

ASGN: An Active Semi-supervised Graph Neural Network for Molecular Property Prediction

Zhongkai Hao¹, Chengqiang Lu¹, Zhenya Huang¹, Hao Wang¹, Zheyuan Hu¹, Qi Liu^{1,*},
Enhong Chen¹, Cheekong Lee²

{hzk171805,lunar,huangzhy,wanghao3,ustc_hzy}@mail.ustc.edu.cn

{qiliuql,cheneh}@ustc.edu.cn, cheekonglee@tencent.com

1: Anhui Province Key Lab of Big Data Analysis and Application, School of Computer Science and Technology, University of Science and Technology of China, 2: Tencent America,

ABSTRACT

Molecular property prediction (e.g., energy) is an essential problem in chemistry and biology. Unfortunately, many supervised learning methods usually suffer from the problem of scarce labeled molecules in the chemical space, where such property labels are generally obtained by Density Functional Theory (DFT) calculation which is extremely computational costly. An effective solution is to incorporate the unlabeled molecules in a semi-supervised fashion. However, learning semi-supervised representation for large amounts of molecules is challenging, including the joint representation issue of both molecular essence and structure, the conflict between representation and property learning. Here we propose a novel framework called Active Semi-supervised Graph Neural Network (ASGN) by incorporating both labeled and unlabeled molecules. Specifically, ASGN adopts a teacher-student framework. In the teacher model, we propose a novel semi-supervised learning method to learn general representation that jointly exploits information from molecular structure and molecular distribution. Then in the student model, we target at property prediction task to deal with the learning loss conflict. At last, we proposed a novel active learning strategy in terms of molecular diversities to select informative data during the whole framework learning. We conduct extensive experiments on several public datasets. Experimental results show the remarkable performance of our ASGN framework.

KEYWORDS

Active Learning; Molecular Property Prediction; Graph Neural Network; Semi-Supervised Learning

ACM Reference Format:

Zhongkai Hao, Chengqiang Lu, Zhenya Huang, Hao Wang, Zheyuan Hu, Qi Liu, Enhong Chen, Cheekong Lee. 2020. ASGN: An Active Semi-supervised Graph Neural Network for Molecular Property Prediction. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*

*Corresponding Author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '20, August 23–27, 2020, Virtual Event, CA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7998-4/20/08...\$15.00

<https://doi.org/10.1145/3394486.3403117>

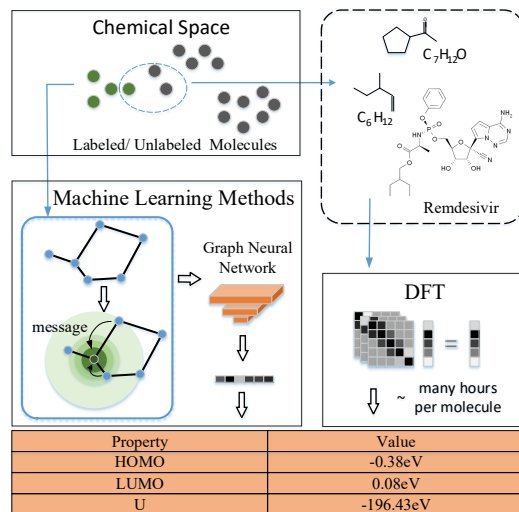


Figure 1: Methods for molecular property prediction. Left: Machine learning methods using message passing graph neural networks. Right: DFT calculation.

(KDD '20), August 23–27, 2020, Virtual Event, CA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394486.3403117>

1 INTRODUCTION

Predicting the property of molecules, such as the energy, is a fundamental issue in many related domains including chemistry, biology and material science, which has led to many significant relevant research and applications. For example, the process of drug discovery [10] can be accelerated if we can accurately predict the properties of molecules in time to help develop specific medicines for the epidemic, such as H1N1 flu, SARS, Covid19.

In chemistry, density functional theory (DFT) is commonly used computational methods for molecular property prediction, which has been studied dating back to the 1970s [4]. It offers accurate and explainable solutions for molecular following complete theory [21]. However, in practice, it suffers from a critical problem of expensive computation cost as it needs to solve many linear equations iteratively for the solutions. For example, experimental results find that it takes an hour to calculate the properties of a molecule with only 20 atoms [13]. Obviously, such low efficiency of DFT has limited its applications when screening from a large set of molecules.

Recently, researchers have attempted to use machine learning methods that are cost-effective for molecular property prediction [15].

Along this line, the most representative methods are graph neural networks (GNN), including MPNN [13], SchNet [30] and MGCN [23], which have shown superior performance. Generally, they treat a molecule as a graph where the nodes denote atoms and the edges represent the interaction between atoms. They design several neural layers to project each node into latent space with a low-dimensional learnable embedding vector and pass its interaction message through the edges iteratively. At last, the node messages can be aggregated to represent the molecule for property prediction.

Though GNNs have achieved great success, they are usually data-hungry, which requires a big amount of labeled data (i.e., molecules whose properties are known) for training [13]. However, the labeled molecules usually take an extreme small portion in the whole chemical space since they can only be provided by expensive experiments or DFT calculation, which restricts GNN based development. To gain further promotion, as shown in top left part of the Figure 1, there are still many valid molecules in the chemical space, though the properties remaining unknown, that have some benefits in terms of their structures. If we can effectively leverage these unlabeled molecules, it could be potentially helpful to improve the performance. Therefore, in this paper, we aim to explore semi-supervised learning (SSL) by fully taking advantage of both labeled molecules and unlabeled ones for property prediction.

However, it is highly challenging due to the following domain-specific characteristics. First, learning molecular graph representation is non-trivial because it involves both the node and the graph level information. Different from traditional applications like social networks since we usually meet a large number of graphs in chemical space rather than a single graph with large number of nodes. Though some existing semi-supervised learning methods, such as Ladder Networks [29], have shown their performance in various domains, such as image and text, they cannot be directly used for molecular graph learning. Second, it is difficult to handle the imbalance between labeled and unlabeled molecules in chemical space since the number of labeled ones generally take extreme small portion. Directly applying previous SSL methods leads to loss conflict caused by large number of unlabeled molecules for their structural representation but ignores our main goal of property prediction. Third, the performance might be still unsatisfactory due to limited labels, we need to find new molecules for labeling to improve the model. To increase the efficiency of labeling, we need a mechanism to find most informative molecules for labeling.

To address these challenges, we design a novel framework called Active Semi-supervised Graph Neural Network (ASGN) for molecular property prediction by taking advantage of both labeled and unlabeled molecules. Generally, ASGN uses a novel teacher-student framework consisting of two models that work alternatively. Specifically, in the teacher model, we propose a novel semi-supervised learning method to learn a general representation that jointly explores molecular features both at a global scale and local scale. The local one represents the essences of molecules, i.e., atoms and bonds while the global one learns the whole molecular graph encoding with respect to the chemical space. Then, to deal with the loss conflict between the unsupervised structure representation and property prediction, we introduce the student model by fine-tuning on property prediction task only on the small labeled molecules. By doing so, the student model can focus on the prediction to achieve

lower error than the teacher model and converge much faster. Additionally, it can alleviate over-fitting than training from scratch only on the labeled dataset. Moreover, to improve labeling efficiency, we propose a novel strategy based on active learning to select new informative molecules. That is, ASGN uses the embeddings by the teacher model to select a diversified subset of molecules in the chemical space and add them to the labeled dataset for finetuning two models repeatedly until the label budget or desired accuracy is reached. We conduct extensive experiments on real-world datasets, where the experimental results demonstrate the effectiveness of our proposed ASGN. To the best of our knowledge, this is the first attempt to incorporate both unlabeled and labeled molecules for property prediction actively in a semi-supervised manner.

2 RELATED WORK

In this section, we summarize the related work with the following three categories.

Molecular Property Prediction. Predicting the properties of molecules is a fundamental task with applications in many areas such as chemistry and biology [3, 25]. According to quantum physics, the states of a molecule are characterized by Schrödinger equation [37]. The first class like Density Functional Theory (DFT) [4] are simulation based methods directly derived or approximated by the Schrödinger equation. However, DFT methods are time-consuming because it solves some big linear equations and the complexity of DFT is $O(N^4)$ where N is the number of atoms.

Another class of molecular properties prediction methods are data-driven [9, 13, 15, 44]. Researchers attempted to use traditional machine learning methods with empirical descriptors or handcraft features to represent a molecule and use them for linear or logistic regression [15, 44]. However, these methods cannot achieve desirable accuracy due to the limited effectiveness of handcrafted features and model capacity [13].

Inspired by the remarkable development of graph neural networks in various domains [13] [40][27][39], researchers have noticed the potentials of them for molecular property prediction. Generally, by treating the molecule as a graph, several graph neural networks have been applied [14, 24, 40] as an architecture that can directly deal with noneuclidean data like graphs. Variants of graph neural networks like MPNN [13], SchNet[30], can be applied for molecular properties prediction where they use nodes to represent atoms, and the edges are weighted by the distances between atoms. Then the node embeddings are propagated and updated using the embeddings of their neighborhood, named message passing. The graph embedding can be pooled from nodes for property prediction.

Semi-supervised Representation learning. Semi-supervised learning is a popular framework to improve model performance by incorporating unlabeled data into training [45]. The main idea is to use the unlabeled data to learn a general and robust representation to improve the performance of the model. On the one hand, methods like ladder network [29] borrow the idea of jointly learning representation for unlabeled data (via generation) and labeled data [20]. On the other hand, a popular fashion is developed recently which uses self-supervised methods that force the networks to be consistent under the handcrafted transformations like image in-painting [26], rotation[12], contrastive loss [16]. Usually, these methods use a pseudo-labeling mechanism to assign each

unlabeled data with a pseudo label and force the neural network to predict these pseudo labels. Then the pre-trained models can be used for downstream tasks like classification or regression. For example, Gidaris et al. uses the rotation degree of an image as a kind of pseudo label. These pseudo labels are often obtained from transformations of data without changing their semantic feature. Deep Clustering [6] shows that the convolutional neural network itself can be viewed as a strong prior to processing image data. Accordingly, they design a self-supervised method based on learning the clustering results of the features by the neural networks.

Active Learning. Active learning is a popular framework to alleviate data deficiency and it has been applied in many tasks [11, 19, 41, 43]. Active learning framework starts with a small set of labeled data and a large set of unlabeled data. In every iteration, it develops a model to select a batch of unlabeled data to be labeled for supplementing the limited labeled data so that it achieves better performance. Generally, the representative methods consider the strategy selection from two perspectives, i.e., uncertainty, and diversity [11] [31]. Specifically, the uncertainty based methods define the model uncertainty for a new unlabeled data leveraged by some statistics properties (e.g., variance) and then select the data with the highest value [11] [38]. Comparatively, the diversity based methods aim to choose a small subset that is the most representative for the whole dataset [31].

As is pointed out in [2], the data selected by the uncertainty strategy are almost identical in batch mode settings, so it might be not suitable for large datasets like our scenarios. In this paper, we propose a novel diversity based active learning strategy for informative molecule selection where the semi-supervised embeddings are used for calculating the distance between molecules.

3 DEFINITIONS AND NOTATIONS

In this section, we will give formal definitions of terminologies and problems in this paper for clarity. Following the previous works [13] [30], we treat each molecule in chemical space as a graph, hence we define a molecular graph as follows:

Definition 3.1. Molecular Graph: A molecule is denoted as a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the vertex set $\mathcal{V} = \{v_i : 1 \leq i \leq |\mathcal{G}|\}$, we use \mathbf{x}_i to represent the feature vector of the node (atom) v_i indicating its type such as Carbon, Nitrogen. $|\mathcal{G}|$ is the total number of atoms. $\mathcal{E} = \{e_{ij} = |\mathbf{r}_i - \mathbf{r}_j| : 1 \leq i, j \leq |\mathcal{G}|\}$ is the set of edges connecting two atoms (nodes) v_i and v_j . Specifically, in a certain molecule, the coordinates of each atom can be represented as $\mathbf{r}_i = (r_i^{(1)}, r_i^{(2)}, r_i^{(3)})$. Therefore, we further denote the edge between two atom nodes e_{ij} as weighted by their coordinate distance $|\mathbf{r}_i - \mathbf{r}_j|$.

Then we give the formal definition of chemical space.

Definition 3.2. Chemical Space: Generally, the whole chemical space consists of a set of molecules, which can be denoted as: $\mathcal{M} = \{\mathcal{G}_i : 1 \leq i \leq N\}$. In practice, only a subset of molecules in the space have been examined to obtain their several properties (e.g., energy) by typical DFT calculation. Therefore, we divide the chemical space \mathcal{M} into two subset $\mathcal{D}_l, \mathcal{D}_u$. Specifically, $\mathcal{D}_l = \{(\mathcal{G}_1, \mathbf{y}_1), \dots, (\mathcal{G}_{N_l}, \mathbf{y}_{N_l})\}$ represents the subset of molecules whose properties have been examined, where $\mathbf{y}_i \in \mathbb{R}^m$ denotes the property vector with real value of molecule \mathcal{G}_i . Comparatively, $\mathcal{D}_u = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_{N_u}\}$ represents the subset of molecules whose

properties remain unknown. Without loss of generality, we call the subset \mathcal{D}_l and \mathcal{D}_u as "labeled set" and "unlabeled set", respectively.

With the above definition, our problem can be formalized as that we want to find a model $f(\mathcal{G}) \rightarrow \mathbf{y}$ using limited labels $|\mathcal{D}_l|$, for precisely predicting the properties of molecules.

4 ASGN: ACTIVE SEMI-SUPERVISED GRAPH NEURAL NETWORK

In this section, we present a description of the framework of ASGN. Then we describe the components of ASGN comprehensively.

4.1 Framework

In this paper, we propose a novel Active Semi-supervised Graph Neural Network (ASGN) for molecular property prediction by incorporating both labeled and unlabeled molecules in chemical space. The general framework is illustrated in Figure 2.

Generally, we use a teacher model and a student model that work iteratively. Each of them is a graph neural network. In the teacher network, we use a semi-supervised fashion to obtain a general representation of molecular graphs. We jointly train the embeddings for unsupervised representation learning and property prediction. Then in the student model, we handle the loss conflict by fine-tuning the parameters transferred from the teacher model for property prediction. After that, we use the student model to assign pseudo labels for the unlabeled dataset. As feedback for the teacher, the teacher model can learn the student's knowledge from these pseudo labels. Also, to improve the labeling efficiency, we propose using active learning to select the new representative unlabeled molecules for labeling. We then add them to the labeled set and finetune two models iteratively until accuracy budget is reached. Specifically, the key idea is to use the embeddings output by the teacher model to find a subset that is most diversified in the whole unlabeled set. We then assign ground truth labels such as using DFT calculation to these molecules. After that, we add them into the labeled set and repeat the iteration to improve performance.

In the following, we will first describe technical details of our teacher model and student model.

4.2 Semi-supervised Teacher Model

In the teacher model, we use semi-supervised learning. We first introduce the network backbone. Then, we introduce the loss for representation learning. Specifically, a property loss on labeled molecule \mathcal{D}_l and two unsupervised loss (from both the graph and the node level) on all molecules $\mathcal{D}_u \cup \mathcal{D}_l$ are designed to guide it.

4.2.1 Message Passing Graph Neural Network. The task of the teacher model is to learn a general representation for molecular graphs from both labeled set and unlabeled set. We first introduce a message passing graph neural network (MPGNN) as the backbone that transforms a molecular graph into a representation vector based on message passing graph neural networks. The graph neural network consists of L message passing layers. At l -th layer, it first embeds each node in a graph to a high dimensional space as their embeddings using $f(v_i) = \mathbf{z}_i \in \mathbb{R}^d$. Then the node embeddings are updated by aggregating node embeddings of its neighbors $\mathcal{N}(v_i)$ along the weighted edges called message passing:

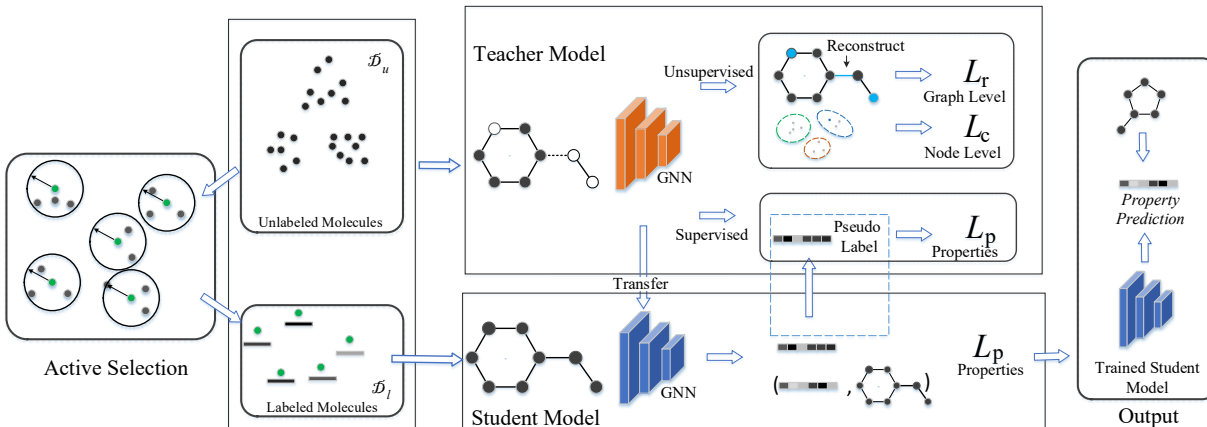


Figure 2: The overall framework of our method. We use a teacher model that jointly learns node embeddings at the local level and the distribution of the data at the global level with the property task. A student network uses the teacher’s weight to fine-tune network parameters on property prediction task. Active learning and pseudo labeling are used to combine these steps effectively into a framework.

$$z_i^{l+1} = \sigma(\mathbf{W}^l \cdot \text{AGG}(z_i^l, \{e(v_i, v_j) : v_j \in \mathcal{N}(v_i)\})), \quad (1)$$

where $\sigma(\cdot)$ is the activation function, \mathbf{W}^l is a learnable weight matrix, AGG is the aggregation function such as sum, mean, max [24]. Here we choose sum as the aggregation type which directly adds the messages from its neighbors as suggested in [42]. $e(v_i, v_j)$ is a vector called message function determined by the node embeddings and edge weights that pass from node v_i to v_j . As the interactions decay with the growth of the distances between two atoms, we use a Gaussian radical basis [30] to embed the edge information that reflects the interaction strength between nodes:

$$e(v_i, v_j)[k] = z_i^l[k] \cdot \exp(-\gamma(\|r_i - r_j\| - d_k)^2), \quad (2)$$

for $1 \leq k \leq N_f$ where $\{d_k : 1 \leq k \leq N_f\}$ is a set of pre-defined filter centers. More intensive centers means higher resolution and can capture minor difference of different bond length. After L layers of message passing and aggregation, we aggregate all node embeddings to get the whole graph embedding:

$$z_{\mathcal{G}} = \text{Pool}(\{z_i^L : v_i \in \mathcal{V}\}). \quad (3)$$

In this paper, we utilize a simple pooling method which directly averages or sums all node embeddings. At last, multi-layer perceptron f_{θ} is used to get the property $f_{\theta}(z_{\mathcal{G}})$.

Traditionally, MPGNN is trained in a supervised manner where all the labels are given and we usually use mean square loss (MSE) between predictions and labels \mathbf{y}_i (i.e. the labeled properties in \mathcal{D}_l) to guide the optimization of the model parameters:

$$\mathcal{L}_p = \sum_{i=1}^{N_l} \|\mathbf{y}_i - f_{\theta}(z_{\mathcal{G}_i})\|^2. \quad (4)$$

However, in practice the training set with small number of labels easily results in an over-fitted model. Additionally, end-to-end training that only learns a high-level representation guided by the property/label is less effective for structural representation. To overcome these challenges, in this paper we propose a semi-supervised representation learning method by considering both local level and

global level unsupervised information to enhance the expressive power of a model for both labeled and unlabeled molecular graphs.

4.2.2 Node Level Representation Learning. In node level representation learning, we learn to capture domain knowledge from geometry information of a molecular graph. The main idea is to use node embeddings to reconstruct the node types and topology (distances between nodes) from the representation. Specifically, we first randomly sample some nodes and edges from the graph as shown in Figure 2, then pass these nodes’ representation to a MLP and use them to reconstruct the node types f_i and distances between nodes e_{ij} . Mathematically, we minimize the following cross-entropy loss:

$$\begin{aligned} \mathcal{L}_r = & -\mathbb{E}_{v_i \sim \mathcal{V}} \left[\sum_{m=1}^{K_n} f_{im} \log(g_{\theta_n}(z_i)) \right] \\ & - \mathbb{E}_{e_{ij} \sim \mathcal{E}} \left[\sum_{m=1}^{K_e} e_{ijm} \log(g_{\theta_e}(z_i, z_j)) \right], \end{aligned} \quad (5)$$

where first term is the loss function for node types reconstruction, and the second term is the edge weights reconstruction. For both terms, we optimize the expectation of the samples. K_n is the number of atom types, we transform the continuous edge weights into a discrete classification problem by dividing the continuous distance into several discrete bins and K_e is the total number of bins. It means that $e_{ijm} = 1$ only if d_m is the nearest to the weight of edge e_{ij} . $g_{\theta_n}, g_{\theta_e}$ is a multi-layer perceptron.

Practically, we randomly sample some nodes and edges to reconstruct their attributes and optimize the expectation of samples. We found such random sampling to be significantly more efficient without sacrificing much performance. We sample $\alpha|\mathcal{G}|$ ($0 < \alpha < 1$) edges from the graph along with the nodes to reconstruct their features. What’s more, we notice that using a fully connected graph to represent a molecule contains redundant information because a molecule contains only $3n$ degrees of freedom since the coordinates of each atom can be decided by 3 numbers as $(r^{(1)}, r^{(2)}, r^{(3)})$. Therefore sampling edges with size $O(|\mathcal{G}|)$ is an efficient trade-off between performance and algorithm complexity. By optimizing the

reconstruction loss (Eq. (5)), we can obtain the node embeddings that contains the topology and features of molecular graphs.

4.2.3 Graph Level Representation Learning. Although node embeddings that can reconstruct the topology of molecules can effectively represent the structure of molecules, a recent study [18] shows that combing graph level representation learning is beneficial for downstream tasks like property prediction. In order to learn a graph level representation, the key insight is to use the mutual relation between molecules within the chemical space, i.e. similar molecules roughly have similar properties. Inspired by this intuition, we propose a method based on learning to cluster to enhance graph level representation. First, we calculate the graph level embedding by the network. Then we use an implicit clustering based method to assign N molecules each with a cluster id which contains M clusters generated by the implicit clustering process. After that we optimize the model with a penalty loss function. The process is iteratively done until at least a local minima is reached.

Next, we introduce the details of graph level representation learning. We denote s as the cluster id in the rest of this section. First we pass the graph level embedding into a multi-layer perceptron and predict the probability distribution $p(s|\mathcal{G})$. We assume there exists a posterior distribution $p(s|\mathcal{G})$ of cluster id. We optimize the cross-entropy loss between p and q as following:

$$H(y, x) = - \sum_{i=1}^N \sum_{j=1}^M p(s_j|\mathcal{G}_i) \log q(s_j|\mathcal{G}_i). \quad (6)$$

However, we easily get a trivial solution if no constraint is applied on $p(s|\mathcal{G})$. The key is to confine these clustering ids to a pre-defined prior distribution $p(s)$ as $\sum_{i=1}^N p(s_j|\mathcal{G}_i) = p(s_j)$ [5] [1]. We choose a uniform distribution with fixed M supports which means that the whole dataset is roughly divided into equally partitioned subsets. Practically, we use hard labeling technique to constraint $p(s|\mathcal{G}_i)$ to be a discrete label by applying the hardmax function. Then we explicitly write the optimization object as:

$$\begin{aligned} \min_{p, q} \mathcal{L}_c &= \sum_{i=1}^N \sum_{j=1}^M p(s_j|\mathcal{G}_i) \log q(s_j|\mathcal{G}_i) \quad (7) \\ \text{s. t. } &p(s_j|\mathcal{G}_i) \in \{0, 1\}, \sum_{i=1}^N p(s_j|\mathcal{G}_i) = p(s_j). \end{aligned}$$

We iteratively optimize predictive distribution $q(s|\mathcal{G})$ by performing gradient descent on the network parameters and the posterior distribution $p(s|\mathcal{G})$ by the following method which can be viewed as an implicit clustering approach. We first rewrite Eq. (7) as:

$$\min \mathcal{L}_c = \min_{Q \in U(p, q)} \langle P, Q \rangle, \quad (8)$$

with $\langle \cdot, \cdot \rangle$ denotes the Frobenius dot-product between two matrices, $P_{ij} = p(s_j|\mathcal{G}_i)$, $Q_{ij} = q(s_j|\mathcal{G}_i)$, and $U(p, q)$ denotes the joint distribution of p and q . This is a typical optimal transport problem and we add an entropy regularization and use Sinkhorn-Knopp algorithm [7] for a better convergence speed:

$$\min \mathcal{L}_c = \min_{Q \in U(p, q)} \langle P, Q \rangle - \frac{1}{\lambda} \text{KL}(Q||pq^T). \quad (9)$$

In fact, this process can be viewed as a type of clustering [8] so we name this loss as clustering loss for self-supervision.

In a nutshell, to train a teacher model under a semi-supervised manner, we need to optimize the following loss jointly combining Eq. (4), Eq. (5) and Eq. (8) as:

$$\mathcal{L}_t = \sum_{\mathcal{G} \in \mathcal{D}_l} \mathcal{L}_p + \sum_{\mathcal{G} \in \mathcal{D}_u \cup \mathcal{D}_l} \mathcal{L}_r + \sum_{\mathcal{G} \in \mathcal{D}_u \cup \mathcal{D}_l} \mathcal{L}_c. \quad (10)$$

4.3 Supervised Student Model

Practically, directly optimizing Eq. (10) of the teacher model yields unsatisfactory results for property prediction. The teacher model will be heavily loaded since it requires to learn several tasks simultaneously. Due to the conflict of these optimization targets, we observe that each target gets worse performance compared with optimizing them separately. Especially, it is also inefficient because if $|\mathcal{D}_l| \ll \mathcal{D}_u$ then little attention will be paid to optimization of \mathcal{L}_p in an epoch, however property prediction is what we care the most. As a result, the property prediction loss is much higher compared with a model that only needs to learn this task. To alleviate this problem, we propose introducing a student model. We use the teacher model to learn such representation by jointly optimizing the objects above. When the teacher’s learning process ends, we transfer the teacher’s weight to the student model, and use the student model to fine-tune only on the labeled dataset to learn the target properties the same as Eq. (4) shown in Figure 2:

$$\mathcal{L}_s = \sum_{\mathcal{G}_i \in \mathcal{D}_l} \|\mathbf{y}_i - f_{\theta_s}(z_{\mathcal{G}_i})\|^2. \quad (11)$$

After fine-tuning, we use the student model to infer the whole unlabeled dataset and assign each unlabeled data a pseudo label indicating the student’s prediction of its properties then the unlabeled dataset is $\mathcal{D}_u = \{(\mathcal{G}_i, f_{\theta_s}(\mathcal{G}_i)) : 1 \leq i \leq |\mathcal{D}_u|\}$ where θ_s is the parameters of student model. In the next iteration, the teacher model also needs to learn such pseudo labels as Eq. (10) becomes:

$$\mathcal{L} = \sum_{\mathcal{G} \in \mathcal{D}_u \cup \mathcal{D}_l} \mathcal{L}_p + \sum_{\mathcal{G} \in \mathcal{D}_u \cup \mathcal{D}_l} \mathcal{L}_r + \sum_{\mathcal{G} \in \mathcal{D}_u \cup \mathcal{D}_l} \mathcal{L}_c. \quad (12)$$

This can be viewed as the teacher learns the knowledge from the students as feedback inspired by the idea of knowledge distillation [17]. In summary, we handle the loss conflict by using two models whose targets are different. The teacher model learns a general representation while the student model aims to learn accurate prediction of molecular graph properties. The pre-training of the teacher provides a warm start for the student model.

4.4 Active Learning for Data Selection

We have incorporated the information in both labeled and unlabeled molecules. However, due to the limited number of labels available, the accuracy might still be unsatisfactory, we need to find new labeled data to improve its performance. Therefore, in each iteration we use the embeddings output by the teacher model to iteratively select a subset of molecules, and the properties (ground truth labels) will be computed (i.e., by DFT). Then we add these molecules output by active learning into the labeled set for finetuning two models iteratively. Along this line, the key strategy of active learning is to find a small batch of most diversified molecules in the chemical

space for labeling. A well-studied method to measure diversity is to sample from k -DPP as [22] suggests. However, the subset selection is NP-hard therefore a greedy approximation is taken advantage of, which is the k -center method. Denoting the unlabeled dataset by \mathcal{D}_u , and the labeled dataset by \mathcal{D}_l , we use a myopic method that in each iteration we choose a subset of data that maximize the distance between labeled set and unlabeled set. Concretely, for every $0 < i < b$ within the k -th batch, we choose the data point that satisfies the following condition:

$$\operatorname{argmax}_{j \in [n] \setminus \mathcal{D}_u^k} \min_{i \in \mathcal{D}_l^k} d(\mathcal{G}_i, \mathcal{G}_j), \quad (13)$$

where $d(\mathcal{G}_i, \mathcal{G}_j) = \|\mathbf{z}_{\mathcal{G}_i} - \mathbf{z}_{\mathcal{G}_j}\|$ is the distance between two molecules. We use $L - 2$ norm on the representations by the teacher model. Since the teacher model learns a general representation we naturally believe that the distance between the representations of two molecules indicates the difference of them. Moreover the features are automatically extracted, we do not need to rely on handcraft distances like graph edit distance which might not suit our problem. Additionally, since the teacher model is trained in a semi-supervised manner, the teacher model only needs to be fine-tuned when new labeled data is added, thus accelerating the training process.

4.5 Method Summary and Discussion

In this subsection, we briefly summarize the framework in Algorithm 1. Given an unlabeled set and a labeled set. In each iteration, we use k -center active learning strategy to get a new batch of data for labeling and add them to the labeled set (Line 4), next we transfer the teacher’s weight to the student network (Line 5) and fine-tune the student network (Line 6), then we use the student model to assign a pseudo label of the property for the rest of the unlabeled dataset (Line 7). After that, we continue to fine-tune the teacher model jointly with three tasks (Line 8). At last, the trained student model will be applied to predict the properties of the molecules.

Algorithm 1 ASGN framework

Input: Unlabeled, labeled, test dataset $\mathcal{D}_u, \mathcal{D}_l, \mathcal{D}_{test}$, error $\epsilon(\cdot, \mathcal{D})$, batch size b , stopping error, label budget ϵ, B

Output: student model θ_s

- 1: Initialize teacher and student θ_t, θ_s , labeled dataset \mathcal{D}_l
- 2: **while** $\epsilon(\theta_s, \mathcal{D}_{test}) > \epsilon$ or $|\mathcal{D}_l| \leq B$ **do**
- 3: Pre-train/finetune the teacher model by minimizing $\mathcal{L} = \mathcal{L}_r + \mathcal{L}_c + \mathcal{L}_p$ to get graph embeddings $\{\mathbf{z}_{\mathcal{G}} : \mathcal{G} \in \mathcal{D}_u\}$.
- 4: Use k -center active learning with $\mathbf{z}_{\mathcal{G}}$ for querying new labeled data s , $\mathcal{D}_l \leftarrow \mathcal{D}_l \cup s$, $|s| = b$.
- 5: Transfer the weights of teacher to student $\theta_s \leftarrow \theta_t$.
- 6: Finetune the student network by minimizing $\mathcal{L}_p = \epsilon(\theta_s, \mathcal{D}_l)$.
- 7: Assign pseudo label for the unlabeled dataset using student model, $\mathbf{y}_i \leftarrow f_{\theta_s}(\mathcal{G}_i), i \leq |\mathcal{D}_u \setminus \mathcal{D}_l|$.
- 8: **end while**

Return: student model θ_s

To summarize, we propose a novel approach to predict the properties of molecules using graph neural networks. First, we use a

multi-level representation learning method to obtain general embeddings for molecular graphs. The node embeddings store essential components of molecular graphs and they are composable to form meaningful graph level embeddings with respect to the whole data distribution. Subsequently, a teacher-student framework is used to effectively combine semi-supervised learning and active learning to deal with label insufficiency. Compared with vanilla semi-supervised learning methods [31], the separation of the two models can alleviate loss conflict. Compared with naive learning methods that re-trains the model from scratch when every new batch data points are selected, the weight transferred from the teacher provides a warm start for the student and avoids overfitting of the small labeled dataset and accelerates training. Besides, the two models communicate via weight transfer and feedback from assigning pseudo labels so that they can be mutually promoted.

5 EXPERIMENTS

In this section, we conduct extensive experiments to show the effectiveness of ASGN on two popular molecular datasets. The code is publicly available ¹.

5.1 Datasets

- **QM9:**² The QM9 dataset [28] is a well-known benchmark datasets that contains the equilibrium coordinates of 130,000 molecules along with their quantum mechanical properties. We use 10,000 molecules for testing and 10,000 for validation. Coordinates and properties for all molecules are calculated using DFT methods. Molecules in QM9 contain no more than 9 heavy atoms (atom heavier than hydrogen).
- **OPV:**³ OPV [34] is a dataset with roughly 100,000 medium size molecules, each contains 20 to 30 heavy atoms. Again the properties and equilibrium coordinates of these molecules are obtained through DFT. We use 5,000 for testing and 5,000 for validation.

5.2 Experiments Setup

We evaluate our method under two experimental settings. We first describe the implementation details and parameters of ASGN. We run all experiments are on one Tesla V100 GPU and 16 Intel CPUs.

Graph Neural Network Hyperparameters. For the network backbone, we use 4 message passing layers and embedding dimension of 96 in Eq. (1). We use Adam optimizer with a learning rate $1e-3$. We use filters from 0 to 3nm with an interval of 0.01nm in Eq. (2).

Semi-Supervised Learning Hyperparameters. The teacher model has an additional linear classifier after the graph neural network. We divide the distance of the edge into 30 bins in Eq. (5) for reconstruction. We use $M = 100$ in Eq. (6). The regularization constant λ is set to be 25 in Eq. (9). We train (fine-tune) the teacher model for 20 epochs in each iteration. We train the student network until the loss does not decrease for about 20 epochs.

Active Learning Hyperparameters. In each iteration, we select 1,000 new unlabeled molecules in Eq. (13) to be labeled and add them into the training dataset.

¹https://github.com/HaoZhongkai/AS_Molecule

²<http://quantum-machine.org/datasets/>


³<https://cscdata.nrel.gov/#/datasets/ad5d2c9a-af0a-4d72-b943-1e433d5750d6>

Table 1: Results on QM9 dataset for effectiveness experiment.

| Properties | U_0 | U | G | H | C_v | HOMO | LUMO | gap | ZPVE | R^2 | μ | α |
|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|----------------|-------------------|---------------|-------------------|
| Unit | eV | eV | eV | eV | Cal/MolK | eV | eV | eV | eV | Bohr ² | Debye | Bohr ³ |
| Supervised | 0.3204 | 0.2934 | 0.2948 | 0.2722 | 0.2368 | 0.1632 | 0.1686 | 0.2475 | 0.0007 | 10.05 | 0.3201 | 0.5792 |
| Mean-Teachers | 0.3717 | 0.2730 | 0.2535 | 0.2150 | 0.2036 | 0.1605 | 0.1686 | 0.2394 | 0.00054 | 5.22 | 0.3488 | 0.5792 |
| InfoGraph | 0.1410 | 0.1702 | 0.1592 | 0.1552 | 0.1965 | 0.1605 | 0.1659 | 0.2421 | 0.00036 | 4.92 | 0.3168 | 0.5444 |
| ASGN (Ours) | 0.0562 | 0.0594 | 0.0560 | 0.0583 | 0.0984 | 0.1190 | 0.1061 | 0.2012 | 0.00017 | 1.38 | 0.1947 | 0.2818 |

Table 2: Results on OPV dataset for effectiveness experiment.

| Property | HOMO | LUMO |
|--------------|--------------|--------------|
| Unit | Hartree | |
| Supervised | 0.080 | 0.078 |
| Mean-Teacher | 0.078 | 0.075 |
| InfoGraph | 0.077 | 0.076 |
| ASGN (Ours) | 0.059 | 0.057 |

**Figure 3: The results of efficiency experiments of property HOMO on QM9 and OPV datasets.**

5.3 Effectiveness Experiment

To demonstrate that our method could achieve lower error with limited labeled data, we first conduct an effectiveness experiment. Under this experimental setting we have a fixed label budget which is the maximum number of labels. Given a fixed label budget, we compare the final Mean-Absolute-Error(MAE) [30] on the test dataset after training. We use a label budget of 5000 for both QM9 and OPV about 5%. Other than these 5,000 labeled data, other labels are not available. We compare our methods with baselines listed below.

5.3.1 Baselines. For accuracy experiments, we mainly compare our method with several semi-supervised learning baselines. To ensure fairness, all baselines are conducted on the same network backbone (i.e. MPGNN). The compared baselines are selected from two perspectives, one is traditional semi-supervised learning, the other is semi-supervised learning baselines for graph data.

- **Supervised** : We train the network backbone using fully supervised manner only on the small labeled dataset.
- **Mean-Teachers** [36]: This is a method for semi-supervised learning by using a consistency regularization and uses moving average for the models' weights as the teacher.
- **InfoGraph** [35]: This is the state-of-the-art method for semi-supervised learning or unsupervised learning on graphs. It maximizes the mutual information between the graph level representations and the substructures of the graphs.

5.3.2 Results. The results are listed in Table 1 on QM9 dataset and Table 2 on OPV dataset.

First, We found that our method is significantly better than baseline methods on all properties. We achieved a reduction of more than 50% on several properties such as U , U_0 , α and C_v compared with the state-of-the-art method. This shows our semi-supervised learning method is effective and incorporating unlabeled data can help the prediction of molecular properties.

Second, the semi-supervised reconstruction captures domain knowledge for molecules and achieves better results than supervised model (i.e. MPGNN) and Mean-Teachers. The global representation learning at graph level is beneficial for molecular property prediction and its performance is better than InfoGraph.

5.4 Efficiency Experiment

To demonstrate ASGN is label efficient, we conduct an efficiency experiment. In this experiment, we start with 5,000 labeled molecules and the rest in the unlabeled set. Then, in each iteration, after the model selects a molecule from \mathcal{D}_u , we add it to \mathcal{D}_l . During this process, we measure the Label Rate-Mean Absolute Error(MAE) curve to show how many labels are saved for a fixed error. For a fixed error, the less labeled data is used, the better the model is.

5.4.1 Baselines. The baselines are selected from active learning methods. We apply these methods on the backbone of ASGN (i.e. MPGNN). We simply omit some methods that cannot be applied to our settings. We use a batch number of 2500 new labeled molecules in every iteration in Eq. (13) for ASGN. The computational cost of QBC method on OPV dataset is unaffordable so we simply omit it.

- **Random**: Choosing data points randomly from the unlabeled dataset in each iteration. The model is re-initialized when a new batch of labeled data is selected. This method equals the passive learning.
- **Query By Committee (QBC)** [32]: We jointly train a group of models named committee initialized in the same method but different parameters. Each iteration we choose a batch of data points with the biggest disagreement of the committee members. We use 8 models as a committee, training 8 models at the same time is time consuming.
- **Deep Bayes Active Learning (BALD)** [11]: This is a method based on uncertainty. We approximate the uncertainty by performing Monte Carlo dropout [33] on layers of the network.
- **Vanilla k -center** [31]: The representation learned by the semi-supervised learning methods actually benefits the selection of new data points. We also compare our method with the vanilla plain k -center active learning strategy.

5.4.2 Results. We plot the results on HOMO (highest occupied molecular orbital) on both QM9 dataset and OPV dataset in Figure

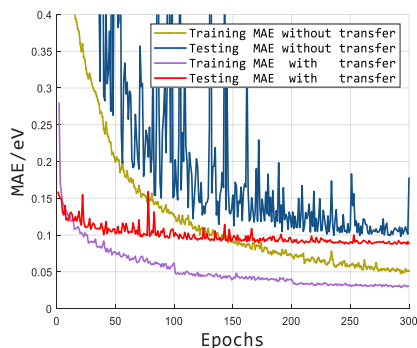


Figure 4: Ablation results on the necessity of weight transfer.

3. "Full" denotes the MAE for a supervised MPGNN using all labeled data. We have the following conclusions.

First, we show that for all datasets and properties, when the label number is fixed, the MAE is much lower than baselines which proves the effectiveness of our model. This shows that the active learning strategy is beneficial for model training. Additionally, the performance is better than a fully supervised model on all labeled data, proving the effectiveness of combining semi-supervised loss as regularization.

Second, when we set a fixed error target, we found that our model is about 2 ~ 3 times label efficient than baselines. This means that if we only need a predictor with given accuracy, we could use only 1/3 ~ 1/2 labels compared with other methods. Specifically, we use 50% labeled data to reach full accuracy on QM9 and 40% for OPV.

Third, we found that some baseline methods that work well in deep learning for image classification like BALD and k -center do not perform well on molecular data. Additionally, since BALD requires dropout, the performance is better when few labels are available but worse when we use all the labels.

5.5 Ablation Experiments

In this section, we conduct more experiments on ASGN including the ablation study to demonstrate how every part of our model affects the performance and a visualization experiment to support the interpretability of our model.

5.5.1 Necessity of Teacher-Student Framework. First, to show the effectiveness of the teacher-student framework in our model, we conduct an ablation study of ASGN without the teacher model or the student model. We denote ASGN with only the teacher model as ASGN-T which means that we jointly learn all tasks without handling the loss conflict. We list the results of HOMO on QM9 and OPV datasets in Table 3. We see that with the student network, the model achieves better performance on property prediction task.

We also study the case without the teacher model as ASGN-S which means no semi-supervised learning is used. Notice that ASGN-S is identical to a vanilla k -center active learning method[31]. Results show that it is necessary using the teacher-student framework.

5.5.2 Necessity of Weight Transfer. The essential step in connecting the student model and teacher model in our method is to transfer the weight of the teacher model to the student model in order to accelerate the training process. Here we use an ablation experiment




Figure 5: Visualization results on QM9 dataset of molecular graph embeddings using t-SNE method.

Table 3: Results of Ablation experiments on the necessity of teacher-student framework.

| Name/Dataset Unit | Homo(QM9) eV | | | Homo(OPV) Hartree | | |
|----------------------|-----------------|---------------|---------------|----------------------|--------------|--------------|
| | 5k | 10k | 50k | 5k | 10k | 50k |
| ASGN-T | 0.1668 | 0.1523 | 0.0682 | 0.080 | 0.053 | 0.020 |
| ASGN-S | 0.1632 | 0.1252 | 0.0653 | 0.076 | 0.049 | 0.019 |
| ASGN | 0.1190 | 0.0951 | 0.0517 | 0.060 | 0.039 | 0.015 |

to demonstrate the effect of the weight transfer. In Figure 4, we plot the MAE of ASGN with weight transfer and without weight transfer on the test dataset of QM9 on LUMO (lowest unoccupied molecular orbital) property when 10,000 labeled data are available. Results show that both training and testing MAE converge faster and are more stable with weight transfer. The final performance is also better using weight transfer.

5.6 Visualization Experiments

Our representation learning has considered the mutual relation between molecules within the chemical space and we use the information mutually for predicting the clustering to enhance the representation. To demonstrate that the distribution of molecules exhibits a clustered structure, we use t-SNE method to visualize the graph level representation of molecules using ASGN, shown in Figure 5. We see after using t-SNE the embedding of molecules can be clustered, and there is obvious distance between the clusters which verifies that we have got discriminative graph level embeddings. Additionally, similar molecules are clustered into the same cluster that means the embeddings can capture structural information.

6 CONCLUSIONS

In this paper, we proposed a novel framework to improve the performance for molecular property prediction with limited labels by incorporating unlabeled molecules. We designed a teacher-student framework consisting of two graph neural networks that work iteratively. Then we introduced the details of our semi-supervised representation learning method for molecular graphs that consider both graph level and node level information. Weight transfer and pseudo labeling are used to optimize two models to balance the loss functions. Furthermore, we used diversity based active learning to select new molecules for labelling. ASGN achieves much better

performance compared with baselines when labels are limited. Additionally, we showed the necessity for components in ASGN using ablation experiments. In future work, we will attempt to extend our model to more general molecular property prediction.

ACKNOWLEDGMENTS. This research was supported by grants from the National Natural Science Foundation of China (Grants No. 61922073, U1605251). Qi Liu gratefully acknowledges the support of the Youth Innovation Promotion Association of CAS (No. 2014299).

REFERENCES

- [1] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. 2019. Self-labelling via simultaneous clustering and representation learning. *arXiv preprint arXiv:1911.05371* (2019).
- [2] Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. 2019. Deep batch active learning by diverse, uncertain gradient lower bounds. *arXiv preprint arXiv:1906.03671* (2019).
- [3] Axel Becke. 2007. *The quantum theory of atoms in molecules: from solid state to DNA and drug design*. John Wiley & Sons.
- [4] Axel D Becke. 2014. Perspective: Fifty years of density-functional theory in chemical physics. *The Journal of chemical physics* 140, 18 (2014), 18A301.
- [5] Piotr Bojanowski and Armand Joulin. 2017. Unsupervised learning by predicting noise. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 517–526.
- [6] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. 2018. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 132–149.
- [7] Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*. 2292–2300.
- [8] Marco Cuturi and Arnaud Doucet. 2014. Fast computation of Wasserstein barycenters. (2014).
- [9] Kien Do, Truyen Tran, and Svetha Venkatesh. 2019. Graph transformation policy network for chemical reaction prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 750–760.
- [10] Sean Ekins, Ana C Puhl, Kimberley M Zorn, Thomas R Lane, Daniel P Russo, Jennifer J Klein, Anthony J Hickey, and Alex M Clark. 2019. Exploiting machine learning for end-to-end drug discovery and development. *Nature materials* 18, 5 (2019), 435.
- [11] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. 2017. Deep bayesian active learning with image data. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 1183–1192.
- [12] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728* (2018).
- [13] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 1263–1272.
- [14] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems*. 1024–1034.
- [15] Katja Hansen, Franziska Biegler, Raghunathan Ramakrishnan, Wiktor Pronobis, O Anatole Von Lilienfeld, Klaus-Robert Muller, and Alexandre Tkatchenko. 2015. Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space. *The journal of physical chemistry letters* 6, 12 (2015), 2326–2331.
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. [n. d.]. *Momentum Contrast for Unsupervised Visual Representation Learning*. Technical Report. arXiv:1911.05722v2
- [17] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [18] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay S Pande, and Jure Leskovec. [n. d.]. *Pre-training Graph Neural Networks*. Technical Report. arXiv:1905.12265v1
- [19] Zhenya Huang, Qi Liu, Yuying Chen, Le Wu, Keli Xiao, Enhong Chen, Haiping Ma, and Guoping Hu. 2020. Learning or Forgetting? A Dynamic Approach for Tracking the Knowledge Proficiency of Students. *ACM Trans. Inf. Syst.* 38, 2 (2020), 19:1–19:33. <https://doi.org/10.1145/3379507>
- [20] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [21] Walter Kohn and Lu Jeu Sham. 1965. Self-consistent equations including exchange and correlation effects. *Physical review* 140, 4A (1965), A1133.
- [22] Alex Kulesza and Ben Taskar. 2011. k-DPPs: Fixed-size determinantal point processes. (2011).
- [23] Chengqiang Lu, Qi Liu, Chao Wang, Zhenya Huang, Peize Lin, and Lixin He. 2019. Molecular property prediction: A multilevel quantum interactions modeling perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1052–1060.
- [24] Yao Ma, Suhang Wang, Charu C Aggarwal, and Jiliang Tang. 2019. Graph convolutional networks with eigenpooling. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 723–731.
- [25] Dino Ogljic, Roman Garnett, and Thomas Gärtner. 2017. Active search in intentionally specified structured spaces. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [26] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2536–2544.
- [27] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. 2020. Geom-GCN: Geometric Graph Convolutional Networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. <https://openreview.net/forum?id=S1e2agrFvS>
- [28] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data* 1 (2014), 140022.
- [29] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. 2015. Semi-supervised learning with ladder networks. In *Advances in neural information processing systems*. 3546–3554.
- [30] Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Sauceda Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. 2017. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. In *Advances in neural information processing systems*. 991–1001.
- [31] Ozan Sener and Silvio Savarese. 2017. Active Learning for Convolutional Neural Networks: A Core-Set Approach. (2017), 1–13. arXiv:1708.00489 <http://arxiv.org/abs/1708.00489>
- [32] H Sebastian Seung, Manfred Opper, and Haim Sompolinsky. 1992. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*. 287–294.
- [33] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [34] Peter C St. John, Caleb Phillips, Travis W Kemper, A Nolan Wilson, Yanfei Guan, Michael F Crowley, Mark R Nimlos, and Ross E Larsen. 2019. Message-passing neural networks for high-throughput polymer screening. *The Journal of chemical physics* 150, 23 (2019), 234111.
- [35] Fan-Yun Sun, Jordan Hoffmann, and Jian Tang. 2019. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. *arXiv preprint arXiv:1908.01000* (2019).
- [36] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*. 1195–1204.
- [37] David J Thouless. 2014. *The quantum mechanics of many-body systems*. Courier Corporation.
- [38] Daniel Ting and Eric Brochu. 2018. Optimal subsampling with influence functions. In *Advances in Neural Information Processing Systems*. 3650–3659.
- [39] Hao Wang, Enhong Chen, Qi Liu, Tong Xu, Dongfang Du, Wen Su, and Xiaopeng Zhang. 2018. A United Approach to Learning Sparse Attributed Network Embedding. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 557–566.
- [40] Hao Wang, Tong Xu, Qi Liu, Defu Lian, Enhong Chen, Dongfang Du, Han Wu, and Wen Su. 2019. MCNE: An End-to-End Framework for Learning Multiple Conditional Network Representations of Social Network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1064–1072.
- [41] Likang Wu, Zhi Li, Hongke Zhao, Zhen Pan, Qi Liu, and Enhong Chen. 2020. Estimating Early Fundraising Performance of Innovations via Graph-Based Market Environment Model. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*. AAAI Press, 6396–6403.
- [42] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks? (2018), 1–17. arXiv:1810.00826 <http://arxiv.org/abs/1810.00826>
- [43] Zhilin Yang, Jie Tang, and Yutao Zhang. 2014. Active learning for streaming networked data. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. 1129–1138.
- [44] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 974–983.
- [45] Xiaojin Jerry Zhu. 2005. *Semi-supervised learning literature survey*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences.