model to enforce chemical valid-

ity while generating molecules significantly faster than the node-per-node generation procedure applied in autoregressive methods (You et al., 2018).

Inter-Message Passing. The idea of inter-message passing between graphs has been already heavily investigated in practice, mostly in the fields of deep graph matching (Wang et al., 2018; Li et al., 2019; Fey et al., 2020) and graph pooling (Ying et al., 2018; Gao & Ji, 2019). For graph pooling, most works focus on *learning* a coarsened version of the input graph. However, due to being learned, the coarsened graphs are unable to strengthen the expressiveness of GNNs by design. For example, DIFFPOOL (Ying et al., 2018) always maps the atoms of two disconnected rings to the *same* cluster, while the top_k pooling approach (Gao & Ji, 2019) either keeps or removes *all* atoms inside those rings (since their node embeddings are shared). The approach that comes closest to ours involves inter-message passing to a “virtual” node that is connected to *all* atoms (Gilmer et al., 2017; Hu et al., 2020a). Our approach can be seen as a simple yet effective extension to this procedure.

5. Experiments

We evaluate our proposed architecture on the ZINC dataset (Kusner et al., 2017) and a subset of datasets stemming from the MOLECULENET benchmark collection (Wu et al., 2018). For all experiments, we make use of the GIN-E operator for learning on the molecular graph (Hu et al., 2020b), and the GIN operator (Xu et al., 2019) for learning on the associated junction tree. GIN-E includes edge features (e.g., bond type, bond stereochemistry) by simply adding them to the incoming node features. All models were trained with the ADAM optimizer (Kingma & Ba, 2015) using a learning rate of 10^{-4} , while other hyperparameters (#epochs, #layers, hidden size, batch size, dropout ratio) are tuned via an additional validation set. Our method is implemented in PYTORCH (Paszke et al.,

Table 1. Results on the ZINC dataset.

Method	Mean Absolute Error (MAE)	
	ZINC (10K)	ZINC (FULL)
GCN	0.367 \pm 0.011	—
GRAPHSAGE	0.398 \pm 0.002	—
GIN	0.408 \pm 0.008	—
GAT	0.384 \pm 0.007	—
MONET	0.292 \pm 0.006	—
GATEDGCN	0.435 \pm 0.011	—
GATEDGCN-E	0.282 \pm 0.015	—
GIN-E	0.252 \pm 0.014	0.088 \pm 0.002
Ours	0.151\pm0.006	0.036\pm0.002

2019) and utilizes the PYTORCH GEOMETRIC (Fey & Lenssen, 2019) library. Our source code is available under <https://github.com/rustyls/himp-gnn>.

ZINC. The ZINC dataset (Kusner et al., 2017) contains about 250 000 molecular graphs and was introduced in Dwivedi et al. (2020) as a benchmark for evaluating GNN performances (using a subset of 10 000 training graphs). Here, the task is to regress the constrained solubility of a molecule. While this is a fairly simple task that can be exactly computed in a short amount of time, it can nonetheless reveal the capabilities across different neural architectures. We compare ourselves to all the baselines presented in Dwivedi et al. (2020), and additionally report results of a GIN-E baseline that does not make use of any additional junction tree information. Furthermore, we also perform experiments on the full dataset.

As shown in Table 1, our method is able to significantly outperform all competing methods. In comparison to GIN-E, its best performing competitor, the additional junction tree extension is able to reduce the error rate about 40–60%.

MoleculeNet Datasets. Following upon Murphy et al. (2019), we evaluate our model on the HIV, MUV and TOX21 datasets from the MOLECULENET benchmark collection (Wu et al., 2018), using a 80%/10%/10% random split. Here, the task is to predict certain molecular properties (cast as binary labels), *e.g.*, whether a molecule inhibits HIV virus replication or not. We compare ourselves to the neural graph fingerprint (NGF) operator (Duvenaud et al., 2015), and its relational pooling variant RP-NGF (Murphy et al., 2019), as well as our own GIN-E baseline.

As the results in Table 2 indicate, our method beats both NGF and GIN-E in test performance. Although RP-NGF is able to distinguish any graph structure by considering permutations of nodes, our approach leads to overall better generalization despite its simplicity.

Table 2. Results on a subset of the MOLECULENET datasets.

Method	ROC-AUC		
	HIV	MUV	Tox21
NGF	81.20 \pm 1.40	79.80 \pm 2.50	79.4 \pm 1.00
RP-NGF	83.20 \pm 1.30	79.40 \pm 0.50	79.9 \pm 0.60
GIN-E	83.83 \pm 0.67	79.57 \pm 1.14	86.68 \pm 0.77
Ours	84.81\pm0.42	81.80\pm2.02	87.36\pm0.50

Table 3. Results on the molhiv and molpcba datasets of OGB.

Method	ROC-AUC	PRC-AUC
	ogbg-molhiv	ogbg-molpcba
GCN-E	76.07 \pm 0.97	19.83 \pm 0.16
GATEDGCN-E	77.65 \pm 0.50	20.77 \pm 0.27
GIN-E	75.58 \pm 1.40	22.17 \pm 0.23
Ours	78.80\pm0.82	27.39\pm0.17

OGB Datasets. We also test the performance of our model on the newly introduced datasets ogbg-molhiv and ogbg-molpcba from the OGB benchmark dataset suite (Hu et al., 2020a), which are adopted from MOLECULENET and enhanced by a more challenging and standardized scaffold splitting procedure. We closely follow the experimental protocol of Hu et al. (2020a) and report ROC-AUC and PRC-AUC for ogbg-molhiv and ogbg-molpcba, respectively. We compare ourselves to three variants that do not make use of additional junction tree information, namely GCN-E, GATEDGCN-E and GIN-E (Kipf & Welling, 2017; Bresson & Laurent, 2017; Dwivedi et al., 2020; Hu et al., 2020b;a).

Results are presented in Table 3. As one can see, our approach is able to outperform all its competitors. Interestingly, our model achieves its best results in combination with a small amount of layers (2 or 3), making its runtime and memory requirements on par with the other baselines (which make use of 5 layers). This can be explained by the fact that the additional coarse-to-fine information flow enhances the receptive field size of a GNN, and therefore omits the need to stack a multitude of layers.

6. Conclusion

We introduced an end-to-end architecture for molecular property prediction that utilizes inter-message passing between graph representations of different hierarchy. Our proposed method can be used as a plug-and-play extension to strengthen the capabilities of a GNN operating on molecular graphs with little to no overhead. In future works, we are interested in studying how the proposed approach can be applied to other domains as well, *e.g.*, social networks.

References

- Albooyeh, M., Bertolini, D., and Ravanbaksh, S. Incidence networks for geometric deep learning. *CoRR*, abs/1905.11460, 2019.
- Bartók, A. P., Kondor, R., and Csányi, G. On representing chemical environments. *Physical Review B*, 87(18), 2013.
- Bresson, X. and Laurent, T. Residual gated graph convnets. *CoRR*, abs/1711.07553, 2017.
- Chen, Z., Li, L., and Bruna, J. Supervised community detection with line graph neural networks. In *ICLR*, 2019.
- Duvenaud, D. K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. In *NIPS*, 2015.
- Dwivedi, V. P., Joshi, C. K., Laurent, T., Bengio, Y., and Bresson, X. Benchmarking graph neural networks. *CoRR*, abs/2003.00982, 2020.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR-W*, 2019.
- Fey, M., Lenssen, J. E., Morris, C., Masci, J., and Kriege, N. M. Deep graph matching consensus. In *ICLR*, 2020.
- Gao, H. and Ji, S. Graph U-Nets. In *ICML*, 2019.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *ICML*, 2017.
- Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., and Aspuru-Guzik, A. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 2018.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open Graph Benchmark: Datasets for machine learning on graphs. *CoRR*, abs/2005.00687, 2020a.
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. Strategies for pre-training graph neural networks. In *ICLR*, 2020b.
- Hy, T. S., Trivedi, S., Pan, H., Anderson, B., and Kondor, R. Predicting molecular properties with covariant compositional networks. *The Journal of Chemical Physics*, 148(24), 2018.
- Jin, W., Barzilay, R., and Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. In *ICML*, 2018.
- Jin, W., Yang, K., Barzilay, R., and Jaakkola, T. Learning multimodal graph-to-graph translation for molecule optimization. In *ICLR*, 2019.
- Jørgensen, P. B., Jacobsen, K. W., and Schmidt, M. N. Neural message passing with edge updates for predicting properties of molecules and materials. In *NIPS*, 2018.
- Kingma, D. P. and Ba, J. L. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- Klicpera, J., Groß, J., and Günnemann, S. Directional message passing for molecular graphs. In *ICLR*, 2020.
- Kusner, M. J., Paige, B., and Hernández-Lobato, J. M. Grammar variational autoencoder. In *ICML*, 2017.
- Li, Y., Gu, C., Dullien, T., Vinyals, O., and Kohli, P. Graph matching networks for learning the similarity of graph structured objects. In *ICML*, 2019.
- Lin, T. Y., Dollar, P., Girshick, R., He, K., Hariharan, B., and Belongie, S. Feature pyramid networks for object detection. In *CVPR*, 2017.
- Loukas, A. What graph neural networks cannot learn: Depth vs width. In *ICLR*, 2020.
- Maron, H., Ben-Hamu, H., Serviansky, H., and Lipman, Y. Provably powerful graph networks. In *NeurIPS*, 2019.
- Morris, C., Ritzert, M., Fey, M., Hamilton, W. L., Lenssen, J. E., Rattan, G., and Grohe, M. Weisfeiler and Leman go neural: Higher-order graph neural networks. In *AAAI*, 2019.
- Murphy, R., Srinivasan, B., Rao, V., and Ribeiro, B. Relational pooling for graph representations. In *ICML*, 2019.
- Newell, A., Yang, K., and Deng, J. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- Rarey, M. and Dixon, J. S. Feature trees: A new molecular similarity measure based on tree matching. *Journal of Computer-aided Molecular Design*, 12(5):471–490, 1998.

- Ronneberger, O., Fischer, P., and Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- Schütt, K., Kindermans, P. J., Sauceda F., H. E., Chmiela, S., Tkatchenko, A., and Müller, K. R. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. In *NIPS*, 2017.
- Unke, O. T. and Meuwly, M. PhysNet: A neural network for predicting energies, force, dipole moments, and partial charges. *Journal of Chemical Theory and Computation*, 15(6):3678–3693, 2019.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *ICLR*, 2018.
- Wang, Z., Lv, Q., Lan, X., and Zhang, Y. Cross-lingual knowledge graph alignment via graph convolutional networks. In *EMNLP*, 2018.
- Weisfeiler, B. and Lehman, A. A. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsia*, 2(9), 1968.
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. MoleculeNet: A benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530, 2018.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *ICLR*, 2019.
- Ying, R., You, J., Morris, C., Ren, X., Hamilton, W. L., and Leskovec, J. Hierarchical graph representation learning with differentiable pooling. In *NIPS*, 2018.
- You, J., Ying, R., Ren, X., Hamilton, W. L., and Leskovec, J. GraphRNN: Generating realistic graphs with deep auto-regressive models. In *ICML*, 2018.